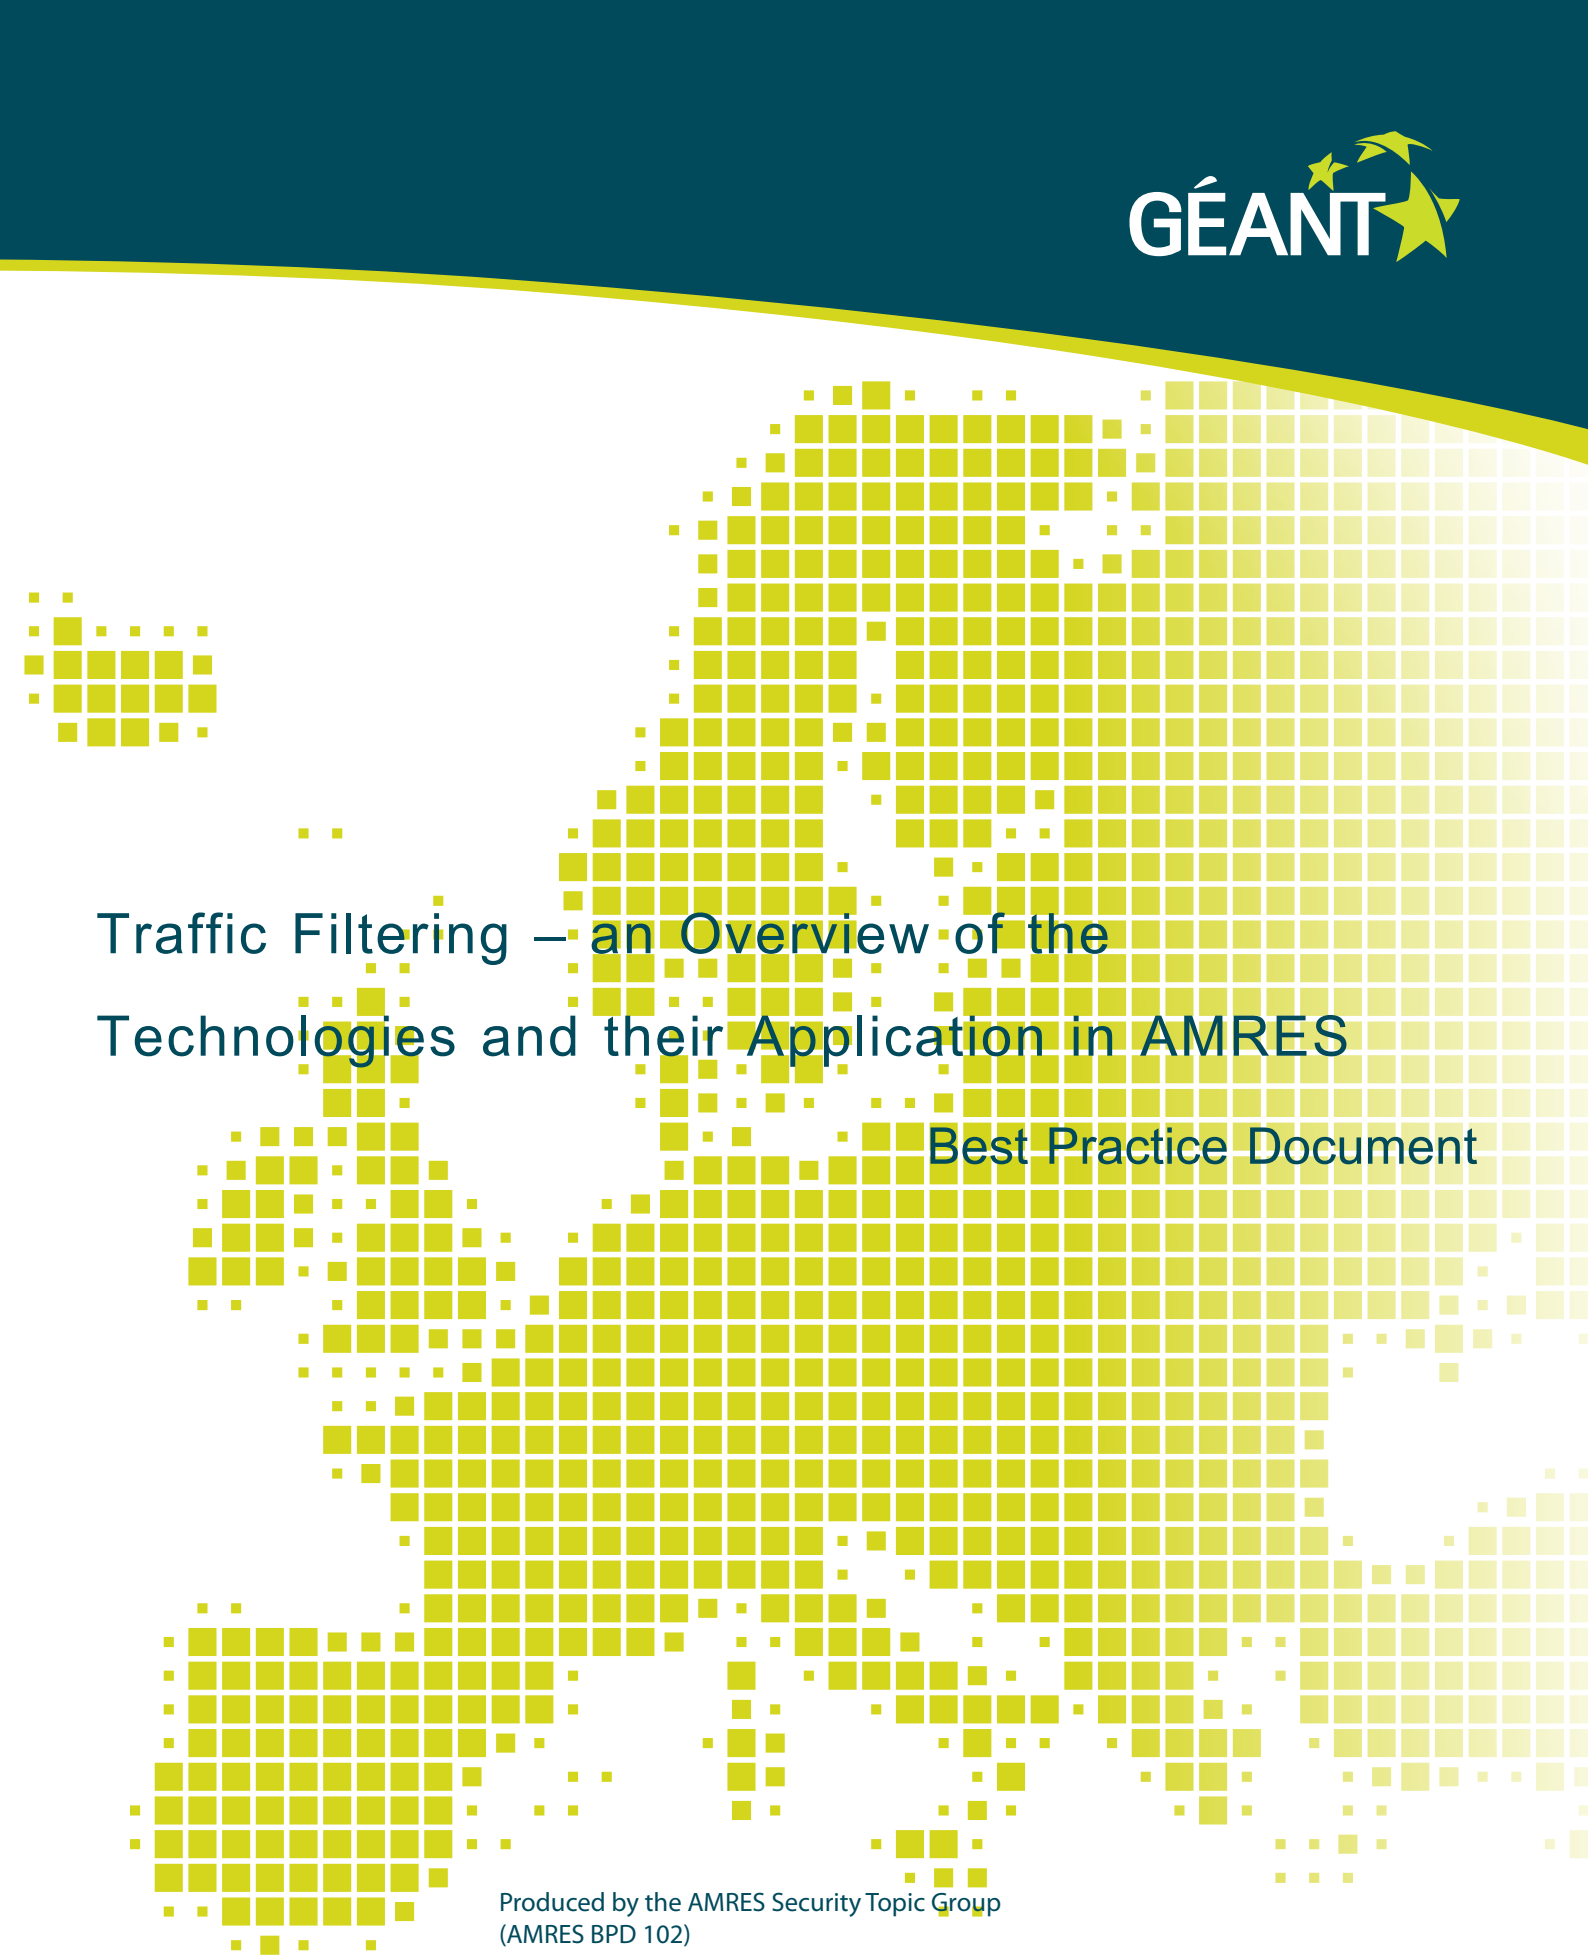

Traffic Filtering – an Overview of the Technologies and their Application in AMRES

Best Practice Document

Produced by the AMRES Security Topic Group
(AMRES BPD 102)

Authors: Zoran Mihailović, Bojan Jakovljević
and Mara Bukvić

July 2011

Document No:	GN3-NA3-T4-AMRES-BPD-102
Version / date:	July 2011
Source language:	Serbian
Original title:	"*Preporuke za filtriranje saobraćaja u krajnjim institucijama*"
Original version / date:	Revision 1 (of a part of the document of May 2010) / 11 July 2011
Contact:	helpdesk@rcub.bg.ac.rs

AMRES/RCUB is responsible for the contents of this document. The document was developed by the Security Topic Group organised at AMRES with the purpose of implementing joint activities on developing and disseminating documents containing technical guidelines and recommendations for network services in higher-education and research institutions in Serbia.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 238875, relating to the project 'Multi-Gigabit European Research and Education Network and Associated Services (GN3)'.

Table of Contents

Executive Summary

The objective of this document is to present the available traffic-filtering technologies and their general application, as well as to indicate the procedures and planned applications in the hierarchical structure of the AMRES network.

In order to reduce security threats, AMRES uses various devices, technologies and techniques for traffic filtering. Depending on its requirements and its needs, each institution/organisation that wishes to improve the efficiency of filtering and security in its network should select a technology and set of rules for traffic filtering to implement in its local network.

A better understanding of the manner in which the traffic is filtered at certain places in the AMRES network will make the process of defining the filtering rules in individual institutions easier and will enable better compliance with the rules and technologies applied in the rest of the network.

This document is primarily intended for managers and administrators who are responsible for the development of filtering rules and the selection of filtering technologies for the network of their institution. The document can also help network administrators in campus networks indirectly with identifying and eliminating the causes of traffic-filtering problems at specific places in their networks.

Introduction

A large number of workstations, servers and other devices at AMRES are daily exposed to threats, both within the AMRES network and from the Internet. For that reason, AMRES administrators are faced with the challenge of blocking threats as much as possible, and at the same time enabling the normal functioning of the network. In addition, it is necessary to prevent the spread of threats from a network for which an administrator is responsible to other networks within AMRES, as well as to the Internet.

In order to address this challenge, administrators filter traffic at various places in the AMRES network. Traffic filtering can be achieved by implementing filtering rules on routers and/or firewall devices, as well as through the application of other technologies described in this document.

This document has been prepared with a view to acquainting the administrators who manage networks, either on campuses, at the level of a service centre or at AMRES itself, with the basic rules and practices of packet filtering at AMRES, including the places at which traffic filtering is implemented in the hierarchical topology of the AMRES network.

A better understanding of the manner in which the traffic is filtered at certain places in the AMRES network will make the process of defining rules for individual institutions easier, and will enable better compliance of those rules with the practice applied in the rest of the network.

The procedure of development and the application of packet filters in campus networks/organisations of AMRES are defined in the document entitled AMRES BPD 110 *Traffic Filtering in Campus Networks*.

This document also contributes to a better understanding of the need to define and publish the existing rules in the form of a policy, as well as to the development and adoption of this policy in AMRES institutions.

1 Four Basic Recommendations for Traffic Filtering

In order to reduce security threats, AMRES uses various devices, technologies and techniques for traffic filtering. Each institution/organisation that wishes to improve the efficiency of filtering and increase the level of security in its network should apply the following recommendations:

1. Define traffic-filtering rules that will determine the manner in which the incoming and outgoing traffic flow in the network will be regulated. A set of traffic-filtering rules can be adopted as an independent packet filtering policy or as a part of the information security policy;
2. Select a traffic-filtering technology that will be implemented depending on the requirements and needs;
3. Implement defined rules on the selected technology and optimise the performance of devices accordingly;
4. Maintain all the components of the solution, including not only devices, but also the policy.

It should be noted that these recommendations can be applied in any organisation, including AMRES as a whole.

The document entitled AMRES BPD 102 *Traffic Filtering – An Overview of Technologies and their Application in AMRES* has been developed in order to support the implementation of Recommendation number 2. The document entitled AMRES BPD 110 *Traffic Filtering in Campus Networks* has been prepared to support the implementation of Recommendations 1 and 3 in campus networks.

2 An Overview of Traffic-filtering Technologies

Traffic-filtering technologies are commonly divided into packet filtering/stateless firewall and stateful firewall technologies.

The packet-filtering functionality (stateless firewall) is built into the majority of operating systems and devices with a traffic routing feature. In most cases, it is a router on which access control lists (ACLs) are applied.

A packet filter implemented on a router is the simplest, but only one of the available traffic-filtering methods. This chapter provides a description of other available firewall technologies and their general application. It also provides recommendations for their application within the hierarchical structure of the AMRES network.

2.1 Packet Filter vs. Stateful Firewall

Packet filtering is the basic feature of all firewall devices. The first firewall devices, with only a packet filter, were also called stateless inspection firewalls. Unlike them, modern firewall devices provide far more possibilities for packet filtering.

A packet filter enables the implementation of control of access to resources by deciding whether a packet should be allowed to pass, based on the information contained in the IP packet header. The packet filter does not analyse the content of the packet (unlike a content filter), nor does it attempt to determine the sessions to which individual packets belong, based on the information contained in the TCP or UPD header, and therefore it does not make any further decisions in that regard. For this reason, the process is also known as **stateless packet inspection**.

Due to its manner of operation, which does not track the information on the state of connections, it is necessary to explicitly allow two-way traffic on the connection when configuring a stateless firewall device.

Stateless firewall devices analyse each packet individually and filter them based on the information contained in Layers 3 and 4 of the OSI reference model.

A filtering decision is made based on the following information:

- source IP address;

- destination IP address;
- protocol;
- source port number;
- destination port number.

They are commonly implemented as a part of the functionality on routers (ACL, firewall filters, etc.), but can also be implemented on servers.

The advantages of applying packet filters:

- simple implementation;
- supported by most routers, so there is no need to invest in new equipment and software;
- rarely cause bottlenecks in the area of their application, even at high speeds in Gigabit networks.

The disadvantages of applying packet filters:

- vulnerability to IP spoofing attacks;
- vulnerability to attacks that exploit problems within the TCP/IP specification and the protocol stack;
- problems with filtering packets that are fragmented (causing interoperability and non-functioning of VPN connections);
- no support for the dynamic filtering of some services (the services that require dynamic negotiation about the ports that will be used in communication – passive FTP).

The places where packet filters are applied in the AMRES network

Bearing in mind the hierarchical structure of AMRES, described in Chapter 3, packet filters are used on the AMRES connections to the Internet (position/layer 1), on the connections between individual institutions/organisations, on the regional service centre to which they belong (position/layer 3) and on the servers dedicated to a specific service (dedicated servers). They can also be found in the internal networks of member institutions/organisations, although rarely.

Stateful packet inspection improves the packet filtering process by monitoring the state of each connection established through a firewall device.

It is known that the TCP protocol allows two-way communication and that TCP traffic is characterised by three phases: establishing the connection, data transfer, and terminating the connection. In the connection establishment phase, stateful packet inspection records each connection in the state-table.

In the data transfer phase, the device monitors certain parameters in the header of the L3 packet and L4 segment and makes a filtering decision depending on their values and the content of the state-table. The state-table contains all currently active connections. As a result, a potential attacker trying to spoof a packet with a header indicating that the packet is a part of an established connection can only be detected by the stateful inspection firewall device, which verifies whether the connection is recorded in the state-table.

The state-table contains the following information:

- source IP address;
- destination IP address;
- source port number;
- destination port number;
- TCP sequence numbers;
- TCP flag values.

The state of the *synchronize* (SYN), *reset* (RST), *acknowledgment* (ACK) and *finish* (FIN) flags are monitored within the TCP header and a conclusion is reached about the state of a specific connection.

The UDP protocol does not have a formal procedure for establishing and terminating a connection. However, devices with stateful inspection can monitor the state of individual flows[1] and match different flows when they logically correspond to each other (e.g., a DNS response from an external server will only be allowed to pass if the corresponding DNS query from the internal source to that server has previously been recorded).

The advantages of applying stateful firewall devices:

- a higher level of protection compared to stateless firewall devices (greater efficiency and more detailed traffic analysis);
- detection of IP spoofing and DoS attacks;
- more log information compared to packet filters.

The disadvantages of applying stateful firewall devices:

- no protection against application layer attacks;
- performance degradation of the router on which they are deployed (this depends on the size of the network and other services run on the router);
- not all of them provide support for UDP, GRE and IPSEC protocols, treating them in the same way as stateless firewall devices;
- no support for user authentication.

The places where stateful firewall devices are applied in the AMRES network

Bearing in mind the hierarchical structure of AMRES, described in Chapter 3, stateful firewall devices may not be applied above position/layer 3. They can also be found in the internal networks of member institutions/organisations.

[1] Flows are defined by specific source and destination IP addresses, a pair of ports (source, destination), the protocol, the TOS field and the device's input interface.

Recommendation

Compare the requirements and the institution's needs for traffic filtering with the features/performance of the device you plan to purchase. Based on this, decide whether the investment in a significantly more expensive firewall device is justified. Bear in mind the experience of AMRES administrators who believe that a packet filter implemented on the router is a sufficient solution for most of the needs of smaller institutions.

2.2 A Description of Other Available Firewall Technologies

Lately, attempts have been made to improve the standard stateful packet inspection technology by adding basic solutions from intrusion detection technology. The improved version is called stateful protocol analysis, also known as DPI (Deep Packet Inspection) analysis of data on the application layer.

The devices resulting from this development trend include Application Firewall, Application Proxy Gateways and Proxy servers.

Unlike stateful firewall devices that filter traffic based on the data on layers 3, 4 and 5 of the OSI reference model, these devices also enable traffic filtering based on the information on the application layer of the OSI reference model (Layer 7).

2.2.1 Application Firewall

Application Firewall (AF) devices perform a stateful protocol analysis of the application layer. They support numerous common protocols, such as HTTP, SQL, e-mail service (SMTP, POP3 and IMAP), VoIP and XML.

Stateful protocol analysis relies on predefined profiles of acceptable operating modes for the selected protocol, enabling the identification of potential deviations and irregularities in the message flow of the protocol through the device. Problems may arise if there is a conflict between the operating mode of a specific protocol, which is defined on the AF device, and the way in which the protocol is implemented in the specific version of the application or of the operating systems used in the network.

The stateful protocol analysis can:

- determine whether an e-mail message contains a type of attachment that is not allowed (e.g., *exec* files);
- determine whether instant messaging is used via an HTTP port;
- block the connection through which an unwanted command is executed (e.g., an FTP *put* command on the FTP server);
- block access to a page with unwanted active content (e.g., Java);
- identify an irregular sequence of commands exchanged in the communication between two hosts (e.g., an unusually large number of repetitions of the same command or the use of a command before using the command it depends on);

- enable the verification of individual commands and the minimum and maximum length of appropriate command-line arguments (e.g., the number of characters used in a username).

An AF device cannot detect attacks that meet the generally acceptable procedures of operation of a specific protocol, such as DoS (Denial of Service) attacks caused by the repetition of a large number of acceptable message sequences in a short time interval.

Due to the complexity of the analysis they perform, and the large number of concurrent sessions they monitor, the main disadvantage of the method of stateful protocol analysis is the intensive use of AF devices.

The places where AF devices are applied in the AMRES network

Bearing in mind the hierarchical structure of AMRES, described in Chapter 3, AF devices are not used frequently, although their application is conditionally justified in the internal networks of the AMRES member institutions. The application of AF devices cannot be expected beyond position 3.

2.2.2 Application Proxy Gateway

Application Proxy Gateway (APG) devices also perform an analysis of the traffic flow on the application layer. Compared to AF devices, APG devices provide a higher level of security for individual applications since they never allow a direct connection between two hosts, and they can perform an inspection of the content of application-layer messages.

APG devices contain so-called proxy agents or “intermediaries” in the communication between two end hosts. In this way, they prevent direct communication between them. Each successful connection between the end hosts consists of two connections – one between the client and the proxy server and the other between the proxy server and the destination device. Based on the filtering rules defined on the APG device, proxy agents decide whether network traffic will be allowed or not. Traffic-filtering decisions can also be made based on the information contained in the header of an application-layer message or even based on the content conveyed by that message. In addition, proxy agents can require user authentication.

There are also APG devices with the capability of packet decryption, analysis and re-encryption, before a packet is forwarded to the destination host. Packets that cannot be decrypted are simply forwarded through the device.

Compared to packet filters and stateful devices, APG devices have numerous deficiencies. The manner of operation of APG devices requires a significantly greater utilisation of resources, i.e., they require more memory and greater utilisation of processor time for analysing and interpreting each packet passing through the device. As a result, APG devices are not suitable for filtering applications that are more demanding in terms of bandwidth or applications that are sensitive to time delays (real-time applications). Another deficiency of these devices is the limitation in the number of services that can be filtered through them. Each type of traffic passing through the device requires a specific proxy agent that acts as an intermediary in the communication. Consequently, APG devices do not always support the filtering of new applications or protocols.

Due to their price, APG devices are commonly used for protecting data centres or other networks containing publicly available servers that are of high importance to an organisation.

In order to reduce the load on APG devices and achieve greater efficiency, modern networks more frequently use proxy servers (dedicated proxy servers) that are dedicated to specific services that are not so sensitive to time delays (e.g., e-mail or web proxy servers).

The places where APG devices are applied in the AMRES network

Bearing in mind the hierarchical structure of AMRES, described in Chapter 3, APG devices are not used frequently, although their application is conditionally justified in the internal networks of the AMRES member institutions. The application of APG devices cannot be expected beyond position 3.

2.2.3 Dedicated Proxy Server

Like APG devices, Dedicated Proxy (DP) servers also have a role as “intermediaries” in the communication between two hosts, although their traffic-filtering capabilities are significantly lower. This type of device is intended for the analysis of the operation of specific services and protocols (e.g., HTTP or SMTP).

Due to their limited traffic-filtering capabilities, DP devices are deployed behind firewall devices in the network architecture. Their main function is to perform specialised filtering of a specific type of traffic (based on a limited set of parameters) and carry out the logging operation. The execution of these specific activities significantly reduces the load on the firewall device itself, which is located in front of the DP server.

The most widely used devices of this type are Web Proxy servers. A common example of their use is an HTTP proxy server (placed behind the firewall device or router), to which users need to connect when they wish to access external web servers. If an institution has an outgoing connection (uplink) of lower bandwidth, the use of the caching function is recommended in order to reduce the level of traffic and improve the response time.

As a result of an increase in the number of available web applications and the number of threats transferred through the HTTP protocol, Web Proxy servers are growing in significance. Consequently, many equipment manufacturers today add the functionality of various firewall technologies to the standard Web Proxy servers, thus increasing their traffic-filtering capabilities.

The places where dedicated proxy servers are applied in the AMRES network

Bearing in mind the hierarchical structure of AMRES, described in Chapter 3, dedicated proxy servers in AMRES service centres (position/layer 2) are configured in such a way that they can support the optional application of a proxy server in the internal networks of AMRES member institutions (position/layer 4). The configuration and application of the proxy technology at various places need to be harmonised.

2.3 Solutions Combining Traffic Filtering with Other Technologies

In addition to their basic purpose of blocking unwanted traffic, firewall devices often combine their filtering functionality with other technologies, primarily routing. It is the other way around with routers. As a result, NAT (Network Address Translation) is sometimes considered to be a firewall technology, although essentially it is a routing technology.

Other related functionalities, such as VPN and IDP, are often available on firewall devices. In order to have a complete overview and due to their frequent use, these technologies are also addressed briefly in this chapter.

2.3.1 NAT (Network Address Translation)

NAT is a technology that enables devices that use private IP addresses to communicate with devices on the Internet. This technology translates private IP addresses, which can be used by devices within a Local Area Network (LAN), into publicly available Internet addresses.

The application of NAT technology may limit (intentionally or unintentionally) the number of available services, i.e., it may disable the functioning of the services that require direct, end-to-end connectivity (e.g., VoIP).

There are three types of NAT translations: dynamic, static and PAT.

Dynamic NAT uses a set of publicly available IP addresses, successively assigning them to hosts with private IP addresses. When a host with a private IP address needs to communicate with a device on the Internet, dynamic NAT translates its private IP address into a publicly available IP address, by taking the first available IP address from a defined pool of publicly available IP addresses. Dynamic NAT is suitable for client computers.

Static NAT provides one-to-one mapping between the private IP address of a host and the public IP address assigned to it. In this manner, the host with a private IP address always appears on the Internet with the same public IP address. This is the main difference between static and dynamic translation. Static NAT is suitable for servers.

In both types of translation mentioned above, each private IP address is translated into a separate, public IP address. In order to support a sufficient number of simultaneous user sessions, an organisation using dynamic and/or static NAT needs to have a sufficient number of public IP addresses.

PAT (Port Address Translation or so-called NAT overload) performs mapping between several private IP addresses and one or more public IP addresses. The mapping of each private IP address is performed by way of the port number of the public IP address. PAT translation ensures that each client on a LAN that establishes a connection with a device on the Internet is assigned a different port number of the public IP address. The response from the Internet, which comes as a result of the request, is sent to the port from which the request was forwarded. In this manner, a device that performs the translation (a router, firewall or server) knows to which host from the LAN it should forward the packet. This feature of PAT increases the level of security of the LAN to a certain degree, since it prevents a connection from the Internet being established directly with the hosts on the LAN. Due to this manner of operation, PAT is sometimes, incorrectly, regarded as a security technology, although it is primarily a routing technology.

2.3.2 VPN (Virtual Private Network)

VPN (Virtual Private Network) technology is used to increase the security of data transfer through a network infrastructure that does not provide a sufficient degree of data security. It enables the encryption and decryption of network traffic between external networks and an internal, protected network.

VPN functionality can be available on firewall devices or implemented on VPN servers that are placed behind firewall devices in the network architecture.

In many cases, the implementation of VPN services on a firewall device itself is the most optimal solution. Placing a VPN server behind the firewall device requires the VPN traffic to pass through the firewall device in an encrypted form. As a result, the firewall device cannot perform an inspection, access control or logging of the network traffic, and therefore cannot scan it for certain security threats.

However, regardless of the place of the implementation, the VPN service requires the application of certain filtering rules of the firewall device in order to enable its uninterrupted operation. Accordingly, special attention should always be paid to making sure that the appropriate protocols and the TCP/UDP services that are necessary for the functioning of the chosen VPN solution are supported.

2.3.3 IDP (Intrusion Detection and Prevention)

Network Intrusion Detection (ID) is based on monitoring the operation of computer systems or networks and analysing the processes they perform, which can point to certain incidents.

Incidents are events posing a threat to or violating defined security policies, violating AUP (Acceptable Use Policy) rules, or generally accepted security norms. They appear as a result of the operation of various malware programmes (e.g., worms, spyware, viruses, and Trojans), as a result of attempts at unauthorised access to a system through public infrastructure (Internet), or as a result of the operation of authorised system users who abuse their privileges.

Network Intrusion Prevention (IP) includes the process of detecting network intrusion events, but also includes the process of preventing and blocking detected or potential network incidents.

Network Intrusion Detection and Prevention systems (IDP) are based on identifying potential incidents, logging information about them, attempting to prevent them and alerting the administrators responsible for security. In addition to this basic function, IDP systems can also be used to identify problems concerning the adopted security policies, to document existing security threats and to discourage individuals from violating security rules.

IDP systems use various incident-detection methods. There are three primary classes of detection-methodology:

a. **Signature-based detection**

Certain security threats can be detected based on the characteristic manner in which they appear. The behaviour of an already detected security threat, described in a form that can be used for the detection of any subsequent appearance of the same threat, is called an attack signature.

This detection method, based on the characteristic signature of an attack, is a process of comparing the known forms in which the threat has appeared with the specific network traffic in order to identify certain incidents.

Although it can be very efficient in detecting the subsequent appearance of known threats, this detection method is extremely inefficient in the detection of completely unknown threats, of threats hidden by using various techniques, and of already known threats that have somehow been modified in the meantime. It is considered the simplest detection method and it cannot be used for monitoring and analysing the state of certain, more complex forms of communication.

b. **Anomaly-based detection**

This method of IDP is based on detecting anomalies in a specific traffic flow in the network. Anomaly detection is performed, based on the defined profile of acceptable traffic and its comparison with the specific traffic in the network.

Acceptable traffic profiles are formed by tracking the typical characteristics of the traffic in the network during a certain period of time (e.g., the number of e-mail messages sent by a user, and the number of attempts to log in to a host, or the level of utilisation of the processor in a given time interval). These characteristics of the behaviour of users, hosts, connections or applications in the same time interval are then considered to be completely acceptable.

However, acceptable-behaviour profiles can unintentionally contain certain security threats, which lead to problems in their application. Likewise, imprecisely defined profiles of acceptable behaviour can cause numerous alarms, generated by the system itself as a reaction to certain (acceptable) activities on the network.

The greatest advantage of this detection method is its exceptional efficiency in detecting previously unknown security threats.

c. **Detection based on stateful protocol analysis**

Stateful protocol analysis is a process of comparing predefined operation profiles with the specific data flow of that protocol on the network. Predefined profiles of operation of a protocol are defined by the manufacturers of IDP devices and they identify everything that is acceptable or not acceptable in the exchange of messages in a protocol. Unlike anomaly-based detection, where profiles are created based on the hosts or specific activities on the network, stateful protocol analysis uses general profiles generated by the equipment manufacturers.

Most IDP systems use several detection methods simultaneously, thus enabling a more comprehensive and precise method of detection.

3 The Application of Traffic-filtering Technologies in the Hierarchical Structure of AMRES

Traffic-filtering rules can be defined and adopted at various places on a network.

Taking into account the topology of the Academic Network of Serbia (AMRES), several places can be identified where traffic filtering is performed:

- Position/Layer 1 – on the AMRES connections to the Internet;
- Position/Layer 2 – in the AMRES service centres (RCUB, CIT-UNS, JUNIS and UNIC);
- Position/Layer 3 – on the connections between certain institutions/organisations and the regional service centre they belong to and are connected to;
- Position/Layer 4 – on the internal network of the AMRES member institutions/organisations.

This chapter describes specific solutions, i.e., the manner in which traffic-filtering technologies have been applied or their application can be expected at certain hierarchical positions in the AMRES network. A better understanding of the manner in which the traffic is filtered at certain places in the AMRES network will make the process of defining policies at individual institutions easier for administrators and will enable better compliance with the policies that are applied in the rest of the network.

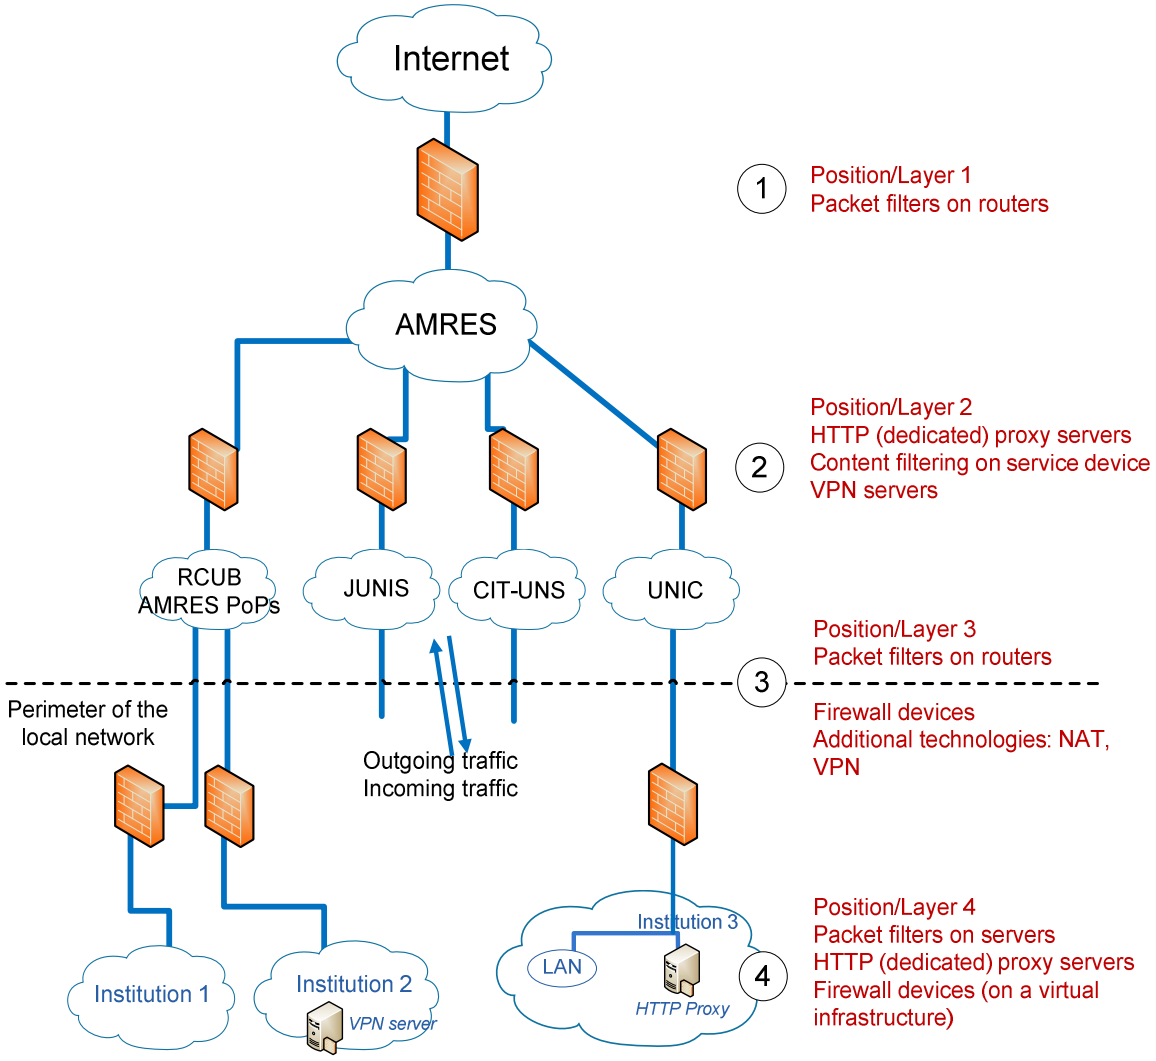

The following are the basic policies and practices applied at the traffic-filtering locations in AMRES:

- Packet filtering of the traffic on routers, whose main purpose is to block disallowed protocols, destination service ports and IP addresses, as well as disallowed source IP addresses and service ports, is applied at all the connections of the AMRES network with the Internet.

  Filtering at this level should be regulated by the *Policy for Traffic Filtering* at AMRES. The *Policy for Traffic Filtering* is established at the level of the NREN, and is adopted by the managing body of the NREN. A draft of the *Policy for Traffic Filtering* has been prepared. Until its adoption, AMRES applies the rules contained in the Code that was established in earlier phases of the development of AMRES.

- On position 2, presented in the above figure, traffic-filtering techniques are applied on the application layer. Dedicated proxy servers and occasionally other special-purpose devices (such as content-filtering devices) are used. The solutions and techniques that are used continually are implemented exclusively in the AMRES services centres: RCUB, JUNIS, CIT-UNS and UNIC. Special-purpose devices (other solutions and techniques), which can be applied occasionally in parts of the network or in the network as a whole, can be implemented as necessary, not only in the AMRES service centres, but also in other larger AMRES nodes (PoP AMRES).

On position 2, the proxy servers are the most important technology for HTTP traffic filtering. According to on-line statistics, HTTP traffic accounts for 80% of the total traffic. Traditional reasons – the mitigation of the lack of resources and maintenance of bandwidth - are no longer the primary reasons for using proxy servers in AMRES. The current reasons are primarily related to the obligation of recording traffic in AMRES, protection against various security threats and controlled support for the KOBSON (Consortium of the Coordinated Acquisition of Serbian Libraries) service.

Among other technologies, content filtering can be applied on layer 2.

The document that should regulate traffic filtering on layer 2 is the *Policy for Traffic Filtering* in the AMRES network, also used for the previous layer (position 1).

- Access to scientific journals is ensured through the service provided by KOBSON, and it may be gained **exclusively** by using designated HTTP proxy servers on position 2. As requested by the National Library and the competent ministry, the KOBSON service needs to be limited and only available to users from member institutions of AMRES. A solution requiring the authentication of each individual user accessing the KOBSON service has been avoided, and the above requirement has been met through the AMRES proxy service. Access is monitored on proxy servers at the AMRES service centres. All HTTP proxy servers on position 2 are used in a non-transparent mode. This means that the computers of end users in the AMRES member institutions need to be configured in order to pass through the AMRES proxy service.

- The above figure shows that position 3 is located between the service centre/AMRES PoP and individual institutions that access the backbone through that service centre/AMRES PoP. The document entitled AMRES BPD 110 *Traffic Filtering in Campus Networks* contains recommendations for the AMRES member institutions concerning the development and application of traffic-filtering rules on position 3.

  It is good practice to have the service centres filter traffic on behalf of their faculties/institutes in one place – at the point of the delivery of traffic to/from the AMRES PoP. Regardless of the fact that this solution is physically implemented in the AMRES service centre, it logically belongs to position 3, so that the recommendations that are valid for any other individual institution can be applied.

- Almost all traffic-filtering technologies presented in the previous chapters, or combinations thereof, can be applied in position 3. In general, packet filters on routers/L3 switches are used, either independently or in combination with the NAT technology. Another popular option is the use of firewall devices, on which a VPN server is occasionally placed.

- AMRES service centres are advised to apply a packet filter on the interface connecting each institution, because its purpose is to defend the rest of the network from unauthorised traffic coming from the network of an AMRES member institution (IP spoofing, etc.). Please note that this recommendation concerns the application of the so-called ***set of general filtering rules of AMRES*** to just one direction of an institution's traffic (the outgoing traffic of the AMRES members' networks).

  The set of general filtering rules of AMRES includes rules concerning anti-spoofing, anti-spam, limitations of the use of private space and address space, limitations to the use of the ICMP protocol, and limitations to the use of protocols used predominantly in LAN networks (NetBIOS, SQL and similar types). For more information on the set of general filtering rules, please consult the document entitled AMRES BPD 110 *Traffic Filtering in Campus Networks*. The set of general filtering rules is included in the Draft of *The Policy for Traffic Filtering at AMRES*.

- Layer 3 is also the perimeter of the network of an AMRES member. The terms ‘incoming’ and ‘outgoing’ traffic are defined in relation to this position. Outgoing traffic is the traffic generated by a resource located inside the internal network of an AMRES member attempting to access an external service. Alternative terms used for outgoing traffic are ‘exit traffic’ or ‘egress traffic’. Incoming traffic is the traffic generated by an external resource attempting to access a service within the internal segment of the network of an AMRES member. Alternative terms for incoming traffic are ‘entrance traffic’ or ‘ingress traffic’.

  Each AMRES member institution should control the incoming and outgoing traffic on the perimeter of its network independently in accordance with the policies applied at AMRES, which may be expanded by defining policies and applying technologies adopted by the institution. This requires certain efforts on the part of the AMRES member institutions/organisations, especially in the beginning. The AMRES BDP 110 document on traffic filtering on campus networks contains recommendations and suggestions (including the commands for configuring the packet filters on the IOS operating system). These recommendations will help the member institutions in the process of taking over filtering activities, although the individual institutions/organisations are not obliged to apply these recommendations. Ultimately, the application of the recommendations contained in this document will increase the level of security and efficiency of traffic filtering across AMRES.

  The document to which an AMRES member institution should refer in order to regulate traffic filtering at the third, fourth and all other layers, if it wishes to do so, is its *Policy for Traffic Filtering* (if that exists as a separate document). Traffic filtering may also be regulated as part of its Information Security Policy, if such a document exists.

- A technical solution currently applied at several nodes in the AMRES network allows the network perimeter of a campus network to be located “inside” a device (most commonly an L3 switch under the administrative control of the AMRES service centre), rather than at the connection between two devices, as illustrated in the figure above. Due responsibility issues, all such cases require more coordination between the service centre and the campus network in terms of implementing and maintaining the access lists. The issue of the transfer of responsibility is resolved through the use of an authorisation method available on the network device. For Cisco devices (which are used in most situations), two solutions have been tested that expand the user’s rights to include the set of commands necessary for the implementation of access lists. Authorisation is achieved through assigned roles in the Role-based CLI environment or by defining different levels of privileges for using the *enable password* command.

- Level 4 has been included in order to describe different needs and methods of implementing traffic-filtering technologies within the network of individual member institutions of AMRES. A particularly important point is the use of a proxy service.

  As noted above, the HTPP proxy server of an institution can be implemented behind the packet filter on position 3 (on position 4 in the internal network of an AMRES member institution). In this case, it is necessary to bear in mind that access to the KOBSON service can only be obtained by using a proxy server at AMRES service centres (layer 2). Therefore, member institutions of AMRES that use their own proxy servers are advised to have them configured in parent mode in relation to the HTTP proxy servers at AMRES service centres (on position 2).

  It is possible to configure an independent proxy server at an AMRES member institution in such a way that it accesses HTTP or HTTPS servers outside AMRES directly, except for the KOBSON service, which can only be accessed through the proxy servers at AMRES service centres. Such a solution requires more skill to configure and more effort to maintain.

All the service centres of AMRES, as well as all of the AMRES members, must use the *X-Forwarded-For* standard for identifying the original IP address of a client who has initiated the traffic passing through the proxy server. Log files must be kept for at least three months.

- AMRES member institutions may use a firewall on position 4, illustrated in the figure above, in order to ensure an additional level of security in their local networks. This is most often done with a view to preventing unauthorised access to the parts of the network where particularly important resources are located or where certain sensitive functions are performed, such as accounting and student services.

- All member institutions of AMRES are advised to use packet filters or various firewall solutions available on their servers that allow access only to the services of that particular server, and to disable all the other ports on the server.

4 References

[1] Karen Scarfone, Paul Hoffman, *Guidelines on Firewalls and Firewall Policy*, Recommendations of the NIST, September 2009.

[2] Eizabet D. Zwicky, Simon Cooper & D. Brent Chapman, *Building Internet Firewalls*, O'Reilly Media, Second Edition, June 2000.

[3] Brian Morgan, Neil Lovering, *CCNP ISCW Official Exam Certification Guide*, Cisco Press, July 2007.

[4] NIST SP 800-95, *Guide to Secure Web Services*, http://csrc.nist.gov/publications/PubsSPs.html.

[5] NIST SP 800-94, *Guide to Intrusion Detection and Prevention Systems (IDPS*,) http://csrc.nist.gov/publications/PubsSPs.html.

[6] *JNCIA-Junos Study Guide—Part 2*, https://learningportal.juniper.net

5 Glossary

ACL	Access Control List
AF	Application Firewall
AMRES	Academic Network of Serbia
APG	Application Proxy Gateway
AUP	Acceptable Use Policy
CIT-UNS	Centre for Information Technology of the University of Novi Sad
DP	Dedicated Proxy
DPI	Deep Packet Inspection
DNS	Domain Name System
DoS	Denial of Service
FTP	File Transfer Protocol
GRE	Generic Routing Encapsulation
HTTP	HyperText Transfer Protocol
HTTPS	HyperText Transfer Protocol Secure
ICMP	Internet Control Message Protocol
IDP	Intrusion Detection and Prevention
ID	Intrusion Detection
IMAP	Internet Message Access Protocol
IP	Internet Protocol
IP	Intrusion Prevention
IPSec	Internet Protocol Security
JUNIS	Unique Research and Education Information System of the University of Niš
KOBSON	Consortium of the Coordinated Acquisition of Serbian Libraries enabling the procurement of scientific magazines
LAN	Local Area Network
L3	OSI layer 3 – Network layer
L4	OSI Layer 4 – Transport layer
NAT	Network Address Translation
NetBIOS	Network Basic Input/Output System
NREN	National Research and Education Network
OSI	Open Systems Interconnection
PAT	Port Address Translation
POP3	Post Office Protocol Version 3
RCUB	Computer Centre of the University of Belgrade
SMTP	Simple Mail Transfer Protocol
SQL	Structured Query Language

TCP	Transmission Control Protocol
UNIC	Computer Centre of the University of Kragujevac
VoIP	Voice over Internet Protocol
VPN	Virtual Private Network
XML	Extensible Markup Language