Centralised web traffic filtering system

Best Practice Document

Produced by the AMRES-led working group
on Security (AMRES BPD 113)

Authors: Ivan Ivanović, Miloš Kukoleča,
Jovana Palibrk

March, 2013

Document No: GN3-NA3-T4-AMRES-BPD-113
Version / date: March 2013
Original language: Serbian
Original title: “Centralizovani sistem za filtriranje web saobraćaja”
Original version / date: Version1 / 19 March 2013
Contact: ivan.ivanovic@rcub.bg.ac.rs, milos.kukoleca@amres.ac.rs, jovana.palibrk@amres.ac.rs

AMRES/RCUB is responsible for the contents of this document. The document was developed by the AMRES-led working group on Security with the purpose of implementing joint activities on the development and dissemination of documents encompassing 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

Summary

The purpose of this document is to introduce the reader to an IronPort firewall technical solution for web traffic filtering that can be used in a campus environment. In addition to describing the design, configuration and positioning of the centralised IronPort firewall system, the document also deals with some important recommendations regarding the mechanisms that ensure a redirection and even distribution of web traffic toward the firewall devices, and the collection and analysis of the data on user activity in the network. The document also looks into the advantages and shortcomings of such a centralised system. Although the document describes the operation of the specialised IronPort firewall equipment, certain ideas and techniques can be also applied to the equipment of other manufacturers.

Introduction

As a result of the increased activity of users on the Internet, and the development of malicious software that enables highly sophisticated attacks, it has become necessary to increase the level of security for end users within the AMRES network. The protection methods that have been used so far for web traffic filtering have relied on commonly used access lists that were placed on router platforms or proxy servers in the network. This method of web traffic filtering has its limitations, as it only enables filtering that is based on protocols, ports and IP addresses.

The lack of traffic filtering and inspection above layer L4 of the OSI model is the main reason why it has become necessary to install a firewall system that provides additional protection, using content filtering.

1 The System Architecture and Design

Figure 1 shows the architecture of the system used within the AMRES network. As presented below, the network part of the system is divided into two segments. The first segment is the network that enables the traffic flow of production data, while the second segment is the network that carries the management data traffic. Since the IronPort devices are equipped with several gigabit Ethernet network interfaces, one of the interfaces is used for the production traffic (P interface) and the other for the management traffic (M interface). In this way, the production traffic is physically separated from the management traffic by way of an OOB (Out-of-Band) network. The M and P ports of all IronPort devices are connected to the rest of the network via gigabit Ethernet switches, in order to establish the high throughput links. The OOB part of the network is reached through the NAT router, to which access is granted only from the IP addresses belonging to the administrators of the IronPort system.

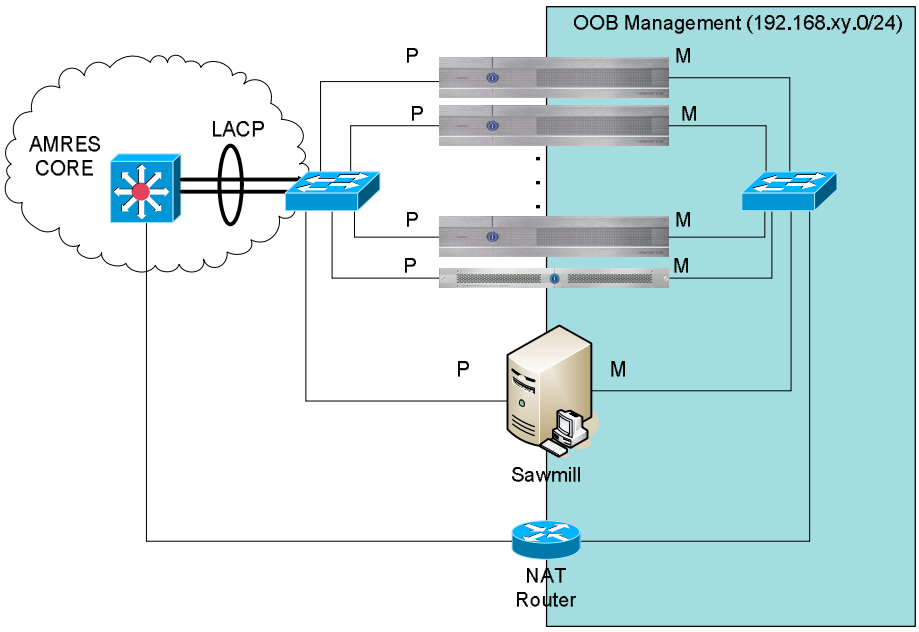

Figure 1 - The Architecture of the System.

P (production) interfaces are used for production web traffic.

M (management) interfaces serve for managing the devices, i.e., they enable:

- access to and administration of the IronPort devices via the WEB and SSH protocols;
- the transport of the user activity logs on the Sawmill server in the OOB part of the network;
- access to and administration of devices through the central IronPort M160 management device;
- monitoring of the devices via the SNMP protocol.

2 Centralised Management

The brain of the web traffic filtering system is comprised of one IronPort M160 centralised management device and five IronPort S670 firewall devices. The IronPort M160 is a centralised management device that provides the administrator with tools to apply the unique configuration simultaneously to all available IronPort S670 firewall devices. The IronPort S670 devices are firewall devices that process web traffic based on a set of rules defined by the administrator. This way, administrator avoids repeating the same configuration on several firewall devices manually and reduces the possibility of making errors in the configuration. Likewise, this ensures the consistency of the configuration on all firewall devices. The IronPort M160 device is also capable of gathering and processing log information, but this solution requires the purchase of an additional licence. In the case of the solution employed within AMRES, the gathering and processing of log information has been taken care of in another way, as described in Chapter 7 of this document.

The term “firewall device” will hereinafter be used to denote all the IronPort S670 devices with a web traffic filtering and blocking functionality, while the term “centralised management device” will be used to refer to the IronPort M160 device used for the centralised management of all the IronPort S670 devices.

Figure 2 describes the principle of OOB management, where separate physical infrastructure is used for configuring the firewall devices. Within the OOB network, a private address range is used that is not redistributed through the rest of the network, thus ensuring an additional level of security and access control for the firewall system. The management device has two ports: the P port, by way of which the configuration options are accessed through the production part of the network; and the M port, the management port, which is isolated within the OOB part of the network and through which the desired configuration is applied to all the firewall devices.

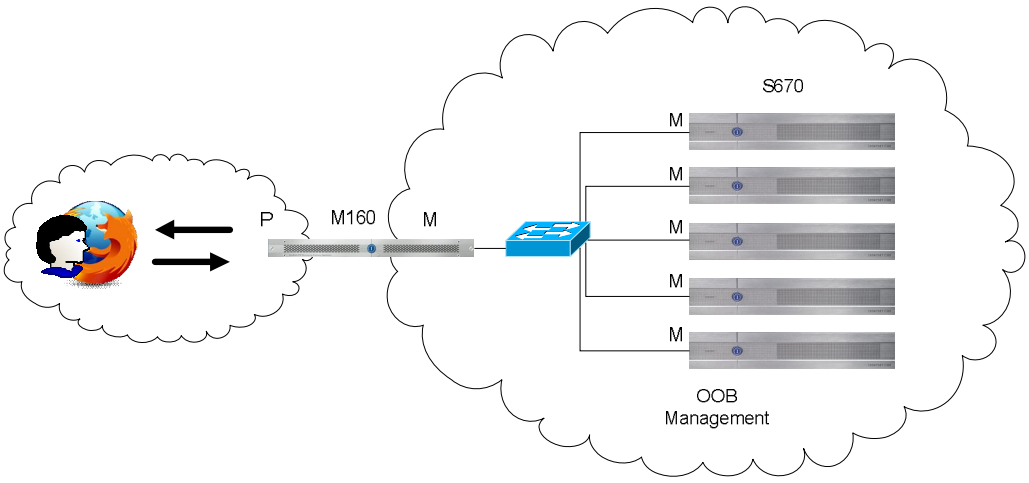

Figure 2 - OOB management by way of the web browser

3 The Position of the Firewall System

The web traffic filtering system is positioned in the core part of the network. Figure 3 illustrates the position of the system in relation to the other devices in the network.

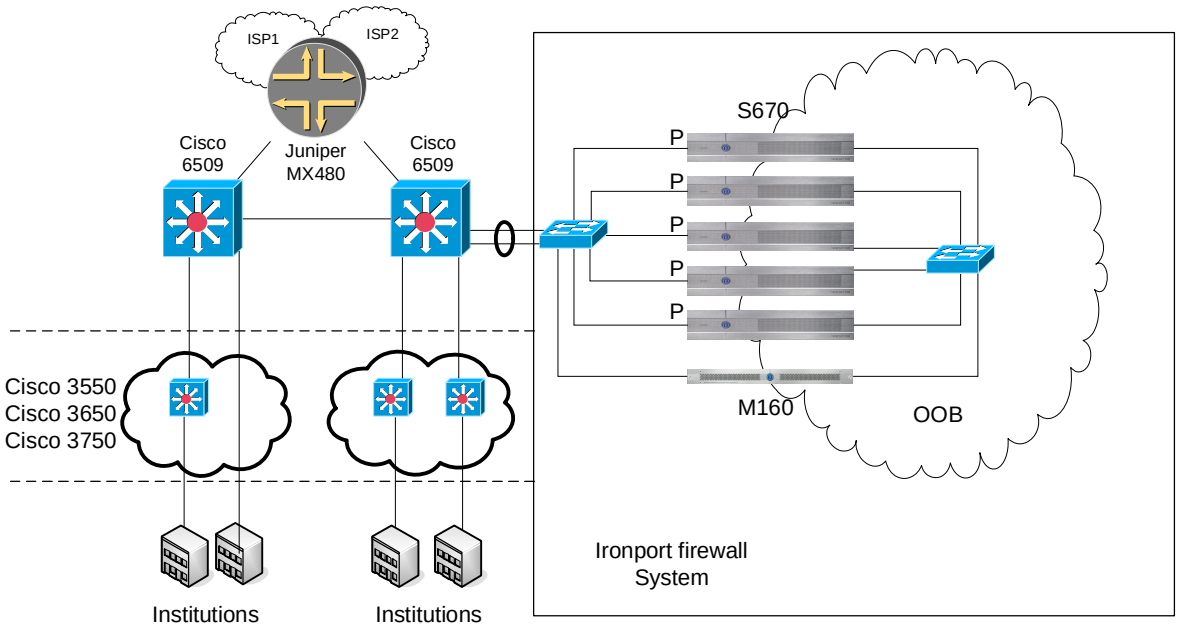

Figure 3 - The position of the firewall system

The firewall devices behave like proxy servers, i.e., all the generated web traffic must first reach the firewall devices, where it is filtered and scanned, and then the processed web traffic is forwarded to its final destination. The returning web traffic is also scanned in accordance with a number of defined parameters. As all the outgoing and incoming traffic must pass through the firewall system, the traffic is duplicated at the links, i.e., the amounts of incoming and outgoing traffic at the firewall devices' ports are approximately the same. An example of statistics of this duplicated incoming and outgoing production web traffic for a period of twenty-four hours is provided in Figure 4, while Figure 5 illustrates an example of such a situation.

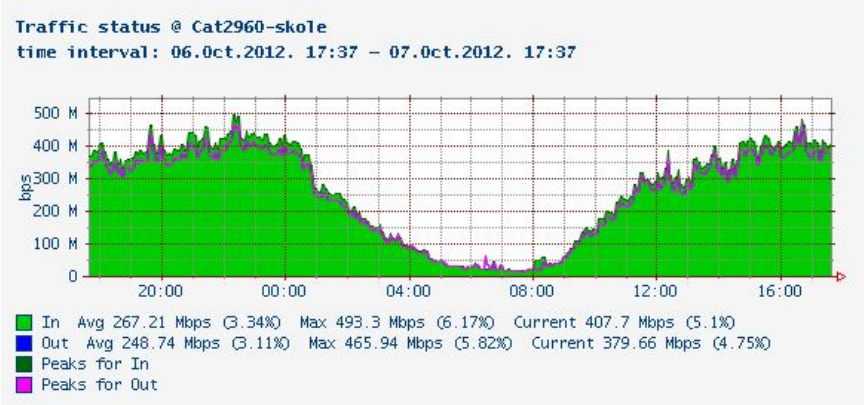

Figure 4 - An example of incoming and outgoing traffic statistics at the main link to the firewall system

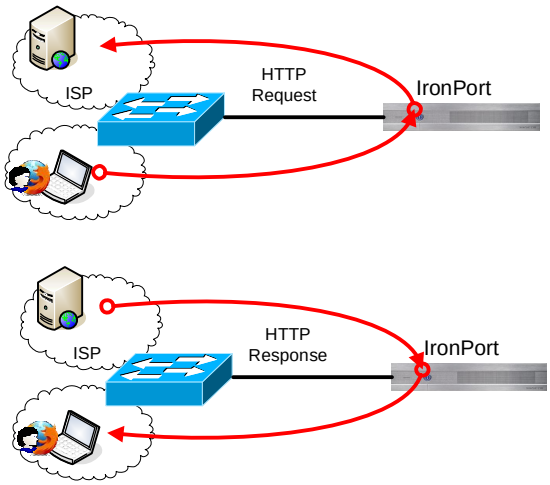

Figure 5 - An example of the duplication of incoming and outgoing traffic

Due to the duplication of traffic, it is necessary to pay attention to the links connecting the firewall system to the rest of the infrastructure. The use of link aggregation is recommended (e.g., LACP protocol) in order to avoid congestion at the links, between the core part of the network and the firewall system.

It is also recommended to position the firewall system inside the core part of the network, since all traffic directed to the Internet from the access part of the network is aggregated here. Thus, this positioning of the firewall system would enable the optimal utilisation of the network resources.

4 Web Traffic Redirection

4.1 Methods of Redirecting Web Traffic

In order to ensure the optimum utilisation of the firewall system resources, it is necessary to evenly redirect the traffic coming from the end users to each of the firewall devices. Below is a series of examples and recommendations as to how this can be achieved without any specialised equipment.

It is possible to redirect the web traffic using one of the following scenarios:

- web traffic redirection based on web browser proxy configuration;
- web traffic redirection by means of policy-based routing;
- web traffic redirection using the WCCP protocol (Web Cache Communication Protocol).

In each of the three possible web redirection techniques, the emphasis is placed on the possibility of the even distribution of web traffic to several firewall devices.

4.2 Configuration within the Internet Browser

Configuration within the Internet browser is performed by setting the appropriate options of the proxy configuration of the browser itself. There are two types of configuration: static configuration, and dynamic auto-configuration.

4.2.1 Static Configuration

4.2.1.1 *Manual Configuration by Entering the IP Address or DNS Name of the Firewall Device*

The traffic redirection is implemented based on the information on the IP address or DNS name of the firewall device, which the end user must enter in the proxy settings of his/her Internet browser. It is recommended to use the DNS name for the following reasons.

- A change of IP address of the firewall device is transparent to the user. In the case of changing the IP address, the user does not need to set the Internet browser again, since the DNS service will always return the right/modified IP address of the firewall device.
- Where several firewall devices are used, it is recommended to set the same DNS name for all the IP addresses of the firewall devices, and to configure the DNS service to resolve the DNS name of firewall devices in a round-robin fashion. In this way, the user is able to evenly use the IP addresses of all the firewall devices.

Figure 6 shows an example of the distribution of production web traffic using the round-robin technique within the AMRES network. It displays the number of established connections to the five different firewall devices. In Figure 6, it can be seen that the collected values regarding the number of established connections are not the same on each firewall device, which is expected, bearing in mind that some AMRES users still use the IP address in the proxy configuration of their web browser instead of the DNS name.

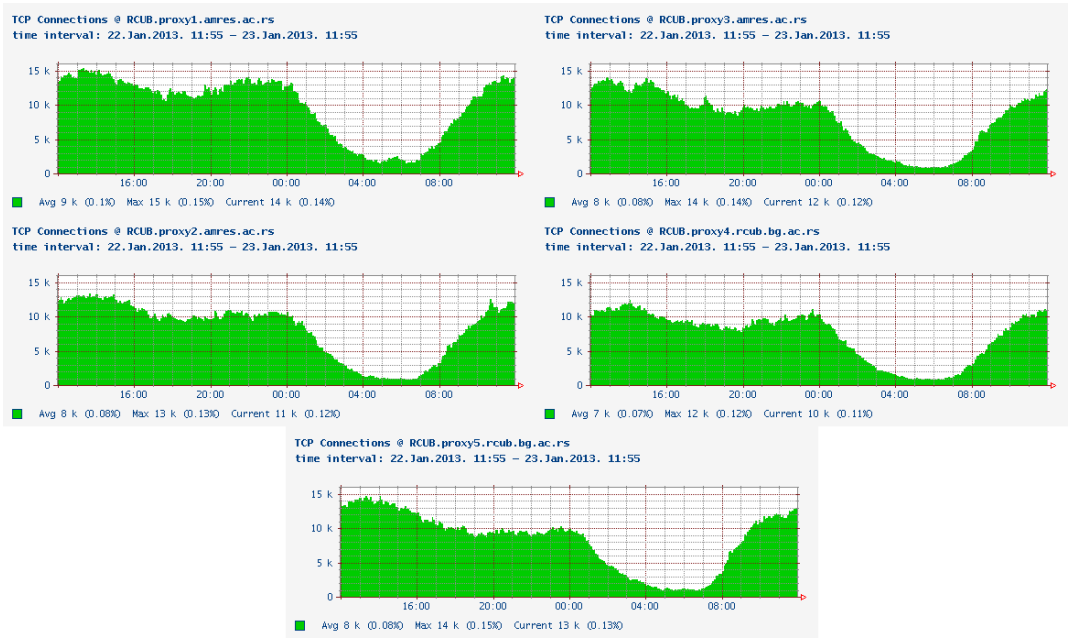

Figure 6 - The distribution of web traffic by means of the round robin DNS resolution technique

4.2.1.2 *Manual Configuration by Entering the PAC File Location*

In this solution, the user must enter the URL path to the PAC (*Proxy Auto Configuration*) file in the proxy settings of his/her Internet browser. The PAC file is written in Java Script and contains the rule by which the Internet browser will conduct the forwarding of web traffic. Figure 7 illustrates an example of the contents of a PAC file. Sections 1 through 4 in Figure 7 define the way the Internet browser will act depending on what has been entered in its URL field.

- Section 1 dictates that the firewall system will not be used if the URL field contains only a name and no dots. In that case, the Internet browser will not use the firewall system because it is assumed that the URL entered is just the name of a device from the local domain in which the computer running the Internet browser is also located.

- Section 2 also checks whether the used URL is a part of the local domain, in which case the traffic is not redirected to the firewall system.

- Section 3 inquires whether the URL is resolved to one of the IP addresses from the local domain, in which case the traffic is not redirected to the firewall system.

- Section 4 resorts to proxying in fail-over mode. This means that if none of Sections 1 through 3 have been applied, the DNS name *proxy.mydomain.com:8080* will be used first (the DNS will resolve *proxy.mydomain.com* URL in a round-robin fashion). If the DNS server is unavailable, then the 172.16.0.1:8080 device will be used, and if this device is also unavailable, 172.16.0.2:8080 and 172.16.0.3:8080 will be used, respectively. If, however, none of the firewall devices is available, no proxying will be implemented and the Internet browser will attempt to connect directly to the Internet.

```
function FindProxyForURL(url, host) {

// Section 1. If URL has no dots in host name, send traffic direct.
        if (isPlainHostName(host))
                return "DIRECT";

// Section 2. If specific URL needs to bypass proxy, send traffic direct.
        if (shExpMatch(url,"*.myfirst.domain1.com*") ||
                shExpMatch(url,"*.mysecond.domain2.com*") ||
                shExpMatch(url,"*localhost*"))
                return "DIRECT";

// Section 3. If IP address is internal or hostname resolves to internal
IP, send direct.

        var resolved_ip = dnsResolve(host);
        if (isInNet(resolved_ip, "10.0.0.0", "255.0.0.0") ||
                isInNet(resolved_ip, "172.16.0.0", "255.255.0.0") ||
                isInNet(resolved_ip, "192.168.0.0", "255.255.0.0") ||
                isInNet(resolved_ip, "127.0.0.0",  "255.255.255.0"))
        return "DIRECT";

// Section 4. All other traffic uses below proxies, in fail-over order.
        return "PROXY proxy.mydomain.com:8080; PROXY 172.16.0.1:8080; PROXY
172.16.0.2:8080; PROXY 172.16.0.3:8080; DIRECT";
```

Figure 7 - An example of PAC file configuration

4.2.2 Dynamic Auto-configuration

Dynamic auto-configuration by means of the WPAD protocol (*Web Proxy Auto-Discovery Protocol*) is initiated in the proxy settings of the Internet browser by selecting the field for the auto-detection of the proxy server.

Below are examples of the options in the most widely used Internet browsers:

- Mozilla Firefox
    - “Autodetect proxy settings for this network”
- Google Chrome and Internet Explorer
    - “Automatically detect settings”

The WPAD protocol functions as follows: it checks the domain of the user’s computer and attempts to find the *wpad.dat* file in that domain. For example, assuming that the computer that is set to automatically detect the proxy server is in the *mydomain.example.com* domain and that the web location of the *wpad.dat* file is *wpad.example.com,* the WPAD protocol operates by taking the following steps:

- the browser will first attempt to find the wpad.dat file *at http://wpad.mydomain.example.com/wpad.dat*;
- once it establishes that there is no wpad.dat file at the above location, it will attempt to find the file in the parent domain, i.e., it will try to access the *http://wpad.example.com/wpad.dat* location where it will find the wpad.dat file.

In this case, the DNS name of the firewall device will also be used in the *wpad.dat* file, i.e., the DNS resolution in round-robin fashion will enable an even distribution of web traffic to the all firewall devices. The syntax of the WPAD file is identical to that of the PAC file (i.e., it is Java script).

4.3 Policy Based Routing Redirection

Traffic can be redirected by means of the Policy Based Routing (*PBR*), i.e., on the router platforms themselves, in the core part of the network. In order to use PBR, it is necessary to select the desired web traffic by access-list(s) and redirect it to the firewall devices. A problem with this approach is reflected in the fact that not all the firewall devices can be used optimally because the PBR routing does not enable an even distribution of traffic. The traffic can be redirected to one firewall device only. This solution requires no configuration on the user side, but it is necessary to run the firewall devices in transparent working mode. Figure 8 is an example of configuration on a Cisco 6500 device. All TCP traffic generated from network 10.20.30.0/24 towards any IP address with destination ports 80 or 443 will be redirected to the firewall device with next hop IP address 172.16.0.1. This principle makes sense if there is only one firewall device to which the web traffic should be redirected. It is not recommended to use PBR with multiple firewall devices.

```
ip access-list extended WEB
permit tcp 10.20.30.0 0.0.0.255 any eq 80
permit tcp 10.20.30.0 0.0.0.255 any eq 443
!
route-map REDIRECT2IRONPORT permit 10
match ip address WEB
set ip next-hop 172.16.0.1
!
interface GigabitEthernet0/0
ip policy route-map REDIRECT2IRONPORT
```

Figure 8 - An example of configuration on a Cisco 6500 device

4.4 WCCP Redirection

WCCP (*Web Cache Communication Protocol*) is a protocol developed by Cisco, and today it is supported by many other manufacturers of network and server equipment. Currently, there are two versions of the protocol – V1 and V2. The use of version 2 is recommended because it contains some new functionalities, such as support for services other than HTTP, the grouping of WCCP routers into clusters, MD5 authentication, and equal distribution of traffic. Unlike PBR routing, WCCP enables an equal distribution of traffic, which is why it is recommended in situations where there are several firewall devices and it is necessary to distribute the traffic between them equally. Depending on the topology of the network and the connections between the network devices, WCCP can be configured in two ways. It is recommended to connect the firewall devices to the core routers directly at level L2, in which case it is possible to forward web traffic using only forwarding plane of the device. Conversely, if the firewall devices are in some separate L3 segment, WCCP will use the GRE (*Generic Routing Encapsulation*) protocol in order to tunnel the redirected traffic.

When configuring the WCCP protocol, it is necessary to pay attention to the type of the device running WCCP. In Cisco series 6500 devices, the WCCP protocol only supports hardware forwarding of packets if the WCCP redirection at the device's interface is run in the *in* direction. If, however, the WCCP redirection is run in the *out* direction, the Cisco 6500 device will implement software-based forwarding (i.e., the device will use *Routing Engine*), which may increase the CPU load and severely affect the performance of the devices. In the case of using the WCCP protocol, all web traffic will be forwarded to the firewall devices and the users will not have to configure any settings in their Internet browsers. Figure 9 shows an example of the WCCP topology, while Figure 10 illustrates an example of the WCCP protocol configuration on a Cisco 6500 device. If several firewall devices were used, the WCCP process on Cisco 6500 would employ the load balancing technique in order to distribute the traffic as evenly as possible among them. Although the WCCP protocol has been developed by Cisco, it can also be used with other vendors, such as BlueCoat, as well as with open source solutions, such as the Squid proxy server.

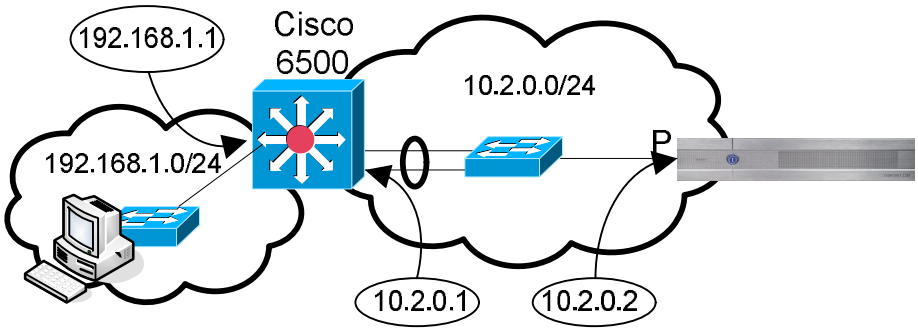

Figure 9 - An example of L2 topology on which the WCCP protocol is run

```
ip wccp 50 redirect-list WCCP-REDIRECT
!
!
interface GigabitEthernet0/1
ip address 192.168.1.1 255.255.255.0
ip wccp 50 redirect in
!
ip access-list extended WCCP-REDIRECT
permit tcp 192.168.1.0 0.0.0.255 any eq 80
permit tcp 192.168.1.0 0.0.0.255 any eq 443
```

Figure 10 - An example of the WCCP protocol configuration on a Cisco 6500 device

5 IronPort Cloud Service

The centralised web firewall system is devised to function as a cloud service. The AMRES staff manages and maintains the firewall system with the aim of providing a cloud firewall service to institutions and end users. The technical contacts of AMRES member institutions have strictly defined privileges in the centralised firewall system, so they can set filtering and scanning parameters that only affect the web traffic originating from their respective institutions. The cloud solution means that a member institution's technical contact can access the remote (centralised) hardware and software resources and adjust them to the needs of the users of his/her institution. The centralised firewall system is used by the all end users of AMRES as protection against malicious traffic on the Internet.

A technical contact of an institution is assigned the right to configure the access policy that conducts scanning and filtering of the traffic coming from the technical contact's home institution. If the institution does not want to administer its access policy, all traffic of that institution will be covered by the default access policy managed by the AMRES staff.

The entire web traffic passing through the firewall devices must be processed within one of the defined access policies, and each access policy relates to a certain address range. In this way, the access policies are tied to the traffic of the individual institutions. When scanning a web transaction, the firewall devices successively check the access policies searching for the one that corresponds to the address range of the web transaction. If the corresponding access policy is found, the web transaction is treated by the firewall parameters defined in the relevant access policy. The last in the sequence of access policies is the default access policy that will process the entire traffic that was not covered by the previous policies.

The ultimate default access policy defines the basic rules to be observed by the all institutions. Additionally, the institutions may make the firewall parameters stricter in their respective access policies, but they cannot relax them in any way. The firewall parameters of an access policy are grouped into the following five sections:

- allowed protocols and Internet browsers;
- URL filtering;
- application control;
- object control;
- web reputation.

5.1 Allowed Protocols and Internet Browsers

In this section, it is possible to block any of the following protocols:

- HTTP
- HTTPS
- FTP
- FTP via HTTP.

Additionally, it is possible to disallow the use of certain Internet browsers for accessing the Internet. If it is found that certain Internet browsers do not meet the security standards specified by the AMRES institution, the technical contact may prohibit its use by the users of the institution. The default access policy does not block any of the above protocols, nor does it prohibit the use of any Internet browsers.

5.2 URL Filtering

This section defines the set of URL categories that are permitted to be accessed via the Internet. Cisco conducts the categorisation of websites based on their content. The URL category information is updated periodically in the databases of the firewall devices on which the website categorisation is conducted. At the time of the preparation of this document, there were 78 different categories of websites, and Cisco periodically adds new categories. With regard to the IronPort URL filtering system, there is the possibility that individual sites are placed in the wrong categories. Therefore, AMRES has created a special category that explicitly allows access to all the sites that have been defined therein. If it is determined that a site has been inaccurately categorised and blocked, the site is placed into the special category of sites to which access is granted directly. Thus, access to the site is temporarily allowed and the error is subsequently reported to Cisco support. In this way, the problem of erroneously categorised sites is temporarily solved until the relevant global changes have been made in the Cisco category database, after which the site is removed from this special category.

The default access policy blocks the following web categories.

- Child Abuse Content
- Filter Avoidance
- Gambling
- Hate Speech
- Illegal Drugs
- Pornography

5.3 Application Control

In this section, it is possible to block the use of certain web applications. Cisco has created a database of the most commonly used web applications that may be blocked. The database is periodically updated to include newly created applications. These applications come mainly from the most popular Internet services, such as facebook, Google+, iTunes, and LinkedIn. The default access policy does not block any web application, while the technical contacts can block access to web applications for the users of their respective institutions if they find it necessary.

5.4 Object Control

The object control section can prevent the transfer of certain objects to the end clients, e.g., files with certain extensions or certain types of traffic, such as video or audio streaming. The technical contacts can also prevent flash or JavaScript functions in the users' Internet browsers. This section also allows the blocking of individual MIME (Multipurpose Internet Mail Extensions) types of data, which ensures additional granularity. It also enables the introduction of restrictions, depending on the size of the object that is being transferred. As a result, this section can prevent downloads of large files in order to save the institutions' uplink capacities. The default access policy does not block any object within the AMRES network, nor does it introduce any restrictions in relation to the size of objects that can be downloaded from the Internet.

5.5 Web Reputation

This section allows the definition of situations in which websites will be blocked and malicious software scanned or in which access to websites will be enabled without further scanning. The decision on these actions is based on the reputation of websites. Cisco has established a system of evaluation of websites based on suspicious and malicious activities, which includes a large number of websites. Websites are examined on a daily basis in order to ensure proper system coverage and fair scores. The score given to a website represents the extent to which the website has created problems for its users in terms of malware, phishing activity, spyware, viruses, or spam messages. The scores range from -10 to +10, where -10 is the worst and +10 the best score. When an end user tries to access a website, the firewall device will check its web reputation and if the score is satisfactory, the firewall will grant access to the website. It is important to determine the optimum threshold values for taking actions in order to make a compromise between the burden of the firewall device and the possibility of detecting malicious websites. Figure 11 shows an example of the default behaviour for AMRES users.

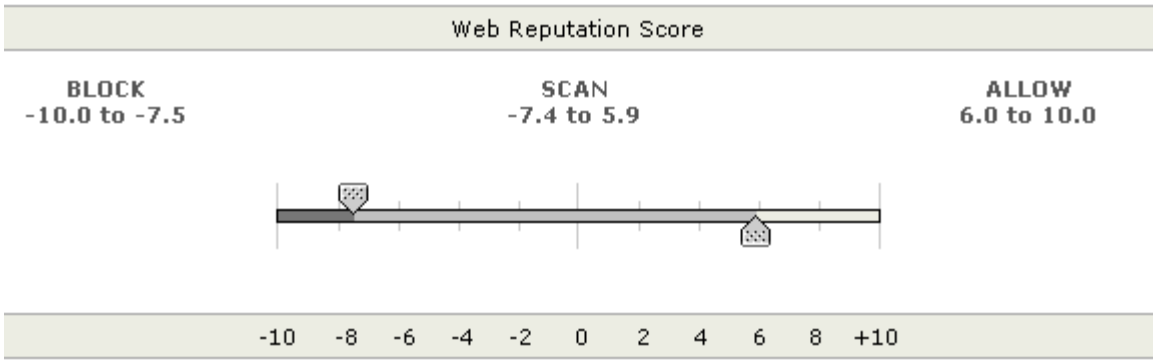

Figure 11 - Filtering and scanning based on the web reputation

Figure 11 shows that the access to a website that has a Web Reputation Score that is less than -7.4 will be blocked. If the Web Reputation Score of a website is between -7.4 and 5.9, the access to the website will only be granted after scanning is performed using the Webroot anti-malware scanning software. If the score of a website is higher than +6, access to the website will be granted without scanning.

6 Firewall Configuration

AMRES has around 100,000 end users, which makes the processing of web traffic filtering very demanding. The number of AMRES network users requires the use of several firewall devices to evenly distribute the processing of web traffic. In order to ensure the proper functioning of the centralised firewall system, it is necessary to maintain a consistent configuration on all the firewall devices. Having purchased the IronPort M160 device for centralised management and the corresponding licence, the AMRES is able to set uniform configurations on all five of its firewall devices using this centralised management device. This adds significantly to the scalability of the system since new firewall devices can be added with the increase in the number of end users and traffic flow, so that the system functions efficiently.

The parameters of the firewall system are configured through the centralised management device. The centralised management device consists of two network interfaces: the P interface with a public IP address and the M interface with a private IP address (Figure 2). By entering the corresponding URL address in his/her Internet browser, the technical contact will be connected to the P interface of the management device. The Internet browser will display a page for registering with the management device. After entering the technical contact's credentials, authentication and authorisation will be performed in the LDAP database. The technical contact will then be granted the right to modify the policy, concerning only the traffic coming from his/her institution.

After configuring the desired parameters and selecting the appropriate commands, the created configuration will be forwarded to all the firewall devices managed by the centralised management device. The configuration will be forwarded via the M interface of the management device, and via the OOB network, it will reach all the firewall devices in the system. This ensures a consistent configuration on all the firewall devices in the centralised web firewall system.

The AMRES administrators also connect to the P interface of the management device, but they have privileges allowing them to edit all the parameters of the firewall system. In addition, the AMRES staff can connect to other firewall devices through their M interfaces. Certain configuration parameters, such as login data collection or software upgrades, need to be set on each firewall device separately and are configured exclusively by the AMRES staff.

7 LDAP Authentication

The implementation of the cloud service in the AMRES network requires a solution for the authentication and authorisation of the member institutions' technical contacts. The optimal solution should take into account the following preconditions:

- one member institution may have several technical contacts who can edit firewall parameters;
- the technical contacts already have user accounts opened for the purposes of another AMRES service;
- through authentication on the centralised web firewall system, the technical contacts will acquire the right to edit firewall parameters affecting only the traffic of their home institution.

Authentication of the technical contacts is performed on the management device. Through the authorisation, the institutions' technical contacts acquire the right to modify the firewall configuration of the management device. The modified firewall configuration is then forwarded to all the firewall devices in the centralised web firewall system through the OOB part of the network. Possible authentication procedures on the management device are as follows:

- authentication using the local database (internal authentication);
- authentication using the RADIUS protocol (external authentication);
- authentication using the LDAP protocol (external authentication).

Taking into account all the above preconditions, the existing authentication solutions in the AMRES network, and the available procedures for authentication on the management device, authentication using the LDAP protocol stands out as the optimal solution.

For the purposes of authentication in the AMRES network, AMRES maintains an LDAP directory storing the data on the technical staff of the AMRES member institutions. A separate LDAP directory contains a branch with the AMRES service accounts. The LDAP directory that contains the AMRES service accounts also contains the administration account for the firewall system. By using this account, the management device connects to the LDAP directory and verifies the information about a user trying to use the firewall service. The branch storing the technical contacts contains all accounts of the member institutions' technical contacts. Each account makes a set of LDAP attributes, where the main attributes for the firewall system are:

- uid – the technical contact's username;
- password – the technical contact's password;
- organisation – the name of the technical contact's home institution;
- eduPersonEntitlement – the attribute that defines whether the technical contact is a user of a specific AMRES service.

Each administrator of the firewall system is uniquely defined by the username, the password, and his/her privileges. The username and the password are stored in the LDAP directory, or more precisely, in the branch containing the technical contacts, while the set of privileges is stored on the centralised management device. There are several predefined sets of privileges (User Roles) on the management device, although it is possible to establish new sets of special privileges (Custom User Roles). By creating a new set of privileges on the management device, a policy can be established that will affect only the web traffic of a specific AMRES member. When a member institution is added to the firewall cloud service, a new set of privileges (Custom User Role) bearing the name of the relevant institution will be created. In this new Custom User Role, it is specified that only the policy concerning the new institution can be changed. The linking of the account of the technical contact from the LDAP database to the User Role on the centralised management system is performed using the following two attributes within the user accounts in the LDAP directory:

- EduPersonEntitlement;
- Organisation.

Figure 12 shows an example of a technical contact's account in the LDAP directory. If the institution's technical contact is authorised to use the firewall cloud service, his/her account in the LDAP directory will contain the attribute *eduPersonEntitlement*, which has the value "Ironport". This attribute allows the authorisation of the technical contact on the centralised management device. The attribute *Organisation* contains the name of the technical contact's home institution and is linked to the corresponding set of privileges of the authenticated technical contact.

```
uid:                          john_doe
Password:                     ********
EduPersonEntitlement:         Ironport
Organization:                 Institution_Name
```

Figure 12 - An example of a technical contact's account in the LDAP directory

The process of authentication of the technical contact is performed in several steps. The technical contact connects to the centralised management device through the P interface (Figure 2) and enters his/her username and password. The centralised management device connects to the LDAP server through port 389 and verifies the user's credentials using a special AMRES service LDAP account. Following the authentication with the AMRES service account, the management device acquires the right to view the branch containing the technical contact's attributes in order to perform their authentication and authorisation. In the branch containing the technical contact's attributes, a check of the user trying to authenticate is performed. If the user exists, his credentials are verified. After successful end-user authentication, the centralised management device checks whether the technical contact has the "eduPersonEntitlement" attribute with the appropriate "IronPort" value. If this attribute exists, the centralised management device can authenticate the technical contact. The verification of the "Organisation" attribute is then performed, which is linked to the Custom User Role bearing the same name on the centralised management device itself. If the appropriate User Role exists, the centralised management device will grant the technical contact the appropriate User Role and authorise him/her. Following the authorisation, the technical contact can modify the configuration on the centralised management device, but only in the part corresponding to the User Role granted.

8 Collecting, Analysing and Storing Logs

Each firewall device registers web connections and collects logs individually, which is why the configuration related to log data collection needs to be set individually on each device. Log data are grouped into log files depending on the type of information registered. Logs can roughly be divided into two groups.

- logs regarding the status, processes and the operation of the device – System Logs
- logs regarding the activity of the end users – Access Logs

The logs regarding the status, processes, and operation of the device record messages about the operation of individual components of the firewall device. There is a large number of these logs and they are useful in situations when certain problems, concerning the operation of the device, need to be examined. The AMRES practice is to store the System Logs on the firewall device itself. The System Logs can be accessed through the web interface of the firewall device. The logs regarding the status and operation of the device are configured to contain information for one whole day. Ten files of each System Log are kept on the device, i.e., the AMRES staff can view log information about the operation of the device components for the last ten days. The System Logs are subsequently deleted so that they do not overload the memory of the device. The practice applied so far has shown that there is no need to store these logs for a longer period of time since they are viewed when a problem arises, and problems are dealt with as soon as possible.

The logs regarding the activity of the end users are recording every web transaction between the AMRES end users and the Internet. These logs are called Access Logs and are based on a form similar to the Squid logs. The Access Logs store the information on a web transaction up to OSI layer 7, which is very useful if forensic analysis is required. These logs record the following information on the web transaction.

- the exact time of the transaction (Unix time)
- the duration of the transaction
- the end user's IP address
- the result of the inquiry into the cache memory of the device
- the response of the remote server
- the size of the transferred data
- the HTTP method
- the URL requested by the end user
- the MIME type
- the user's home institution/network
- the access policy on the firewall device that processed the transaction

- the web reputation of the remote server
- information on any malware
- information on whether the transaction was successful and if not, the reason for the blockade

Access Logs can also be adjusted to meet the needs of the organisation by adding the desired fields at the end of the log entry. AMRES adds two fields to its Access Logs: the “X Forwarded For” field, and the exact time of the transaction in HH:MM:SS format. The “X Forwarded For” information is convenient since certain member institutions have their proxy servers so this field enables the detection of clients hiding behind the proxy servers of their home institutions. The exact time of the transaction in HH:MM:SS format is convenient in situations requiring the analysis of a log file since it is easier to follow the transactions in this way than by using Unix epoch time.

Since all web transactions passing through the device are recorded, the size of the log files is quite large, so their storage on the firewall device itself is not practical. In the AMRES network, Access Logs are transferred from the firewall device to a special server for processing and then to a remote server for storing. A log file can contain information on an arbitrary time interval – half an hour, one or more hours, one or several days, etc. The AMRES network’s practice is that one Access Log file contains log information for one hour. The log-creation process has several steps. At the beginning of each hour, the firewall device creates a new log file and begins recording the web transaction information. At the end of the hour, the firewall device closes the file and transfers it to the processing server. The transfer to the processing server can be done using one of the following methods.

- SCP (Secure Copy Protocol)
- FTP (File Transfer Protocol)
- Syslog push method

The log file is transferred through the OOB network, which ensures the security of the entire process. Initially, the log files were transferred using the SCP method. Occasionally, it would happen that the firewall device and the log-processing server were not able to establish the SCP connection, and for this reason, this type of transfer was eventually abandoned. The FTP transfer method is now used.

There are two problems with regard to the transfer and handling of these logs. The first problem is that each firewall device creates its own log file for the previous hour. During one hour of operation of the entire firewall system, five separate log files are transferred to the log-processing server (Sawmill), i.e., one log file from each firewall device. The review, analysis, and storage of the log information for each hour of operation of the firewall system are therefore very inefficient. For this reason, the server that stores log files every hour runs a script that merges all five of the log files into one log file containing the information on all the transactions that have passed through the firewall system during the last hour. Figure 13 below shows the process of storing Access Logs.

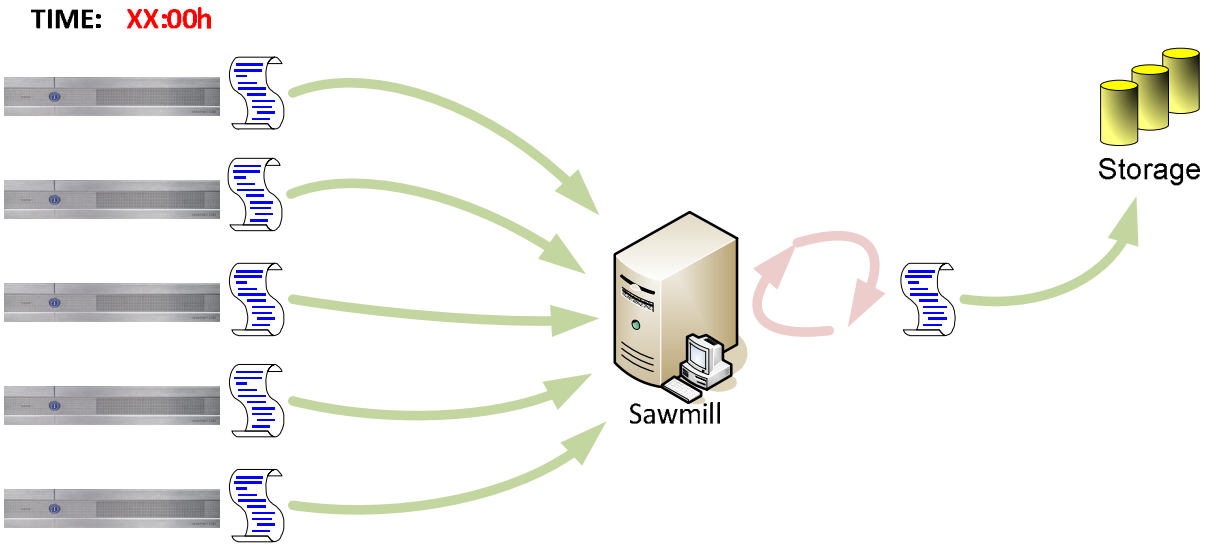

Figure 13 - Collecting, analysing, and storing logs in the AMRES network

A unique one-hour log file is given the name in the time format *YYYY-MM-DD-HH*, where *HH* is the hour in the 24-hour format when the file was created. All transactions within the one-hour log file are sorted by the time of their occurrence, which facilitates the reviewing and analysis of the log information. The size of the unique one-hour Access Log file varies between 200MB to 3GB depending on the time of day. The storage of log files in their original form is inefficient from the point of view of the available storage memory. For this reason, all log files are compressed and stored on the remote storage server. The compressed Access Log files are reduced to a size of 50 to 500 MB, which saves the memory space significantly. One month of log information on the end user activity in the compressed format takes up 100 to 200 GB of memory, while annually around 1.5 TB is required for storing log files. If it is required to analyse the web traffic from several months ago, the one-hour log files of interest can be found easily by reviewing the compressed log files.

The other problem with the collection and storage of logs is related to the system of creating the logs on the firewall devices. The creation of a log file and the beginning of recording of the log information in the file has to be started manually on the firewall device itself using a mouse click selection. The option defining the time after which the existing file will be closed and a new one opened can be selected (thirty minutes, one or more hours). After that, the process is automated, i.e., after the defined time interval, the existing file will be closed and sent to the server, and a new log file will be created and will begin recording log information. Therefore, the first definition of the log configuration is very important. The problem arises with the synchronisation of the creation of log files on all five firewall devices, as it is practically impossible to manually set (mouse click selection) the Access Log configuration on all firewall devices at the same time. A more flexible solution would be to have an option in the log configuration enabling the creation of a new log file at an exact time, e.g., at the zero minute of every hour, similar to the CRON function in Linux. Unfortunately, the firewall devices do not have this option so the Access Log configuration needs to be set in parallel on all firewall devices at the desired time. Special attention should also be paid to the fact that any change in the Access Log configuration will lead to automatic creation of a new file. For this reason, AMRES takes care that all changes in the Access Log configuration are performed simultaneously on all firewall devices at the zero minute of an hour.

For the purposes of traffic analysis and web statistics, the Academic Network purchased the *Sawmill for IronPort 7.3.3* application, which is located on the log file processing server. All one-hour Access Logs are kept on the server during the day. When the last log file arrives at midnight, the Sawmill application begins reading the log files from the monitored day and creates its database of web transactions. After the Sawmill application has read all the files, they are compressed and stored on the remote storage server where they are kept for twelve months. The Sawmill application database contains parsed log information from the Access Log files of the firewall devices. Based on the information from the database, Sawmill creates several parsed reports providing a general picture of the amount of web traffic, trends in the institutions' networks, patterns of behaviour of end users, and potential security threats to which the end users are exposed. These reports can help locate computers infected by malware or web servers posing a security threat to the AMRES users. The Sawmill application also enables the reviewing of its database, and can handle queries concerning any transaction of interest. The Sawmill application database stores the information on the activity of the end users in the last two weeks and also allows the possibility of recording and storing predefined reports in the PDF format, which later can be used for the preparation of reports on the operation of the service and the network.

Based on the statistical analysis of the log information, conclusions can be reached on the effects of the set firewall parameters.

9 Monitoring of the IronPort Firewall System

As most of the AMRES network web traffic passes through the IronPort firewall system, it is very important that this system functions with a high level of reliability. The real-time monitoring of the working parameters of the firewall devices is one of the most important tasks of the AMRES staff. The web interface of each firewall device contains the “Reporting” section, which provides the working parameters of the device, and statistical data on the users, websites, malware, etc. This section also shows data on the CPU load of the device, the RAM memory usage, and the free memory space for log file storage. It also shows the average network flow, the average response time of a device, and the total number of current connections on the device. The active monitoring of the operation of the entire firewall system requires the AMRES staff to connect to each individual device and to simultaneously monitor the Reporting sections. This solution is impractical since it does not provide an insight into the working parameters of all the firewall devices in one place. Therefore, AMRES has decided to integrate the monitoring of the operation of the firewall system into its own system for the monitoring and supervision of network devices – NetIIS.

NetIIS is a computer network monitoring system developed in the Computer Centre of the University of Belgrade. NetIIS actively monitors the operation of elements of the computer network, such as routers, switches, severs, and connections between devices. The system is based on the SNMP (Simple Network Management Protocol) protocol so that it can collect data on the operation of the network devices and the status of the links in the network. The AMRES staff actively monitoring the performance of entire Academic Network, as well as the performance of individual services via NetIIS. Each service in the Academic Network has a special section in NetIIS, which enables the real time monitoring of their performance. The integration of the monitoring firewall system into NetIIS enables an overview of all the working parameters of the firewall devices in one place. Practically, the AMRES staff can monitor the performance of all the firewall devices in real time on one web page and promptly respond if the working parameters of the firewall system are not satisfactory.

The IronPort firewall devices support SNMP protocol versions 1, 2, and 3, and for the purposes of monitoring within AMRES, SNMP version 2 is used. NetIIS itself generates alarms based on the results of the SNMP requests and the optimal threshold values set in NetIIS.

Figure 14 shows a section in the NetIIS system that actively monitors the parameters of a firewall device. The following working parameters are monitored for each firewall device:

- the status and current network flow on the production P interface;
- the status of the current network flow on the management M interface;
- the CPU load and the usage of the RAM memory on the device;
- the average number of TCP connections in the last minute;
- the average number of active TCP connections in the last minute;
- the average number of passive TCP connections in the last minute;
- proper operation of traffic proxying and the response time;
- ping monitor-packet loss and delay.

proxy2.amres.ac.rs	
P&T@Cat2960-skole.Gi0/26 [IW2-D1]	Up / Up , 184.14 Mbps / 165.90 Mbps
P&T@proxy2.amres.ac.rs.P1 [P1]	Up / Up , 155.15 Mbps / 171.84 Mbps
P&T@proxy2.amres.ac.rs.Management [Management]	Up / Up , 1.63 Kbps / 11.33 Kbps
CPU & Mem Monitor@proxy2.amres.ac.rs	16 % / 79 %
TCP Connections@proxy2.amres.ac.rs	11856
TCP Active Opens@proxy2.amres.ac.rs	111.01
TCP Passive Opens@proxy2.amres.ac.rs	267.97
facebook.com monitor@proxy2.amres.ac.rs	HTTP OK: HTTP/1.1 302 Found - 322 bytes in 0.207 second response time / time=0.206528s;;;0.000000 size=322B;;;0 / 0 (OK) / 1 / 0.206528
Ping monitor@proxy2.amres.ac.rs	0.33 ms / 0.49 ms / 0.37 ms / 10 / 10 / 0 %

Figure 14 - Working parameters of a firewall device

For the purposes of monitoring the firewall devices, AMRES only uses two OID identifiers taken over from the official Cisco MIB model:

- CPU load of the device (IronPort CPU);
- RAM memory usage (IronPort Memory).

These OID identifiers are placed in one SNMP monitor – the CPU & Mem Monitor. The TCP OID identifiers are taken over from the standard TCP MIB model, which are supported by most of the vendors present on the market.

The verification of the proper functioning of the firewall devices is crucial in system monitoring. The website monitor checks the traffic proxying function on the firewall device (Figure 15). The website monitor in the NetIIS system is implemented with the Nagios plug-in – “check_http”. The proxying towards at least two websites on the Internet is examined on each firewall device. Checking various websites enables the AMRES staff to determine if a problem in proxying is occurring on a firewall device and not on the website that is being checked. It is essential to check whether the end users can access the websites on the Internet through the firewall system and to determine the system response time. Figure 15 shows a check of the proxying function of a firewall device, where the *facebook.com* website can be accessed successfully, and where the response time is 0.207 seconds, which is an acceptable result. If the proxying were unsuccessful, that would mean that there is probably a problem with the functioning of the firewall device or the AMRES network. The response time is carefully monitored, since it can point to delays in the AMRES network or to problems concerning the proxy functionality of the firewall device.

facebook.com monitor@proxy2.amres.ac.rs	HTTP OK: HTTP/1.1 302 Found - 322 bytes in 0.207 second response time / time=0.206528s;;;0.000000 size=322B;;;0 / 0 (OK) / 1 / 0.206528

Figure 15 - The website monitor that checks the proxy functionality of the firewall device

In addition to actively monitoring the firewall devices, AMRES also performs passive monitoring. Firewall devices can send alarm messages in case any component of the system does not function properly. These alarm messages are actually e-mail messages containing notifications on the performance of specific functions of the firewall devices. The alarm messages may be:

- system messages – transmitting information on the functionality of the firewall device itself;
- hardware messages – transmitting information on the performance of hardware components;
- software update messages – transmitting information on the update of certain software parts;
- web proxying messages – transmitting information on the proxy functionality of the device.

In addition to various types of messages, each message is marked by a criticality measure. There are messages containing only information on an event, messages containing warnings concerning the operation of the device, and messages pointing to critical events threatening the functionality of the device. In order to properly monitor the operation of the firewall system, a mailing list has been established to which the firewall

devices send all alarm messages. Members of this mailing list include the AMRES engineers in charge of the maintenance of the firewall system. Through these alarm messages, the AMRES staff is notified about any problem on a device the moment it occurs. Previous experience in the operation of the firewall devices shows that these messages can be extremely efficient in the event of the failure of a hardware component or in the event of an unsuccessful transfer of the log information to the remote server. Thanks to timely notifications on the problem, the AMRES staff has been able to react quickly and restore the full functionality of the firewall system.

10 Conclusion

Although the IronPort firewall security system is designed to meet a large number of the requirements of closed organisations, such as banks and companies whose business requires a high level of security, certain components of the system can be used effectively in other organisations where the level of security of the organisation is not so important, but where the focus is placed on the protection of individual end users. Primary and secondary schools, faculties, and libraries are examples of such organisations. In addition to innovations, current trends in technological development come with certain bad aspects, such as identity theft, spreading destructive viruses, and the abuse of the anonymity provided by the Internet. A large number of parameters suggest that more attention should be paid to the security of Internet users, especially among the younger population. This security system can be widely used in the protection of the end users of academic institutions and it is expected that its use will increase significantly in the near future.

Glossary

AMRES	Acedemic Network of Serbia
FTP	File Transfer Protocol
LACP	Link Aggregation Control Protocol
LDAP	Lightweight Directory Access Protocol
MIB	Management Information Base
MIME	Multipurpose Internet Mail Extensions
NetIIS	Network Informational System
OID	Object identifier
OOB	Out Of Band
PAC	Proxy Auto-Config
PBR	Policy Based Routing
RADIUS	Remote Authentication Dial In User Service
SCP	Secure Copy
SNMP	Simple Network Management Protocol
SSH	Secure Shell
WCCP	Web Cache Communication Protocol
WPAD	Web Proxy Autodiscovery Protocol

Complete BPDs available at www.terena.org/campus-bp/

campus-bp-announcements@terena.org