

A large, stylized map of Europe is the central background element. It is composed of a grid of small squares in various shades of yellow and light green, creating a pixelated or mosaic effect. The map covers most of the page's width and height.

Network Monitoring and Management Recommendations

Best Practice Document

Produced by AMRES led working group
on network monitoring
(AMRES BPD 101)

Authors: Esad Saitović and Ivan Ivanović
February 2011

© TERENA 2010. All rights reserved.

Document No: GN3-NA3-T4-AMRES-BPD-101
Version / date: February 2011
Source language: Serbian
Original title: "Preporuke za monitoring i menadžment mreže"
Original version / date: Revision 1 (of the document dated 24 October 2009) / 2 February 2011
Contact: esad.saitovic@rcub.bg.ac.rs, ivan.ivanovic@rcub.bg.ac.rs

AMRES/RCUB is responsible for the contents of this document. The document was developed by AMRES led working group on network monitoring 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.

Parts of this document may be freely copied, unaltered, provided that their original source is acknowledged and copyright preserved.

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	5	
Introduction	6	
1	Preparing the Network for NMS Implementation	7
1.1	In-Band vs. Out-of-Band	7
1.1.1	In-Band Environment	7
1.1.2	Out-of-Band Environment	8
1.2	Logical Network Segmentation in an In-Band Management Environment	8
2	Interfaces for Accessing Devices	9
2.1	Network Devices	9
2.2	Servers	10
2.3	Other Devices	11
3	Access to the Management Network	12
3.1	IP Addressing	12
3.2	Access by Administrators to the Management Parts of the Network	12
3.3	Mutual Isolation between Devices in the Management Network	13
3.4	Examples of Topologies	13
4	Logical Access to Devices (Device Access Protocols)	16
4.1	Protocols for Controlling and Configuring Devices	16
4.2	Access to Network Devices	17
4.3	Access to Servers	18
4.4	Access to Other Devices	19
4.5	Device Monitoring Protocols	19
4.5.1	SNMP v2c	20
4.5.2	SNMP v3	21
5	Configuration Maintenance	22
5.1	Configuration Backup	22
6	NMS Server	23
6.1	The Position of the NMS Server in the Network	23
6.2	The Recommended Version of SNMP on Network Devices and Servers	23
6.3	Recommended Variables to Monitor	24
6.3.1	Network Devices	24
6.3.2	Servers	24
6.3.3	UPSs	25
6.4	MIB Variables	25

6.4.1	Standard MIB Variables	25
6.4.2	Proprietary MIB Variables	25
6.5	Trap Mode	26
6.6	Examples of SNMP Configurations on Devices	26
6.6.1	Cisco Router	26
6.6.2	Linux Server	28
7	Saving Syslogs	31
7.1	SysLog Protocol	31
7.2	Location of the SysLog Server	32
7.3	Installation	33
7.3.1	Network Devices	33
7.3.2	Servers	34
8	Traffic Analysis Protocol	36
8.1	NetFlow Protocol	36
8.2	How the System Operates	37
8.3	The Location of Collectors in the Network	37
8.4	Configuring the NetFlow Exporter	38
8.5	Indirect Solutions for Collecting NetFlow Statistics	39
References	42	
Glossary	43	

Summary

The purpose of this document is to provide an insight into basic NMS (Network Management System) activities, along with recommendations for administrators of campus and/or local networks intending to apply the NMS tools within their networks.

The document begins with a discussion of network topology. Changes in topology are recommended based on the idea that a majority of NMS activities should be conducted through the management segment of the network. The discussion focuses on two alternatives: the management network and the production network may either be separated physically (out-band management segment) or they may share the same physical infrastructure (the VLAN segment of the network).

The document further identifies the three components that should, as a minimum, be covered by a Network Management System. They are the configuration management and log management, along with the already recognised Network Monitoring component implemented using one of the NMS software packages.

Finally, the document briefly describes the most frequently used management protocols and their application in different environments and on different types of devices within a network (such as network devices, servers, UPS devices and A/C), provided they do not jeopardise the security of the network.

Introduction

Network management systems are nowadays among the most important elements of a successfully functioning computer network. The maintenance and configuration of network devices, servers and services, as well as continual monitoring of the operation of all the devices within the network, are the key elements of a network management system. In order to ensure the reliable and secure management of devices and services, it is necessary to design a network in such a manner that it provides the highest level of security isolation of management traffic from production traffic. Another aspect of the successful management of computer networks concerns the network protocols used for this purpose, as well as their implementation, i.e., the manner in which they are used. The first part of this document deals with the design of the management part of the network (out-of-band and in-band), and the recommendations for the use of device access protocols, including the methods of backing up configurations. The second part of the document describes the monitoring protocols within computer networks and the traffic analysis protocol (NetFlow), as well as their implementation and manner of use.

1 Preparing the Network for NMS Implementation

1.1 In-Band vs. Out-of-Band

In-band management involves the simultaneous use of interfaces and network equipment for production traffic and management purposes.

Out-of-band management (OOBM) involves the use of a separate network infrastructure and a separate interface for management purposes, different from the network devices and interfaces used for production traffic.

Out-of-band management is recommended for campus networks or parts of campus networks where the network equipment and servers are in a single room (machine room / server room / single hub).

1.1.1 In-Band Environment

Advantages:

- It does not require additional physical infrastructure (such as network interfaces on servers, separate network devices and passive infrastructure).

Disadvantages:

- A lower level of security, bearing in mind that the sensitive content of management traffic passes through the same infrastructure as the production traffic, i.e., the traffic intended for end-users;
- In cases of congestion (during regular operation or due to a denial-of-service attack), it may be difficult (or even impossible) to access the devices in order to perform interventions that would help eliminate the problem.

1.1.2 Out-of-Band Environment

Advantages:

- A physically separate infrastructure provides more security for sensitive management information.
- Access is enabled even in case of problems with the production links (e.g., link disconnection, congestion, etc.).

Disadvantages:

- A separate network infrastructure involves expenses related to the procurement of equipment.
- It requires a higher initial engagement of administrator staff and implementation costs.

1.2 Logical Network Segmentation in an In-Band Management Environment

One of the key elements in network organisation is the logical network segmentation. This is achieved through defining groups and creating VLANs for each group. It is also necessary to define a VLAN for management purposes.

VLANs can be defined in the following manner:

- VLAN-MGMT – Management VLAN. Although this VLAN is usually VLAN 1, it is recommended to define a different VLAN for management purposes, a VLAN that will be used for management traffic only, in order to increase security.
- VLAN-SERVER-ENT – a VLAN for enterprise servers (DNS, proxy, e-mail, web ...).
- VLAN-SERVER-<workgroup> – a VLAN for workgroup servers (various application servers, database servers, etc.)
- VLAN-ADMIN – Administrator VLAN.
- VLAN-USER-<workgroup> – User VLANs depending on the workgroup.

2 Interfaces for Accessing Devices

It is necessary to consider the options available for accessing various types of devices, as well as recommendations for applying certain solutions depending on the network topology. Three possible groups of devices are defined as follows:

- Network devices – routers, layer-2 and layer-3 switches
- Servers
- Other devices – UPS, air conditioners, printers, etc.

2.1 Network Devices

Network devices can be accessed in one of the following ways:

- Console port – this port enables access to the Command Line Interface (CLI). This type of access does not require network communication but rather a direct connection between the serial (COM) port of the computer used to access the device and the Console port itself. The access enabled through the Console port is used for initial configuration of the devices, for updating software on the devices, for resetting passwords and when the device cannot be accessed via network interfaces. The access is enabled by terminal access software (Hyper Terminal, Putty, SecureCRT, etc.).
- AUX – this is a port that can be used for remote connection to the device using a dial-in connection. Because this document aims to formulate a set of recommendations for campus and local networks, there are no situations that are sufficiently serious to justify the implementation of this access solution.
- OOBM – some network equipment vendors build in out-of-band management ports that have an assigned IP address and enable access in the same way as enabling access to any other network interface, although the access is limited to supported management protocols (telnet, ssh, http, etc.).
- A separate Ethernet interface – when it comes to network devices, having a single network interface intended exclusively for management purposes is not common (and it can be said that it is not recommended), taking into account the fact that the price of an interface on a router is significant. For switches with a large number of ports, having a separate interface is a recommended solution, but only within data centres, i.e., unless the installation of additional passive infrastructure is required (the cost of OOBM, see Section 1.1.2).
- VLAN-MGMT access – the access to network devices by way of logical separation of traffic in a separate VLAN intended for management purposes is recommended since the implementation of this

solution requires only configuration of the existing equipment. The configuration requires the following for each kind of device:

Switches – all connections between switches should be in the mode for transferring multiple VLANs (generally the IEEE 802.1Q standard, such as the trunk mode in Cisco devices, or tagging mode in other vendors' devices). The management VLAN (VLAN-MGMT) needs to go through this link as well.

Routers – sub-interfaces with an IP address within the range defined for the management VLAN need to be defined on routers (also by using the IEEE 802.1Q standard).

When frame tagging is configured (802.1Q), Cisco devices (and perhaps some other vendors' devices) send preconfigured frames from VLAN1 as 802.3 frames, i.e., they send the frames untagged.

For other vendors' devices it is necessary to define which VLAN will be sent untagged.

In order to increase security, it is recommended that traffic for all VLANs is tagged, if that can be configured on the devices. If it is not possible to send all VLANs tagged on some of the devices, it is recommended that a VLAN is defined that is isolated from the rest of the network (a black hole) and that the traffic from this VLAN goes untagged.

2.2 Servers

For management purposes, access to servers can be enabled in one of the following ways:

- KVM switch – Servers are usually located in one place – in rack cabinets, on server room desks, etc. In order to make access to the servers easier, the KVM switch is usually used to enable a connection between the keyboard-video-mouse ports on all servers and a single physical set of these components (keyboard, video and mouse). The server can thus only be managed if the administrator is physically present in the server room where the KVM switch is located.
- OOBM – Recognised vendors incorporate out-of-band management ports in servers. This type of port represents another network interface and it has a special purpose.

Advantages:

This port has a separate controller on the motherboard and it enables TCP/IP access to the server independently of the operating system on the server. This enables remote review of the events on the server itself, which is the same as direct access to the server (keyboard-video-mouse), including the server booting process.

Remote access to the server BIOS settings.

Remote control – turning the server on and off, restarting the server.

Disadvantages:

The implementation of this solution varies depending on the vendor, i.e., there is no uniform solution.

In the basic (free of charge) version of server access via this port, the functionalities are limited. In order to get a wider set of functionalities, some vendors require a payment for licences.

- Network interface – it is recommended that servers have at least two network interfaces, which is usually the case with the servers that are currently on the market. On the other hand, it is not

uncommon to use higher-performance client computers with a single network interface for server functions.

Access to servers in these various circumstances is treated in the following manner:

- Servers with at least two network interfaces:

It is recommended that one of the server interfaces is configured as a management interface, i.e., it should be located in the management network (the OOB part of the network or the management VLAN). The other server interface (the rest of the server interfaces) should not be used for management purposes.

The other server interface (the rest of the server interfaces) should be used for production access to the services offered by the server.

- Servers that have a single network interface:

As it is impossible to physically separate the management traffic from the production traffic on these servers, two approaches are recommended depending on the hardware. Regardless of the approach used, it is recommended that access protocol encrypting traffic is used.

If the network card supports the IEEE 802.1Q standard, it is recommended that the logical interfaces assigned to various VLANs are defined on the card itself, including a logical interface assigned to the management VLAN.

If the network card does not support the IEEE 802.1Q standard, the management traffic cannot be separated from the production traffic by way of a logical interface. In this case, the server should not be linked to the management part of the network.

2.3 Other Devices

“Other devices” are defined as all devices whose primary purpose does not require network communication (e.g., uninterruptible power supply devices, air conditioners or humidity sensors). Management access is enabled through the following interfaces:

- Serial port – management through a serial port is enabled on the devices by way of specifically designed manufacturer’s software. Usually, if there is a network interface on the device, the serial port serves only for the purpose of initial configuration of IP parameters in order to go to the network interface access.
- Network interface – recognised vendors also offer network interface cards in their product portfolios, which are built in their devices for management purposes.

3 Access to the Management Network

In order to define the manner in which the management network will be accessed, the following need to be defined:

- IP addresses for the devices in the management network,
- methods of accessing the management network,
- the level of isolation between the devices in the management network.

3.1 IP Addressing

In order to enhance the security of the management parts of the network, it is recommended that separate IP address ranges are defined for these parts of the network. This applies to the out-of-band management network, as well as to the management VLAN segment. These ranges should not be routed to the rest of the network.

3.2 Access by Administrators to the Management Parts of the Network

It is necessary to define the methods of accessing the management parts of the network in different environments and situations.

We have defined the following three methods of access:

Accessing devices from the management part of the network:

- For this method of access, it is necessary to have a computer within the management part of the network whose only network interface is connected to the management network itself.
- Access to this computer is only allowed from within the management part of the network.

Accessing devices from the administrator VLAN segment:

- It is necessary to use the NAT functionality for administrator computers accessing devices in the management part of the network.
- The addresses into which the administrator computers' addresses are translated are in the address range of the management part of the network. This achieves the same effect as access from the management network.

Accessing devices from remote locations:

- Using VPN access technology is mandatory.
- The NAT functionality is applied here in the same way as when accessing devices from the administrator VLAN segment.

3.3 Mutual Isolation between Devices in the Management Network

In order to ensure a higher level of security, it is necessary to limit communication between devices in the management network. Below are the recommendations for such limitations:

- The management servers are able to communicate with all other devices in the management network, but only for management protocol purposes.
- Other devices cannot communicate with each other through the management network.

This level of isolation can be achieved on Cisco devices (which may also be the case with some other vendors' devices) by using the following functionalities on the OOBM switch:

- Private VLAN
The ports on the OOBM switch to which all the devices are connected (except the management server) are in an isolated mode of operation. Ports set to work in isolated mode can only communicate with ports working in promiscuous mode.
The ports on the OOBM switch to which the management servers are connected (as well as the ports used by the administrator to access the OOBM network) are in promiscuous mode. Ports set to promiscuous mode can communicate with all ports, regardless of their mode of operation.
- MAC-Access Control List
Enables filtering within a broadcast domain, based on the MAC addresses of the devices.

3.4 Examples of Topologies

A comprehensive example of the topology of a management network is shown in Figure 3.1 below. The example includes the following:

- Out-of-band management network

- Three methods of accessing the management network (a computer in OOBM, access from the administrator VLAN and remote access through a VPN)
- Access to the management VLAN network from the OOBM network
- The connections of the network devices with the OOBM network established via the console port and via a special OOBM port
- The connection of the network devices with the management network established through the management VLAN
- Additional interface of the server providing connection with the OOBM or the management VLAN part of the network
- The connections of other devices with the OOBM or the management VLAN part of the network
- The locations of the management server and the connections with the management part of the network, and with the rest of the network for the monitoring system.

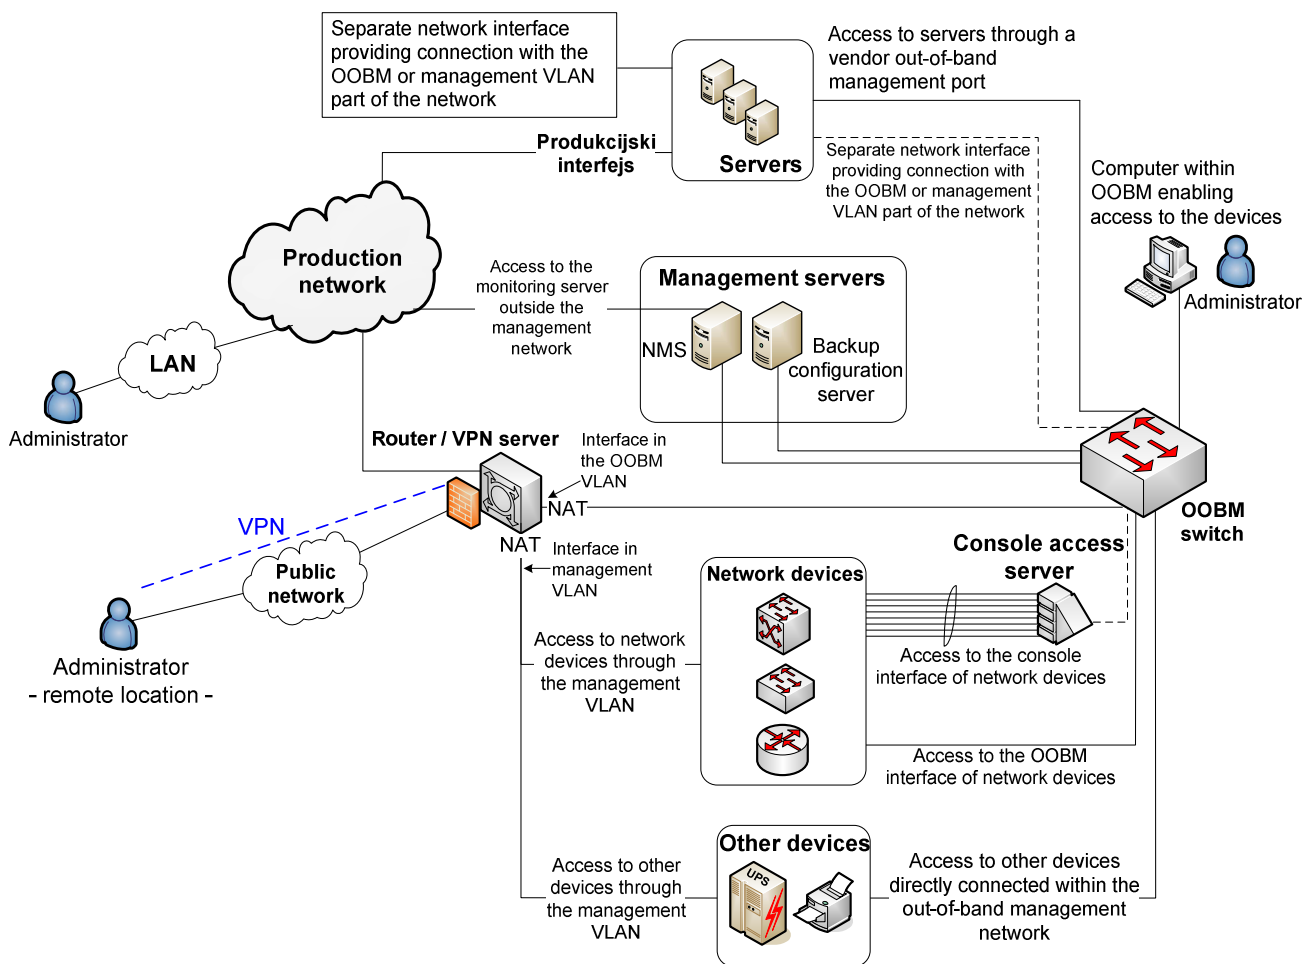


Figure 3.1: An example of the comprehensive topology of the management network

Figure 3.2 shows an example of network segmentation and the connections of devices with the management VLAN:

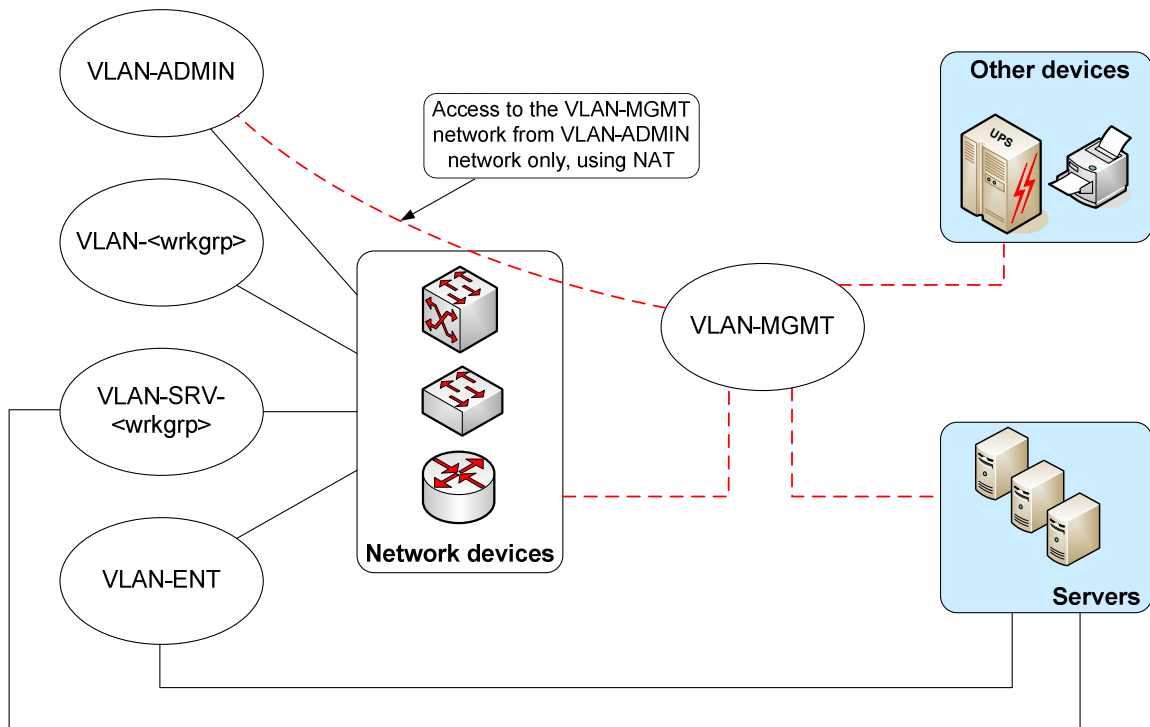


Figure 3.2: The topology of a segmented network and the connections of the devices with the management VLAN

4 Logical Access to Devices (Device Access Protocols)

4.1 Protocols for Controlling and Configuring Devices

Upon defining the ways of physical access to devices, it is also necessary to define the communication protocols that are recommended for use under different circumstances.

The description of protocols and their basic characteristics are provided first, followed by recommendations for use that are given separately for communication with different types of devices.

TTY:

- The protocol for asynchronous serial communication. This is used for communication with the Console port of network devices.

Telnet:

- Enables access to the command line (Command Line Interface - CLI).
- This protocol is supported by most devices.
- Communication between a Telnet client (the device from which the access is established) and a Telnet server (the device being accessed) is performed in non-encrypted mode (clear text), which is the main disadvantage of this protocol.

SSH (Secure Shell):

- Enables access to the command line (it is visually identical to the Telnet protocol).
- Communication between the SSH client and the SSH server is performed in encrypted form, which is why wherever possible the application of this protocol is recommended instead of the Telnet protocol.
- In order to work with this protocol, devices must have encryption software support.

RDP (Remote Desktop Protocol):

- Enables graphical access to the server (desktop access).
- Developed by Microsoft, although it is possible to have it implemented on other OS platforms.
- The level of encryption of the data being transferred during communication between an RDP client and the RDP server can be set on the RDP server (Windows Server platforms). It uses the RSA RC4 algorithm, which has shortcomings and cannot be considered completely secure.
- It is recommended that this protocol is used through VPN networks.

VNC (Virtual Network Computing):

- It is an application based on the RFB protocol (Remote Frame Buffering).
- Like RDP, it enables graphical access to servers.
- There are several versions of VNC applications, some of which are free while some require a licence.
- In its basic (free) version, VNC does not support any traffic encryption mechanisms.
- It is recommended to be used through VPN networks.

HTTP(S) (HyperText Transfer Protocol (Secure)):

- Enables web access to devices.
- In its basic version, this protocol does not support encryption. HTTPS (HTTP Secure) represents a combination of the HTTP and the SSL/TLS protocol, which is considered sufficiently secure for data transfer.

4.2 Access to Network Devices

For network devices, the use of certain protocols is recommended under the following circumstances:

The **TTY** protocol, i.e., access to the console port of a device, is recommended:

- when it is necessary to reset the device passwords;
- when no access to the device is provided via the network through one of the protocols described below;
- if it is difficult or impossible to access the device through the network due to a DoS attack, deleted configuration and/or incorrect configuration that blocks access to the device, etc.

The **Telnet** protocol is recommended:

- where there is no direct connection to the console port of the device;
- when the device does not support SSH access, which depends on the version and supported functions of the device's operating system;

- if there is a VPN network through which the connection is established with the secure part of the network, which in turn enables access to the devices (out-of-band segment). Also, if there is a direct connection with the management VLAN.

The **SSH** protocol is recommended:

- if the device's operating system supports traffic encryption, i.e., the protocol itself;
- where there is no direct connection to a network from which the devices are accessed securely (out-of-band segment), or to the management VLAN.

The **HTTP(S)** protocol is recommended:

- if the HTTP protocol is supported and under the following conditions:
for the basic version of the protocol (HTTP) the same recommendations apply as in the case of the Telnet protocol;
for a combination of HTTP + SSL/TLS (HTTPS) protocols the same recommendations apply as in the case of the SSH protocol.

4.3 Access to Servers

Access to servers for management purposes can be set in a more flexible manner, bearing in mind that virtually all access protocols can be implemented via software on most operating systems. Below are the environments in which the use of specific protocols is recommended.

SSH access recommendations:

- For Unix-based operating systems it is recommended to stick to the usual practice of accessing servers via the SSH protocol.

Graphical access recommendations:

- It is recommended that graphical access protocols (RDP, VNC, etc.) are applied exclusively through a management network (out-of-band and/or management VLAN).
- The use of these protocols outside the management part of the network is only recommended if the client and the management part of the network are connected via a VPN connection.

Telnet access recommendations:

- The Telnet protocol is rarely used to access servers. It is recommended that this protocol is used only if the servers are accessed through an interface located within the management network (out-of-band or management VLAN). This requires the existence of at least two network interfaces on the server (see Section 2.2).

- Telnet access is recommended only if there is a VPN connection between the client (the computer accessing the server) and the management part of the network.

4.4 Access to Other Devices

Access to “other” devices (UPS devices, air conditioners, etc.) depends on the type of specific device, as well as on the hardware access to the device – i.e., it depends on the type of network interface and supported protocols. The usual access methods and recommendations for their use are presented below:

Telnet access:

- The Telnet protocol is only recommended if the network interface of the device is in the management part of the network (out-of-band and/or management VLAN).
- If it is not possible to connect the network interface of the device to the protected management part of the network, the Telnet protocol is not recommended.

Web access:

- If web (HTTP) access is supported, it is recommended to use it only if the network interface is within the protected management part of the network. Web access via the HTTPS protocol is recommended if the network interface of the device is not connected to the protected management part of the network.

4.5 Device Monitoring Protocols

Many of today’s network systems rely on the SNMP protocol for monitoring purposes. The SNMP protocol itself is designed to place very little load on the network. It is called a *simple protocol* because it uses simple (unstructured) types of data. This OSI application layer protocol is an integral part of the TCP/IP protocol stack. It consists of a set of standards that define the manner of network management, information storage databases and the structures of the data in use. It relies on UDP as a transport protocol, although it is possible to set it to work via TCP. However, such a setting is not recommended for larger networks due to the large number of connections, which may overload the devices, and the size of the TCP protocol header, which may increase the amount of traffic on the link. Currently, there are two versions of the SNMP protocol in use: v2c and v3. Their security characteristics are given in Table 4.5.1 below.

SNMP Security Models and Levels				
Model	Level	Authentication	Encryption	Method of operation
v1	noAuthNoPriv	Community String	-	Uses community string for authentication.
v2c	noAuthNoPriv	Community String	-	Uses community string for authentication.
v3	noAuthNoPriv	Username	-	Uses username for authentication.
v3	authNoPriv	MD5 or SHA	-	Authentication is based on HMAC-MD5 or HMAC-SHA algorithm. MD5 or SHA hash is sent instead of password.
v3	authPriv	MD5 or SHA	DES/AES	Authentication is based on MD5 or SHA algorithm. Enables DES/AES encryption during data transfer.

Table 4.5.1 – The characteristics of SNMP protocols

There are two basic operation modes of the SNMP protocol: READ and READ/WRITE. The READ mode enables only reading the SNMP variables from a remote device, while the READ/WRITE mode enables setting certain variables on the remote device, i.e., controlling the device (router restart, backup of the current configuration, etc.). While configuring an agent on the remote device, it is possible to set limitations in the Management Information Base (MIB). If the READ/WRITE option is used for setting only one of the OID variables, the limitations on the SNMP agent in the MIB base should be set to include only a specific OID value. In this case, WRITE access to other OID values would be forbidden. Relevant examples will be provided in the section dealing with setting the agents on specific types of devices.

4.5.1 SNMP v2c

The most widely used version of the SNMP protocol is currently v2c (RFC 1901-1908). SNMP v2c performs authentication using a community string that is sent via the network as clear text. If this traffic is intercepted by a sniffing application, it could easily discover the community string and jeopardise the proper operation of the network. The use of SNMP v2c is only recommended if the devices do not support SNMP v3, and in such situations other mechanisms are introduced for protection during data transfer (ACL, firewalls, etc.). The READ mode is recommended for monitoring the network and, if it is necessary to control the network (READ/WRITE) via the SNMP protocol, it is recommended to introduce limitations in the MIB base. When starting the SNMP v2c agent, the devices usually have a predefined community string value that is set to “public”. This predefined

and widely known community string should be set to some other value, if possible to a combination of digits and letters.

4.5.2 SNMP v3

The need for security within the network brought about SNMP v3.

SNMP v3 introduces important security aspects:

1. message integrity, which prevents packet alteration during transfer;
2. authentication - a confirmation that the message came from the right sources;
3. packet encryption, which prevents unauthorised sources from reading the messages.

Table 4.5.1 shows that SNMP v3 introduces three different levels of security. The most secure level uses authentication based on the SHA algorithm and DES or AES encryption. Before introducing any level of security, it should be checked if the network device supports encryption of traffic. If a device does not support encryption, a lower level of security may be used to include authentication only. The third and lowest level of security practically acts like SNMP v2c and uses only a username, in the same way as SNMP uses a community string to access the device. When starting the SNMP v3 agent, it is recommended to introduce limitations in the MIB base on certain OID values for READ and READ/WRITE modes.

5 Configuration Maintenance

5.1 Configuration Backup

Because it is always possible for a device or a part of its functionality to fail unexpectedly, it is desirable to set the functionality of a new device to the status that the device being replaced was in before its failure, and to do that with minimum down-time. This is possible if the configurations of all devices in the network are properly and regularly maintained and backed up.

Configuration backup is of importance primarily for network devices. Configuration backup of network devices is usually performed using a simple data transfer protocol, such as TFTP and RCP.

A common disadvantage of these protocols is that they transfer data in an unencrypted form. This poses a problem because the configurations being backed up often contain confidential information (username/password, versions of the operating system, etc.). TFTP is considered an even less secure protocol because it does not have a server access authentication function.

Due to these issues, it is recommended to perform the configuration backup in such a manner that the backup servers (TFTP and/or RCP) are **connected only** within the protected management part of the network and are isolated from the rest of the network.

For servers it is necessary to back up all configuration parameters. This can be done if the configuration parameters are located in a configuration file, which can be transferred to the backup server via a data transfer protocol such as FTP or SCP.

It is recommended to use the SCP transfer protocol because, unlike the FTP protocol, it encrypts data during the transfer.

It is also possible to perform a backup if the server (client) application itself has a built-in backup mechanism.

6 NMS Server

6.1 The Position of the NMS Server in the Network

When defining the position of the NMS server in the network, it is necessary to strictly define the policy of access to the NMS server. Recommendations for campus networks are as follows:

- The NMS server must have one network interface located inside the management network (OOBM or management VLAN). This interface serves for both managing the NMS server itself and for the communication between the NMS tools and the other devices in the network.
- If necessary, the NMS may have an additional network interface in the production part of the network. This interface would allow access to the monitoring system in order to monitor the current status of devices and perform alarm detection. The access enabled through this interface should be limited to read-only mode, with the possibility of performing certain predefined diagnostic actions. Likewise, it is necessary to limit access through this interface to intended users only (administrators and helpdesk service).

6.2 The Recommended Version of SNMP on Network Devices and Servers

Most devices support SNMP v2c, while only newer devices support SNMP v3. The version of the SNMP protocol on the server depends exclusively on the type of operating system in use. Since Windows OS does not support SNMP v3, it is necessary to find an SNMP agent that supports it and have that installed. One of the free Windows versions of the SNMP v3 agent, which is based on the Linux “NET-SNMP” package, can be downloaded from the website <http://marksw.com/snmpv3agent/windowsagent.html>. There are a number of free versions of Windows SNMP agents that can be downloaded from the Internet. The “NET-SNMP” package that comes as part of Linux OS supports SNMP v3 and may be installed together with the operating system. The implementation of SNMP v3 on the server will not have any major impact on the performance of devices, so SNMP v3 with the highest level of security is recommended for use on the server. For switches and routers a problem may arise if the memory and CPU are already loaded by other processes, so starting SNMP v3 may affect the performance of the devices, especially if the entire MIB base is being read and a security level employing encryption is used. In that case, it is not recommended to start encrypting the SNMP traffic, despite the fact that the devices support SNMP v3, in order to avoid a degradation of performance of the devices; rather a level that uses authentication only should be started.

6.3 Recommended Variables to Monitor

Before implementing the monitoring system in the network, it is necessary to define the parameters that will be monitored. As the MIB base offers a large number of OID parameters, the question is how to choose the values that provide the most important information about the status of network devices and links. A prevailing tendency in the IT world is to use standard IETF MIB bases, which should be supported by any device vendor.

6.3.1 Network Devices

The most frequently monitored parameters in network devices, such as routers and switches, include:

1. The state of the interface (layer-2 and layer-3 connections)
2. Interface data flow (obtained indirectly, by reading the counter consecutively and dividing the value by the interval between the readouts)
 - i. Standard In/Out traffic (bits/s)
 - ii. Rejected In/Out traffic (bits/s)
 - iii. Transferred traffic per In/Out packet (packets/s)
3. Processor load
4. Memory load
 - i. I/O memory
 - ii. CPU memory

If it is necessary to monitor functions that are not available on every device, i.e., functions that are vendor-specific, it is required to analyse the MIB bases of the vendor who manufactured the device.

6.3.2 Servers

The OID variables that can be read out on the server depend on the operating system. In general, all operating systems support the standard IETF MIB bases, so most of the OID values are universal for all devices that support SNMP. It is recommended to monitor the following values:

1. The state of the interface (layer-2 and layer-3 connections)
2. Interface statistics (obtained indirectly)
 - i. Standard In/Out traffic (bits/s)
 - ii. Rejected In/Out traffic (bits/s)
 - iii. Flow per In/Out packet (packets/s)
 - iv. How long the interface is active
3. Processor load
4. Memory load
 - i. HDD memory
 - ii. RAM memory
5. Swap space memory
6. Number of system processes
7. List of started services on the server
8. Number of established TCP connections
9. Number of currently logged system users

6.3.3 UPSs

In the case of monitoring the SNMP variables on UPS devices, most of the OID values can be found in the vendors' MIB bases. It is recommended to monitor the following variables:

1. The current state of the UPS, i.e., its work mode (battery mode, online mode, malfunction, etc.)
2. UPS battery capacity
3. How long the UPS can work in battery mode
4. The temperature of the battery
5. UPS output load
6. Incoming voltage
7. Output voltage
8. Input current
9. Output current

6.4 MIB Variables

Variables in the MIB base are divided into two groups. The first group comprises a set of variables that can be found on any device (standard variables), while the second group consists of variables that are specific for certain vendors of network devices (proprietary variables).

6.4.1 Standard MIB Variables

The standard IETF MIB bases are located under the MIB-2 node (.1.3.6.1.2.1) of the MIB tree. Some of the most frequently used variables in this node include the following:

1. interfaces (.1.3.6.1.2.1.2) – All information about the state of the interfaces on the device is located here.
2. ifMIB (.1.3.6.1.2.1.31) - ifMIB is an expansion of the MIB interfaces from 32-bit counters to 64-bit counters.
3. tcp (.1.3.6.1.2.1.6) – Here you can find the parameters describing the tcp connections.
4. host (.1.3.6.1.2.1.25) – The Host table contains information about the state of the processor and memory on the servers.

6.4.2 Proprietary MIB Variables

Proprietary MIB variables are defined and implemented by the vendor of the network devices and they can only be used on the devices produced by that vendor. All proprietary MIB variables are found under the enterprise node (.1.3.6.1.4.1) in the MIB base. Below are some examples of the MIB variables of certain vendors of network devices.

1. Cisco(.1.3.6.1.4.1.9) – Contains all proprietary MIB variables supported by various types of Cisco devices.
2. APC(.1.3.6.1.4.1.318) – Contains all proprietary MIB variables supported by various types of APC devices.
3. juniperMIB(.1.3.6.1.4.1.2636) – Contains all proprietary MIB variables supported by various Juniper devices.

6.5 Trap Mode

The SNMP protocol is used for periodic readouts of data on remote devices. If a change occurs on a remote device, it will only be detected when the NMS server reads it, and this interval may be very long. That is why the concept of trap messages has been introduced. If a change occurs on a remote device, the remote device itself will generate an SNMP trap message defining the change and it will send it to the NMS server. The SNMP trap mode has been designed in such a way that it sends SNMP trap messages via UDP port 162, and it does it in one direction only, without requesting confirmation as to whether the trap was received. SNMP v2c sends the trap message using a community string in clear text. SNMP v3 version sends information to a specific user with a specific password and a specific engineID, and the information can be encrypted depending on the security model. This implies that the NMS server must know the username, password and engineID configured on the remote device in order to be able to decrypt the SNMP v3 trap received. The settings used for SNMP v3 on remote devices must be set in the same way on the NMS server. Instead of creating a large number of different users on the NMS, the same user may be generated on all remote devices as the user that sends trap messages only. Since the engineID is automatically generated on remote devices when a user is created, it is necessary to manually change the engineID so that it is the same for the trap user on all devices. The advantages of using SNMP v3 for trap messages include the increased security when receiving trap messages, although the SNMP v2c trap mode may be used if the network is designed in such a way that it is not easy to perform any DoS attacks on the NMS. The complexity of configuring the SNMP v3 trap is one of the main reasons why SNMP v2c is more commonly used for trap mode.

6.6 Examples of SNMP Configurations on Devices

The parameter values (e.g., community string) in the configurations used in the examples below are shown in italics.

6.6.1 Cisco Router

6.6.1.1 SNMP v2c

This example shows the commands for setting the SNMP v2c version of the protocol.

The following command, which is used in the configuration mode, starts the SNMP agent on the router:

```
1. SNMPTEST(config)#snmp-server community donotusepublic ro acl10
```

The string used as authentication – *donotusepublic* – is a form of protection, so the router will only respond to the device that sends a request containing this string. The **ro** option emphasises that it is only possible to read the data without changing it (ro - read only). It is also possible to change certain variables (wr - write command), which may cause a change in the workings of the router (restart the router), so it is very important not to use the factory predefined values for the community string and to limit SNMP queries to readouts only, rather than changing the variables. Finally, the end of the command defines the access list – *acl10*, which can be used to allow access to the SNMP agent on the device from specific IP addresses only.

In order to properly set the SNMP trap mode, it is necessary to define a community string for the trap mode, start the SNMP trap and define the destination address to which the trap messages will be sent.

```
2. SNMPTEST(config)#snmp server enable traps snmp linkup likdown
3. SNMPTEST(config)#snmp server host 192.168.10.1 version 2c donotusepublic
```

The first command defines the type of action for the trap. Thus, if the link goes down or goes back up, a trap message will be generated.

The second command defines the IP address to which the traps will be sent, the SNMP version used and the community string.

6.6.1.2 SNMP v3

The following commands start the SNMP v3 protocol on Cisco devices:

```
1. SNMPTEST(config)#snmp-server view MYGROUPV interfaces included
2. SNMPTEST(config)#snmp server group MYGROUP v3 auth read MYGROUPV
3. SNMPTEST(config)#snmp server user pera MYGROUP v3 auth md5 perapass priv
des56 pera1234
4. SNMPTEST(config)#snmp server enable traps linkup linkdown
5. SNMPTEST(config)#snmp server host 192.168.10.1 traps version 3 priv MYGROUP
```

The first command defines the values in the MIB base of OIDs that can be read on the device. In this case, the reading of OIDs (**interface**) that describe the state of the interface on the device is enabled. If no group is defined, it is assumed that the reading of all values in the MIB base is allowed.

The second command defines the group – **MYGROUP**, which uses the SNMP v3 protocol and authentication. This group is allowed to read the data from the MIB base, and only the data defined under the “**view**” of **MYGROUPV**.

The third command defines the user – **pera**, who belongs to **MYGROUP**, uses SNMP v3, the MD5 based authentication algorithm and has the **perapass** password. The last option in the third command - **des56 pera1234** defines the passphrase used for the **des56** encryption of the SNMP traffic.

The fourth command starts the SNMP trap mode.

The fifth command defines the NMS server that will collect the trap messages. In this case, the SNMP v3 protocol will be used for communication between the NMS and the Cisco device, along with the rules defined in **MYGROUP**. This command shows that the engineID was generated automatically and if we want to enable the NMS server to receive trap messages from the **pera** user, it is necessary to find out the engineID of the **pera** user and configure it in the SNMP agent on the NMS server.

Another way of doing this is to manually set the engineID for the local or remote user, by using the following command:

```
6. snmp server engineID [local engineid-string] | [remote ip-address udp-port  
port- number engineid-string]
```

6.6.2 Linux Server

6.6.2.1 Linux Server - SNMP v2c

When setting the SNMP protocol on Linux operating systems, it is necessary to first install the SNMP daemon on the server. The following example describes the installation on the CentOS 5.X operating system using the YUM command. The following command enables the automated installation of the SNMP daemon and some useful commands that control the work of SNMP.

```
1. yum install net-snmp net-snmp-utils
```

The next step sets the service to start automatically each time the server is started. It is necessary to enter the following command:

```
2. chkconfig snmpd on
```

The next step is to set the community string and OID objects that can be read from the server. It is necessary to edit the snmpd.conf file, which is usually located in the /etc/snmp/ directory, and change the following lines.

```
com2sec notConfigUser default public
```

In this line, it is necessary to change the default community string public into a desired community string.

```
view systemview included .1.3.6.1.2.1.1
```

This line shows that all the OIDs located under the .1.3.6.2.1.1 node in the MIB tree are available. By using the *excluded* command, it is possible to exclude certain OID values, i.e., to introduce limitations in the manner the MIB base is displayed. At this point, it is necessary to define the OID values that will be returned by the server in response to SNMP queries. If the NMS requires an OID that is not defined here, the server will not respond to the NMS.

It is now necessary to start the service using the following command:

```
3. service snmpd start
```

A check can be performed using the following command:

```
4. snmpwalk -v 2c -c mojcommunity 127.0.0.1
```

As a result, the whole MIB table (tree) will be shown, or a part of the MIB table defined by the previous commands imposing limitations in reading the MIB base.

6.6.2.2 Linux Server - SNMP v3

The installation of the SNMP v3 agent is performed in the same manner as the installation of the previous v2c version, but this time it is necessary to start the SNMP v3 version. It is required to edit the `snmpd.conf` file and add the following commands:

```
syslocation MojGradiliLokacija
syscontact mojemail@provajder.com
view mojpogled included .1.3.6.1.2.1.2.2
createUser john MD5 john1234 DES john5678
rouser john priv -V mojpogled
```

The first two parameters, *Syslocation* and *Syscontact*, are the data used to provide general information about the server. This data is not essential for the proper functioning of the SNMP protocol and it can be added regardless of the version of SNMP. It is important for server administrators, who will thus have data about the location of the server and the contact person they can approach in case a problem emerges, in addition to the basic information on the status of the server. *Syslocation* and *Syscontact* are some of the parameters that can be found on all devices that support the SNMP protocol.

The third line defines the view – *mojpogled*, i.e., the set of OID values from the MIB tree. In this case, the server interface table is defined. It is possible to add more tables or to exclude some OID values by using the *exclude* command. These limitations are very important because it is possible to set certain parameters through SNMP that may directly affect the proper functioning of the device.

The fourth command creates the user – *john*, whose password is *john1234*. The *MD5* algorithm is used for authentication and the traffic is encrypted using the *DES* algorithm. The passphrase used during encryption is *john5678*.

The fifth command gives the user – *john* read-only (*rouser*) privileges, but only regarding the *mojpogled* view of the MIB base as defined in the third command.

After editing the SNMP configuration files, it is necessary to restart the service so that the changes entered in the files can take effect.

A shortcoming of this type of configuration is reflected in the fact that the values entered in the configuration files are kept as clear text.

The restart command is shown below:

```
1. service snmpd restart
```

It is possible to check the settings of SNMP v3 on the server by using the following command:

```
2. snmpwalk -v 3 -u john -l authPriv -a MD5 -A john1234 -x DES -X john5678
192.168.1.1
```

6.6.2.3 *Configuring the SNMP Protocol Using a Perl Script*

Perl scripts are automatically installed along with the Net-SNMP package, and they enable the basic configuring of the SNMP agent. The three scripts used for configuring the SNMP agent are `snmpconf`, `snmpusm` and `net-snmp-conf`. The first script enables the interactive configuration of the basic functions of the SNMP agent, while the second and third ones may be used to create SNMP v3 users.

An advantage of configuring the protocol using Perl scripts is reflected by the fact that after restarting the server the sensitive parameters are kept in an encrypted form. This is a recommended way of configuring the SNMP agent.

It is possible to change the settings of the SNMP agent in the following two ways: by manually editing the configuration files `snmpd.conf` and `snmptrapd.conf` (as shown in Section 6.2.2.2) or by using the Perl scripts. More detailed information about the Net-SNMP package can be found at <http://net-snmp.sourceforge.net>.

7 Saving Syslogs

7.1 SysLog Protocol

The SysLog (System Logger) protocol has been developed as a mechanism for gathering information about the changes and events in Unix operating systems, and one of its very useful features is the possibility of sending the information via the network. This has enabled collecting the messages on a central server, which in turn allows for the quicker and easier detection of problems and their solutions. SysLog uses the UDP protocol (port 514) on the transport layer, while the application layer offers no mechanism for providing information on whether the message was properly transferred to its destination, and therefore the protocol has been placed in the category of unreliable protocols. Despite these shortcomings, SysLog is one of the commonly used protocols for gathering information about the state of the system.

The message format carries the following information describing the state of the system:

- Facility – identifies the object that generated the message. It can be the operating system, a process or an application. The facility is presented as an integer, and the values ranging from 0 to 15 are reserved for Unix operating systems, while values ranging from 16 to 23 are traditionally used for network devices (routers, switches, etc.).

Table 7.1.1 below contains an overview of Facility values.

Integer	Facility
0	Kernel messages
1	User-level messages
2	Mail system
3	System daemons
4	Security/authorization messages
5	Messages generated internally by SysLog
6	Line printer subsystem
7	Network news subsystem
8	UUCP subsystem
9	Clock daemon
10	Security/authorization messages
11	FTP daemon
12	NTP subsystem

13	Log audit
14	Log alert
15	Clock daemon
16	Local use 0 (local0)
17	Local use 1 (local1)
18	Local use 2 (local2)
19	Local use 3 (local3)
20	Local use 4 (local4)
21	Local use 5 (local5)
22	Local use 6 (local6)
23	Local use 7 (local7)

Table 7.1.1 – Facility values

- Severity – may assume one of 8 integer values describing the current severity of the problem. Possible values are provided in Table 7.1.2.:

Integer	Severity
0	Emergency: System is unusable.
1	Alert: Action must be taken immediately.
2	Critical: Critical conditions.
3	Error: Error conditions.
4	Warning: Warning conditions.
5	Notice: Normal but significant condition.
6	Informational: Informational messages.
7	Debug: Debug-level messages.

Table 7.1.2 – Severity values

- Hostname – contains the IP address of the device sending the SysLog data, which usually includes the IP address of the interface from which the messages are sent.
- Timestamp – provides information about the time the SysLog message was generated. It is recommended to use the NTP protocol in order to ensure proper time synchronisation of the devices. It is very important to have a correct time sequence for all SysLog messages.
- Message – contains the SysLog message generated by the device and some additional information about the process that generated the message.

7.2 Location of the SysLog Server

The application that gathers SysLog data is often installed on the server collecting the SNMP data. It is a simple and acceptable solution for systems whose server performance characteristics are good enough to support simultaneous processing of SNMP and SysLog data. The critical point is the operation of the database in which the messages are usually stored. A large number of generated SysLog messages will cause a situation where messages are frequently written in the database and thereby onto the hard disc as well. With a suitable choice of hardware configuration, the database can be divided onto separate hard discs in order to

increase the performance of the system. Another solution for the massive export of SysLog messages is to designate a separate server for the SysLog application.

For the location of the SysLog server in the network the same recommendations as for the NMS server should be followed.

7.3 Installation

Before starting the SysLog service in the network, it is very important to properly start the time synchronisation on the network devices and the SysLog server itself. This is important in order to ensure that the messages are saved in the correct time sequence in the database of SysLog collector. It is recommended to export all SysLog messages and then filter the most important ones, or to set the service to export only the SysLog messages related to certain important functions of the devices or services. The manner in which the SysLog agent is started on specific devices is described below.

7.3.1 Network Devices

The SysLog agent on routers and switches is started in the following manner on Cisco devices.

The SysLog service is started using the following commands:

```
1. ABC(config)#logging on
2. ABC(config)#logging host 10.10.5.1
3. ABC(config)#logging trap informational
4. ABC(config)#logging source-interface Loopback0
5. ABC(config)#logging buffered 100000
6. ABC(config)#logging buffered debug
7. ABC(config)#logging monitor informational
8. ABC(config)#no logging console
```

The first command will start the data logging process.

The second command defines the IP address of the SysLog server (collector) to which the data will be exported.

The third command defines the severity level of the information that will be exported to the server. The informational level is level 6, which means that all messages ranging from severity level 0 to severity level 5 will be collected.

The fourth command defines the source address that will show up in the logs, and it is recommended to set it to the loopback address because this address is always active. If the SysLog server is unavailable or if we want to read the messages directly on the device, it is recommended to separate some system memory in which the data will be kept.

The fifth command defines the buffer size in bytes.

The sixth command enables the logging of debug commands.

The seventh command defines the severity level for the terminal lines, while the last command sets out that no logs will be sent to the console line.

7.3.2 Servers

7.3.2.1 Windows

Windows operating systems have no SysLog agent installed, but they rely on its agent called Event Logger. If we want to collect the SysLog data generated by a server using a Windows platform, it is necessary to install a separate SysLog agent that will translate the messages generated by the Event Logger into SysLog format and send these messages to a remote SysLog server. Two free versions of such agents may be downloaded from the following websites: <https://engineering.purdue.edu/ECN/Resources/Documents/UNIX/evtsys/> and <http://ntsyslog.sourceforge.net/>.

7.3.2.2 Linux

When installing the Linux operating system, the SysLog agent is installed automatically. The agent is preconfigured to collect all messages generated by the system in the files located in the `/var/log/` directory. If we want to set the facility and severity parameters of certain parts of the system, including the location where the messages will be stored (a local file or a remote server), it is necessary to edit and configure the `/etc/syslog.conf` file.

An example of the `/etc/syslog.conf` file is provided below:

```
*.info;mail.none;authpriv.none;cron.none /var/log/messages

authpriv.* /var/log/secure

mail.* -/var/log/maillog

cron.* /var/log/cron

*.emerg *

local7.* /var/log/boot.log
```

The first command shows that all the messages whose severity level is higher than or equal to the info value (informational) will be exported to the `/var/log/messages` folder, except for messages whose facility value is mail, authpriv or cron.

The second command exports all messages whose facility values equal authpriv to the `/var/log/secure` folder.

The third command exports the messages whose facility values are equal to mail, while the severity may assume any of the values, to the `/var/log/maillog` folder.

The fourth command exports the messages whose facility values are equal to cron, while the severity may assume any of the values, to the `/var/log/cron` folder.

The fifth command exports all the messages whose severity values equal emergency to the console of all logged-in users.

The sixth command exports all messages whose facility values equal local7 to the `/var/log/boot.log` folder.

If we want to export messages to the configuration file `/etc/syslog.conf` on a remote centralised SysLog server (192.168.1.1), it is necessary to add the following command:

```
*.* @192.168.1.1
```

This command in `/etc/syslog.conf` enables the export of all messages to the remote SysLog collector.

Finally, it is necessary to restart the SysLog agent using the following command:

```
service syslog restart
```

If we want to collect SysLog messages from an application operating on the Linux server, it is necessary to set the application to export its logs to any local facility and to set the option in the `syslog.conf` file so that the SysLog agent exports the local facility to the remote SysLog server.

It is recommended to keep the logs in the files at the same time, so that the administrator may view the changes on the server registered by the SysLog even when the central SysLog server is unavailable.

If a process that generates a lot of messages is started on the server, the files will be filled rather quickly and take up memory. In that case, it is necessary to start the `logrotate` option that keeps the files for a specified period and subsequently erases them.

The relevant commands are entered into the configuration file `/etc/logrotate.conf` to define the time period, compression, privileges and the size of the files.

Nowadays, Linux platforms often use a new version of SysLog called *rsyslog*. This is an improved version compared to the standard SysLog, but it cannot be found as a standard package in all Linux distributions.

The standard SysLog agent can also be configured to receive messages from other devices and to act as a collector of SysLog messages. One of the available collectors of SysLog messages supported by Linux operating systems can be downloaded from the following website: <http://code.google.com/p/php-syslog-ng/downloads/list>.

8 Traffic Analysis Protocol

In today's networks, it is very important to perform traffic analysis/inspection, not only up to layer-2 (SNMP - the amount of bytes/packages transferred via a device's interface), but also on layer-3 and layer-4. This provides insight into the nature of the traffic and services that are most frequently used (network load), including information about the amount of traffic.

8.1 NetFlow Protocol

One of the most widely used protocols for gathering network statistics is the NetFlow protocol developed by Cisco Systems, Inc. Versions 5 and 9 are the most commonly used versions. Unlike version 5, version 9 offers a flexible message format and support for MPLS and IPv6. Other vendors also offer this protocol, but under a different name (such as Juniper J-Flow, Huawei NetStream, etc.). All of these variants of the NetFlow protocol are mutually compatible. Figure 8.1 below shows the format of a NetFlow v5 message.

source IP address			
destination IP address			
next hop IP address			
input interface index		output interface index	
packets			
bytes			
start time of flow			
end time of flow			
source port		destination port	
pad	TCP flags	IP protocol	TOS
source AS		destination AS	
src netmask length	dst netmask length	padding	

Figure 8.1: An example of the NetFlow v5 message format

Due to the need of having a universal standard, the IETF defined the IPFIX protocol as the industry standard protocol used for exporting traffic flow information.

8.2 How the System Operates

When the NetFlow protocol is started on a device, it begins to collect statistics for all the traffic passing through that device. The information is then periodically exported to the server (collector) on which an application is started to receive and process the data according to the set criteria. The results of the analysis are usually presented in the form of graphics and tables that enable easy identification of problems that emerge. Some applications offer automatic detection of problems (attacks) in the network.

The results of this analysis contain the following information:

- information on the total amount of traffic transferred between certain subnets (bytes, packets, flows);
- information on the total amount of in/out traffic on specific interfaces of the exporter;
- information on the total amount of traffic at the levels of protocol, service and host;
- information on the hosts accessed from the external network;
- detection of rejected traffic (the traffic rejected by ACL, bad routing, etc.);
- prediction of future traffic behaviour.

In this manner, the following can be detected:

- the existence of viruses in the network (a large volume of traffic is generated in the OUT direction or towards the DNS or MAIL servers);
- DoS attacks;
- abuse of traffic (YouTube, Facebook, Torrent, etc.);
- access to forbidden sites;
- attempts to attack/access protected network devices;
- the existence of open ports in the network;
- *Top Talker* users.

8.3 The Location of Collectors in the Network

The location of the collector that gathers the NetFlow statistics depends on the network architecture. The volume of NetFlow data exported by network devices is directly dependent on the amount of traffic passing through a specific device (exporter). It has been empirically demonstrated that the percentage of NetFlow traffic does not exceed 1% of the total traffic in the network, so the “distance” between the server (collector) and the network device exporting the data (exporter) is not relevant. More important parameters are the accessibility and security of the server. The physical location of the server is usually connected to the main node, because the majority of main traffic passes through there. It is recommended that the server is positioned in a separate VLAN (management VLAN) and that it is protected by a firewall. In the event of failure of certain network devices it is important to make sure that the NetFlow server is available for traffic collection and analysis.

8.4 Configuring the NetFlow Exporter

The configuration of the NetFlow statistics exporter on devices depends on the characteristics of the devices themselves, and on the network architecture. Figure 8.2 below shows an example of an incorrect configuration of the NetFlow exporter:

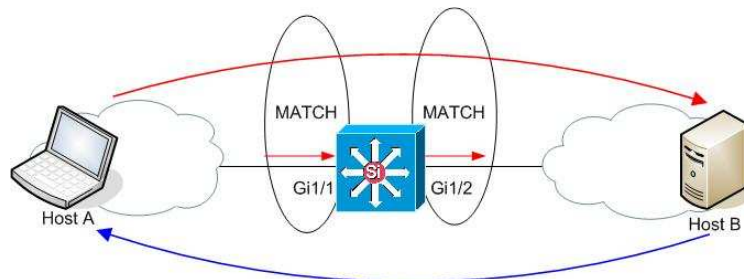


Figure 8.2: Improperly configured export of the NetFlow statistics

The figure shows that the collection of NetFlow statistics is configured in the IN direction on interface Gi1/1 and in the OUT direction on interface Gi1/2. In this case, the flow passing from point A to point B will be processed twice and exported to the NetFlow collector. However, the flow passing from point B to point A will not be included in the statistics. As a result, the statistics will be incorrect (the values will be doubled and show only one-way flow). The devices usually support the export of NetFlow statistics, which can be set at the level of the interface, although only in one of the two directions – in/ingress or out/egress. Some devices enable the export of both ingress and egress at the same time. Figure 8.3 below shows a properly configured export on the central device. On both the Gi1/1 and Gi1/2 interfaces, the NetFlow export is configured in the IN direction only.

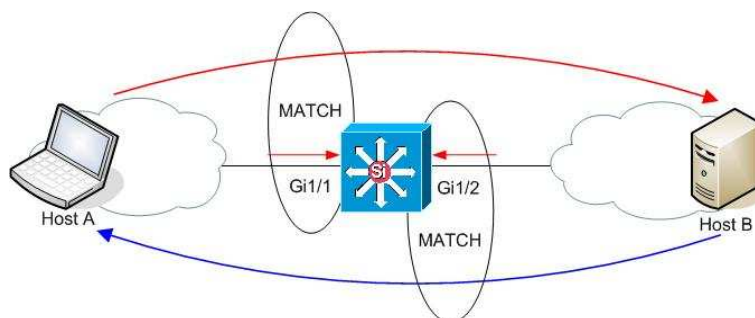


Figure 8.3: Properly configured NetFlow export

In this case, the traffic information will be exported in the direction from point A to point B (export from interface Gi1/1 for the traffic conducted in the IN direction), and the exported statistics will not be doubled. The figure also shows that it is possible to start the NetFlow export on the central device and thus cover the entire network. All the traffic passing through the central router will be exported to the NetFlow collector. The only traffic that will not be detected is traffic not passing through the central router. Figure 8.4 shows an example of a situation where the statistics of a part of the traffic is not exported because the flow is not passing through the central core device .

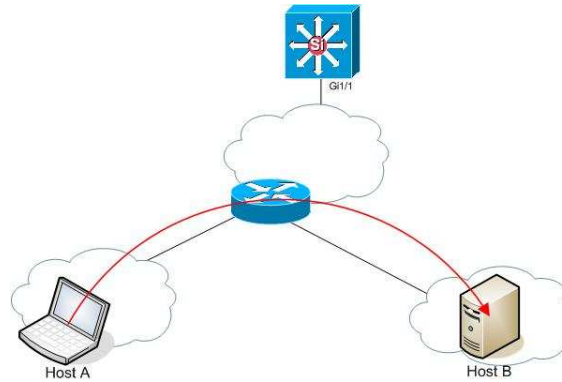


Figure 8.4: Traffic traversing from host A to host B is not exported via Netflow

8.5 Indirect Solutions for Collecting NetFlow Statistics

Support for the NetFlow protocol is usually implemented in router platforms. However, most switches do not support the NetFlow protocol, except for certain layer-3 switches (such as Cisco 6500, Cisco 4500, etc.).

If it is necessary to analyse traffic when network devices do not support the NetFlow protocol, it is possible to collect the relevant information in an indirect manner.

Figure 8.5 below shows a situation like that.

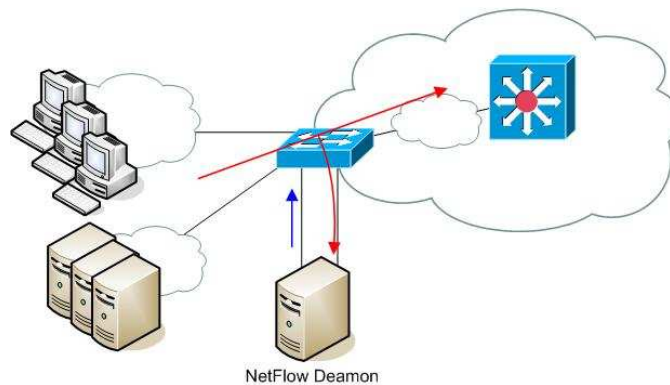


Figure 8.5: Redirecting the traffic towards the NetFlow daemon server

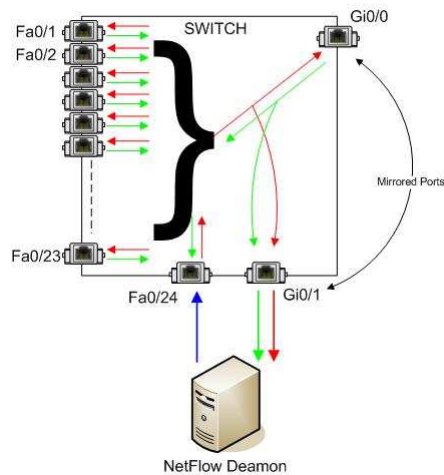


Figure 8.6: A detailed description of positioning the server and connecting the ports (port mirroring)

Figures 8.4 and 8.5 show that the traffic is redirected (port mirroring) towards the server on which the NetFlow daemon is started. When port mirroring is started on a switch, the interface towards which the complete traffic is forwarded becomes unusable for normal communication between the devices. Instead, it only forwards the traffic (IN/OUT) from the interface on which the port mirroring is set up. The problem is how to export the statistics if the interface on the server with the NetFlow daemon is unusable for normal communication. The problem can be solved by adding another network card on the server and connecting it to the switch. The blue arrow in Figure 8.5 shows the export of NetFlow statistics from the second network card on the server. Such a configuration also enables the export of NetFlow statistics from layer-2 devices, although a disadvantage may be that it takes up additional ports on the switch and requires an additional server. One port on the switch is used for receiving the mirrored IN/OUT traffic and the other port is used to send the NetFlow statistics. At this point, the application that conducts NetFlow analysis of collected traffic can be started on the server. One of the tools for creating NetFlow statistics of IN/OUT traffic on the server is the **softflow** application. It enables the export of the statistics locally (127.0.0.1) to a collector on the server itself, or to a collector on a remote server. Figures 8.7 and 8.8 show examples of NetFlow statistics being exported locally and to a remote location containing the NetFlow collector, respectively.

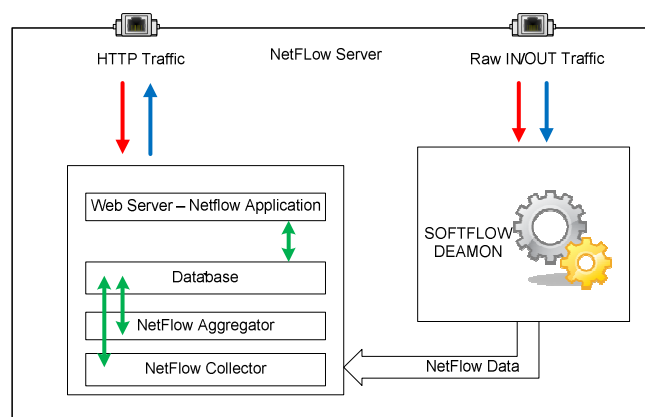


Figure 8.7: An example of NetFlow statistics being locally exported to the NetFlow collector

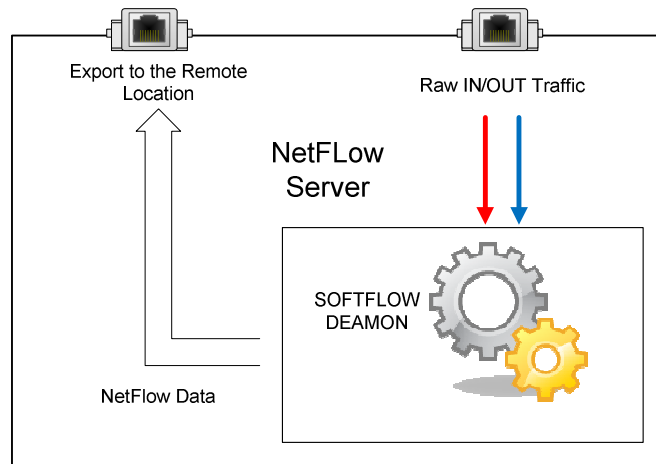


Figure 8.8: An example of NetFlow statistics being exported to a remote server on which the NetFlow collector is located

In this way, it is possible to collect NetFlow statistics even when the devices do not support the NetFlow protocol.

References

- [1] <http://www.net-snmp.org/>
- [2] <http://www.mindrot.org/projects/softflowd/>
- [3] http://www.cisco.com/en/US/docs/ios/solutions_docs/netflow/nfwhite.html
- [4] RFC-1901-1908, SNMP v2c
- [5] RFC-3411-3418, SNMP v3
- [6] Mauro D., Schmidt K.(July 2001), Essential SNMP.
- [7] CentOS SNMP and Syslog manuals

Glossary

ACL	Access List
AES	Advanced Encryption Standard
CLI	Command Line Interface
CPU	Central Processing Unit
DES	Data Encryption Standard
DNS	Domain Name System
DoS	Denial of Service
FTP	File Transfer Protocol
HMAC-MD5	Hashed Message Authentication Code, Message Digest 5
HMAC-SHA	Hashed Message Authentication Code, Secure Hash Algorithm
HTTP	HyperText Transfer Protocol
HTTPS	HyperText Transfer Protocol Secure
IPFIX	Internet Protocol Flow Information Export
IPv6	Internet Protocol version 6
MIB	Management Information Base
MPLS	Multiprotocol Label Switching
NAT	Network Address Translation
NMS	Network Management System
NTP	Network Time Protocol
OID	Object Identifier
OOBM	Out of Band Management
RCP	Remote Copy Protocol
RDP	Remote Desktop Protocol
RFB	Remote Frame Buffering
SCP	Secure Copy
SNMP	Simple Network Management Protocol
SSH	Secure Shell
SSL	Secure Socket Layer
SYSLOG	System Logger
TCP	Transmission Control Protocol
TFTP	Trivial File Transfer Protocol
TLS	Transport Layer Security
UPS	Uninterruptible Power Supply
VLAN	Virtual Local Area Network
VNC	Virtual Network Computing
VPN	Virtual Private Network

