

NEXCOM International Co., Ltd.

Industrial Firewall Solutions

Henge™ Industrial Firewall, 5-port VPN Router

IFA 3610

User Manual

Contents

Introduction: Getting Started

- Getting Started 1
- About this Reference Manual 1
- Conventions Used in This Document 1
- The Zones 2
- The IFA 3610 Appliance Management Interface 3
- Accessing the IFA 3610 Appliance 7

Chapter 1: The System Menu

- 1.1 Dashboard 9
 - 1.1.1 System Information Plugin 9
 - 1.1.2 Hardware Information Plugin 10
 - 1.1.3 Service Information Plugin 10
 - 1.1.4 Network Information Plugin 11
 - 1.1.5 Signatures Information Plugin 11
 - 1.1.6 Uplink Information Plugin 11
- 1.2 Network Configuration 12
 - 1.2.1 1/8 - Choose Type of RED Interface 12
 - 1.2.2 2/8 - Choose Network Zones 12
 - 1.2.3 3/8 - Network Preferences 13
 - 1.2.4 4/8 - Internet Access Preferences 14
 - 1.2.5 5/8 - Configure DNS resolver 16
 - 1.2.6 6/8 - Configure Default Admin Mail 17
 - 1.2.7 7/8 - Apply configuration 17
 - 1.2.8 8/8 - End 17
- 1.3 Event Notifications 17
 - 1.3.1 Settings 18
 - 1.3.2 Events 18
- 1.4 Updates 19
 - 1.4.1 Status 19
 - 1.4.2 Schedule for Retrieving the Update List 19
- 1.5 Support 19
- 1.6 Henge™ Network 20
 - 1.6.1 Subscription 20
 - 1.6.2 Remote Access 20
- 1.7 Passwords 21
- 1.8 Web Console 21
- 1.9 SSH access 21
 - 1.9.1 Secure Shell Access Settings 21
 - 1.9.2 SSH host keys 22
- 1.10 GUI Settings 22
- 1.11 Backup 22
 - 1.11.1 Backup 23
 - 1.11.2 Scheduled backups 24
- 1.12 Shutdown 25
- 1.13 License Agreement 25

Chapter 2: The Status Menu

- 2.1 System status 26
- 2.2 Network status 27
- 2.3 System graphs 27
- 2.4 Traffic graphs 28
- 2.5 Proxy graphs 29

- 2.6 Connections 29
- 2.7 VPN Connections 30

Chapter 3: The Network Menu

- 3.1 Edit hosts 31
- 3.2 Routing 32
 - 3.2.1 Static routing 32
 - 3.2.2 Policy routing 33
- 3.3 Interfaces 34
 - 3.3.1 Uplink editor 34
 - 3.3.2 VLANs 35

Chapter 4: The Services Menu

- 4.1 DHCP Server 36
- 4.2 Dynamic DNS 39
- 4.3 Time server 40
- 4.4 Intrusion Prevention 40
 - 4.4.1 Intrusion Prevention System 40
 - 4.4.2 Rules 41
 - 4.4.3 Editor 41
- 4.5 High Availability 41
- 4.6 Traffic Monitoring 43
- 4.7 SNMP Server 44
- 4.8 Quality of Service 44
 - 4.8.1 Devices 44
 - 4.8.2 Classes 45
 - 4.8.3 Rules 45

Chapter 5: The Firewall Menu

- 5.1 Common Configuration Items 47
- 5.2 Port Forwarding / NAT 48
 - 5.2.1 Port forwarding / Destination NAT 48
 - 5.2.2 Source NAT 50
 - 5.2.3 Incoming Routed Traffic 50
- 5.3 Outgoing Traffic 51
- 5.4 Inter-Zone Traffic 52
- 5.5 VPN traffic 53
- 5.6 System access 53
- 5.7 Firewall Diagrams 54

Chapter 6: Proxy

- 6.1 DNS 55
 - 6.1.1 DNS Proxy 55
 - 6.1.2 DNS Routing 55
 - 6.1.3 Anti-Spyware 56

Chapter 7: The VPN Menu

- 7.1 OpenVPN Server 57
 - 7.1.1 Server Configuration 57
 - 7.1.2 OpenVPN Settings 57
 - 7.1.3 OpenVPN Server Instances 58
 - 7.1.4 VPN Client Download 60
- 7.2 OpenVPN Client (Gw2Gw) 61



- 7.2.1 Add Tunnel Configuration 61
- 7.2.2 Advanced Tunnel Configuration 61
- 7.2.3 Import Profile from OpenVPN Access Server..... 63
- 7.3 IPsec 63
 - 7.3.1 IPsec..... 63
 - 7.3.2 IPsec Settings..... 64
 - 7.3.3 Debug Options..... 64
 - 7.3.4 Connections..... 64
- 7.4 L2TP 66
- 7.5 Authentication..... 68
 - 7.5.1 Users..... 68
 - 7.5.2 Groups..... 70
 - 7.5.3 Settings..... 70
- 7.6 Certificates 72
 - 7.6.1 Certificates..... 72
- 7.7 Certificate Authority 73
- 7.8 Revoked Certificates 74
- 7.9 Certificate Revocation List..... 74

Chapter 8: The Logs and Reports Menu

- 8.1 Dashboard 75
 - 8.1.1 Common elements..... 75
 - 8.1.2 Summary..... 76
 - 8.1.3 System 76
 - 8.1.4 Web 77
 - 8.1.5 Mail..... 77
 - 8.1.6 Intrusion Attempts..... 77
 - 8.1.7 Connections..... 77
- 8.2 Traffic Monitoring 78
 - 8.2.1 Dashboard..... 78
 - 8.2.2 Flows..... 78
 - 8.2.3 Hosts..... 79
 - 8.2.4 Interfaces 80
- 8.3 Live..... 80
- 8.4 Common actions 81
- 8.5 Summary 82
- 8.6 System..... 82
- 8.7 Service..... 82
- 8.8 Firewall..... 83
- 8.9 Proxy 83
 - 8.9.1 HTTP and Content filter..... 83
 - 8.9.2 HTTP Report 83
 - 8.9.3 SMTP..... 83
- 8.10 Settings 84
- 8.11 Trusted Timestamping..... 85



Introduction: Getting Started

Getting Started

This section presents the conventions used in the remainder of the manual, then provides introductory notions about the concept of zones, and finally describes the GUI of the Henge™ products and the possible ways to access the IFA 3610 appliance.

About this Reference Manual

This manual has been written for the 1.0 release, with a Software 1.0 as guide, but it is intended for all types of the Henge™ series. Since the functionalities and abilities may differ between the various appliances, the description of some of the displayed data or configuration options may slightly vary for some appliances or not being present at all. This guide is intended both as a contextual help and an user manual, as well as providing quick introductory descriptions to some of the concepts that lay behind the various functionalities provided by the IFA 3610 appliance.

The remainder of this section contains some basic information about this guide and how to move your first steps within the IFA 3610 appliance, introducing some important concepts and describing the most significant parts of GUI.

Conventions Used in This Document

To improve the readability and clarity of this document, several conventions are used:

Besides for emphasis, italics is used to denote non-interactive objects or labels within the web GUI, while a bolded word(s) indicates objects that require user interaction, i.e., clicking on a button or to open a hyperlink.

Admonitions are employed to mark items, actions, or tasks that require special attention:

Warning: Changing this value will cause the service to restart!

Note: Remember that you can modify this later.

Hint: Tips about configuration of options

This is an example box.

Boxes like this one contain example of configurations or short how-tos for the quick setup of some feature or service described in the main document.

A relevant subject or an example

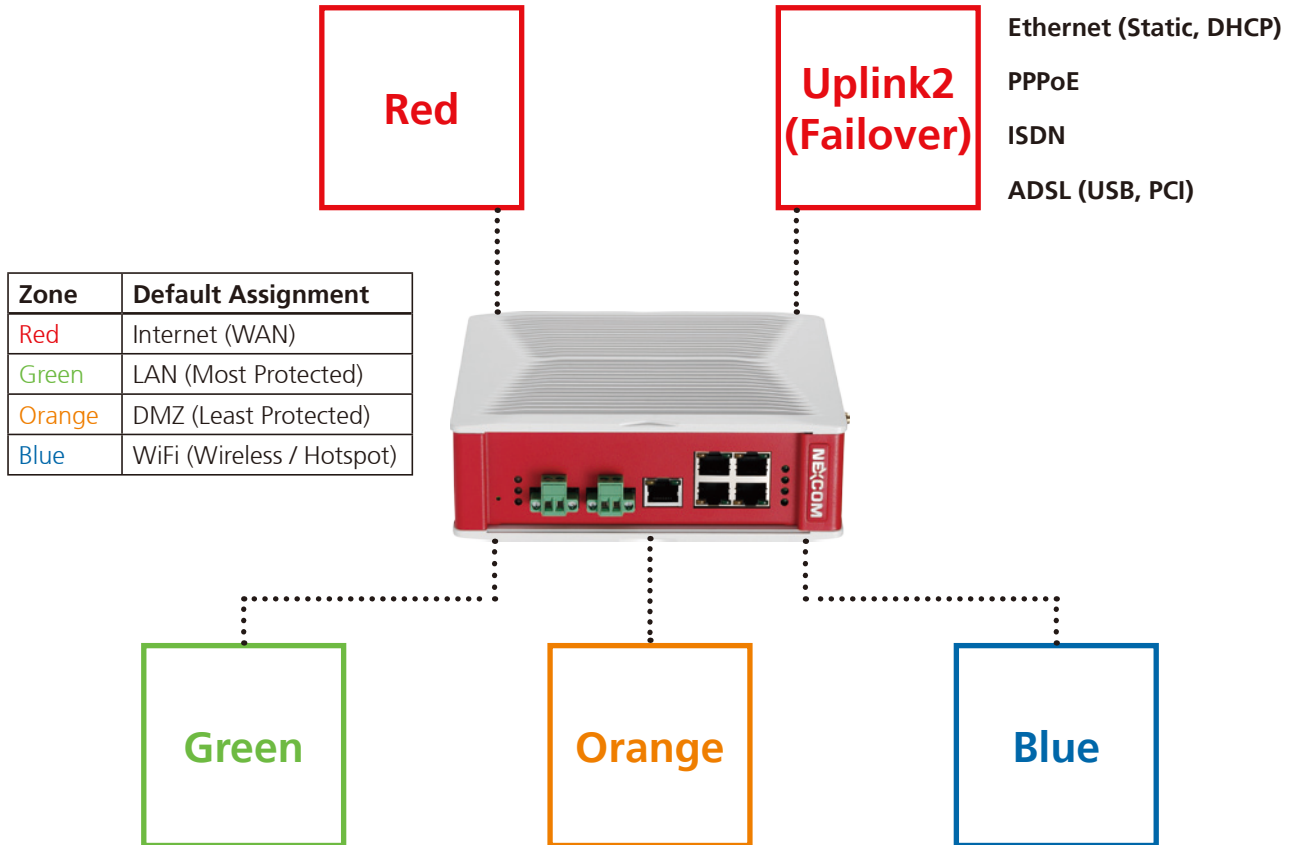
In boxes like this one ("topic"), you can find the explanation of some subject that requires a not-so-short explanation and is relevant to the topic of the section or to the configuration of some setting. Also, quick how-tos or examples may appear in it. At their bottom there might be present one or more hyperlinks to online resources.

A sequence like *Menubar ► Firewall ► Port forwarding/DNAT ► Show system rules* requires to click on each of the items, in the sequence shown, to reach a particular page or configuration item. This example shows how to reach the page that shows the configuration of the system rules for the firewall's DNAT.

Alternatively, in a sequence like *Menubar ► Firewall ► Port forwarding/DNAT ► [Rule list] ► Edit*, the [...] means that there is a large number of objects (in this case there is a list of firewall's rules) from which one should be chosen to carry out on it the action (*Edit*).

The Zones

One of the most important concepts on which the IFA 3610 appliance is grounded, the *Zone*, finds its root in IPCOP's idea to protect the networks it can reach by grouping them into different segments -the *zone*, indeed- and allowing the traffic to be exchanged only in certain directions among these segments. The four main zones are identified by a color and may group together a number of servers or workstation that have a same purpose.



- RED, this is the so-called Untrusted segment, i.e., the WAN: It encompasses all the networks outside the IFA 3610 appliance or, broadly speaking, the Internet, and is the source of incoming connections. This is the only zone that can not be managed: but only access to and from it can be granted or limited.
- GREEN, the internal network, i.e., the LAN. This zone is the most protected one and is dedicated to the workstations and should never be directly accessed from the RED zone. It is also the only zone that by default can access the *management interface*.
- ORANGE, The DMZ. This zone should host the servers that need to access the Internet to provide services (e.g., SMTP/POP, SVN and HTTP and so on). It is a good practice that the ORANGE zone be the only zone directly accessible from the RED zone. Indeed, if an attacker manages to break into one of the servers, she will be trapped within the DMZ and will not be able reach the GREEN zone, making impossible for her to gain sensitive information from local machines in the GREEN zone.
- BLUE, the WiFi zone, i.e., the zone that should be used by wireless clients to access the Internet. Wireless networks are often not secure, so the idea is to trap by default all the wireless connected clients into their own zone without access to any other zone except RED.

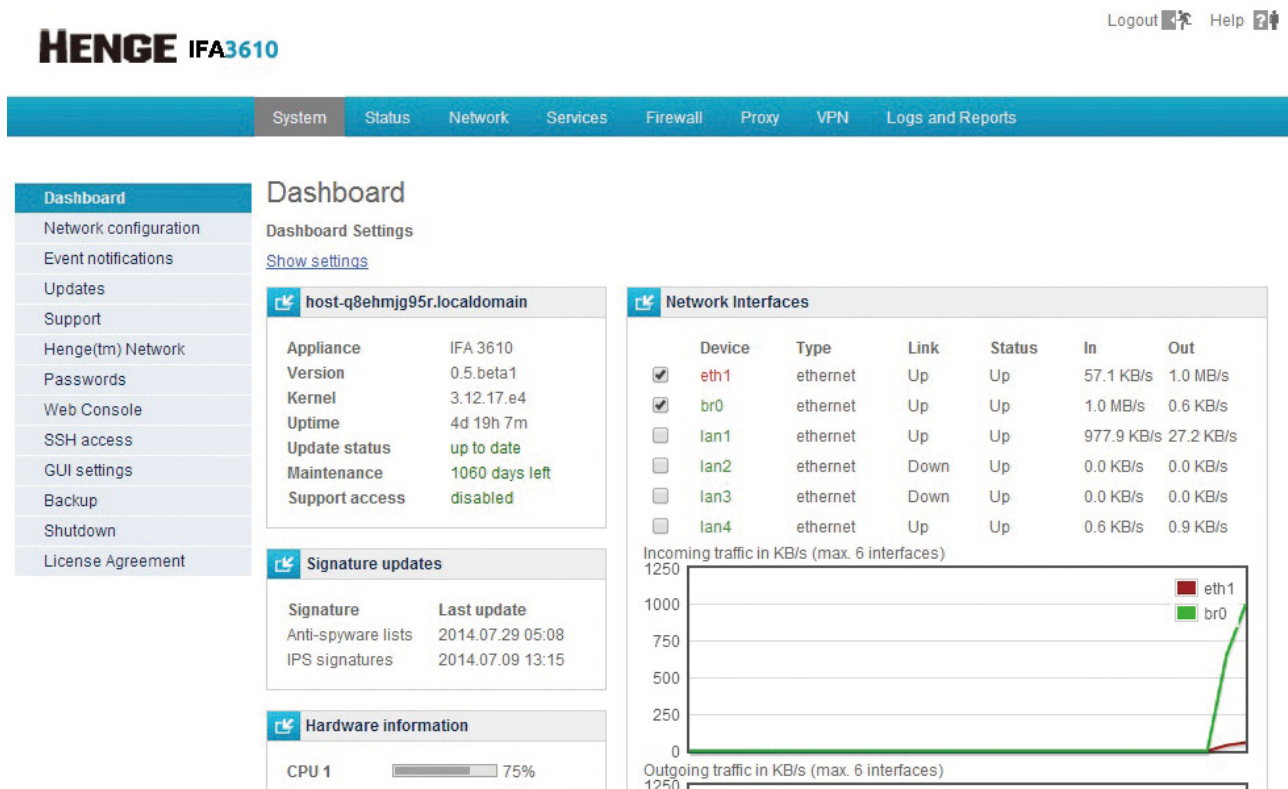
For the appliance to correctly operate, it is not necessary to configure the ORANGE and BLUE zones. Indeed, it suffices to define the GREEN zone, since also the RED zone can be in some cases left unconfigured.

The appliance has pre-defined firewall rules that forbid the network traffic to flow between some of the zones. Besides the four main zones, two more zones are available, but are used only in advanced setups: The OpenVPN clients zone (sometimes called PURPLE), and the HA zone. These are two special zones that are used as networks for the OpenVPN remote users that should connect to the IFA 3610 appliance and for the HA service. By default, they use the 192.168.15.0/24 and 192.168.177.0/24 networks respectively, so those networks ranges should not be used in the main zones, especially when planning to use either of these services. Indeed, those networks would overlap, possibly causing undesirable effects. The IP ranges of these two zones can however be modified during the set up of the OpenVPN or HA services.

To each zone corresponds an (*network*) interface and an IP address. The *interface* is the (ethernet or wireless) port through which the network traffic flows to the zone, so *RED interface* it the port through which you can reach the RED zone and the Internet. The IP address of the interface is the *<Zone>IP*. For example, the factory setting for the GREEN zone is the 192.168.0.15/24 network, hence the GREEN interface will have IP 192.168.0.15, which is referenced to as the GREENIP.

The IFA 3610 Appliance Management Interface

The GUI of the IFA 3610 appliance has been designed to be easy to use, and consists of five main parts: The header, the main menubar, the sub-menu, the main area, and the footer. A sample screenshot of the *Service module* can be seen below.



The screenshot displays the HENGE IFA3610 management interface. At the top right, there are links for "Logout" and "Help". Below the header is a navigation menu with options: System, Status, Network, Services, Firewall, Proxy, VPN, and Logs and Reports. The main content area is divided into several sections:

- Dashboard Settings:** Shows system information for host-q8ehmjg95r.localdomain, including Appliance (IFA 3610), Version (0.5.beta1), Kernel (3.12.17.e4), Uptime (4d 19h 7m), Update status (up to date), Maintenance (1060 days left), and Support access (disabled).
- Signature updates:** Lists signature updates for Anti-spyware lists (last updated 2014.07.29 05:08) and IPS signatures (last updated 2014.07.09 13:15).
- Hardware information:** Shows CPU 1 usage at 75%.
- Network Interfaces:** A table listing network interfaces with their status and traffic rates.

| Device | Type | Link | Status | In | Out |
|--|----------|------|--------|------------|-----------|
| <input checked="" type="checkbox"/> eth1 | ethernet | Up | Up | 57.1 KB/s | 1.0 MB/s |
| <input checked="" type="checkbox"/> br0 | ethernet | Up | Up | 1.0 MB/s | 0.6 KB/s |
| <input type="checkbox"/> lan1 | ethernet | Up | Up | 977.9 KB/s | 27.2 KB/s |
| <input type="checkbox"/> lan2 | ethernet | Down | Up | 0.0 KB/s | 0.0 KB/s |
| <input type="checkbox"/> lan3 | ethernet | Down | Up | 0.0 KB/s | 0.0 KB/s |
| <input type="checkbox"/> lan4 | ethernet | Up | Up | 0.6 KB/s | 0.9 KB/s |
- Traffic Graphs:** Two line graphs showing incoming and outgoing traffic in KB/s for the last 6 interfaces. The incoming traffic graph shows a sharp increase for eth1 and br0, while the outgoing traffic graph shows minimal activity.

The Footer

Status: Connecting... main **Uptime:** 00:06:53 up 6 min, 0 users, load average: 1.73, 1.28, 0.71

Nexcom Appliance release 0.5.alpha1 (Deployset #0) (c) NEXCOM International Co., Ltd.

The footer is placed at the very bottom of the page. It consists of two lines of text with a few information on the running IFA 3610 appliance. The top line shows (**Status:**) whether an uplink is connected or connecting and which one (if there are more than one uplinks defined) and the time elapsed (**Uptime:**) since the last time the connection was established and the uptime of the machine, which is reported as the output of the **uptime** command, i.e., the time since last boot, the number of users and the load average. When you change page, the information are updated. The bottom line shows the version of the appliance with the deployset, and the copyright.

The Main Navigation Bar



The main navigation bar, situated right below the header, is a menu bar with a black background and a green bottom line that displays all the available sections of the IFA 3610 appliance. When clicking on one of the modules (e.g., *Services*), its background becomes green, to emphasise the current open module. Upon clicking on a menu item, the sub-menu on the left of the page and the title at the top of the main area change, since they are context-dependant. By default, the GUI opens on the *System* menu.

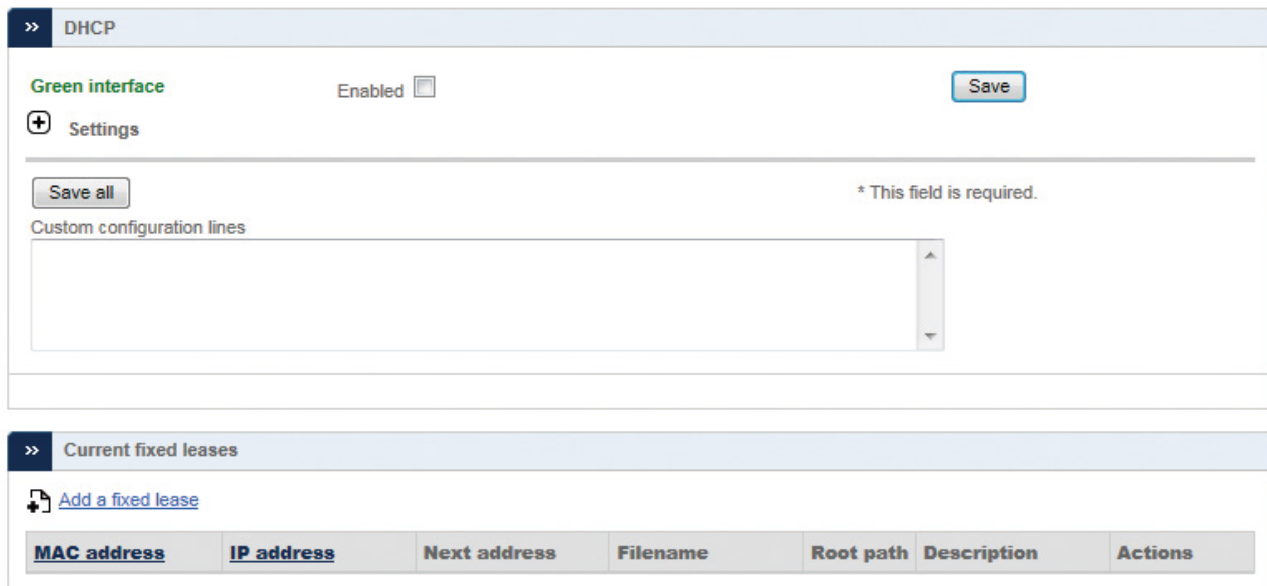
The Sub-Menu

| |
|----------------------|
| DHCP server |
| Dynamic DNS |
| Time server |
| Intrusion Prevention |
| High availability |
| Traffic Monitoring |
| SNMP Server |
| Quality of Service |

The sub-menu appears on the left-hand side of the GUI and changes depending on the module selected on the menubar. It appears as a vertical list of items that can be clicked to change the content of the main area and to access all the functionalities included in that IFA 3610 appliance's module.

The Main Area

DHCP configuration



>> DHCP

Green interface Enabled Save

+ Settings

Save all * This field is required.

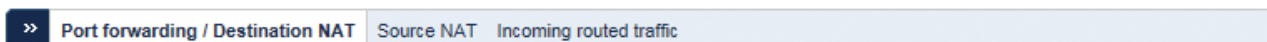
Custom configuration lines

>> Current fixed leases

+ Add a fixed lease

| MAC address | IP address | Next address | Filename | Root path | Description | Actions |
|-------------|------------|--------------|----------|-----------|-------------|---------|
|-------------|------------|--------------|----------|-----------|-------------|---------|

The main area contains all the information and settings encompassed by the current selection of the menu/sub-menu combination. Some of the pages (e.g., the Dashboard or parts of the *Service* and *Logs* modules) are simply informative, showing the current status of the appliance either graphically or textually, in the latter case conveying the output of linux commands on the screen. The vast majority of the pages, however, shows a table containing various information about the current configured settings, allowing to modify or delete existing items and settings and to add new ones. Particularly elaborate services like e.g., the HTTP proxy or the firewall, contain so many configuration options that a single page does not suffice to present them all, so the available settings are grouped together and organised in tabs.



>> Port forwarding / Destination NAT Source NAT Incoming routed traffic

Within tabs, often the configuration options are packed in one or more boxes, that gather together settings that refer to a common part of the overall configuration.

The Icons

Many icons are used throughout the pages served by the IFA 3610 appliance to denote either an action that can be quickly carried out, or convey some meaning to the settings shown.

Switches




Switches are used to entirely enable or disable a service and are present on the top of the main area. The gray switch suggests that the service is disabled and inactive, with the main area showing no settings or configuration options. Upon clicking on it, the service and the daemons that are necessary for its proper functioning are started and initialised. After a few seconds, the switch's color turns azure and all the configuration options available will appear. To disable the service, click again on the switch: This causes all the daemons to be stopped, the switch to turn grey, and the settings to disappear.

Policies

These icons are found in those services that require some form of access policies or traffic control, like, e.g., firewall rules or proxy specifications. Whenever a packet matches a rule, the policy specified for that rule is applied, determining if and how the packet can pass or not.



 accept the access with no restriction.


 allow the access but only after the packets have positively passed the IPS. This policy is only available in firewall rules.

-  blocks the packets and discards it.
-  blocks the packets, but a notification is sent to the source.
-  partial accept the rules. This is only found on the heading of a list of policies, to give at a glance the idea that some of the policies in the list are accepted and some are rejected, like e.g., in *Menubar* ► *Proxy* ► *HTTP* ► *Contentfilter*.

Other Icons

Additional icons that can be found on the appliance.

-  expands a panel, revealing its content.
-  closes a panel, hiding its content.



Navigation bar






In most places where a long list of item appears, a navigation bar appears to ease the listing of the items, which is composed of several cells: First | ◀ and Previous ◀ on the left, Next ▶ and Last ▶ | on the right, which enclose the page number. Clicking on these buttons will lead to either the first or last page, or to the previous or next page.

Common Actions and Tasks

There are two types of actions that can be performed within the GUI: Actions on a single item in a list of configuration settings (i.e., one firewall rule), and 'global' actions to save, store, and apply all the settings in a list, a box, or a page.

Actions and icons

These icons are placed in the *Actions* column on the right of the various tables that appear on the pages and usually show a list of the items defined, like e.g., the firewall rules or the OpenVPN users. The actions' icons allow to execute one task on the element of the list to which they correspond. Some action is only available on some type of lists:

- and indicate the status of an item, enabled and disabled respectively. You can change the status by clicking on the icon. After that, a callout may notify you to restart service, if this is needed, to let the daemons reload the configuration and activate the changes.
- ↑ and ↓ are available only in lists where the order is important, e.g., firewall rules, and allow to modify the order by moving up ↑ or down ↓ the corresponding item.
-  allows to modify the current item. Clicking on this icon will open the appropriate editor for that item.
-  causes the selected item to be removed from the list and from the configuration. A message will appear, asking for confirmation before the item is definitely deleted.
-  allows to download the item (usually an archive).
-  is used in limited locations, e.g., in *Menubar* ► *Services* ► *Spam Training* to test the connection of an item to a remote servers.
-  appears in the IPS (*Menubar* ► *Services* ► *Intrusion Prevention*) and allows to log the packets that are allowed to pass or are blocked after they have matched a rule.

'Global' Actions

At the bottom of every page that allows the customisation of one or more options, there is the option to **Save** and store the new configuration on disk or to **cancel** the customisation done so far. In the latter case, no further action is required, since the configuration did actually not change. In the former case, however, it proves necessary to restart the service just modified, and perhaps also a few other related or dependant services, for the new settings to be reloaded and used in the running configuration. For the sake of convenience, when this action is required, a callout is displayed after the settings have been saved, with an **Apply** button, to be clicked to restart the service.

Whenever a *Multiselect box* is used (e.g., in *Menubar* ► *Hotspot Settings*), **Add all** and **Remove all** can be clicked as shortcut to add or remove all the available entries from the list of the available items or the selected and active items, respectively.

Multiple entries in one configuration option

In several places, several values can be entered for a single configuration item, for example the source or destination of a firewall rule. In those cases, either a textarea or a drop-down menu is shown. In the former case it is possible to enter one value per line, like e.g., it a MAC address, a network range (in CIDR notation), or an OpenVPN user. In the latter case, the choice is limited among a number of predefined values, that can be selected by holding the Control key on the keyboard and clicking on the values to be selected.

IPv4 and CIDR notation.

An IPv4 address is a network address whose length is 32 bits, divided in four, 8-bits long octets. In decimal, each octet can assume any value between 0 and 255 ($2^8 = 256$).

When specifying a network range, the IP address of the first host on the network along with the subnet mask, or netmask for short, is given, which defines the number of hosts available in that network. The subnet is defined as the length of the network prefix, i.e., that part of the address shared by all the hosts in a network.

There are two possibilities to denote the network/netmask pair:

- Explicitly, i.e., both are given in quad dotted notation. For example:

```
network 192.168.0.0
netmask 255.255.255.0
```

This is a network starting at the address 192.168.0.0 with 256 host available, i.e., the network range from 192.168.0.0 to 192.168.0.255. The first three octet in the netmask are 255, showing that there are no free host (or that this part of the address is the network prefix), while the fourth is 0, meaning that all hosts ($256 - 0 = 0$) are available.

- In CIDR notation, a more compact way to show the network range, in which the free bits instead of the free hosts are given. The same network range as above is expressed as:

```
192.168.0.0/24
```

This notation shows the length in bits of the shared part of the IP address. 24 means that the first three octets (each consisting of 8 bits) are shared, while the fourth octet is free, giving a number of free hosts that is equivalent to $32 - 24 = 8$ bits, i.e., 256 hosts.

The same line of reasoning can apply to an IPv6 address, with the only difference that IPv6 addresses are 128 bits long.

Accessing the IFA 3610 Appliance

There are several ways to access the IFA 3610 appliance: The most intuitive and straightforward one is from the web-based GUI. There are also console-based access via SSH and serial console, although they are suggested to advanced users only.

There are several ways to access the IFA 3610 appliance: The most intuitive and straightforward one is from the web-based GUI. There are also console-based access via SSH and serial console, although they are suggested to advanced users only.

The IFA 3610 Appliance GUI

Hint: The default IP address of the IFA 3610 appliance is 192.168.0.15.

The recommended access to the IFA 3610 appliance GUI is very simple: Start the browser and enter the GREENIP address, whether or not this is the first time the IFA 3610 appliance is used.

The browser will be redirected to a secure HTTPS connection on port 10443. Since IFA 3610 appliance uses a self-signed HTTPS certificate, the browser might ask to accept the certificate during the first connection. The system will then ask for username and password. Specify "admin" as the username and provide the password received from the reseller or, if the IFA 3610 appliance has already been customised, insert the password that provided during the installation.

Console-based access

Console-based access to the IFA 3610 appliance is suggested only to users that are acquainted with the Linux command line.

Two possibilities are available to reach the **CLI**: Using SSH access or via serial console. SSH access is by default disabled, but can be activated under *Menubar* ► *System* ► *SSH access*, while Serial Console access is enabled by default on all appliances with the following parameters:

- port: ttyS0
- bit, parity bit, stop bit: 8, N, 1
- speed: 115200 baud in newer appliances.

The connection using the serial console requires:

- A suitable terminal program like minicom for Unix/Linux boxes or *putty* for MS Windows.
- A workstation with a serial interface
- A nullmodem cable to connect a workstation to the appliance

or

- Terminal program.
- Networked Serial-to-Ethernet adapter.
- Serial-to-Ethernet cable to connect the appliance to the adapter.

Note: In case the network is not configured properly, the serial console may represent the only way to access the IFA 3610 appliance.

Chapter 1: The System Menu

The System menu provides several information about the IFA 3610 appliance and its status, and allows to define the network setup and some access modalities (e.g., via SSH or for the NEXCOM support).

The sub-menu on the left-hand side contains the following items, which allow for some basic administration tasks and to monitor the running activities of the IFA 3610 appliance.

- Dashboard - overview of the system and of the connections status
- Network configuration - network and network interface configuration
- Event notifications - set up of notification via e-mail
- Updates - management of system updates
- Support - support request form
- Henge™ Network - Henge™ Network registration information
- Passwords - set system passwords
- Web console - a console shell on the browser
- SSH access - enable/configure SSH access to the IFA 3610 appliance
- GUI settings - web interface language settings
- Backup - backup or restore IFA 3610 appliance settings as well as reset to factory defaults
- Shutdown - shutdown or reboot the IFA 3610 appliance
- Credits - acknowledgement to all contributors
- License Agreement - a copy of the User License Agreement

The remainder of this section will describe the various parts that compose the System menu items.

1.1 Dashboard

The Dashboard is the default page, the one that is displayed upon every login. It encompasses several boxes (“plugins”) organised in two columns that provide a complete overview of the running system and of its health. The top of each box reports the name of the box. The Dashboard has lately undergone some changes in its usability and new features have been added to improve the interaction with the user. The information visible on screen are updated at regular intervals.

The available plugins and the information they display are described here.

1.1.1 System Information Plugin

It shows several information about the installed system. It usually presents the hostname and domain name of the appliance in the title.

Appliance: The appliance type.

Version: The version of the firmware.

Kernel: The current running kernel.

Uptime: The time since the last reboot.

Update status: A message depending on the appliance status:

- “up to date”. No updates are available.
- “update required”. New packages can be installed: A click on the message leads to the *Updates* page where it is possible to review the list of new packages.
- “Register for enterprise”. The system has not yet been registered to *Henge™ Network*: A click on the message will open the *Henge™ Network* page, in which to compile a form to complete the registration.

Maintenance: The remaining days of validity of the maintenance support.

Support access: Whether the support team can access the IFA 3610 appliance or not. In the former case, it is also shows the date until the access is granted.

This plugin also shows the remaining days of validity of the additional modules *Panda Antivirus* and *CommTouch*, if purchased.

1.1.2 Hardware Information Plugin

It shows the main hardware information of the appliance and the resource availability. All the information are provided with the absolute value (graphically with a small bar and in number at the end of a line) and the percentage of their use. The only exception is the CPU load, which shows only the percentage of use, in graphic and numbers.

CPU x: The load of the CPU, where x represents the CPU number, for those appliance that have more than one CPU.

Memory: The amount of the RAM memory used.

Swap: How much swap disk space is used. A high percentage here usually means there is something not working correctly.

Main disk: The usage of the root partition.

Temp: The space used in the /tmp partition.

Data disk: the usage of the /var partition.

Configuration disk: The space occupied by the partition containing all the IFA 3610 appliance services and settings.

Log disk: The amount of space used in the partition containing the logs.

The latter values, showing disk space availability, can vary depending on the appliance, since the data, system, and log partitions may be located in different places.

Warning: A partition on the hard disk (e.g., main disk, data disk, /var/log) shall never have a usage of 95% or more, as this can cause malfunctioning and data loss.

1.1.3 Service Information Plugin

Information about the most important services installed on the IFA 3610 appliance, along with their actual status, are displayed by this plugin. For each service is shown the status, either ON or OFF, and a summary of the tasks accomplished during the last hour and the last days. A click on the service's name expands or collapses additional information on the tasks carried out by the service. For running services, there is the possibility to open in a new window the respective *Live Logs*. Hence, if some number in the summaries sounds strange (e.g., a number of email rejected that is twice as normal) or not common compared to the normal activities (e.g., the IDS has detected some attack), the logs can be controlled to search for some useful message that has been recorded. The services currently supported by this plugin are:

Intrusion Detection: The number of attacks logged by snort.

SMTP Proxy: Statistics about the e-mails processed. The number of e-mail currently in the postfix-queue, of the received e-mails and how many of them were clean, the number of viruses found, and how many e-mails were blocked.

HTTP Proxy: The numbers of cache misses and hits of squid and of the viruses found.

POP3 Proxy: Statistics about the received, blocked, and virus-containing e-mails that went through the POP3 Proxy.

Hint: Inactive services are marked with a red *OFF* message.

1.1.4 Network Information Plugin

It shows information about the network interfaces of the firewall and the traffic. The upper part of this plugin shows several data about the network interfaces of the appliance: Their name, type, link (*Up* if a connection is established, *Down* otherwise) and status (*Up* if the device is activated, *Down* if not), and the In- and Outgoing traffic. The latter two data are updated in real-time. When ticking the checkbox near the device name, that device is shown in the graphs underneath. The devices' name is coloured according to the zone they serve.


The lower part of the plugin contains two charts: The first one shows the incoming traffic, while the second one the outgoing traffic on each of the interfaces chosen. The traffic of each interface is coloured according to the zone it belongs to, different interfaces serving the same zone have different nuances. Bridges built on one device are shown in the same colour as the device. Like the traffic data in the upper part, both charts are updated in real-time.

Hint: Up to six interfaces can be selected and shown in the charts.

1.1.5 Signatures Information Plugin

This plugin shows information about the actual status of those services requiring the download of signatures that are installed and enabled on the appliance. In case no signature has been downloaded and no service has already been enabled, the message *No recent signature updates found* is displayed, otherwise the plugin presents the signatures installed for the various daemons and the timestamp (date and time) of the last download. The list includes the signatures for the anti-spyware, antivirus, contentfilter, and intrusion prevention services.

1.1.6 Uplink Information Plugin

This plugin shows a table detailing the uplinks' connection status. For each defined uplink are shown name, IP address, status, uptime, whether it is active or not , managed or manual . The circular arrow , when clicked, allows to immediately reconnect the corresponding uplink. Of particular interest is the *Status* field of each individual uplink, which can be:

Stopped: Not connected.

Inactive: Not connected.

Connecting: Not yet connected, but a connection is ongoing.

Connected or UP: The connection has been established and it is fully operational.

Disconnecting: The uplink is closing the connection. The appliance keeps pinging the gateway and announces when it becomes available.

Failure: There was a failure while connecting to the uplink.

Failure, reconnecting: There was a failure while connecting to the uplink, but the IFA 3610 appliance is now trying again.

Dead link: The uplink is connected, but the hosts that were defined in the uplink configuration (*Menubar ► Network ► Interfaces*, option *Check if these hosts are reachable* in the *Uplink editor*) to check the connection could not be reached. In other words, the uplink is not operational.

Managed and manual uplink.

Each uplink can be operated in either managed mode, which is the default, or manual mode. In managed mode, the appliance monitors and restarts the uplink automatically when needed. If managed mode is disabled, the uplink has to be activated or deactivated manually: This implies that there will be no automatic reconnection attempt if the connection is lost, but clicking on **Reconnect** is required to restart a nonoperational uplink. The management mode of an uplink can be selected under *Menubar ► Network ► Interfaces*.

While an uplink should always be managed to allow for a quick reconnection in case of a connection loss, the manual mode proves useful for troubleshooting or testing connections before actually establishing them.

1.2 Network Configuration

The configuration of the networks and of the network interfaces serving the zones is fast and easy with this 8-step wizard. It is possible to freely navigate back and forth the step, using the <<< and >>> buttons and even decide at any moment to cancel the actions done so far. Only at the last step it is required to confirm the new settings: In that case, all the changes made will be applied. Note that while applying the new settings, the web interface might not respond for a short period.

The 8 steps in which the wizard is divided are:

1.2.1 1/8 - Choose Type of RED Interface

At installation time, the appliance receives a default GREEN IP. This screen allows to choose the type of the RED interface (i.e., the type of uplink) among those supported by the appliance.

ETHERNET STATIC

The RED interface is in a LAN and has fixed IP address and netmask, for example when connecting the RED interface to a simple router but with the convenience that the appliance be always reachable at the same IP address.

ETHERNET DHCP

The RED interface receives its network configuration via (dynamic) DHCP from a local server, router, or modem, i.e., the RED interface is connected to a simple router but without the need to have a fixed address.

PPPoE

The RED interface is connected to an ADSL modem. This option is only needed when the modem uses bridging mode and requires to use PPPoE to connect to the provider. This option should not be confused with the *ETHERNET STATIC* or *ETHERNET DHCP* options, used to connect to ADSL routers that handle the PPPoE themselves.

ADSL (USB, PCI)

The RED interface connects to an ADSL modem via a USB or PCI cable, not via an Ethernet one.

ISDN

The RED interface is an ISDN connection.

ANALOG/UMTS Modem

The RED interface is an analog (dial-up) or UMTS (cell-phone) modem.

GATEWAY

The appliance has no RED interface. While this represents an unusual situation, since a firewall normally should have at least two interfaces, this configuration may be suitable for some special scenarios, like for example when only some specific services of the appliance are needed. Another, more sophisticated example is a scenario in which the BLUE zone of an appliance is connected through a VPN to the GREEN interface of a second appliance. The second firewall's GREEN IP address can then be used as a backup uplink on the first firewall. If this is the case, a default gateway shall be configured later on.

A small box recalling the number of network interfaces available on the system is shown to the right of the available choices. The RED interface can be fully configured during step 4.

1.2.2 2/8 - Choose Network Zones

The appliance separates the networks connected to it into four main zones, as described in this section. At this point the two most important zones - *GREEN* and *RED* - have already been encountered during the installation: This step allows to enable one or two additional zones, depending on the services that should be provided by the appliance: *ORANGE* -used as the DMZ network portion- and *BLUE* -used as segment for wireless clients. Their full configuration will be possible in the next step.

1.2.3 3/8 - Network Preferences

This step concerns the configuration of the GREEN zone, if needed, and of any zone chosen in the previous step. For each of the zones enabled, the following options can be configured:

IP Address

The IP address (such as 192.168.0.1) of the interface, which should not be already in use in the network.

Hint: Good practice suggest that the last octet be 1, since the interface will gather the traffic of the whole subnet.

Remember also that a change in the IP addresses of an appliance, especially in a production environment, might require to adjust additional settings elsewhere, for example the HTTP proxy configuration in the workstations, otherwise the web browsers will not work correctly.

Warning: When configuring the interfaces of the GREEN zone, make sure to not remain locked out of the web interface! This situation may occur for example when changing the GREEN IP address into one that is not reachable from the current GREEN segment and then saving the settings. In this case the only access to the appliance is via serial console.

Network Mask

Define the network mask from a drop-down menu containing the possible masks (e.g., /24 - 255.255.255.0).

Hint: All the devices connected to the same subnet shall have the same netmask to communicate properly.

Additional Addresses

Additional IP addresses for different subnets can be added to the interface here.

Interfaces

Map a network interface to a zone, with the following rules:

1. Each interface can be mapped to only one zone and each zone must have at least one interface.
2. When more than one interface is assigned to a zone, these interfaces will be bridged together and act as if they were part of a switch.

For each available interface these information are shown:

- A colored checkbox, showing which zone the interface serves. No color means that the interface is not assigned to any zone.
- *Port*, the number of the port.
- *Link*, shows the current status by means of icons: link ✓ -the link is active, ✘ -no link or no cable plugged in, ? -no information from the driver.
- *Description*, the interface's PCI identification string, as returned by **lspci**. The string is trimmed, but it can be shown by moving the mouse on the ?.
- *MAC*, the interface's MAC address.
- *Device*, the logical name of the device.

Note: Internally, the appliance handles all zones as bridges, regardless of the number of the assigned interfaces. Therefore, the Linux name of the interfaces is **brX**, not **ethX**.

Finally, the system's host name and domain name can be set in the two text boxes at the bottom of the screen.

Private IP Addresses

It is suggested to follow the standard described in RFC 1918 (which has been recently been updated by RFC 6761) and to use for the zone's setup only the IP addresses contained in the network segments reserved for private use by the IANA, which are:

10.0.0.0 to 10.255.255.255 (10.0.0.0/8, 16,777,216 addresses)
 172.16.0.0 to 172.31.255.255 (172.16.0.0/12, 1,048,576 addresses)
 192.168.0.0 to 192.168.255.255 (192.168.0.0/16, 65,536 addresses)

This choice avoids incurring in DNS resolution errors, as IP addresses not falling within these ranges are likely to have been reserved by other organisations as their public IPs. Moreover, different IP ranges must be used in the different network segments for each interface, for example:

IP = 192.168.0.1, network mask = /24 - 255.255.255.0 for GREEN
 IP = 192.168.10.1, network mask = /24 - 255.255.255.0 for ORANGE
 IP = 10.0.0.1, network mask = /24 - 255.255.255.0 for BLUE

Note also the first and the last IP address of a network segment (which are usually .0 and .255) are reserved as the network address and the broadcast address respectively, and must not be assigned to any device.

1.2.4 4/8 - Internet Access Preferences

This step allows the configuration of the RED interface chosen in step 1, that connects to the Internet or to any other untrusted network outside the appliance.

Depending on the type of the selected RED interface, different configuration options will be available, as required by each interface type. At the bottom of the page appear two options that are commonly available, namely *MTU* and *Spoof MAC address with*, described below, and the choice of the DNS resolver, available for almost all interface types, which is wither *Dynamic* or *Manual*: In the latter case, one valid IP address of a DNS server must be provided manually in the next step. The other configuration options are:

ETHERNET STATIC

The IP address and network mask of the RED interface, as well as the IP address of the default gateway, that is, the IP address of the gateway that connects the appliance to the Internet or to another untrusted network. Optionally, the Ethernet hardware address (MAC address) of the interface can be specified.

ETHERNET DHCP

Only one available option, namely the DNS choice.

PPPoE

To configure PPPoE, fill in the form with the username and password assigned by the provider, and the authentication method. Optionally, the provider's service and concentrator name can be configured, though this is usually not needed.

Hint: If unsure whether to select PAP or CHAP authentication, keep the default option.

ADSL (USB, PCI)

There are 3 sub-screens for this choice.

- In the first one, select from the drop-down menu the appropriate driver for the modem, among the possibilities offered.
- In the second one, choose the *ADSL type* from the drop-down menu among the four choices: PPPoA, PPPoE, static IP, or DHCP.
- Finally, depending on the selection made in the previous two steps, some of the following settings are required, which can be asked to the ADSL provider:
 - *VPI/VCI numbers* and the encapsulation type
 - the *username* and *password* assigned by the provider and the *authentication method* (if unsure, keep the default PAP or CHAP)
 - the IP address and network mask of the *RED* interface,
 - the IP address of the *default gateway* (required for static IP only);

Note: If PPPoE was chosen at point 2. above, then the configuration is exactly like explained in the previous paragraph, PPPoE.

ISDN

To configure the ISDN connection, the modem driver, phone numbers (the provider's number and the number used to dial out), as well as the username and password that have been assigned by the provider, and the authentication method are needed (if unsure, keep the default PAP or CHAP). Also specify whether the IP address of the DNS should be assigned automatically or set manually.

ANALOG/UMTS Modem

While the appliance supports most modern UMTS modems, some care is required when using them in conjunction with the appliance. On one side, some UMTS modems are USB mass storage devices as well and usually register two devices (e.g., `/dev/ttyUSB0`, `/dev/ttyUSB1`): In this case the first device `/dev/ttyUSB0` is the modem, the second one is the storage. These types of modem can cause problems when restarting the firewall because the appliance tries to boot from the USB mass storage device. On the other side, some SIM cards require a personal identification number (PIN) to work, but this is not supported. To allow those cards to work with the appliance, the PIN should be removed from the card.

There are 2 sub-screens for this choice.

1. In the first one, specify to which serial port the modem is connected to and whether it is an analog modem or an UMTS/HSDPA modem.

Hint: The `/dev/ttyS0` device is reserved for the serial console and is therefore not available as port for modems.

2. In the second one, configure the modem's bit-rate, the dial-up phone number or access point name, the username and password that have been assigned by the provider and the authentication method (if unsure, keep the default PAP or CHAP). For UMTS modems it is also necessary to specify the access point name.

GATEWAY

The IP address of the default gateway - that is, the IP address of the gateway that connects the appliance to the Internet or another untrusted network.

The common options are:

MTU

The MTU size of the packets send over the network.

Spoof MAC address with

Specify a custom MAC address for the RED interface. This setting is required for the proper failover of slave devices in an HA setup. See *High availability* for more information about the RED address in HA setups.

The MTU Size

While the vast majority of the ISPs uses a standard value of 1500 bytes, in some circumstances the standard MTU size results too high. If that happens, some strange network behaviours will be noticed, like e.g., downloads which always stop after a while or connections which will not work at all.

If the ISP does not use a standard MTU size, it is easy to discover the correct one, by sending special ICMP packets with a specific value, that can be lowered until no errors are encountered: At theist point, the MTU size is correct and this value should be entered in the configuration options.

In order to send the icmp packets do the following:

Log in to the EFW and choose a host which can be actually reached (e.g., the ISP's DNS, which should always be reachable) and ping that host with the following command:

ping -c1 -M do -s 1460 <host> (please refer to the **ping(8)** manpage for more info).

If the MTU size 1460 is correct, ping replies like the following one are received:

```
PING 10.10.10.10 (10.10.10.10) 1460(1488) bytes of data.
1468 bytes from 10.10.10.10: icmp_seq=1 ttl=49 time=75.2 ms
```

If however the current MTU size is still too big for packets of the size 1460, an error message like this will appear:

```
PING 10.10.10.10 (62.116.64.82) 1461(1489) bytes of data.
ping: sendmsg: Message too long
```

Retry with different packet sizes (i.e., the value after the -s option), until the correct size has found and no error is displayed. The value shown within brackets in the ping command's output is the MTU size. In this example the output is 1460(1488), therefore 1488 is the value to select for the MTU size.

An MTU value lower than 1500 may cause problems also in the *OpenVPN* setup and require to adjust some setting there.

1.2.5 5/8 - Configure DNS resolver

This step allows to define up to two IP addresses for the DNS server, unless they are assigned automatically: In this case, no configuration option can be set and it is safe to move to the next one. If only one DNS server should be used, the same IP address must be entered twice. The IP address(es) of the DNS must be accessible from the IFA 3610 appliance, otherwise URL and domain resolution will not work.

See also: Changes to the RED interface, i.e., the uplink, and the DNS server can be modified later, separately from the other network configuration:

Uplink Editor

Menubar ► Network ► Interfaces ► [edit uplink]

1.2.6 6/8 - Configure Default Admin Mail

The configuration of a global administrator e-mail address that will be used by all services to send e-mails, is done here. The administrator e-mail address is then used for notifications, in case of problems or emergencies. These email addresses will be used by the *Event notifications*.

There are three fields to configure.

Admin email address

A valid e-mail address to which the system e-mails should be sent.

Sender email address

A valid e-mail address that appears as the sender address. A custom sender address proves useful if the recipient wants to filter messages sent by the appliance.

Address of smarthost

The SMTP server through which the email should be sent.

Hint: Although all the fields may be left blank, it is suggested to supply at least one valid *Admin e-mail address*.

1.2.7 7/8 - Apply configuration

This step informs that the network setup is now finished and all the new settings have been gathered. Clicking on the **OK, apply configuration** button will save the settings and apply the configuration by restarting all the necessary services and daemons.

1.2.8 8/8 - End

In the last step, all the configuration files are written to the disk, all the devices are reconfigured and the network-dependent services and daemons (e.g., the firewall and ntpd) are restarted as necessary. The whole process may take up to 20 seconds, during which the connection to the administration interface and through the appliance may not be possible.

The administration interface will then reload automatically. If the GREENIP address has changed, the GUI will be reloaded at the new IP address. In this case or in case the hostname changed, a new SSL certificate is generated to identify the new host.

Note: To change later only some of the settings in the network configuration (e.g., the hostname or the network range of a zone), simply start the network configuration, skip all the steps until the one in which to make the desired changes, edit the appropriate values, then proceed to the last step and finally save.

1.3 Event Notifications

Whenever some critical event takes place on the appliance (e.g., a partition is filling up, or there are updates available), there is the option to be immediately informed by e-mail about it and to promptly take some actions to solve a problem, if required.

1.3.1 Settings

The default tab serves for the configuration of the email notification:

Email notifications

Select from a drop-down menu how to use the notification system. Available options are:

- notify using default email address: the default administrator e-mail address (as specified in the Installation wizard or in step 6 of *Menubar ► System ► Network configuration*)
- notify using custom email address: an alternate e-mail address to which the notification e-mail shall be sent. In this case, three more options must be configured, namely:

Mail sender address

The e-mail address that appear as the sender of the e-mail.

Mail recipient address

The e-mail address to which the e-mail will be delivered.

Mail smarthost

The SMTP server that will be used to send the notification e-mail.

- do not notify: no notifications will be sent

1.3.2 Events

This tab shows a list of all the events that can produce a notification message and allows to configure the actions to be done when each of the events takes place. Right above the list there is a small navigation bar and a search field: The latter can be used to filter only the relevant items.

The list contains three columns:

ID



The 8-digit ID ABBCCCCD code of the event, which is built as follows:

- A represents the layer number, i.e., in which system's component the event has taken place: 1 means kernel, 2 the system itself, 3 services, 4 configlayer, and 5 the GUI.
- BB is the module number
- CCCC is a sequential number assigned to the event
- D is the severity of the event, i.e., the degree of badness of the event. The lower the number, the worst the severity: 0 is a critical event, 4-5 neutral, 9 is a positive event.

Description

A short description of the event.

Actions

The actions that can be performed for each event. All e-mail notifications are enabled by default (this is shown by the  icon), but to disable notifications for one event, click on the mail icon in that event's row (this causes also the icon to change into ). To later re-activate the notification, it suffices to click again on the icon. After changing an action, remember to click on the **Apply** button that appears within the green callout above the events' list.

1.4 Updates

The management of the software updates is done from here. It is possible at any time to manually check for available updated packages, or to schedule a periodic check.

In this page there are two boxes: One with the current status of the system and one to schedule a routine check for updates.

1.4.1 Status

The Status box informs whether the system needs updates or not. In the former case, a list of available packages is presented, while in the latter the message "Your appliance is **up to date!**" is displayed. Moreover, additional messages inform of the last date and time when a check for updates and the last upgrade have been carried out. These options are available:

Check for new updates

A manual check for updated packages is started, and any upgradable package found is listed here. Individual packages can be chosen from the list and installed.

Start update process NOW

The update process is launched: The system downloads the updated packages which are then installed, replacing the old ones.

Note: In order to check for updates, a valid maintenance is required, otherwise no update will show up, even if available.

1.4.2 Schedule for Retrieving the Update List

The *Schedule* box allow to set up a periodic job, governed by the **cron** daemon, that retrieves the list of updated packages. The available, mutually exclusive, options are *Hourly*, *Daily*, *Weekly*, and *Monthly*. Moving the mouse over the small ? next to each option shows a tool-tip with the exact time at which the job will run.

1.5 Support

In this page it is possible to manage requests for assistance to the Henge™ support.

Note: To be able to submit a support request, the system must be registered to the Henge™ Network. If not, the "Currently no running maintenance available." message will be displayed.

If the system is not registered, support request can be made to one of the several forums or mailing lists enumerated in the *NEXCOM web sites* section.

The page is divided in two boxes with different purposes: The first one contains a link to open the support's home page, while in the second one it is possible to grant SSH access to the support team.

Visit Support Web Site

This box contains only a hyperlink to the home page of the support.

Please visit our Support Web Site

By clicking on this link, a new tab in the browser will open, where it is possible to find directions on how to fill in an assistance request to the support team.

Access for the Henge™ Support Team

Optionally, access to the firewall can be granted via SSH, a secure, encrypted connection that allows a member of the support staff to log in to the IFA 3610 appliance, verify its configuration and inspect it to find out where the problem lies. The box contains an informative message, the status of the access, which is either DENIED or ALLOWED. When the status is DENIED a button appears at the bottom of the box:

Allow access

Click on this button to grant 4 days of access to the appliance to the support team.

When the support team access is allowed, a new message appears under the status message: Access allowed until: followed by the date and time when access to the appliance will be revoked. Moreover, there are two buttons at the bottom of the box.

Deny access

Immediately revoke the grant to access the appliance.

Extend access for 4 more days

If the support team needs more time to inspect the appliance, a click on this button extends the access grant by four more days.

Note: When enabled, the support team's public SSH key is copied to the system and access is granted via that key. The support team will not authenticate with username/password to the appliance. The root password of the appliance is never disclosed in any way to the support team.

1.6 Henge™ Network

If the appliance has been purchased with a maintenance package, it can be registered and connected to the Henge™ Network, the Henge™ solution for an easy and centralised monitoring, managing, and upgrading of all the registered appliance systems, with just a few clicks. Note that many functionalities of the appliance (e.g., support, sms notification, and so on) require that the appliance be registered to the Henge™ Network.

This page is organised into two tabs, namely *Subscription* and *Remote Access*.

1.6.1 Subscription

If the firewall has not yet been registered to the Henge™ Network, the registration form is shown, that can be filled in before submitting the request for registration. After the registration has been completed, the Subscriptions tab shows three boxes:

System information

Basic data about the appliance: Serial number, activation code, model of the appliance, and the maintenance package chosen.

Registration Status

A summary of the Henge™ Network support status: System name, organisation for which the appliance is registered, system ID, and the date of the last update.

Your Activation Keys

To receive updates from and to participate in the Henge™ Network, at least one valid (i.e., not expired) activation key is required. There is a key for each support channel, but typically just one, shown with the validity time and the days of maintenance left. An expired key is shown by its channel name stricken-through and by the *expired* string in the corresponding *Days left* column.

1.6.2 Remote Access

The Remote Access tab allows to choose whether the appliance can be reached through the Henge™ Network and by which protocol. To allow access, click on the grey switch on  the top of the page: Its color will turn azure, and two access options can be chosen, by ticking the checkbox:

Enable HTTPS access ...

The IFA 3610 appliance can be reached via the web interface.

Enable SSH Access ...

Login via a secure shell to the IFA 3610 appliance is allowed. Activating this option automatically activates the SSH access.

1.7 Passwords

In this page passwords can be changed for each of three default users, by writing each new password twice and then by pressing the corresponding **Change Password** button:

Admin

The user that can connect to the web interface for administration.

Dial

A special user that can only manage uplinks, with a limited interface access. It is not present in recent versions of the IFA 3610 appliance.

Root

The user that can login to the shell for administration. Logins can be made either via the serial console, or remotely with an SSH client.

Hint: Passwords need to be at least 8 characters long.

1.8 Web Console

The web console provides an applet which emulates a terminal within the browser window, that serves as a CLI to carry out administrative tasks.

The functionalities of the web console are the same found upon logging in via serial console or SSH. On the bottom left of the applet, a message shows the status of the console: *Connected* or *Disconnected*. It is possible to exit at any time by typing `exit` in the console and then pressing `Enter` on the keyboard, like in any normal console.

When disconnected, click again on the **Web console** sub-menu item to reconnect. On the bottom right of the applet, two hyperlinks show up:

Enable virtual keyboard.

When clicking on this link, a keyboard applet appears below the console, that can be used to type and execute commands by clicking the mouse on the various *keys*.

Note: When the web console is disconnected, this applet does not communicate with the console.

Disable input

This link toggles the possibility to send input from the keyboard to the web console.

Hint: This option has no effect on the virtual keyboard.

1.9 SSH access

This screens allows to enable remote SSH access to the appliance. This is disabled by default and it is the recommended setting. There are two boxes in the page: *Secure Shell Access Settings* and *SSH host keys*.

1.9.1 Secure Shell Access Settings

The SSH access is activated by clicking on the grey switch . The SSH service is started, and after a few seconds, some configuration options are displayed:

SSH protocol version 1

This is only needed for old SSH clients that do not support newer versions of the SSH protocol.

Warning: The activation of the SSH version 1 is strongly discouraged, since this version is not maintained anymore, deprecated, and contains well known vulnerabilities that could be exploited by malicious users. SSH clients nowadays shall always use version 2 of SSH, which is more secure and reliable.

Allow TCP forwarding

Ticking this option lets other protocols be tunneled through SSH. See *SYS-1* example for a sample use case.

Allow password based authentication

Permit logins using password authentication.

Allow public key based authentication

Logins with public keys are allowed. The public keys of the clients that can login using key authentication must be added to the file `/root/.ssh/authorized_keys`.

Save

Click on this button at the bottom of the box to save the setting of the above four options.

Note: The SSH access is automatically activated when at least one of the following options is true:

- Henge™ support team access is allowed in *Menubar ▶ System ▶ Support*.
- High availability is enabled in *Menubar ▶ Services ▶ High Availability*.
- SSH access is enabled in *Menubar ▶ System ▶ Henge™ Network ▶ Remote Access*.

1.9.2 SSH host keys

At the bottom of the page, a box details the public SSH host keys of the appliance, that have been generated during the first start of the openSSH server, along with their fingerprints and their size in bits.

Example SYS-1 - Traffic Tunnelling over SSH.

Assume that a service such as telnet (or any other service that can be tunneled through SSH) is running on a computer inside the GREEN zone, say port 23 on host *myhost* with IP address 10.0.0.20. To setup a SSH tunnel through the IFA 3610 appliance to access the service securely from outside the LAN, i.e., from the RED zone. While GREEN access from the RED interface is in general not recommended, it might prove useful in some cases, for example during the testing phase of a service.

1. Enable SSH and make sure the host can be accessed, i.e., configure the firewall in *Menubar ▶ Firewall ▶ System* access for *myhost* to be reachable from the outside.
2. From an external system connect to the appliance using the command `ssh -N -f -L 12345:10.0.0.20:23 root@appliance` where `-N` tells SSH not to execute commands, but just to forward traffic, `-f` makes SSH run in the background and `-L 12345:10.0.0.20:23` maps the external system's port 12345 to port 23 on *myhost*, as it can be seen from the appliance.
3. The SSH tunnel from port 12345 of the external system to port 23 on *myhost* is now established. On the external system now it suffices to telnet to port 12345 on localhost to reach *myhost*.

1.10 GUI Settings

Two configuration options for the GUI are present here. The first option is the language that will be used for the section names, the labels, and all the strings used in the web interface and can be selected from a drop-down menu. The languages currently supported are: English, German, Italian, Simplified Chinese, Japanese, Portuguese, Russian, Spanish, and Turkish.

The second option is to display the hostname of the appliance in the browser's window title, activated by ticking the checkbox *Display hostname in window title*.

In the Community release it is also possible to click on the **Help translating this project** link, which will open the appliance translation page. Any help is appreciated!

1.11 Backup

In this section the management of the backups can be carried out: Creation of backups of the current appliance configuration and system rollback to one of these backups when needed. Backups can be saved locally on the appliance host, on a USB stick, or downloaded to a workstation.

It is also possible to reset the configuration to factory defaults, to create fully automated backups, and to carry out various other administrative tasks concerning backups.

This section is organised into two tabs, *Backup* and *Scheduled* backups: The former is used to manage manual backups, while the latter to set up automatic, scheduled backups.

1.11.1 Backup

In the *Backup* tab there are four boxes, that allow to manage the manual backups.

Backup sets

The first box contains a list of the backups stored on the appliance - both manually and scheduled ones, an option to create a new backup, and the legend of the symbols that accompany each backup. If a USB stick is plugged in in the appliance and detected, also backups stored on it are displayed.

When clicking on the **Create new Backup** button, a dialogue box opens up in which to select the data to be included in the backup.

Current configuration

The backup contains all the configuration settings, including all the changes and customisation done so far, or, in other words, all the content of the `/var/efw` directory.

Include database dumps

The content of the database will also be backed up.

Warning: The database dumps may contain sensitive data, so whenever a backup contains a database dump, make sure that it is stored in a safe place.

Include log files

Include the current log files (e.g., `/var/log/messages`, but not log files of the previous days.)

Include log archives

Include also older log files, that have been rotated, like e.g., `/var/log/messages.YYYYMMDD.gz`, etc. Backups created with this option may become very big after some time.

Remark

A comment about the backup, that will appear in the *Remark* column of the table. Hence, it should be meaningful enough to allow a quick recall of the content.

At least one of the checkbox must be ticked to create a new backup.

The format and name of the backup files.

Backup files are created as **tar.gz** archives, using standard Linux's tools **tar** and **gzip**. The files stored in the archive can be extracted using the **tar xzf archivename.tar.gz** or **tar vzxzf archivename.tar.gz** to see all the file processed and extracted and see some informative message on the screen, the `v` option meaning verbose. The name of the backup file is created to be unique and it conveys the maximum information possible about its content, therefore it can become quite a long string, like e.g., **backup-20130208093337-myappliance.mydomain-settings-db-logs-logarchive.tar.gz**, in which `20130208093337` is the timestamp of the backup's creation, in the form `YYYYMMDDHHMMSS` -in this example, 8th of February 2013 at 9:33:37 AM. This choice allows the backups to be lexicographically ordered from the oldest one to the most recent one; `myappliance.mydomain` are the IFA 3610 appliance's hostname and domainname as set in Step 3 of the *Network configuration (Menubar ► System ► Network configuration)*, and `settings-db-logs-logarchive` represent the content of the backup. In this case it is a full backup, since all four parts appear in the name. For example, a backup containing only settings and logs will be identified by the string `settings-logs`.

In order to create a backup on a USB external drive, a USB drive (even a stick) must be plugged in the appliance. It is suggested to use a FAT32/VFAT filesystem, as this maximises portability to other systems. When the stick is detected, the message *USB stick detected* will appear on the right-hand side of the box, along with a new option *Create backup on USB stick*. The checkbox next to this option must be ticked for the backup to be stored on the stick.

Click on the **Create Backup** button to create the backup. After a short time, during which the files required by the backup are gathered and assembled into the archive, the new backup appears in the list. The end of the backup process is marked by a yellow callout that appears above the box, showing the message *Backup completed successfully*.

The list of available backups, which is initially empty, presents for every backup the creation date, the content shown by a set of letters, the remark, and the list of actions available on each backup file. Automatic backups are marked with the string *Auto - backup before upgrade*.

The content of each backup is marked by at least one of the following letters or symbols, corresponding to the option specified during its creation:

A, Archive. The backup contains archived log files.

C, Cron. The backup has been created automatically by a scheduled backup job.

D, Database dumps. The backup contains a database dump.

E, Encrypted. The backup file is encrypted.

L, Log files. The backup contains today's log files.

S, Settings. The backup contains the configurations and settings.

U, USB. The backup has been saved to a USB stick.

!, Error. Something did not succeed while sending the backup file by email.


The available actions are to export  an archive to the local workstation, to delete it , or to restore it  on the appliance.

Encrypt backup archives

The second box makes available the option to encrypt all the backups by providing a GPG public key. Select the GPG public key by clicking on the **Choose file** button to upload the key file from the local file system. Make sure the checkbox *Encrypt backup archives* is ticked, then upload the key file by clicking on **Save**.

Hint: Encrypt backup archives whenever saving sensible data in the backup file, like for example the passwords of users stored in the database or hotspot's users data and billing information.

Import backup archive

The third box lets a previously saved backup archive be uploaded to the appliance. The backup file can be selected by clicking on the **Choose file** button and then choosing the backup file from the local file system. Optionally, some note to the backup can be added in the *Remark* field. Finally, the backup is uploaded by clicking on the **Import** button. The backup appears after a short period in the backup list at the top of the page, and can be restored by clicking on the restore icon .

Note: It is not possible to import encrypted backups on the appliance: Any encrypted backup must be decrypted before being uploaded.

Reset configuration to factory defaults and reboot

The fourth box allows to wipe out all configurations and settings done so far and reboot the system with the default configuration. This result is achieved by clicking on the **Factory defaults** button: The configuration of the appliance is reset to the factory defaults and rebooted immediately, right after a backup copy of the current settings has automatically been saved.

1.11.2 Scheduled backups

Automated backups of the system can be enabled and configured in the *Scheduled backups* tab, which contains two boxes.

Scheduled automatic backups

In the first box, automatic backups are enabled and configured. When enabled, the elements of the IFA 3610 appliance to be included in the backup can be chosen as seen in the *Backup Sets* box in the other tab. The only difference is that for scheduled backups there is no possibility to specify a remark. Additional options are:

Enabled

Enable scheduled backups.

Keep # of archives

Choose from the drop-down how many backups to keep on the IFA 3610 appliance (from 2 up to 10, but they can be exported to save space).

Schedule for automatic backups

The frequency between backups, either hourly, daily, weekly, or monthly.

Send backups via email

In the second box, the system can be configured to send or not the backups by e-mail. The following options are available.

Enabled

Allows backup archives to be sent via e-mail.

Email address of recipient

The e-mail address to which to send the e-mail with the backup.

Email address of sender

The e-mail address that will appear as the sender's e-mail address, which proves useful when backups should appear to have been sent from a special address (say, backups@myappliance.mydomain), and must be provided if the domain or hostname are not resolvable by the DNS.

Address of smarthost to be used

The address of a smarthost to be used to send the e-mails, which is needed in case the outgoing e-mails should go through a SMTP server, like, e.g., the Company's SMTP server, rather than to be sent directly by the appliance.

Hint: The explicit address of a smarthost is needed if the SMTP proxy (*Menubar -> Proxy -> is not enabled. SMTP*) is disabled.

Send a backup now

A click on this button will save the settings and immediately try to send an e-mail with the backup's archive as attachment, an action that serves also as a test for the correctness of the data supplied.

See also: A guide to create a backup on a USB stick.

1.12 Shutdown

Option to either shutdown or reboot the appliance, by clicking on the **Shutdown** or the **Reboot** button respectively, are provided in this page.

Warning: The shutdown or reboot process starts immediately after clicking on the respective button, with no further confirmation request.

After a reboot, it is possible to continue to use the GUI without the necessity of an authentication.

1.13 License Agreement

This section displays the license agreement between NEXCOM and the owner of the appliance.

Note: After an upgrade, if the license agreement changes, at the first login it is necessary to accept the new license agreement before accessing the upgraded system and being allowed to use the appliance.

Chapter 2: The Status Menu

The status menu provides a set of pages that display information in both textual and graphic views about various daemons and services running on the IFA 3610 appliance. No configuration option is available in this module, which only shows the current and recent status of the IFA 3610 appliance.

The following items appear in the sub-menu on the left-hand side of the screen, each giving detailed status information on some functionalities of the IFA 3610 appliance:

- System status - services, resources, uptime, kernel
- Network status - configuration of network interfaces, routing table, ARP cache
- System graphs - graphs of resource usage
- Traffic Graphs - graphs of bandwidth usage
- Proxy graphs - graph of HTTP proxy access statistics in the last 24 hours (week, month, and year)
- Connections - list of all open TCP/IP connections
- OpenVPN connections - list of all OpenVPN connections
- SMTP mail statistics - graphs about the SMTP service
- Mail queue - SMTP server's mail queue

The remainder of this section will describe the various parts that compose the System menu items.

2.1 System status

The default page that opens when clicking on *Menubar* ► *Status* is the *System status* page, which gives a quick overview of the running services, memory, disk usage, uptime and users, loaded modules, and the kernel version, each in its own box. At the top of the page, there are hyperlinks to each box. In more details, these are the information presented in each box, which are usually the output of some Linux command.

It shows several information about the installed system. It usually presents the hostname and domain name of the appliance in the title.

Services

The status -marked as either *Stopped* or *Running* by a red or green square- of each service installed on the IFA 3610 appliance is shown here. A service might appear as stopped because the corresponding daemon or script is not enabled.

Memory

The output of the Linux **free** command supplies the data shown here. All data are represented with the real amount in kilobytes, and with a bar to ease the visualisation of the memory used. The first line shows the total used RAM memory, for which is normal to be close to 100% for a long time running system, since the Linux kernel uses all available RAM as disk cache to speed up I/O operations. The second line shows the memory actually used by processes: Ideally this value should be below 80% to keep some memory available for disk caching. If this value approaches 100%, the system will slow down because active processes are swapped to disk. If the memory usage remains for long periods of time over 80%, RAM should be added to improve performances. The third bar indicates the swap usage. For a long running system it is normal to see moderate swap usage (the value should be below 20%), especially if not all the services are used all the time.

Disk usage

The output of the Linux **df** command shows the disk devices -physical disks and partitions, their mount point and the space of each disk partition. Depending on the type of the IFA 3610 appliance, the data displayed in this box differ. Usually, they are:

- The main disk `/dev/hda1`.
- The data disk `/dev/mapper/local-var`.
- The configuration disk, where all the IFA 3610 appliance settings are stored `/dev/mapper/local-config`.
- The log disk `/dev/mapper/local-log`.
- The shared memory, `/dev/shm/`.

Uptime and users

This box shows the output of the Linux `w` command, which reports the current time, information about how long the system has been running since last reboot, the number of console users that are currently logged into the system (though normally there should be none) and the system load average for the past 1, 5, and 15 minutes. Additionally, if any console user is logged into the system, some information about the user is displayed (like the remote host from which she is logged in or what is she doing). More details can be found on the `w(1)` manual page.

Loaded modules

The output of the Linux `lsmod` command. It shows the kernel modules currently loaded into memory. This information should be useful to advanced users only.

Kernel version

The output of the Linux `uname -r` command, which shows the current kernel version.

2.2 Network status

This page contains several information about the running state of the network interfaces. Four boxes are present on the page. The boxes contain the following information, representing the output of different shell commands.

Interfaces

The first box reports the output of the `ip addr show` command which provides for each network interface the associated MAC address, IP address, and additional communication parameters. The active interfaces are highlighted with the colour of the zone they are serving. The interface can be an ethernet interfaces, a bridge, or a virtual device.

NIC status

The running configuration and capabilities of each of the NIC are shown here. Each interface is highlighted with the colour of the zone it is serving and is labelled as **[Link OK]** to indicate that it is working. Interfaces that are not used are labelled with '[NO Link]'. The command providing the output is `ip link show`.

Routing table entries

The kernel routing table, as provided by the `route -n` command. Typically, there should be one line per active interface, which correctly routes the traffic within the zones served by the IFA 3610 appliance, plus a default route (recognisable by the 0.0.0.0 Destination field) that allows the traffic to reach the Internet.

ARP table entries

The last box shows the output of the `arp -n` command and shows the ARP table, i.e., a table containing the MAC address associated to each known IP address in the local network.

2.3 System graphs

The graphs displayed in this page present the usage of resources during the last 24 hours: CPU, memory, swap, and disk usage, each accompanied with a legend of the data included in the graph, their associated colour, and a summary of the maximum, average, and current percentage of use. Moreover, a message informs of the time and date of the last update to the graphs, which matches the last access to the page.

When clicking on one of the graphs, a new page will open, with summaries of the usage graphs for the last day, week, month, and year. In these pages, a click on the **BACK** button allows to return to the previous page.

Note: The nan (short for "Not A Number") string that may appear in the summaries designate that there are not enough data to calculate the usage of the selected resource. It can appear for example in the "per year usage" when the IFA 3610 appliance is used for only a few weeks.

CPU graph

In this box is shown the CPU usage per day of the IFA 3610 appliance, measured in percentage of the CPU time used by the various processes. The output is provided by the `top` command. Different colors are used to denote the type of running processes:

- White - idle, i.e., time the CPU is not used by any process.
- Green - nice processes, i.e., user processes which have changed their default priority.
- Blue - user processes with default priority.

- Orange - time spent by the CPU waiting for I/O tasks to complete.
- Red - system (kernel) processes
- Pink - softirq, i.e., the time spent for software interrupts
- Brown - interrupt, i.e., is the time spent for hardware interrupts
- Black - steal meaningful only if running as a virtual machine, is the time used by the hypervisor to run the VM.

Memory graph

This graph shows the memory usage during the last 24 hours. The following colours are used to denote the types of memory:

- Green - unallocated memory, that can be allocated to new processes.
- Blue - cache memory, copy of recent data used by processes.
- Orange - buffer memory, a temporary portion of memory that stores data to be sent to -or received from- external devices.
- Red - used memory.

Swap graph

The usage of the swap area, located on the hard disk, is displayed in this box.

- Green - unallocated swap.
- Blue - cached swap.
- Red - swap space used.

Disk usage graphs

Graphs showing the usage of the disk are split into four boxes, each showing the usage of a partition. In each of them, the green colour shows the free space, while the red colour shows the disk space used.

The four boxes show the free and allocated space in the four hard disk's partitions: Main disk, configuration disk, log disk, and data disk. For each of them, those colours are used.

- Green - space available.
- Red - space used.

In older 2.5 IFA 3610 appliance, the disk graph was slightly different, since the data shown was not the disk usage, but the hard disk accesses, denoted by two colours:

- Blue is used to show sectors read per second on the hard disk.
- Green is used to show the sectors written per second on disk.

2.4 Traffic graphs

This page contains the traffic graphs for the last 24 hours, divided by zone. Hence, depending on the zones enabled and configured, this page will contain 2, 3, or 4 boxes, each with one graphs. Like the System graphs, the graphs are accompanied with a legend of the data displayed:

- Green - the outgoing traffic.
- Blue - the incoming traffic

Below the graphs, also the summary of the average, maximum, and current amount of data transmitted and received is displayed and updated in real time.

When clicking on one of the graphs, a new page will open, with summaries of the data flown through the IFA 3610 appliance for the last day, week, month, and year. The data shown are the same in all the graphs: Incoming and outgoing traffic in blue and green respectively. In

Hint: To go back to the page with all the zone's graphs, click on the **BACK** hyperlink on the bottom of the page.

2.5 Proxy graphs

The access statistics of the HTTP proxy during the last 24 hours are shown here. There are no graphs in this page if the HTTP proxy service is not active and has never been enabled. However, if the service has been running even for a short period during the last year, the data produced are still accessible by clicking on the graph. Similarly to the other graphs, older statistics are shown for the last day, week, month, and year. In this page, a click on the **BACK** hyperlink on the bottom allows to go back to the main page.

Note: To show the proxy graphs, HTTP proxy logging must be enabled under *Proxy ► HTTP ► Configuration ► Log settings*, by ticking the *Enable logging* checkbox. Also *queried terms* and *useragents* can be logged to produce more detailed logs and graphs.

After the HTTP proxy has been enabled, the four boxes show the following data:

- *Total traffic per day:* the amount of data flown through the IFA 3610 appliance's proxy service. In green is shown the outgoing traffic, while in blue the incoming traffic.
- *Total Accesses per Day.* The number of HTTP requests, depicted in blue, received by the appliance.
- *Cache hits per day.* The number of cache data requested
- *Cache hits ratio over 5 minutes per day.* The number of cache data requested during a five minutes period.

2.6 Connections

This page shows a table containing the list of current connections from, to, or going through the appliance. The data shown here are devised by the kernel conntrack table. The following colours are employed in the table and used as the background of the cells in the table to denote the source and destination of the connection.

- Green, red, orange, and blue are the zones governed by the appliance.
- Black is used for connections involving the firewall, including daemons and services, like e.g., SSH or web accesses).
- Purple shows connections using VPN or IPsec.

The data displayed in the table are the following.

Source IP

The IP from which the connection has originated.

Source port

The port from which the connection has originated.

Destination IP

The IP to which the connection is directed.

Destination port

The port to which the connection is directed.

Protocol

The protocol used in the connection, which is typically tcp or udp.

Status

The current status of the connection, meaningful only for TCP connections. They are defined in RFC 793, significant states are ESTABLISHED (connection is active) and CLOSE (no connection).

Expires

How long will the connection remain in that particular status.

Hint: The page refreshes automatically every 5 seconds.

Each IP address and each IP port in the table can be clicked to obtain useful information. Clicking on the IP address will launch a whois query that will display who the owner of the IP address is and where it is located. Clicking on the port number will open the Internet Storm Center web page, with information about the port (i.e., the purpose for which it is used) and about which services or malware (e.g., Trojans, viruses) may exploit that port and the number of attacks received on those ports by various servers worldwide.



2.7 VPN Connections

When there are OpenVPN or IPsec servers running, this page shows the connected users, along with the service they rely on for the connection (OpenVPN, L2TP, IPsec Xauth), the time stamp since they are connected, and the possible actions that can be carried out. Currently, only to disconnect the user.


Chapter 3: The Network Menu

The network menu can be used to tweak the networking configuration by adding specific hosts and routes, or configuring the uplink and adding VLANs. This menu should not be confused with the *Network configuration* wizard available at *Menubar ► System ► Network Configuration*, that allows to configure interfaces, zones, and to define uplinks, although many settings and configuration options, especially in the Interfaces menu item are the same found there.

The sub-menu on the left-hand side of the screen contains these items, each of which groups several configuration options:

- Edit hosts - define hosts for local domain name resolution.
- Routing - set up static routes and policy routing.
- Interfaces - edit the uplinks or create VLANs.
- Wireless - set up wireless network connections.

3.1 Edit hosts

The page contains the list of hosts previously defined. Each line contains an IP address, the associated hostname, and the domain name, if specified. Two available actions are available for each entry: To edit it edit or to delete it. 

Warning: Deleting an host entry by clicking on the small delete icon does not require any confirmation and is not reversible. If deleted by mistake, an entry must be re-added manually.

A new entry in the file can be added by clicking on the **Add a host** link right above the table. A simple form will replace the table, in which to enter the following options:

IP address

The IP address of the remote host.

Hostname

The hostname associated to the IP address.

Domain name

An optional domain name.

Note: Unlike in the `/etc/hosts` file (see below), each IP address added here corresponds to one hostname and vice versa. To add two hostnames to a same IP, add two entries with the same IP address.

The choice can be confirmed by clicking on the **Add Host** button. To associate more hostnames to the same IP address, repeat the procedure by inserting the same IP address but a different name.

Hosts management, dnsmasq and `/etc/hosts`.

The dnsmasq application is used in small networks as DNS server for local hosts and as a DNS forwarder and caching server for worldwide DNS servers. The IFA 3610 appliance uses dnsmasq to be able to correctly resolve and answer DNS requests coming from the GREEN, ORANGE, and BLUE zones. It is sometimes desirable (e.g., for testing purposes on a remote website) to override some entries in dnsmasq, or to add some local server to dnsmasq's cache, for local clients to be able to connect to it.

The hosts added in this page are stored in a dnsmasq's settings file and merged with the `/etc/hosts` file at every restart of the daemon. Host added to that files directly via CLI will not persist after a reboot of the IFA 3610 appliance or a restart of dnsmasq.

The `/etc/hosts` file contains the so-called static lookup table, in the form:

```
IP1 hostname1 [hostname2]
IP2 hostname3 [hostname4] [hostname5]
```

Here, *IP1* and *IP2* are unique (numerical) IP addresses and *hostname1*, *hostname2*, *hostname3*, *hostname4*, and *hostname5* are custom names given to those IPs. Names within square brackets are optional: In other words, each IP address can be associated with one or more names of known hosts. Custom host entries can be added to the file, that will then be resolved for all the clients connecting through the appliance. On a typical appliance, the `/etc/hosts` file contains at least the following entries:

```
127.0.0.1    localhost.localhost localhost
172.20.0.21  myappliance.localdomain myappliance
172.20.0.21  spam.spam spam
172.20.0.21  ham.ham ham
172.20.0.21  wpad.localdomain wpad
```

Here, 127.0.0.1 is the IP address of the loopback device, *localhost*, which is a mandatory entry for the correct working of any Linux system; while 172.20.0.21 is the IP address of the GREEN interface. The entries listed for that IP have the following meaning and purposes:

myappliance.localdomain

The hostname and domainname of the IFA 3610 appliance, as set up during the *Network configuration*.

spam.spam spam and ham.ham ham

These two entries combined are used for the training of the spamassassin e-mail filter.

wpad.localdomain wpad

A facility for some browsers to detect and apply proxy settings automatically without the user's interaction when the proxy is not transparent.

3.2 Routing

Besides the default routing table, that can be seen in *Menu*bar ► *Status* ► *Network status*, the routing on the appliance can be improved with static and policy routing rules. This page displays a unique table that contains all the custom routings, although new rules are added from the two different tabs that present on this page. Indeed, static and policy routing rules require slight different settings. The table contains a summary of the rule: the source and destination networks or zones, the gateway, a remark, and the list of available actions: Enable or disable, edit, and delete a rule.

Whenever a modification is carried out on the routing table, it is required that the changes be saved and the service be restarted.

3.2.1 Static routing

A static route allows to associate specific source and destination networks with a given gateway or uplink. A click on the **Add a new route** link above the table allows create new routes by defining the following fields in the form that will appear:

Source Network

The source network, in *CIDR notation*.

Destination Network

The destination network, in *CIDR notation*.

Route Via

Four options are available to define through which means should the traffic be channeled: *Static Gateway*, *Uplink*, *OpenVPN User*, or *L2TP User*. In the case the Static Gateway is selected, the IP address of a gateway should be provided in the text box on the right. Otherwise, a drop-down will appear, proposing the choice among the available uplinks, OpenVPN users, or L2TP users.

Enabled

A ticked checkbox means that the rule is enabled (default). If unchecked, then the rule is only created but not activated: It can always be enabled later.

Remark

A remark or comment to explain the purpose of this rule.

A click on one of the icons will trigger an action on the respective item:



- toggle the status of the item, enabled or disabled.



- modify the item's property.



- remove the item

3.2.2 Policy routing

A policy route rule allows to associate specific network addresses, zones, or services (expressed as port and protocol) with a given uplink.

The table shows all the already defined rules, with some of their properties, and the following actions for each item:



- toggle the status of the item, enabled or disabled.



- modify the item's property.



- remove the item

Policy routing, HTTP proxy, and uplink.

The interaction between these three components of the appliance might produce some behaviour that may appear strange or even wrong when clients in the zones try to access the Internet. There are indeed three steps to highlight, for a correct understanding how traffic flows to the Internet when both HTTP proxy is enabled and there are policy routing rules defined:

1. An HTTP proxy uses the *main* uplink (i.e., it accesses the RED zone and the Internet using the main uplink).
2. An HTTP proxy "breaks" a connection from a client to a remote server in two connections: One from the client to the appliance and one from the appliance to the remote server.
3. Policy routing rules are taken into account *after* the traffic goes through the HTTP proxy.

When clicking on the **Create a policy routing rule** link, a form will open, which seems rather more complicated than the one for static routes and very similar to the firewall rule's editor. However, this policy rule editor is much like the previous one, but gives more control over the definition of the rule. Additionally, the setup of the rule is guided by several drop-down menus, to simplify entering the data in the following fields:

Source

The first drop-down menu allows to choose the source of the traffic. More entries, one per line, are accepted, but all must belong to the same type, either: A zone or interface, OpenVPN or L2TP users, IPs or networks, or MAC addresses. Depending on the choice, different values shall be supplied. To apply the rule to all sources, select <ANY>.

Destination

The second drop-down menu permits the choice of the destination of the traffic, in form of a list of IPs, networks, OpenVPN or L2TP users. Again, by selecting <ANY> the rule will match every destination.

Service/Port

The next two drop-down menus allow to specify the service, protocol, and a destination port for the rule when the TCP, UDP, or TCP + UDP protocols are selected. Some predefined combinations service/protocol/port exists, like HTTP/TCP/80, <ALL>/TCP+UDP/0:65535, or <ANY>, which is a shortcut for all *services, protocols, and ports*. User defined permits to specify a custom protocol and the ports to block, an option that proves useful when running services on ports different from the standard ones.

Protocol

The type of traffic that is interested by the rule: *TCP, UDP, TCP+UDP, ESP, GRE, and ICMP*. TCP and UDP are the most used, GRE is used by tunnels, ESP by IPsec, and ICMP by the **ping** and **traceroute** commands.

Route Via

How the traffic should be routed for this rule. Four options are available:

1. Static gateway: In this case an IP Address shall be provided
2. Uplink: The uplink that should be used for this rule. There is the option, when the uplink becomes unavailable, that the routing be carried over to the backup link corresponding to the selected uplink. This option is enabled when the checkbox next to the drop-down menu is ticked.
3. OpenVPN user: An OpenVPN user, chosen from those available in the drop-down menu.
4. L2TP user: An L2TP user, chosen from those available in the drop-down menu.

Type Of Service

The type of service (TOS) can be chosen here. Four values can be chosen, depending on what is the most important characteristic of the traffic interested by that rule: *default*, *lowdelay*, *reliability*, or *throughput*.

Remark

A remark or comment to explain the purpose of this rule.

Position

The position in which to insert the rule (relative position in the list of rules).

Enabled

Tick this checkbox to enable the rule (default). If unchecked, the rule is created but not active: A rule can be enabled later.

Log all accepted packets

This checkbox must be ticked to log all the packets affected by this rule.

Warning: The activation of this option may cause the size of the log files to dramatically improve.

3.3 Interfaces

The uplinks manager allows to carry out a number of tasks that are related with the uplink and the interfaces, and in particular to define custom VLANs on the network interfaces.

3.3.1 Uplink editor

By default, the uplink editor shows the available uplinks that have been created and the actions that can be executed on each of them, by clicking on the icons in the last column, Actions:



- toggle the status of the item, enabled or disabled.



- modify the item's property.



- remove the item

Hint: The main uplink can not be deleted.

Additional uplinks can be defined by clicking on the **Create an uplink** hyperlink above the list of uplinks. A rather long page, full of configurable options will open, that should be filled with appropriate values -very similar to those in the network configuration. Depending on the type of uplink chosen, the available settings will differ.

Note: Not all the available options are described here: They are the same that are present in the *network configuration wizard* and depend on the type of the uplink chosen, so please refer to that section for the full explanation of each option.

Description

A description of the uplink.

Type

The selection of the type of RED connection includes one additional protocol, compared to those available in the network configuration wizard: **PPTP**. PPTP can be configured to work in static or in DHCP mode, selectable from the respective value from the "PPTP method" drop-down. The IP address and netmask must be defined in the appropriate textfields if the static method has been chosen, in which case additional IP/netmask or IP/CIDR combinations can be added in the field below if the checkbox is ticked. Phone number, username, and password are not required but may be needed for some configurations to work, depending on the provider's settings. The authentication method can be either PAP or CHAP: if unsure, keep the default value "PAP or CHAP".

Uplink is enabled

Tick this checkbox to enable the uplink.

Start uplink on boot

This checkbox specifies whether an uplink should be enabled at boot time or not. This option proves useful for backup uplinks which are managed but do not need to be started during the boot procedure.

Uplink is managed

Tick this checkbox for the uplink to be managed. See the *Uplink Information Plugin* under *Menu* ► *System* ► *Dashboard* for a discussion about managed and manual modes.

If this uplink fails activate

If enabled, an alternative connection can be chosen from a drop-down menu, which will be activated when this uplink fails.

Check if these hosts are reachable

Tick this option to enter a list of IP or hostnames that will be **ping**-ed when the uplink fails, to check whether it has reconnected.

Hint: One of those hosts could be the provider's DNS server or gateway.

In the advanced settings panel, two other options can be customised:

Reconnection timeout


The time interval (in seconds) after which an uplink tries to reconnect if it fails. This value depends on the provider's settings. If unsure, leave this field empty.

MTU

A custom value for the MTU size.

3.3.2 VLANs

The idea behind offering VLAN support in IFA 3610 appliance is to allow arbitrary associations of VLAN IDs to the zones and to provide an additional level of separation (and therefore adding another level of security) between the zones. The existing VLANs are shown in the table, if any had already been created. The only action available is:

-  - remove the VLAN. A pop-up window will open, that requires a confirmation for the deletion.

A new VLAN can be defined by clicking on the **Add new VLAN** hyperlink above the VLAN list. In the form that will open a few click suffice to create an association between an interface and a VLAN, by specifying a few values:

Interface

The physical interface to which the VLAN is connected to. Only the available interfaces can be chosen from the drop-down menu. The menu also shows the status of the link of the interface.

VLAN ID

The VLAN ID, which must be an integer number between 0 and 4095.

Zone

The zone to which the VLAN is associated with. Only the zones that have been defined in the network configuration wizard can be selected. The option "NONE" can be chosen, if that interface is used as a High Availability management port.

Warning: It is not possible to define a VLAN that serves one zone (e.g., a VLAN on BLUE) on an interface that already serves another zone (e.g., eth1 serving GREEN). When trying to do so, the form closes and a red callout appears, informing that the VLAN can not be created.

Whenever a virtual LAN is created, a new interface is created and named as **ethX.y** where X is the number of the interface and y is the VLAN ID. This interface is then assigned to the chosen zone and will show up as a regular interface in the various sections that report network information, like *Menu* ► *Status* ► *Network Configuration* or in the Dashboard, where it can be selected to be drawn in the graph.

Chapter 4: The Services Menu

The appliance includes many useful services to prevent threats and to monitor the networks and the running daemons, whose activation and set up is explained in this section. In particular, among them, we highlight the various proxy services, such as the antivirus engine, as well as the intrusion detection system, high availability, and traffic monitoring. The available services appear as items in the sub-menu list on the left-hand side of the screen.

- DHCP server - DHCP server for automatic IP assignment
- Dynamic DNS - Client for dynamic DNS providers such as DynDNS (for home / small office use)
- Time server - enable and configure the NTP time server, set the time zone, or update the time manually
- Intrusion Prevention - configure snort, the IPS availability setup
- High availability - configure the IFA 3610 appliance in a high availability setup
- Traffic Monitoring - monitor network traffic and flows with ntopng
- SNMP Server - enable or disable support for the Simple Network Management Protocol
- Quality of Service - IP traffic prioritisation.

4.1 DHCP Server


The DHCP server is used by the clients (workstations and servers) in the zones controlled by the appliance to receive an IP address ("lease"), and allows to control the IP address assigned to them in a centralised way. Two types of leases can be assigned to clients: Dynamic and fixed. The DHCP server page is divided into two or three boxes, namely *DHCP*, in which to configure the DHCP server, *Current fixed leases*, showing the fixed leases, and *Current dynamic leases* that shows up only if at least one client has obtained a dynamic lease. Dynamic leases are assigned on a network basis within a given range that is configured in the first box, whereas fixed leases are assigned on a per-host basis and are configured in the second box.

DHCP

When a client (be it either a host or another device such as networked printer) joins the network it will automatically get a valid IP address from a range of addresses and other settings from the DHCP service. The client must be configured to use DHCP, which is sometimes called "automatic network configuration", and is often the default setting on most workstations. Dynamic leases are configured on a zone basis: for example, it is possible to enable them only for clients in the GREEN zone, while the other active zones receive only fixed leases.

It is however possible to let also devices in the ORANGE (DMZ) or BLUE (WLAN) zone to receive dynamic leases.

Note: If the BLUE zone is enabled but managed by the hotspot, the message *DHCP configuration is managed by hotspot* appears, preventing to configure it here.

To customise the DHCP parameter for each zone, click on the small icon  next to the *Settings* label. These are the available options:

Enabled

Enable the DHCP server in the zone.

Start address, End address

The range of IP addresses to be supplied to the clients. These addresses have to be within the subnet that has been assigned to the corresponding zone. If some hosts should receive a fixed lease, (see below), make sure their IP addresses are included neither in this range nor in the range of the OpenVPN address pool (see *Menubar ► VPN ► OpenVPN server*) to avoid conflicts.

Leaving these two fields blank will use the whole IP range of the zone for dynamic leases.

Allow only fixed leases

Tick this checkbox to use fixed leases *only*. No dynamic lease will be assigned.

Default lease time, Max lease time

The default and the maximum time in minutes before the assignment of each lease expires and the client requests a new lease from the DHCP server.

Domain name suffix

The default domain name suffix that is passed to the clients and that will be used for local domain searches.

Default Gateway

The default gateway that the clients in the zone will use. If left blank, the default gateway is the appliance itself.

Primary DNS, Secondary DNS

The DNS used by the clients. Since the appliance contains a caching DNS server, the default value is the firewall's own IP address in the respective zone, though a second server or even the primary value can be changed.

Primary NTP server, Secondary NTP server

The NTP servers used by the clients, to keep the clocks synchronised.

Primary WINS server, Secondary WINS server

The WINS servers used by the clients. This option is only needed for the Microsoft Windows networks that use WINS.

Advanced users might want to add some custom configuration lines to be added to the `dhcpd.conf` file (e.g., custom routes to subnets) by writing them in the text area at the bottom, marked with the *Custom configuration* lines label.

Warning: No syntax check on these lines is carried out: the lines are appended to the configuration file. Any mistake here might inhibit the DHCP server from starting!

Example SRV-1 - PXE boot and dhcpd.conf configuration.

The customisation of the DHCP server proves useful in different networks configuration.

One common use case is for VoIP telephones that need to retrieve their configuration files from an HTTP server at boot time. In this case, the files may also reside on the IFA 3610 appliance, so the configuration of the tftp server can be passed as extra lines like the following:

```
option tftp-server-name "http://$GREEN_ADDRESS";
option bootfile-name "download/voip/{mac}.html";
```

Note the use of `$GREEN_ADDRESS` which is a macro that is replaced in the `dhcpd.conf` file with the GREENIP of the IFA 3610 appliance.

Current fixed leases

It is sometimes necessary or desirable for certain devices to always use the same IP address while still using DHCP, for example servers that provide services like a VoIP box, a SVN repository, a file server, or devices like printers or scanners. A fixed lease is usually referred to as *Static IP Address*, since a device will always receive the same IP address when requesting a lease from the DHCP server.

This box reports the list of all the fixed leases currently active in the local network, providing several information about that lease. By clicking on the **Add a fixed lease** link, new fixed leases can be assigned to a device and insert all the information that will be displayed in the list. The devices are identified by their MAC addresses.

Note: Assigning a fixed lease from the DHCP server is very different from setting up the IP address manually on a device. Indeed, in the latter case, the device will still contact the DHCP server to receive its address and to announce its presence on the network. When the IP address required by the device has already been assigned, however, a dynamic lease will be given to the device.

The following parameters can be set for fixed leases:

MAC address

The client's MAC address.

IP address

The IP address that will always be assigned to the client.

Description

An optional description of the device receiving the lease.

Next address

The address of the TFTP server. This and the next two options are useful only in a few cases (see below for an example).

Filename

The boot image file name. Option needed only for thin clients or network boot.



Root path

The path of the boot image file.

Enabled

If this checkbox is not ticked, the fixed lease will be stored but not written down to the file `dhcpd.conf`.

The actions available for each fixed lease in the table are:

- - toggle the status of the lease, enabled or disabled.
-  - modify the property of the lease.
-  - remove the lease.

A use case for a fixed lease.

A use case that shows the usefulness of a fixed lease is the case of thin clients or disk-less workstations on the network that use PXE, i.e., boot the operating system from an image supplied by a networked tftp server. If the tftp server is hosted on the same server with the DHCP, the thin client receives both the lease and the image from the same server. More often, however, the tftp server is hosted on another server on the network, hence the client must be redirected to this server by the DHCP server, an operation that can be done easily adding a fixed lease on the DHCP server for the thin client, adding a next-address and the filename of the image to boot.

Besides the information supplied during the fixed lease creation, the list allow each lease to be enabled or disabled (by ticking the checkbox), edited, or deleted, by clicking on the icons in the *Actions* column. Editing a lease will open the same form as the creation of a new lease, whereas deleting a lease will immediately remove it from the configuration.

Note: All leases assigned by the DHCP server are stored by default in the `/var/lib/dhcp/dhcpd.leases` file. Although the DHCP daemon takes care of cleaning that file, it may happen that the file stores lease that have already been expired and are quite old. This is not a problem and does not interfere with the normal DHCP server working. A typical entry in that file is:

```
lease 192.168.58.157 {
starts 2 2013/06/11 13:00:21;
ends 5 2013/06/14 01:00:21;
binding state active;
next binding state free;
hardware ethernet 00:14:22:b1:09:9b;
}
```

Current dynamic leases

When the DHCP server is active, and at least one client has received a (dynamic) IP address, a third box appears at the bottom of the page, containing the list of the currently assigned dynamic IP addresses. This list report the IP address, the MAC address, the hostname, and the expiry time of the lease associated to each client.

4.2 Dynamic DNS



A DNS server provides a service that allows to resolve the (numeric) IP address of a host, given its hostname, and vice versa, and works perfectly for hosts with *fixed IP* address and hostname.

DDNS providers, like DynDNS or no-IP, offer a similar service when the IP addresses is dynamic, which is normally the case when using residential ADSL connections: Any domain name can be registered and associated to a server with a dynamic IP address, which communicates any IP address change to the DDNS provider. To be compatible and to integrate with the root DNS servers, each time IP address changes, the update must then be actively propagated from the DDNS provider.

The appliance includes a dynamic DNS client for 14 different providers and if enabled, it will automatically connect to the dynamic DNS provider to communicate the new IP address whenever it changes.

Note: If no dynamic DNS account has been set up, detailed instruction to register a new one, detailed online helps and how tos are available on the web site of the providers.

This page displays the list of the Dynamic DNS accounts. Indeed, more than one DDNS provider can be used. For each account, the list shows information about the service used, the hostname and domain name registered, if the anonymous proxy and the wildcards are active, if it is enabled, and the possible actions:

- - toggle the status of the lease, enabled or disabled.
-  - modify the property of the lease.
-  - remove the lease.

New accounts can be created by clicking on the **Add a host** link, providing the following parameters:

Service

The drop-down menu shows the available DDNS providers.

Behind a proxy

This option only applies to the no-ip.com provider. The checkbox must be ticked if the appliance is connecting to the Internet through a proxy.

Enable wildcards

Some dynamic DNS providers allow *all* the sub-domains of a domain point to the same IP address. This is a situation in which two hosts like www.example.myddns.org and second.example.myddns.org are both located on the same IP address. Ticking this box enables the feature, making all the possible sub-domains redirect on the same IP address. The feature must be configured also in the account on the DDNS provider server, if available.

Hostname and Domain

The hostname and domain as registered with the DDNS provider, for instance "example" and "myddns.org"

Username and Password

The credentials given from dynamic DNS provider to access the service.

Behind Router (NAT)

Activate this option if the appliance is not directly connected to the Internet, i.e., there is another router or gateway before accessing the Internet. In this case, the service at <http://checkip.dyndns.org> can be used to find the IP address of the router.

Enabled

Tick this checkbox to enable the account, which is the default.

Note: It is still necessary to export a service to the RED zone to be able to use the domain name to connect to the appliance from the Internet using its dynamic IP address, since the dynamic DNS provider only resolves the domain name and not the associated services. Exporting a service might typically involve setting up port forwarding (see [Menubar ► Firewall ► Port forwarding / NAT](#)).

After making a change in the configuration or to immediately update the dynamic DNS for all the defined accounts, click on the **Force update** button. This proves useful for example when the uplink has been disconnected and the REDIP has changed: When this happens, updating all the DDNS accounts is required, otherwise the services offered via DDNS will be unreachable.

4.3 Time server

The appliance uses NTP to keep its system time synchronised with time servers on the Internet. The settings available are grouped into two boxes.

Use a network time server

A number of time server hosts on the Internet are preconfigured and used by the system, but custom time servers can be specified after ticking the *Override default NTP servers* checkbox. This might prove necessary when running a setup that does not allow the appliance to reach the Internet. Several time servers addresses can be supplied, one per line, in the small form that will show up.

This box also shows the current time zone setting, that can also be changed by choosing a different one from the drop-down menu. An immediate synchronisation can be done by clicking on the **Synchronize now** button.

Adjust manually

The second box gives the possibility to manually change the system time. While this is not recommended, this action proves useful when the system clock is way off and an immediate update of the appliance's clock to the correct time is needed.

Automatic synchronisation using time servers is not done instantly, but the clock is "slow down" or "speed up" a bit to recover and align to the correct time, hence a system with a significant error in its time may require a long period to be corrected. In those cases, forcing a manual synchronisation represents a more drastic but immediate solution.

4.4 Intrusion Prevention

The appliance includes the well known intrusion detection (IDS) and prevention (IPS) system snort, which is directly built into iptables, to intercept and drop connections from unwanted or distrusted sources.

The page contains three tabs, *Intrusion Prevention System*, *Rules*, and *Editor*.

4.4.1 Intrusion Prevention System

If snort is not active, a grey switch next to the *Enable Intrusion Prevention System* label appears on the page and can be clicked on to start the service. A message appears, informing that the service is being restarted and after a short interval, the box will contain some options to configure the service.

Automatically fetch SNORT Rules

Ticking this box will let the appliance automatically download the snort rules from the Henge™ Network.

Note: If the appliance is not registered, rules are downloaded from the Emerging Threats web page. An informative message is also shown at the bottom of the page.

Choose update schedule







The frequency of download of the rules: A drop-down menu allows to choose one of the *hourly*, *daily*, *weekly*, or *monthly* options. This option appears only if the previous option has been activated.

Custom SNORT Rules

A file containing custom SNORT rules that should be uploaded. Pick one file from the file selection window that opens upon clicking the **Browse** button, and upload it by clicking on the **Upload custom rules** button.

4.4.2 Rules

On the *Rules* tab appears the list of rule sets that are stored on the appliance, along with the number of rules they contain and the actions that can be done on them:

-   - toggle the status of the rule set, enabled or disabled.
-   - The policy applied to packets, either they are allowed to pass or not.
-  - modify the property of the rule set.
-  - remove the rule set.

Note: When editing a ruleset in the Rules tab, the Editor page (see below) will open with that ruleset already selected.

All the actions, except for editing, can be carried out on more than one rulesets at once, by selecting them (tick the checkbox on the left of their filename) and pressing one of the button underneath the list.

By default, the policy for all the rulesets is set to *alert*. This behaviour can be changed by clicking on the alert icon to toggle the policy into *block* and the icon into a red shield. After clicking on the **Apply** button, that ruleset will not cause alerts anymore, but all the traffic that matches its rules will be blocked.

A ruleset can be deleted by clicking on the trash can icon, while a click on the pencil icon redirects to the *Editor* page in which to edit each rule independently.

4.4.3 Editor

At the top of the *Editor* page are shown the rulesets that can be edited. To chose more than one ruleset at once, hold the **CTRL** key and click on the rulesets.

After selecting and clicking on the **Edit** button, the list of the rules included in the selected ruleset(s) is shown. The list can be narrowed down by entering some terms in the text box next to the *Search* label. Like in the *Rules* page, the policy of every entry can be changed.

Warning: Turning on the IPS only implies that snort is running, but it does not yet filter the traffic. For snort to filter packets, the *Allow with IPS Filter* policy must be selected for the rules defined in the various *Firewall* configuration pages.

4.5 High Availability

The appliance can be run in an HA mode, that can easily be setup using at least two IFA 3610 appliances, one of which assumes the role of the active (i.e., *master*) firewall, while the remaining are standby (i.e., *slave*) firewalls.

If the master firewall fails, an election among the slaves takes place and one of them will become the new master, providing for transparent failover. If there is only one slave, though, it will immediately take over the master's duties and allows a seamless failover transition to the secondary appliance in the event of a hardware failure on the primary appliance. This provides unparalleled hardware availability and redundancy for critical network operations and security.

In order to start up the HA service, at least one master and one slave appliances must be configured according to the following guidelines.

Note: The Henge™ HA system is supported on both Henge™ hardware and software appliances. Regardless of choosing hardware or software, the high availability module requires at least two completely identical hardware platforms (e.g. 2 Minis, 2 Macros, 2 x86 systems, etc.).

An important point to focus on when deploying high availability is that a duplication method for each and every connection to the IFA 3610 appliance must be provided. Every connection of the primary unit (e.g., WAN, LAN, etc.) must be replicated across the standby unit(s) to ensure that complete replication capabilities exist.

In this scenario, each network on the appliance (WAN, LAN, etc.) is connected to an external managed switch which has a unique VLAN assigned to each network. This deployment option consumes the least amount of network ports and provides for enhanced extensibility. Another option is to replace a single managed (VLAN capable) switch with smaller, separate switches for each network (WAN, LAN, etc.). This setup however may not be cost-effective and could be less reliable since the failure of any switch could break failover partially or completely.

Warning: Since the HA runs automatically over the GREEN network, the heartbeat can be configured to run over the switch connection or alternatively, an additional Ethernet port can be assigned to the GREEN network to directly connect the master device to the slave unit. The advantage of adding a direct connection is that it removes the switch (and thus possible sources of problems, improving the overall reliability) from the failover equation. The decision on whether to implement this setup may largely depend on the overall reliability of the managed switch (dual power, port failure rate, warranty terms, etc.) – so the more reliable/redundant is the switch configuration, the less critical having a direct connection can become.

In this page, there is only one box, which initially contains only one option:

Enable High Availability

Enable HA on the appliance, by default it is disabled.

After enabled, a second drop-down menu appears, *High Availability side*, that allows to configure the appliance as master or slave. Depending on this choice, different configuration options are available. Configuring a slave unit, however, requires that a Master unit have already been set up.

For the **master side**, the following options are available:

Management network

The special subnet to which all appliance that are part of a same HA setup must be connected and defaults to 192.168.177.0/24. Unless this subnet is already used for other purposes, there is no need to change it.

Master IP Address

The first IP address of the management network. It is automatically set to 1 on the network chosen, and defaults to 192.168.177.1.

Notification: recipient email address, Notification: sender email address, Notification: email subject, Notification: SMTP server to be used

These options can be filled in to be notified by e-mail when a failover event occurs. They are configured the same way as they were configured for other *event notifications* in *MenuBar ► System ► Event notification*: A custom sender, recipient, and subject of the email and the SMTP smarthost used to send the email.

Enable STP

Choose from the drop-down menu whether to enable or not the spanning tree protocol, STP. This option and the next one are important when the appliance is in gateway mode.

STP Bridge Priority

The priority of the bridge. It must be 1 on the master side.

The HA management network.

The appliance uses a special network to connect the master to slave unit(s): 192.168.177.0/24. If this network has already been used in other zones, none of the already defined network(s) is deleted nor any change should be made to them. Indeed, in such a case, simply assign to the HA management network a different range of IP addresses, like, e.g., 172.19.253.0/24 or 10.123.234.0/28. It is important to note that the only requirement of the management network is that it must be large enough to accommodate the master and all the slaves, so if there are only a master and a slave devices, even a network as small as 192.168.177.0/29 should suffice. The management network will be created as an interface on the GREEN network, and it will show up as such on the device or when viewing the network status.

Warning: Make sure that the management network can be reached from the current LAN setup, or it will not be possible to login to the master unit!

After the master unit has been configured, the second appliance, that is going to be the slave, can be set up. The same procedure shall be followed for every additional slave to configure.

Warning: It is strongly suggested to make a backup of the slave unit before configuring it and saving it on a safe place, since it may become useful to restore a slave unit after it is removed from its role.

For the **slave side** the following are the available options.

Master IP address

The IP address of the master unit, which defaults to 192.168.177.1/24 if the management network had not be changed. This value must match the one that appears as value of the *Master IP Address* option on the master unit.

Master root password

The password of the console *root* user (not the graphic administration interface!) on the master. These data will be used by the slave to retrieve from the master all the information needed and to keep the synchronisation.

Enable STP

Choose from the drop-down menu whether to enable or not the spanning tree protocol, STP. On the slave side, this option must have the same value as in the master side.

STP Bridge Priority

The priority of the bridge. On the slave side, it must be a digit or number higher than the one on the master side. Upon saving the setup, the connection to the device will be temporarily lost, since the management network is created and then the two devices (the master and the currently defined slave) begin to synchronise.

After the synchronisation process is complete, the slave itself cannot be reached anymore via its old IP address (be it its factory default or its previous GREENIP address), since it has gone in standby mode and is connected to the master only through the management network. Any change made on the primary unit (the activation of a service, the change of one setting, the deletion of a VPN user, and so on) will automatically be synced to the slave unit(s) with the exception of updates, upgrades, or device backups (these have to be performed manually on the slave unit).

Moreover, the slave IFA 3610 unit will automatically appear on the master's list of slaves and switch to an informational-only web interface that can be accessed from the master, by following the **Go to Management** GUI link next to each of the entries of the list of slaves.

The RED MAC Address

During the HA failover, the RED interface MAC address is not replicated onto the slave unit. This can represent a problem if the ISP requires to use the *Sticky IP* setup. In this situation, the IP address assigned from the ISP is determined from the MAC address of the client's network interface, similarly to a fixed IP assigned from a DHCP server to a client. It may not be possible to reconnect with the slave unit. To avoid this situation, it is necessary to utilise the spoofed MAC address feature on the RED interface in order for HA to work properly. This will ensure that when the HA is activated the MAC address will carry over to the standby unit and will not require manual intervention. This can be achieved on the slave, before activating it, by ticking the option *Use custom MAC address* under *MenuBar ► Network ► Interfaces ► Edit main uplink ► Advanced settings* and specifying the MAC address of the RED interface on the Master. Alternatively, the MAC address can be entered in the step 4 of the network installation wizard, writing the master's MAC address in the *Spoof MAC address* with option.

4.6 Traffic Monitoring

Note: The ntop service is not available on the Mini appliances, due to its limited available resources.

Traffic monitoring is done by ntopng and can be enabled or disabled by clicking on the main switch on this page. Once traffic monitoring is enabled a link to its new **administration interface** appears in the lower section of the page. The traffic can be visualised and analysed by host, protocol, local network interface and many other types of information: All these operations can be carried out directly from the *Traffic Monitoring* module in *The Logs and Reports Menu*.

4.7 SNMP Server

The SNMP is used to monitor network-attached devices, and can be used e.g., to control the status of the internal infrastructure.

To enable the SNMP Server is sufficient to click on the grey switch next to the *Enable SNMP server label*: Once done so, a few options will appear in the *Settings* box.

Community String

A key that is needed to read the data with an SNMP client.

Location

An identification string that can be set to anything, but it is suggested that it describe the location of the appliance.

Override global notification email address

The SNMP Server requires to configure an e-mail address as the system contact, and the global e-mail address provided during the installation procedure is used by default. Nonetheless, to use a custom e-mail address, tick the checkbox and supply it into the *System contact email address* field that will activate right below.

4.8 Quality of Service

The purpose of the QoS module is to prioritise the IP traffic that is flowing through the appliance depending on the service. In other words, the QoS is a convenient way to reserve a given amount of the available bandwidth (both incoming and outgoing) for a given service. Applications that typically need to be prioritised over bulk traffic are interactive services such as SSH or VoIP.

The QoS configuration options are arranged into three tabs: *Devices*, *Classes*, and *Rules*.

4.8.1 Devices

The *Device* tab is also the starting page for the QoS and is initially empty. Once populated, a table showing a list of all the Quality of Service devices appears and for each device, some parameters and the available actions are displayed.

New QoS devices can be added by clicking on the **Add Quality of Service Device** link above the list and by configuring a few options.

Target Device

The network interface that will be used by this device. Choices are among the network interfaces or zones enabled on the system and can be selected from a drop-down menu.

Downstream Bandwidth (kbit/s)

The downstream speed of the interface.



Upstream Bandwidth (kbit/s)

The upstream speed of the interface.

Enabled

Enable the QoS (default) or not.

The actions available on the devices are:

- - enable or disable the device.
-  - modify the properties of the device.
-  - remove the device.

When editing a device, the same form opens as when adding a new device, in which to modify the current device's parameters.

For every device added, four items will appear under the *Classes* tab: Three for high, medium, and low priority, respectively, and one for bulk traffic (see below).

4.8.2 Classes

This tab shows a list of all Quality of Service classes that have been created, if any. For each entry, several data are shown. New items can be added by clicking on the **Add Quality of Service Class** link above the list of classes. The parameters to configure are the same shown in the list:

Name

The name of the Quality of Service class.

QOS Device

The drop down menu allows to choose the Quality of Service device for which the class was created.

Hint: At least one QoS device must have been created before defining a QoS class.

Reserved

The percentage of bandwidth that has been reserved for this class from the device's overall available bandwidth.



Limit

The maximum percentage of bandwidth this class may use.

Priority

The priority of the class, from 0 (low) to 10 (high), selected from a dropdown menu.

Note: The sum of reserved percentages can not be greater than 100 per device. Moreover, the *reserved* bandwidth can not be higher than the *limit* bandwidth.

The actions available are: *  - modify the properties of the device. *   - move the class in the list. *  - remove the device.

Classes can be moved up or down the list: Items closer to the top of the list are the first to be processed when the bandwidth does not suffice for all the traffic and the appliance needs to choose which traffic should be prioritised.

4.8.3 Rules

The third tab displays a list of the already defined Quality of Service Rules and allows to specify which type of traffic should belong to each of the classes. To add a new Quality of Service rule click on the **Add Quality of Service Rule** link. In the form that will open, which is very similar to the one used to define firewall rules, several values should be configured. Many drop-down menus are employed here to ease the choices and guide through the configuration.

Source

Choose from the drop-down menu the traffic source, either a Zone or interface, a network, an IP or MAC address. Depending on this choice, different values can be specified: A zone or interface from the available ones from those that will be displayed, or one or more IP addresses, networks, or MAC addresses.

Destination Device/Traffic Class

Choose the device/class from the drop-down menu and then the destination IP addresses or networks, to be written in the text-area on the right-hand side.

Service/Port, Protocol

The next two drop-down menus are used to define the service, protocol, and a destination port for the rule (when choosing one of TCP, UDP, or TCP + UDP protocols). Some predefined combinations Service/Protocol/Port exists, like HTTP/TCP/80, <ALL>/TCP+UDP/0:65535, or <ANY>, which is a shortcut for all services, protocols, and ports. Finally, in the *Destination port*, one or more custom port number can be supplied (this proves useful when some service does not run on a standard port).

TOS/DSCP

The type of TOS or DSCP value to match.

Match Traffic

Choosing *TOS* or *DSCP* class in the previous drop-down menu allows to choose a suitable value for the traffic to match from another drop-down menu. Otherwise, the choice *DSCP Value* allows to enter a custom value that should match the rule.

**Enabled**



Tick the checkbox to enable the rule.

Comment

A comment to identify the rule.

Note: If there is more than one service in a Quality of Service class, then all these services together will share the reserved bandwidth.

The actions available on the rules are:

- - enable or disable the rule.
-  - modify the properties of the rule.
-  - remove the rule.

Chapter 5: The Firewall Menu

This section allows to set up rules that specify if and how the network traffic flows through the appliance. The firewall on the appliance is divided in different modules, each monitoring and allowing or blocking one specific type of traffic. The modules available are the following:

- Port forwarding / NAT - port forwarding and abbr: *NAT (Network Address Translation)*.
- Outgoing traffic - outgoing traffic, i.e., towards the RED interface
- Inter-Zone traffic - traffic between zones.
- VPN traffic - traffic generated by VPN users.
- System access - grant access to the IFA 3610 appliance host itself.
- Firewall diagrams - pictures that show which traffic is intercepted by each type of firewall.

Within each of the sub-menus, in which all the corresponding existing rules are listed, any customised rules can be added, for any type of service or every port/protocol. The various parts of which the firewall is composed refer to different types of traffic (e.g., OpenVPN governs the traffic from/to the VPN users, inter-zone traffic the one flowing from zone to zone) and are designed to avoid any overlapping or even contrasting rules. In other words, there is no way to write two rules in two different firewall modules whose combined effect causes an unwanted block or access of packets.

The choice to separate the networks controlled by the appliance allows also for an easier management of the firewall, whose configuration may become very complex. Indeed, each of the modules can be considered as an independent firewall, and their combined effect covers all possible packet flows through the appliance.

Additionally, for any of the modules listed above, one or more rule may exist, that can neither be disabled nor removed. These are the so-called *Rules of system services* (or *System rules*) whose purpose is to allow the correct interoperability of the services running on the appliance with the network infrastructure.

The rules that are defined here will be transformed into **iptables** commands, the standard Linux firewall tool since the 2.4 kernel, and therefore organised into tables, chains, and rules. For a more in-depth description of the various elements that compose a firewall rule, or even to learn how to fine-tune and to manage a complex firewall, it is suggested to read either the **iptables(8)** manual page on any Linux box, or some of the countless online resources or tutorials available on the Internet.

5.1 Common Configuration Items

When adding a rule, most of the values to configure in the various modules are of the same type (e.g., the source or destination interfaces), since in the end they are all setup with **iptables**. Therefore, in order to keep this section short and readable, all the configuration items that are common to all modules of the firewall are grouped here and defined only once. There will be some more explanation only in case of significant differences with the descriptions given here.



- *Source or Incoming IP*. Usually in the form of a drop-down menu, this setting is the type of the source or incoming connection that should be matched. Depending on the type chosen, the selection of different connections from the small box underneath the menu will be possible: *Zone/VPN/Uplink* is either the source zone, VPN client, or uplink to which this rule should be applied, *Network/IP/Range* the IP address or range or the network addresses, *OpenVPN User* and *L2TP User* the OpenVPN or 2TP users, respectively.
- *Destination or Target*. Also this setting comes in the form of a drop-down menu and allows the choice among three types of destination that should be matched, which are the same as in the *Source* drop-down menu: *A Zone/VPN/Uplink, Network/IP, OpenVPN User* or *L2TP user*, except for some small change (e.g., for some type of rules, the target can not be an OpenVPN or L2TP user).

Service, Port, and Protocol. A service is usually defined as a combination of a port and a protocol. For example, the SSH service runs by default on port 22 and uses the TCP protocol. These three options control the port and protocol to which to apply the rule and consist of two drop-down menus, from which to choose either a pre-defined *Service*, that will also set the protocol and the port range in the text area, or one Protocol and optionally a port or a port range. Available protocols are: TCP and UDP - the most used, GRE - used by tunnels, ESP - used by IPsec, and ICMP - used by the **ping** and **traceroute** commands.



Note: There exist dozens predefined services that can be chosen from the drop down menus and should suffice to allow the most common services to access the Internet. An user defined combination of port and protocol should be used only if a service is not running on a standard port (e.g., an SSH server listens to port 2345 or a web server runs on port 7981) or if a service is using a particular port (e.g., a multiplayer game on the Internet).

- *'Access from'* sub-rule. Almost every rule can be further detailed by adding *several* Access from rules to it, for example to limit access to a client depending on the zone from which it connects to the appliance. *Access from rules* can be configured when the advanced mode is selected (see below). As a consequence, a rule can appear split on two or more lines, depending on the number of access policies defined. Each *access from* sub-rule can be deleted individually, without changing the main rule. Each of the sub-rules can even have a different filter policy.
- *Policy, Filter Policy.* The action to carry out on the packets that match the current rule. The drop-down menu allows to select among four options: *Allow with IPS* - let the packet pass but analyse it with the *Intrusion Prevention System*, *Allow* - let the packets pass without any check, *Drop* - discard the packet, and *Reject* - discard the packet and send an error packet in response.
- *Enabled.* Every rule created is by default enabled, but it can be saved and not activated by unticking the checkbox, i.e., it will not be taken into account for packet filtering. Disabling a rule may prove useful for troubleshooting connections' problems.
- *Log, Log all accepted packets.* By default, no log entries is written when traffic is filtered. To enable logging for a rule, tick the box.

Warning: If there is a lot of traffic and packets to be analysed, the size of the log files will likely grow rapidly, so in this case remember to check the log directory regularly to avoid running out of space!

- *Remark.* A description or a remark about the rule, to remember the purpose of the rule.
- *Position.* Recall that the iptables rules are processed in the order they appear on the list and that some is a "terminating" rule, i.e., it may drop or reject a packet and stop the processing of the subsequent rules. This drop-down menu allows to choose in which position this rule should be saved.
- *Actions.* On all rules several actions can be carried out:
-   - move the rule upwards or downwards in the list.

Hint: Remember that the ordering matters! The firewall rules are processed in the order they appear in the page, top to bottom.

- - enable or disable the rule.
-  - modify the rule.
-  - remove the rule.

Finally, after every change has been saved in the firewall rules, the firewall should be restarted to reload the configuration. A callout with a clickable **Apply** button will appear to recall this necessity.

5.2 Port Forwarding / NAT

The Port forwarding / NAT module is composed by three tabs: Port forwarding / DNAT, Source NAT, and Incoming routed traffic. Its purpose is to manage all the traffic that flows through the uplink, from the RED zone to the appliance and the NAT-ed traffic, both incoming and outgoing.

5.2.1 Port forwarding / Destination NAT

Destination NAT is usually employed to limit network accesses from an untrusted network or to redirect the traffic coming from the untrusted network and directed to a given port or address-port combination. It is possible to define which port on which interface should be forwarded to which host and port.

The list of the configured rules shows several information: The ID (#) showing the order in which the rules are matched against the traffic, the *Incoming IP* address, the service (i.e., port and protocol) to which the traffic is directed, the *Policy* applied to the traffic, the *Translate to* address (i.e., the host and port where to redirect the traffic), a custom *Remark*, and the available *Actions*.

When editing a rule, the same form open as when adding a new rule, by clicking on the **Add a new Port forwarding / Destination NAT rule**. A link on the top right of the form allows to chose between a **Simple mode** or an **Advanced mode**. The latter mode allows also to fine-tune the *Access from*, the policy, and the type of *Translate to*.

Besides the *common options*, these other settings can be configured:

Translate to

This part of the form changes depending on the current active editing *mode*, simple or advanced. If the mode is set to *advanced*, besides adding *Access from sub-rules*, there is an additional *Type* drop-down menu that allows to chose among different types of translations.

1. The first one is *IP* and corresponds to the only one available in *simple mode*. Here should be written the destination IP address (besides port and NAT), the port or port range to forward to and if to apply NAT or not to the incoming packets.
2. OpenVPN User: choose one OpenVPN user as the destination target for the traffic.
3. Load Balancing: specify a range of IP addresses to which traffic will be split, to avoid bottlenecks or the overloading of a single IP.
4. Map the network. Insert a sub-network to which translate the incoming traffic

Note: The *Map network* translation statically maps a whole network of addresses onto another network of addresses. This can be useful for companies whose subsidiaries all use the same internal network. Indeed, in this case all these networks can be connected to each other through network mapping. An example would be:

```
original network 1: 192.168.0.0/24
mapped network 1: 192.168.1.0/24
original network 2: 192.168.0.0/24
mapped network 2: 192.168.2.0/24
```

5. L2TP User: choose one L2TP user as the destination target for the traffic.

Except when selecting the *Map the network* option, it is always possible to define the port or port range to which the traffic should be sent to, and if to apply NAT on the traffic or not. If *Do not NAT* is chosen, it is not allowed to define a *Filter policy* under the *Access From* (advanced mode).

Warning: When selecting *IP*, *OpenVPN User*, *L2TP User* or *Load balancing*, keep in mind that port ranges will not be mapped 1 to 1, but rather a round robin balancing is performed. For example, mapping incoming ports 137:139 to destination ports 137:139 will result in these ports being used randomly: The incoming traffic to port 138 can unpredictably be redirect to either 137, 138, or 139. Leave the translation *Port/Range* field empty to avoid such occurrences!

Troubleshooting port-forwarding.

There are mainly two reasons why port-forwarding may not work.

1. The appliance is behind a NAT device.

In this case there is a device like a router or like another firewall between the appliance and the Internet, which disallows direct incoming connections. The solution is to configure a port forwarding also on that device to the RED IP of the appliance, if this is possible.

2. The destination server has wrong default gateway.

The server set as the destination of a port-forwarding rule is configured with a wrong or no default gateway. Connections will be directed to the target IP address but due to a wrong default gateway, packets will not be directed through the appliance. The solution is to correct the server's gateway.

5.2.2 Source NAT

In this page can be defined rules that apply SNAT to outgoing connections. The list of already defined rules is also displayed, for each of which the source and destination IP addresses, the service, the NAT status, a custom description of the rule, and the available actions are shown.

Source NAT can be useful if a server behind the IFA 3610 appliance has an own external IP and the outgoing packets should therefore not use the RED IP address of the firewall, but the one of the server. To add a new rule, click on **Add a new source NAT rule** and proceed like in the case of adding a port forwarding rule. Besides the *common options*, only one other setting can be configured:

NAT

Select to either apply *NAT*, *No NAT*, or *Map Network*. The choice to use SNAT allows the selection of the IP address that should be used among those presented in the drop-down menu. The *Auto* entries will automatically choose the IP address corresponding to the outgoing interface.

SNAT and a SMTP server in the orange zone.

In certain cases it is preferable to explicitly declare that *no Source NAT* be performed. An example would be a SMTP server in the DMZ, configured with an external IP, but whose outgoing connections should have the REDIP as the source. Configuring an SMTP server running on the IP 123.123.123.123 (assuming that 123.123.123.123 is an additional IP address of the uplink) in the DMZ with Source NAT can be done as follows:

1. Configure the ORANGE zone with any subnet (e.g., 192.168.100.0).
2. Setup the SMTP server to listen on port 25 on an IP in the ORANGE zone (e.g., 129.168.100.13).
3. In the *Menubar* ► *Network* ► *Interfaces* section, add a static Ethernet uplink with IP 123.123.123.123 to the appliance.
4. Add a source NAT rule and specify the ORANGE IP of the SMTP server as source address. Be sure to use NAT and set the NAT-ed source IP address to 123.123.123.123.

5.2.3 Incoming Routed Traffic

This tab allows to redirect traffic that has been routed through the appliance. This is very useful when having more than one external IP addresses and some of them should be used in the DMZ without the necessity to use NAT. The fields shown for every rule in the list are the traffic source and destination, the service, the policy to apply, a remark, and the available actions.

No other setting can be configured besides the *common options*.

5.3 Outgoing Traffic

The appliance comes with a pre-configured set of rules for outgoing traffic, i.e., to allow traffic flow of specific services, ports, and applications from the various zones to the RED interface and therefore the Internet. These rules are needed to ensure that the most common services always be able to access the Internet and work correctly. Two boxes are present on this page, one that shows the current rules and allows to add new ones, and one that allows to set the outgoing firewall options.

IFA 3610 appliance and Application Firewall.

Application firewalls are a recent development and improvement to stateful firewalls, that combine the ability of the latter to keep track of the connection's origin and path with those of Intrusion Prevention Systems to inspect packets' content, with the purpose to provide higher security from worm, viruses, malware, and all types of threats. The final result from the user experience point-of-view is that firewalls can block not only traffic between ports and IP addresses, but also traffic generated by single applications. This requires however, more efforts from the firewall: While traffic between IP addresses only needs that the first packet be inspected to block or allow the whole flow, to correctly recognise traffic generated by application, it is sometimes necessary the analysis of a few packets -usually not more than 3- of the flow.

Starting with version 3.0, every appliance is equipped with nDPI, an open source library implementing Deep Packet Inspection, thus allowing the deployment of rules for application firewalling. nDPI is deployed as a kernel module and interacts with iptables for the packet analysis.

Hence, there are now two different types of rules that can be defined on the outgoing firewall:

- *Stateful firewall rules*, that filter traffic between IP addresses and ports.
- *Application Rules*, i.e., rules that filter traffic generated by application.

When no application rules have been defined, the behaviour of the firewall is exactly the same as in previous version. Whenever an application rule has been defined, however, the stateful rules preceding it behave normally, while all the rules after undergo nDPI.

It is worth noting that the use of nDPI might present some subtleties, illustrated by the following example, and therefore might produce some unwanted side effect.

Suppose that a company wants to allow all HTTP traffic, except for youtube and gmail. The first default rule defined in appliance is to allow all HTTP traffic, with no restriction. This rule must therefore be disabled as first step. Then, two rules must be defined:

1. An application rule blocking the Gmail and Youtube protocols.
2. A stateful rule allowing all http traffic.

If rule 2. were an application rule with protocol HTTP, then only traffic recognised as HTTP by nDPI would be allowed, but other protocols using HTTP, like e.g., Yahoo and FaceBook would pass, since nDPI does not consider them as being HTTP, but independent protocols.

Current Rules

In detail, these are the services and protocols allowed by default to access the REDIP from the zones and shown in the top box:

GREEN: HTTP, HTTPS, FTP, SMTP, POP, IMAP, POP3s, IMAPs, DNS, ICMP

BLUE: HTTP, HTTPS, DNS, ICMP

ORANGE: DNS, ICMP

Everything else is forbidden by default except for the System rules which allow access to the services in the Henge™ Network. The system rules are defined even if the corresponding zones are not enabled.

Note: Access to Henge™ Network is not permitted to Community Edition appliances.

Possible actions on each rule are to enable or disable it, to edit it or delete it. Additional rules can be added by clicking on the **Add a new firewall rule** link at the top of the page. Please remember that the order of rules is important: the first matching rule decides whether a packet is allowed or denied, regardless of how many matching rules follow. The order of the rules can be changed by using the up and down arrow icons next to each rule.

The following settings differ from the default *common options*.

Source

It can be one or more Zone/Interfaces, Network/IP, or MAC addresses.

Destination

Can be the RED zone, one or more uplinks, or one or more network/host addresses accessible outside the RED interface.

Application

This search widget allows to select the applications that should be part of the rule. Applications are dividend into categories (e.g., Database, filesharing, and so on).

Hint: Enter at least one letter to show all applications whose name starts with that letter.

Outgoing Firewall Settings

It is possible to disable or enable the whole outgoing firewall by clicking on the *Enable Outgoing firewall* switch. When disabled, all outgoing traffic is allowed and no packet is filtered: This setting is however strongly discouraged and the recommendation is to keep the outgoing firewall enabled.

Log accepted outgoing connections


Ticking this checkbox causes all the accepted connections to the RED interface to be logged.

Proxy and outgoing firewall.

Whenever the proxy is activated for a given service (e.g., HTTP, POP, SMTP, DNS), the firewall rules in the outgoing firewall will take no effect, because of the nature of the proxy.

With the proxy activated, whenever a connection starts from a client to the Internet, it will either be intercepted by the proxy on the appliance (in transparent mode) or go directly to the firewall, but never go *through* the firewall. The proxy then starts a new connection to the real destination, gets the data and sends it to the client. Those connections to the Internet always start from the appliance, which hides the clients internal IP address. Therefore, such connections never go through the outgoing firewall, since in fact they are local connections.

5.4 Inter-Zone Traffic

This module permits to set up rules that determine how traffic can flow between the local network zones, excluding therefore the RED zone (traffic through the RED zone can be filtered in *Outgoing traffic* and *Port forwarding / NAT*). To activate the inter-zone firewall, click on the grey switch . Two boxes are present on this page, one that shows the current rules and allow to add new ones, and one that allows to set the inter-zone firewall options.

Current Rules

The appliance comes with a simple set of pre-configured rules: traffic is allowed from the GREEN zone to any other zone (ORANGE and BLUE) and within each zone, with everything else forbidden by default.

Analogously to the outgoing traffic firewall, rules can be disabled/enabled, edited or deleted by clicking on the appropriate icon on the right side of the table. New rules can be added by clicking on the **Add a new inter-zone firewall rule** link at the top of the page. Only the *common options* can be configured.

Inter-Zone Firewall Settings

The inter-zone firewall can be disabled or enabled by using the *Enable Inter-Zone firewall* switch. When disabled, all traffic is allowed among all the BLUE, GREEN, and ORANGE zones. Disabling the inter-zone firewall is strongly discouraged.

Log accepted Inter-Zone connections

Ticking this checkbox causes all the accepted connections among the zones to be logged.

5.5 VPN traffic

The VPN traffic firewall allows to add firewall rules applied to the users and hosts that are connected via OpenVPN.

The VPN traffic firewall is normally not active, which means that, on the one side, the traffic can freely flow between the VPN hosts and the hosts in the GREEN zone, and on the other side, VPN hosts can access all other zones. Please note that VPN hosts are *not subject* to the outgoing traffic firewall or the Inter-Zone traffic firewall. Two boxes are present on this page, one that shows the current rules and allow to add new ones, and one that allows to set the VPN firewall options.

Current Rules

The handling and definition of the rules is identical to the outgoing traffic firewall, so please refer to that section and to the *common options* for directions on the definition and handling of the firewall rules in this module.

VPN Firewall Settings

The VPN firewall can be disabled or enabled using the *Enable VPN firewall* switch.

Log accepted VPN connections

Ticking this checkbox causes all the accepted connections from the VPN users to be logged.

5.6 System access

This section governs the rules that grant or deny access to the appliance itself.

There is a list of pre-configured rules that cannot be changed, whose purpose is to guarantee the proper working of the firewall. Indeed, there are services, among those supplied by the appliance, that require to be accessed from clients in the various local zones. Examples include using the DNS (which requires that the port 53 be open) to resolve remote hostnames or using the administration web interfaces (which uses port 10443): Whenever one of these services is activated, one or more rules are automatically created to allow the proper efficiency of the service itself.

The list of the pre-defined rules is shown when clicking on the **Show rules of system services** button at the bottom of the page.

More system access rules can be added by clicking on the **Add a new system access rule** link. The setting specific to this module of the firewall are:

Log packets

All packets that access or try to access the appliance are logged when this checkbox is ticked. This option proves useful to know who accessed -or tried to access- the system.

Source address

The MAC addresses of the incoming connection.

Source interface

The interface from which the system can be accessed

Note: There is no *Destination* address, as it is the IP address of the interface from which the access is granted or attempted.

Actions are to disable/enable, edit, or delete a rule from the list of rules.



5.7 Firewall Diagrams

This page shows, for each of the modules described in this page, a diagram that shows how the traffic flows among the zones, and which is the firewall module that takes charge of the various flows. The green arrowed lines show which traffic is allowed in each zone and in which directions. In the case of VPN, the arrows from/to the RED interface are marked with a red 'X', meaning that the traffic is not possible between them.

When an image is clicked, it will be opened into a gallery that allows to browse all of them like in a slide show.

Chapter 6: Proxy

6.1 DNS

The DNS proxy is a proxy server that intercepts DNS queries and answers them, without the need to contact a remote DNS server each time it is necessary to resolve an IP address or a hostname. When a same query is often repeated, caching its results locally may sensibly improve performances. The available settings for the DNS proxy are grouped into three tabs.

6.1.1 DNS Proxy

A few options for the DNS proxy can be configured in this page.

Transparent on Green, Transparent on Blue, Transparent on Orange

Enable the DNS proxy as transparent on the GREEN, BLUE, and ORANGE zone, respectively. They appear only if the corresponding zones are enabled.

Specific sources and destinations can be set up to bypass the proxy by filling in their values in the two text areas.

Which sources may bypass the transparent Proxy

Allow the sources under the corresponding text area not to be subject to the DNS scanning. The sources can be specified as IP addresses, networks, or MAC addresses.

Destinations to which bypass the transparent Proxy

Allow the destinations under the corresponding text area not to be subject to the DNS proxy scanning. The destinations can be specified as IP addresses or networks.

6.1.2 DNS Routing

This page allows the management of custom domain - nameservers pairs. In a nutshell, whenever a sub-domain of a domain is queried, the corresponding nameserver in the list will be used to resolve the domain into the correct IP address.

A new domain - nameserver combination can be added by clicking on the **Add new custom nameserver for a domain** link. When adding an entry, a few values can be entered for the various options available:

Domain

The domain for which to use the custom nameserver.



DNS Server

The IP address of the nameserver.

Remark

An additional comment.

On each domain in the list, these actions can be carried out:

-  - edit the rule.
-  - delete the rule.

6.1.3 Anti-Spyware

This page presents configuration options about the reaction of the appliance when asked to resolve a domain name that is known to be either used to propagate spyware or that serves as phishing site. The options that can be set are:

Enabled

The requests are redirected to localhost. In other words, the remote site will neither be contacted nor reachable.

Whitelist domains

Domain names that are entered here are not treated as spyware targets, regardless of the list's content.

Blacklist domains

Domain names that are entered here are always treated as spyware targets, regardless of the list's content.

Spyware domain list update schedule

The update frequency of the spyware domain list. Possible choices are *Daily*, *Weekly*, and *Monthly*. By moving the mouse cursor over the respective question mark, the exact time of the update execution is shown.

Hint: To download updated signatures, the system must be registered to Henge™ Network.

Chapter 7: The VPN Menu

7.1 OpenVPN Server



When configured as an OpenVPN server, the appliance can accept remote connections from the uplink and allow a VPN client to be set up and work as if it were a local workstation or server.

Starting with version 3.0, the OpenVPN server deployed on the appliance allows the simultaneous presence of several instances. Each server will listen to one different port, accepting incoming connections on that port only. Moreover, when the hardware on which appliance is installed has multiple CPU cores, every instance may be assigned more than one core, thus resulting in an increase of the throughput and data processing of that instance. It is nevertheless also possible to have multiple instances of OpenVPN running on a device equipped with a single-core CPU, though this results in the CPU carrying the load of all instances.

The OpenVPN server settings page is composed of two tabs: *Server configuration* and *VPN client download*.

7.1.1 Server Configuration

This page shows a switch *Enable OpenVPN server* , that will start the OpenVPN server and all services related to it (like e.g., the *VPN firewall* if enabled) once clicked. Below, there is one box, *OpenVPN settings*, that allows to set up some global settings. Right below, a link allows to define a new server instance while at the bottom of the page there's the list of the available OpenVPN servers running on the appliance, if any has already been defined. The list shows the following data about each OpenVPN server instance defined: The name, remark, and details about the configuration, namely: The port on which it is listening, the protocol, the type of device, and the type of network. Finally, the actions available are:

- - the server is active or stopped.
-  - modify the server's configuration
-  - remove the configuration and the server.

Note: When starting the OpenVPN server for the first time, the root and host certificates are generated automatically.

7.1.2 OpenVPN Settings

The box on the top shows the current OpenVPN settings, which concern the authentication method, and are:

Authentication type

There are three available authentication methods to connect clients to the OpenVPN server running on the appliance:

- *PSK (username and password)*. Connection is established after providing correct username and password.
- *X.509 certificate*. A valid certificate only is needed to connect.
- *X.509 certificate & PSK (two factor)*. Besides a valid certificate, username and passwords are needed.

Warning: When employing certificate-only authentication, a client with a valid certificate will be granted access to the OpenVPN server even if it has no valid account!

The appliance's default method is *PSK (username/password)*: The client authenticates using username and password. To use this method, no additional change is needed, while the other two methods are described below.

Certificate configuration

This drop-down menu is used to select the method of creation of a new certificate. The available options are:

- *Generate a new certificate*. Create a new certificate from scratch. This option is only available if no host certificate has already been generated. A form will open where to specify all options necessary to create a new certificate. These are the same found in the *new certificates generation* editor, with two slight changes: *Common name* becomes *System hostname* and *Organizational unit name* becomes *Department name*.

- *Use selected certificate.* Select one certificate from those available, shown on the right-hand side of the drop-down menu. It is possible to see the full details of this certificate by clicking on the **View details** hyperlink.

Hint: The name of the certificate selected appears right above the hyperlink.

- *Use an existing certificate.* A second drop-down menu on the left allows to select a certificate that has already been created and stored on the appliance.
- *Upload a certificate.* By clicking on the **Browse...** button that appears underneath the drop-down menu it will be possible to select from the workstation and to upload an existing certificate. The password for the certificate, if needed, can be provided in the textfield on the right-hand side.
- *Upload a certificate signing request.* The **Browse...** button that appears underneath the drop-down menu can be clicked to select from the workstation and upload an existing certificate signing request. The validity of the certificate in days can be provided in the textfield on the right-hand side.

7.1.3 OpenVPN Server Instances

The list of already defined OpenVPN instances is shown in this panel, above which is present the **Add new OpenVPN server instance** hyperlink. A click on this link will open an editor in which to provide all the necessary configuration values for a new VPN instance.

Note: When the number of OpenVPN instances is greater than the cores, a yellow callout informs that the performances may degrade.

In the editor, the following configuration options are shown.

Name

The name given to the OpenVPN server instance.

Remark

A comment for this instance.

Bind only to

The IP address to which the instance should listen to.

Port

The port on which the instance waits for incoming connections.

Device type

The device used by the instance, chosen between TUN and TAP from the drop-down menu. TUN devices require that the traffic be routed, hence the option *Bridged* below is not available for TUN devices.

Protocol

The protocol used, chosen between TCP and UDP from the drop-down menu.

Bridged

Tick this option to run the OpenVPN server in bridged mode, i.e., within one of the existing zones.

Note: If the OpenVPN server is not bridged (i.e., it is routed), the clients will receive their IP addresses from a dedicated subnet. In this case, appropriate firewall rules in the *VPN firewall* should be created, to make sure the clients can access any zone, or some server/resource (e.g., a source code repository). If the OpenVPN server is bridged, it inherits the firewall settings of the zone it is defined in.

VPN subnet

This option is the only available if bridged mode is disabled. It allows the OpenVPN server to run in its own, dedicated subnet, that can be specified in the text box and should be different from the subnets of the other zones.

Bridge to

The zone to which the OpenVPN server should be bridged. The drop-down menu shows only the available zones.

Dynamic IP Pool Start address

The first possible IP address in the network of the selected zone that should be used for the OpenVPN clients.

Dynamic IP Pool End address

The last possible IP address in the network of the selected zone that should be used for the OpenVPN clients.

Routed and bridged OpenVPN server, static and dynamic.

When configuring a pool of IP addresses to be reserved for clients connecting via OpenVPN, it is necessary to keep in mind a few guidelines that help both the prevention of future malfunctioning and the cleaner and easier design and set up.

Before starting the configuration of the server, there is a golden rule to remember, concerning the implementation of the VPN multicore architecture: Regardless of the bridged or routed mode used for a multicore VPN server instance, the reservation of static IP addresses is neglected. In other words, a client connecting to this VPN server, will receive a dynamic IP address, even though in her configuration there is a static IP assignment.

The first choice is to define whether the OpenVPN server should act in routed or bridged mode. In the former case, it is necessary to define a suitable *VPN subnet* that will provide the IP addresses for the clients. The traffic directed to this subnet has to be filtered, if necessary, using the *VPN firewall*. In the latter case, the OpenVPN server is configured to consider the clients, upon connecting, as they were physically connected to that zone, i.e., the server *bridges* the client to one of the zones. In this case, a pool of IP addresses must be defined within that zone using the two option that appear right before this box. This pool must be entirely contained in the zone's subnet and smaller than that one. It is also important to make sure that this pool does not conflict with other pools defined in that zone, like e.g., a DHCP server.

In a bridged OpenVPN server it is possible to assign to some (or even to all) user a static IP address. When planning this possibility, it is a good practice that these static IP addresses do not belong to any of the IP pools defined in that zone, to prevent any conflicts of address and wrong routing. Traffic to this particular client can then be filtered using the VPN (or IPsec) user as source or destination of traffic in the Firewall rules.

In the **Advanced options** box, additional options can be configured.

Number of cores

The drop-down menu allows to choose how many CPUs of the appliance can be used by the instance, hence the options in the drop-down menu may vary.

Allow multiple connections from one account:

Normally, one client is allowed to connect from one location at a time. Selecting this option permits multiple client logins, even from different locations. However, when the same client is connect twice or more, the VPN firewall rules do not apply anymore.

Block DHCP responses coming from tunnel

Tick this checkbox when receiving DHCP responses from the LAN at the other side of the VPN tunnel that conflict with the local DHCP server.

Client to client connections

Select from the drop-down menu the modalities of the communications between clients of the OpenVPN server. This option is only available on single-process servers, i.e., on servers running only one instance of the OpenVPN server.

- **Not allowed:** The clients can not communicate one to the other.
- **Allow direct connections:** The clients can communicate directly with each other but filtering is not possible.
- **Filter connections in the VPN firewall:** The clients can communicate with each other, but their traffic is redirected to the VPN Firewall and can be filtered using suitable rules there.

Note: In case of appliances having multi-core CPUs, there is no selection possible and the option *Filter connections in the VPN firewall* is automatically activated.

Push these nameservers

By ticking this checkbox, the nameserver specified in the textfield below are sent to the clients upon connection.

Nameservers

The nameservers specified in this textfield are sent to the connected clients, when the previous checkbox has been ticked.

Push these networks

By ticking this checkbox, the routes to the networks defined in the textfield below are sent to the connected clients.

Networks

The networks specified in this textfield are sent to the connected clients, when the previous checkbox has been ticked.

Push this domain

By ticking this checkbox, the search domain defined in the textfield on the right-hand side, is added to those of the connected clients.

Domain

The domain that will be used to identify the servers and network resources in the VPN network (i.e., the *search domain*).

Note: The options *Push these nameservers* and *Push domain* only work for clients running the Microsoft Windows operating system.

The first time the service is started a new, self-signed CA certificate for this OpenVPN server is generated, an operation that may take a long time. After the certificate has been generated, it can be downloaded by clicking on the **Download CA certificate** link. This certificate must be used by all the clients that want to connect to this OpenVPN server, otherwise they will not be able to access.

After the server has been set up, it is possible to create and configure accounts for clients that can connect to the appliance in the *Authentication* tab.

Enabled

Tick this checkbox to make sure the OpenVPN server is started.

Troubleshooting VPN connections.

While several problem with VPN connections can be easily spotted by looking at the configuration, one subtle source of connections hiccups is a wrong value of the MTU size. The appliance sets a limit of 1450 bytes to the size of the VPN's MTU, to prevent problems with the common MTU value used by the ISP, which is 1500. However, some ISP may use a MTU value lower than the commonly used value, making the appliance's MTU value too large and causing connection issues (the most visible one is probably the impossibility to download large files). This value can be modified by accessing the appliance from the CLI and following these guidelines:



1. Write down the MTU size used by the ISP (see link below).
2. Login to the CLI, either from a shell or from *Menu* ► *System* ► *Web Console*.
3. Edit the OpenVPN template with an editor of choice: `nano /etc/openvpn/openvpn.conf.tmpl`.
4. Search for the string **mssfix 1450**.
5. Replace 1450 with a lower value, for example 1200.
3. Restart OpenVPN by calling: `jobcontrol restart openvpnjob`.

7.1.4 VPN Client Download

Click on the link to download the VPN client for Microsoft Windows, MacOS X, and Linux from the Henge™ Network. A valid account is needed to download the client.

7.2 OpenVPN Client (Gw2Gw)

In this page appears the list of the appliance's connections as OpenVPN clients, i.e., all tunnelled connections to remote OpenVPN servers. For every connection, the list reports the status, the name, any additional option, a remark, and the actions available:

- - the server is active or stopped.
-  - modify the server's configuration
-  - remove the configuration and the server.

The status is *closed* when the connection is disabled, *established* when the connection is enabled, and *connecting...* while the connection is being established. Beside to enable and to disable a connection, the available actions are to edit or delete it. In the former case, a form will open, that is the same as the one that opens when adding a connection (see below) in which to see and modify the current settings, whereas in the latter case only deletion of that profile from the appliance is permitted.

The creation of a new OpenVPN client connections is straightforward and can be done in two ways: Either click on the **Add tunnel configuration** button and enter the necessary information about the OpenVPN server to which to connect (there can be more than one) or import the client settings from the OpenVPN Access Server by clicking on **Import profile from OpenVPN Access Server**.

7.2.1 Add Tunnel Configuration

There are two types of settings that can be configured for each tunnel configuration: The basic one includes mandatory options for the tunnel to be established, while the advanced one is optional and normally should be changed only if the OpenVPN server has a non-standard setup. To access the advanced settings, click on the >> button next to the *Advanced* tunnel configuration label. The basic settings are:

Connection name

A label to identify the connection.

Connect to

The remote OpenVPN server's FQDN, port, and protocol in the form `myvpn.example.com:port:protocol`. The port and protocol are optional and left on their default values which are *1194* and *udp* respectively when not specified. The protocol must be specified in lowercase letters.

Upload certificate

The server certificate needed for the tunnel connection. Browsing the local filesystem is admitted, to search for the file, of the path and filename can be entered. If the server is configured to use PSK authentication (password/username), the server's host certificate (i.e., the one downloaded from the **Download CA certificate** link in the server's *Menubar ► VPN ► OpenVPN server* section) must be uploaded to the appliance. Otherwise, to use certificate-based authentication, the server's PKCS#12 file (i.e., the one downloaded from the **Export CA as PKCS#12 file** link on the server's *Menubar ► VPN ► OpenVPN server ► Advanced section*) must be uploaded.

PKCS#12 challenge password

Insert here the *Challenge* password, if one was supplied to the CA before or during the creation of the certificate. This is only needed when uploading a PKCS#12 certificate.

Username, Password

If the server is configured to use PSK authentication (password/username) or certificate plus password authentication, provide here the username and password of the account on the OpenVPN server.

Remark

A comment on the connection.

7.2.2 Advanced Tunnel Configuration

In this box, that appears when clicking on the >> button in the previous box, additional options can be modified, though the values in this box should be modified only if the server side has not been configured with standard values.

Fallback VPN servers

One or more (one per line) fallback OpenVPN servers in the same format used for the primary server, i.e., `myvpn.example.com:port:protocol`. The port and protocol values default to `1194` and `udp` respectively when omitted. If the connection to the main server fails, one of these fallback servers will take over.

Hint: The protocol must be written in lowercase letters.

Device type

The device used by the server, which is either TAP or TUN.

Connection type

This drop-down menu is not available if TUN has been selected as *Device type*, because in this case the connection type is always *routed*. Available options are *routed* (i.e., the client acts as a gateway to the remote LAN) or *bridged* (i.e., the client firewall appears as part of the remote LAN). Default is *routed*.

Bridge to

This field is only available if TAP has been selected as *Device type* and the connection type is *bridged*. From this drop-down menu, select the zone to which this client connection should be bridged.

NAT

This option is only available if the *Connection type* is *routed*. Tick this checkbox to hide the clients connected through this appliance behind the firewall's VPN IP address. This configuration will prevent incoming connections requests to the clients. In other words, incoming connections will not see the clients in the local network.

Block DHCP responses coming from tunnel

Tick this checkbox to avoid receiving DHCP responses from the LAN at the other side of the VPN tunnel that conflict with a local DHCP server.

Use LZO compression

Compress the traffic passing through the tunnel, enabled by default.

Protocol

The protocol used by the server: UDP (default) or TCP. Set to TCP only if an HTTP proxy should be used: In this case, a form will show up to configure it.

If the appliance can access the Internet only through an upstream HTTP proxy, it can still be used as an OpenVPN client in a Gateway-to-Gateway setup, but the *TCP* protocol for OpenVPN must be selected on both sides. Moreover, the account information for the HTTP upstream proxy must be provided in the text fields:

HTTP proxy

The HTTP proxy host, e.g., `proxy.example.com:port`, with the port defaulting to 8080 if not entered.

Proxy username, Proxy password

The proxy account information: The username and the password.

Forge proxy user-agent

A forged *user agent* string can be used in some cases to disguise the appliance as a regular web browser, i.e., to contact the proxy as a browser. This operation may prove useful if the proxy accepts connections only for some type of browsers.

Once the connection has been configured, a new box at the bottom of the page will appear, called *TLS authentication*, from which to upload a TLS key file to be used for the connection. These options are available:

TLS key file

The key file to upload, searchable on the local workstation.

MD5

The MD5 checksum of the uploaded file, which will appear as soon as the file has been stored on the appliance.

Direction

This value is set to 0 on servers and to 1 on clients.

7.2.3 Import Profile from OpenVPN Access Server

The second possibility to add an account is to directly import the profile from an OpenVPN Access Server: In this case, the following information must be provided:

Connection name

A custom name for the connection.

Access Server URL

The URL of the OpenVPN Access Server.

Note: Note that the appliance only supports XML-RPC configuration of the OpenVPN Access Server, therefore a URL input here has the form: `https://<SERVERNAME>/RPC2`.

Username, Password

The username and password on the Access Server.

Verify SSL certificate

If this checkbox is ticked and the server is running on an SSL encrypted connection, then the SSL certificate will be checked for validity. Should the certificate not be valid then the connection will be immediately closed. This feature might be disabled when using a self-signed certificate.



Remark

A comment to recall the purpose of the connection.

7.3 IPsec

The IPsec page contains two tabs (IPsec and L2TP), that allow to set up and configure the IPsec tunnels and to enable the L2TP support, respectively.

7.3.1 IPsec

To enable L2TP on the appliance, the switch next to the *Enable L2TP* label should be blue . If it is grey , click on it to start the service.

The IPsec tab contains two boxes: The first one is *IPsec settings*, which concerns the certificate choice and various options, also for debugging purposes. The second one is *Connections*, which shows all the connections and allows to manage them.

IPsec, L2TP, and XAuth in a nutshell.

IPsec is a generic standardised VPN solution, in which the encryption and the authentication tasks are carried out on the OSI layer 3 as an extension to the IP protocol. Therefore, IPsec must be implemented in the kernel's IP stack. Although IPsec is a standardised protocol and it is compatible to most vendors that implement IPsec solutions, the actual implementation may be very different from vendor to vendor, sometimes causing interoperability issues.

Moreover, the configuration and administration of IPsec may become quite difficult due to its complexity and design, while some particular situations might even be impossible to handle, for example when there is the necessity to cope with NAT.

Compared to IPsec, OpenVPN is easier to install, configure, and manage. However, mobile devices rely on IPsec, thus the appliance implements an easy-to-use administration interface for IPsec, that supports different authentication methods and also two-factor authentication when used together with L2TP or XAuth.

Indeed, IPsec is used to authenticate clients (i.e., tunnels) but not users, so one tunnel can be used by only one client at a time.

L2TP and XAuth add user authentication to IPsec, therefore many clients can connect to the server using the same encrypted tunnel and each client is authenticated by either L2TP or XAuth.

An additional option is available when using XAuth and is called *XAuth hybrid mode*, which only authenticates the user.

7.3.2 IPsec Settings

In this box a few global IPsec options can be set, namely two for Dead peer detection, and quite a lot debugging options. Additionally, configuration of certificates used in IPsec tunnelled connections is also carried out here.

Roadwarriors virtual IP pool

The IP interval from which all roadwarrior connections receive their IP address.

Ping delay (in seconds)

The amount of seconds between two successive pings, used to detect whether the connection is still active.

Timeout interval (in seconds) - IKEv1 only

The maximum amount in seconds of the exchange interval for the IKEv1 protocol.

Hint: IKEv2 does not need a timeout interval, as it is capable of detecting when the other endpoint does not reply and which actions to take.

Certificate configuration

Certificate configuration and management is carried out exactly like in the case of *OpenVPN server* (in *MenuBar* ► *VPN* ► *OpenVPN server*), in which all the various management modalities are explained.





7.3.3 Debug Options

Debug options are rather advanced settings and usually not needed, as they only will increase the number of events and messages recorded in the log file.

The activation of all those options proves useful when issues are experienced during the establishment of a connection or to produce more precise and technical messages about the normal operations of a tunnel. This way, the log file will contain very detailed options.

7.3.4 Connections

In this table are shown all the already configured IPsec connection, with the following information:

- Name. The name given to the connection.
- Type. What kind of tunnel is used.
- Common Name. The name of the certificate used to authenticate the connection.
- Remark. A comment about the connection.
- Status. Whether the connection is either *Closed*, *Connecting* or *Established*.
- Actions. The possible operations that can be made on each tunnel:
 - - the connection is active or not.
 -  - modify the connection's configuration
 -  - restart the connection.
 -  - display detailed information about the connection.
 -  - remove the connection.

Hint: When a connection is reset from the appliance, it is necessary for the client to reconnect in order to establish the connection.

Upon clicking on **Add new Connection**, a panel will appear, which contains all options needed to set up a new IPsec connection.

Name

The name of the connection.

Remark

A comment for the connection.

Connection type

There are four different connection modalities can be chosen for the IPsec tunnel:

- *Host-to-Net*. The client is connecting to the IPsec server on the IFA 3610 appliance is a single remote workstation, server, or resource.
- *Net-to-Net*. The client is an entire subnet. In other words, the IPsec connection is established between remote subnets.
- *L2TP Host-to-Net*. The client is a single device, using also L2TP.
- *XAuth Host-to-Net*. The client is a single device and authentication is carried out by XAuth.

Hint: Linux users can read more about XAuth by reading the **Xsecurity(7)** manpage, also available online for everyone.

The options available for each of them are basically same, with only one more option available for Net-to-Net connections.

Authentication Type

The option selected from the drop-down menu determines how the client's authentication is carried out. Available values are:

- *Password (PSK)*. The client shall supply the password specified in the Use a *pre-shared key* textfield situated on the right.
- *Peer is identified by either IPV4_ADDR, FQDN, USER_FQDN or DER_ASN1_DN string in remote ID field*. The client is authenticated by its IP Address, domain name, or by other unique information of the IPsec tunnel.
- *Use an existing certificate*. The certificate chosen from the drop-down menu on the right shall be used.
- *Generate a new certificate*. Additional options will be shown to create a new certificate.
- *Upload a certificate*. Select from the local workstation a certificate to use.
- *Upload a certificate request*. Select from the local workstation a certificate request to obtain a new certificate.
- *XAUTH hybrid*. Only available for *XAuth Host-to-Net* connections: The user will authenticate, while the encryption tunnel must not.

Local ID

A string that identifies the client within the local network.

Interface

The interface through which the host is connecting.

Local subnets

The local subnets that will be accessible from the client.

Note: Mobile devices running iOS can not properly connect via XAuth to the appliance if this value is not set, therefore the special subnet *0.0.0.0/0* is automatically added when the *Connection type* is set to XAuth.

Hint: Only when using IKEv2 it is possible to add more than one subnet, one per line, since IKEv1 only supports one subnet.

Remote ID

The ID that identifies the remote host of the connection.

Remote subnet

Only available for Net-to-Net connections, it specifies the remote subnet.

Hint: When using IKEv2 it is possible to add more than one subnet.

Remote host/IP

The IP or FQDN of the remote host.

Note: When a hostname is supplied in this option, it must match the *local ID* of the remote side.

Roadwarrior virtual IP

The IP Address specified in the textfield will be assigned to the remote client.

Hint: This IP Address must fall within the pool defined in the *IPsec settings* below.

Note: This option is available neither for L2TP Host-to-Net connections, as it is L2TP that takes charge of IP address assignment to clients, nor for Net-to-Net connections.

Dead peer detection action

The action to perform if a peer disconnects. Available choices from the drop-down menu are to *Clear*, to *Hold*, or to *Restart* the peer.

By clicking on the **Advanced** label, additional options are available, to choose and configure different types of encryption algorithm. For every option, many types of algorithm can be chosen.

Note: It is necessary to change algorithm only in case some remote client uses a given algorithm and can not change it.

IKE encryption

The encryption methods that should be supported by IKE.

IKE integrity

The algorithms that should be supported to verify the integrity of packets.

IKE group type

The IKE group type.

IKE lifetime

How many hours are the IKE packets valid.

ESP encryption

The encryption methods that should be supported by the ESP.

ESP integrity

The algorithms that should be supported to verify the integrity of packets.

ESP group type

The ESP group type.

ESP lifetime

How many hours should an ESP key be valid.

Negotiate payload compression

Tick the checkbox to allow payload compression.

7.4 L2TP

L2TP, the Layer 2 Tunnelling Protocol, is described in RFC 2661.

To enable L2TP on the appliance, the switch next to the *Enable L2TP* label should be green. If it is grey, click on it to start the service.

The following options are available to configure L2TP.

Zone

The zone to which the L2TP connections are directed. Only the activated zones can be chosen from the drop-down menu.

L2TP IP pool start address, L2TP IP pool end address

The IP range from which L2TP users will receive an IP address when connecting to the appliance.

Enable debug

Tick this checkbox to let L2TP produce more verbose logs.

How to create a Net-To-Net VPN with IPsec using certificate authentication.

Scenario:

- Firewall CoreFW - REDIP: 100.100.100.100, GREENIP: 10.10.10.1/24
- Firewall LocalFW - REDIP: 200.200.200.200, GREENIP: 192.168.0.1/24

Problem: Connect LocalFW to CoreFW using IPsec.

Solution:

- The following steps have to be performed on CoreFW:

1. Go to *Menu* ► *VPN* ► *IPsec*, enable IPsec, and specify *100.100.100.100* as Local VPN hostname/IP.
2. After saving, click on the **Generate host/root certificate** button, unless they have already been generated, and compile the form.
3. Download the host certificate and save it as **fw_a_cert.pem**.
4. In the *Connection status and control* box click on the *Add* button, then select *Net-to-Net*. In the page that opens, two box will appear.
5. In *Connection configuration* enter *200.200.200.200* in the *Remote host/IP* field, *10.10.10.0/24* as *Local subnet* and *192.168.0.0/24* as *Remote subnet*.
6. In the *Authentication box* select *Generate a certificate* and compile the form. Make sure to set a password.
7. After saving, download the PKCS12 file and save it as **fw_a.p12**.

- The following steps have to be performed on LocalFW:

1. Go to *Menu* ► *VPN* ► *IPsec*, enable IPsec, and specify *200.200.200.200* as Local VPN hostname/IP.
2. After saving click on the **Generate host/root certificate** button. If they had already been generated, **Reset** the previous certificates.
3. In the *Generate host/root certificate*, **Do not** fill in any field in the first section! Instead, upload the **fw_a.p12** file saved from CoreFW, enter the password, and click on the **Upload PKCS12 file**.
4. Click on **Add** in the *Connection status and control* box, then select *Net-to-Net*. In the page that opens, two box will appear.
5. In *Connection configuration* enter *100.100.100.100* in the *Remote host/IP* field, *192.168.0.0/24* as *Local subnet* and *10.10.10.0/24* as *Remote subnet*.
6. In the *Authentication box* select *Upload a certificate* and upload the **fw_a_cert.pem** that have created on MainFW.

7.5 Authentication

This page shows three tabs, which allow to manage local *Users*, local *Groups*, and *Settings* for remote authentication

7.5.1 Users

In this page, all users that have an account on the appliance's VPN server are displayed in the table, and for each the following information are shown:

- Name. The name of the user.
- Remark. A comment.
- Authentication server. The server used for the user authentication, which is either *local* (the appliance itself) or *LDAP* (an external LDAP server, configurable in the *Settings* tab).
- Actions. The available operation that can be carried out on the account. For LDAP users they are *Enable/Disable* and *Edit*, for local users, there is also the possibility to *Delete*. Editing an LDAP user only allows to modify its local options, not of other data like username or password, which are entirely managed by the LDAP server.

Click on **Add new local user** above the table to add a new local account. In the form that will show up, the following options can be specified for each user.

Add new local user

Username

The login name of the user.

Remark

An additional comment.

Password, Confirm password

The password for the user, to be entered twice. The passwords are actually not shown: To see them, tick the two checkboxes on their right.

Certificate configuration

Select the mode to assign a certificate to the user. The available modes are selectable from the drop-down menu: *Generate a new certificate*, *Upload a certificate*, and *Upload a Certificate signing request*. Upon selection, below the drop-down menu appear the available options for each mode, which are described in the Certificates page.

Organizational unit name

The Organisation Unit to which the user belongs to, i.e., the company, enterprise, or institution department identified with the certificate.

Organization name

The organisation to which the user belongs to.

City

The city (L) in which the organisation is located.

State or province

The state or province (ST) in which the organisation is located.

Country

The Country (C) in which the organisation is located, chosen from those in the selection menu. By typing one or more letters, matching countries are searched for and displayed.

Email address

The e-mail address of the user.

Group membership

In this part of the panel it is possible to assign membership to one or more groups to the user. In the search widget it is possible to filter existing groups to find matching groups. Group membership is added by clicking on the **+** on the right of the group name. Groups to which the user belongs are shown in the textfield below. There are also shortcuts to **Add all** and to **Remove all** groups memberships at once.

Override OpenVPN options

Tick this checkbox to allow the OpenVPN protocol to be used. This option will reveal a box in which to specify custom option for the account, see below.

Override L2TP options

Tick this checkbox to show a box in which to choose the L2TP tunnel to be used.

Note: This option can not be selected if no L2TP tunnel has yet been configured. In such a case, an informative message appears as a hyperlink: Upon clicking on it, the IPsec connection editor opens. Once done, it will be possible to allow a VPN user to connect using the L2TP Protocol.

Hint: The box for L2TP options will appear below the *OpenVPN options* box, if also OpenVPN option are to be overridden.

Enabled

Tick the checkbox to enable the user, i.e., to allow her to connect to the OpenVPN server on the appliance.

OpenVPN Options

direct all client traffic through the VPN server

If this option is checked, all the traffic from the connecting client, regardless of the destination, is routed through the uplink of the appliance. The default is to route all the traffic whose destination is outside any of the internal zones (such as Internet hosts) through the client's uplink.

Push only global options to this client

For advanced users only. Normally, when a client connects, tunnelled routes to networks that are accessible via VPN are added to the client's routing table, to allow it to connect to the various local networks reachable from the appliance. This option should be enabled if this behaviour is not wanted, but the client's routing tables (especially those for the internal zones) should be modified manually.

Push route to GREEN [BLUE, ORANGE] zone,

When this option is active, the client will have access to the GREEN, BLUE, or ORANGE zone. These options have no effect if the corresponding zones are not enabled.

Networks behind client

This option is only needed if this account is used as a client in a Gateway-to-Gateway setup. In the box should be written the networks laying behind this client that should be pushed to the other clients. In other words, these networks will be available to the other clients.

Static IP addresses

Dynamic IP addresses are assigned to clients, but a static IP address provided here will be assigned to the client whenever it connects.

Note: If the client connects to a multicore VPN server running on the appliance, this assignment will not be taken into account.

Push these nameservers

Assign custom nameservers on a per-client basis here. This setting (and the next one) can be defined, but enabled or disabled at will.

Push these domains

Assign custom search domains on a per-client basis here.

Note: When planning to have two or more branch offices connected through a Gateway-to-Gateway VPN, it is good practice to choose different subnets for the LANs in the different branches. For example, one branch might have a GREEN zone with the 192.168.1.0/24 subnet while the other branch uses 192.168.2.0/24. Using this solution, several possible sources for errors and conflicts will be avoided. Indeed, several advantages come for free, including: The automatic assignment of correct routes, without the need for pushing custom routes, no warning messages about possibly conflicting routes, correct local name resolution, and easier WAN network setup.

L2TP Options

IPsec Tunnel

This drop-down menu allows to choose the tunnel that will be employed by the user, among those already defined.

7.5.2 Groups

In this page a table is displayed, which shows all the groups that are either defined on the appliance or on an external LDAP server. For each group the following information are shown:

- Groupname. The name of the group.
- Remark. A comment.
- Authentication server. The server used for the user authentication, which is either *local* (the appliance itself) or *LDAP* (an external LDAP server, configurable in the *vpnauthsettings* tab).
- Actions. The available operation that can be carried out on the account. For LDAP servers the only action is to *Edit* the local properties, while for local groups there is also the possibility to *Delete* the group.

Click on **Add new local groups** above the table to add a new local group. In the form that will show up, the following options can be specified for each group.

Group Name

The name given to the group.

Remark

A comment.

Users

In this part of the panel it is possible to assign users to the group. in the search widget it is possible to filter existing local users to find matching users. Users are added to the group by clicking on the **+** on the right of the username. Users in the Group are shown in the textfield below. There are also shortcuts to **Add all** and to **Remove all** users to/from a group.

Override OpenVPN options

Tick this checkbox to allow the OpenVPN protocol to be used. This option will reveal a box in which to specify custom option for the account, which are the same as those specified for the *local users*.

Override L2TP options

Tick this checkbox to show a box in which to choose the L2TP tunnel to be used from a drop-down menu.

Note: This option can not be selected if no L2TP tunnel has yet been configured. In such a case, an informative message appears as a hyperlink: Upon clicking on it, the IPsec connection editor opens. Once created a new L2TP tunnel, it will be possible to associate it to a user.

Hint: The box for L2TP options will appear below the *OpenVPN options* box, if also OpenVPN option are to be overridden.

Enabled

Tick the checkbox to enable the user, i.e., to allow her to connect to the OpenVPN server on the appliance.

Warning: While the same user can be legally part of one or more groups, care must be taken that the groups the user belongs to do not define contrasting *override* options. As an example, consider a user member of two groups, one allowing access only to the GREEN zone, and one only to the BLUE. In this case, it is not easy to predict whether that user will be granted or not access to the BLUE or GREEN zone. The management of these issues is left to the manager of the OpenVPN server.

7.5.3 Settings

This page contains the current configuration of the authentication servers on which the appliance relies and allows for their management. Currently, only local and LDAP / Active Directory are supported, though in future releases additional types of authentication server might be added, like e.g. Radius servers.

There are two tables in this page, one displaying information about *Authentication servers*, and one showing *Authentication server mappings*. In the former, those information is shown:

- Name. The name given to the server
- Type. Whether the server is a local or an external LDAP one.
- Service. Which authentication is available for that server.
- Actions. For local authentication, it is possible to *enable/disable* the server, to *edit it*, or to delete it. For LDAP servers there is also the ability to *refresh* the connection, for synchronising the users and groups.

The table at the bottom shows the correspondences between a service (IPsec XAuth, OpenVPN, and L2TP) and the type of authentication allowed. The only *Actions* for the mappings is to *Edit* them. By clicking on **Edit**, a form will appear, in which a selector allows to select which authentication backends will be used for that service.

A click on the **Add new authentication server** link above the tables opens a form in which to supply all data to set up a new authentication server.

This form replaces the tables displaying the already defined authentication servers and allows to configure a new one, by specifying appropriate values for the following configuration options.

Name

The name given to the authentication server.

Enabled

Tick the checkbox to enable the server.

Type

Select from the drop-down menu whether the server shall be *LDAP / Active directory* or *local*. All the next options, except for the last one, are available only for the configuration of LDAP servers.

LDAP server URI

The URI of the LDAP server.

LDAP server type

This drop-down menu allows the choice of the type of the authentication server among *Generic*, *Active Directory*, *Novell eDirectory*, or *OpenLDAP*.

LDAP bind DN username

The fully distinguished name of the bind DN user, which must have the permission to read user attributes.

LDAP bind DN password

The password of the bind DN user.

The following options depend on the server's setup and are used to identify which users and groups shall be granted access to appliance's OpenVPN server: *LDAP user base DN*, *LDAP user search filter*, *LDAP user unique ID attribute*, *LDAP group base DN*, *LDAP group unique ID attribute*, *LDAP group member attribute* and *LDAP group search filter*.

Limit to specified groups






This option allows to select which groups on the LDAP server are allowed to connect to the appliance's OpenVPN server.

7.6 Certificates

The *Certificates* page allows the management of the certificates that are needed by the various OpenVPN server instances running on the appliance and is composed of three tabs: *Certificates*, *Certificate Authority*, and *Revoked Certificates*.

7.6.1 Certificates

Here it is possible to manage all the certificates stored on the appliance. The table, initially empty, shows all certificates along with the following details, one per each column:

- *Serial*. A unique number identifying the certificate.
- *Name*. The name assigned to the certificate.
- *Subject*. The collection of information that identify the certificate. itself. See the options below.
- *Expiration Date*. The final date of validity of the certificate.
- *Actions*. What can be done with the certificate:
 -  - to show all its details.
 -  - to download it in PEM format.
 -  - to download it in PKCS12 format.
 -  - to delete the private key associated to it.
 -  - revoke the certificate.

Above the list, a link can be clicked to **Add new certificate**. Upon clicking, the page will be replaced by a form that allows to provide all data necessary to the generation of a new certificate.

At the bottom of the table, on the left-hand side there is a navigation widget, that allows to navigate among the various pages composing the table, if there are many certificates, whereas on the right-hand side there is a reload widget, used to refresh the list of certificates.

Add new certificate

Three alternatives are available to store a new certificate on the appliance, selectable from this drop-down menu: *Generate a new certificate*, *Upload a certificate*, and *Upload a Certificate signing request*.

Generate a new certificate

The first alternative allows to create a new certificate directly on the appliance, by providing the following information. The capital letters in parentheses show the field of the certificate that will be filled by the value supplied and form the *Subject* of the certificate.

Note: A Root Certificate Authority is needed to create certificates, so create the Root CA before creating certificates.

Common name

The common name (CN) of the certificate's owner, i.e., the name with which the owner will be identified.

Email address

The e-mail address of the certificate's owner.

Organizational unit name

The Organisation Unit (OU) to which the owner belongs to, i.e., the company, enterprise, or institution department identified with the certificate.

Organization name

The organisation (O) to which the owner belongs to.

City

The city (L) in which the organisation is located.

State or province

The state or province (ST) in which the organisation is located.

Country

The Country (C) in which the organisation is located, chosen from those in the selection menu. By typing one or more letters, matching countries are searched for and displayed.

Subject alt name (subjectAltName=email:*,URI:*,DNS:*,RID:*)

An alternate name for the subject, i.e., the certificate.

Certificate type

The type of the certificate, chosen between *Client* and *Server* from the drop-down menu.

Validity (days)

The number of days before the certificate expires.

PKCS12 file password

The password for the certificate, if needed.

PKCS12 file password (Confirmation)

Type once more the certificate's password for confirmation.

Upload a certificate

The next alternative is to upload an existing certificate from the local workstation to the appliance.

Certificate (PKCS12/PEM)

By clicking on the **Browse** button or on the textfield, a file chooser will open, in which to supply the path to the certificate to be uploaded.

PKCS12 file password

The password for the certificate, if needed.

Upload a certificate signing request

The third alternative is to upload a CSR from the local workstation to the appliance, i.e., an encrypted text file containing all necessary information to generate a new certificate, recognised by the server.

Certificate Signing Request (CSR)

By clicking on the **Browse** button or on the textfield, a file chooser will open, in which to supply the path to the CSR to be uploaded.




Validity (days)

How many days shall the certificate be valid.

7.7 Certificate Authority

This page allows to manage the CA, which are necessary for the correct working of an OpenVPN encrypted connection. There are two ways to add a CA: Either by clicking on the link above the table of already existent certificates to generate a new certificate, or by uploading one using the widgets below the table.

The table, once populated, shows the same information as in the *Certificates* tab, with the only difference in the *Actions* available, which are:

-  - to show all CA details.
-  - to download it in PEM format.
-  - to delete the certificate.

To upload a certificate, supply the following information:

CA name

The name of the **Authority** who created the certificate.

Certificate (PEM)

By clicking on the **Browse** button or on the textfield, a file chooser will open, in which to supply the path to the certificate to be uploaded.

Clicking on the **Upload CA certificate** will start the upload process.

Generate new root/host certificates

This procedure can be applied only once and will generate two certificates: A root certificate authority and a host certificate, with the latter that shall appear in the lint shown in the *Certificates* tab. When clicking on the link, a form will replace the list, in which to supply the following data, that will be used in the new root and host certificates.

Note: The only way to generate a new root certificate is to delete the existing one.

System hostname

The name of the system, that will be used as the certificate's Common Name.

Email address

The e-mail address of the system's owner or responsible.

Organizational unit name

The Organisation Unit (OU) to which the system belongs to.

Organization name

The organisation (O) to which the system belongs to.

City

The city (L) in which the organisation is located.

State or province

The state or province (ST) in which the organisation is located.

Country

The Country (C) in which the organisation is located, chosen from those in the selection menu. By typing one or more letters, matching countries are searched for and displayed.

Subject alt name (subjectAltName=email:*,URI:*,DNS:*,RID:*)

An alternate name for the subject, i.e., the certificate.

Validity (days)

The number of days before the certificate expires.

7.8 Revoked Certificates

The certificates that have been revoked are listed in the table, that show the serial number and the subject of the certificate.



Download the Certificate Revocation List

A click on this link will allow to download the on a local workstation the Certificate Revocation List.

7.9 Certificate Revocation List

In this page can be managed all the Certificate Revocation lists that have been uploaded.

The table shows all the Certificate Revocation Lists that have been uploaded and for each item in the table are show the name of the certificate, the issuer, and the issued date. Available actions are:

-  - display the certificate details
-  - download the certificate on the local workstation.

Chapter 8: The Logs and Reports Menu

In the logs and reports section of the appliance there are different possibilities to look at and to analyse the log files.

The sub-menu on the left-hand side of the screen contains the following items:

- Dashboard - the brand new reporting module.
- Traffic monitoring - the ntopng graphic interface gives a real time overview of the network traffic using charts.
- Live Logs - get quick, live view of the latest log entries as they are being generated.
- Summary - get daily summaries of all logs.
- System - system logs (`/var/log/messages`) filtered by source and date.
- Service - logs from the intrusion detection system (IDS), OpenVPN, and antivirus.
- Firewall - logs from iptables rules.
- Proxy - logs from the HTTP, SMTP, and content filter proxies.
- Settings - customise all the log options.
- Trusted Timestamping - securely time stamp the log files to verify they have not been altered.

In a nutshell, there are two modalities to access the log from the GUI: Live and “by-service”: In the live mode the log files are visualised as soon as they are created, while in the “by-service” mode only the logs produced by one daemon or service are displayed.

8.1 Dashboard

The reporting GUI is a new module, introduced in version 3.0, whose purpose is to graphically show the occurrence of various types of event on the system.

In a nutshell, the reporting module shows events happened on the appliance using different widgets and graphs. All events occurring on the system and the information concerning them recorded by the syslog daemon are parsed and used to populate a sqlite3 database. From here, data are gathered according to the options and filters applied in the GUI and are displayed by the widgets.

Note: This module is loosely coupled with the *Event notifications* located in *Menubar ► System ► Event notifications*. All events recorded there, and for which email or SMS alerts are sent, appear also here, but the vice-versa is not true.

This page is divided into six tabs: *Summary*, *System*, *Web*, *Spam*, *Attacks*, and *Virus*. Except for the first tab, which shows an overview of all events, each of them is dedicated to a precise service running on the appliance.

8.1.1 Common elements

All the tabs share the same design: Below the tabs, on the left-hand side there are a date selector on the left-hand side and a **Print** button on the right-hand side. Then, a line chart at with an horizontal slider right below, atop one informative boxes (*Summary Grid*) and a pie-chart. At the bottom, there are one or more tables, depending on the tab and the data shown. The table that is always present is the one displaying the syslog messages related to the events shown.

More in detail, here is a description of all the widget present in the reporting module.

Date selector

At the top left-hand side of the GUI there is an hyperlink that shows the interval within which occurred those events that have been considered for the charts. By clicking on it, a small panel gives access to other choices of intervals. There are two types of choices, the first one concerns events that took place in the *last ... days*, namely events from the last day, week month, quarter, or year; the second one selects all the events occurred in one of the last 12 months. Upon selecting a new time span, the other widgets are also updated. There is also the possibility to not change the interval shown, by clicking on **Cancel**.

Print

A click on this button shows a print preview of the current page, in which the **Back** button replaces **Print**, and open a pop up window in which to choose the printing device.

Line Chart and Time Slider

The line chart shows the event happened on the appliance during the selected time span in a two dimensional graph, in which the x-axis shows the time interval and the y-axis shows the number of occurrences. A coloured line connects events of the same type.

Hint: Different types of event are denoted with different colours.

The time slider is located underneath the chart and allows, within the selected time span, a more fine-grained view of the events, depicted here as histograms. Indeed, the two grey handles on the left and right limits of the slider can be clicked and dragged to reduce the time span shown in the line chart. When reduced, the slider can also be moved by clicking in its middle and dragging it to the left or the right.

Summary Grid

The summary grid has a twofold purpose: On the one hand to show the number of occurrences of the various types of events that took place on the appliance in the selected period, whereas on the other end to filter which type of events are shown in the line chart. Its content changes according to the tabs it is located, i.e., to the types of events logged. The summary grid is not present in the *Mail*, *Attacks*, and *Virus* tabs, in which is replaced by a number of tables with details about the events.

Pie Chart

The pie chart diagram shows graphically the number of event that took place in the selected time span. When in the *Summary* tab, each slice can be clicked, to open the tab corresponding to the type of event and show a more detailed representation.

Syslog Table

A table that shows the syslog messages extracted from the log files and related to the events shown in the charts. When the table carries lot of messages, these are divided into many pages and can be browsed using the buttons and number at its left bottom. At the right bottom there is an icon that allows to refresh the table's content.

8.1.2 Summary

The *Summary* tab gives an overview of all categories of events recorded on the appliance. The summary grid allows to filter the following types of events:

- System (Green). Number of Log ins and other events connected with system administration tasks (e.g., uplinks change of status, start and stop of logging, and so on).
- Mail (Dark Grey). Number of spam e-mail received.
- Web (Blue). Number of pages blocked by the content filter.
- Virus (Red). Number of viruses found.
- Intrusion attempts (Yellow). Events recorded by the IPS.

Each category can be shown separately, with more information and a higher level of details in the other tabs of the page, see further on.

8.1.3 System

The *System* tab displays all events that are related to the system efficiency and to system administration. These are all the events shown:

- Uplink (Red). The times the uplink(s) went online or offline.
- Status (Dark Grey). Changes in the state of the appliance.
- Login (Blue). The number of logins, both successful and not.
- Disk (Yellow). Events involving disk I/O.
- Upgrade (Green). Events involving upgrade of system or of packages.
- Support (Light Grey). Number of accesses and operations done by the Support Team.

A click on the small icon on the left-hand side of each event category causes the other categories to not be shown, while the current is further detailed and the pie chart is updated.

8.1.4 Web

The *Web* tab displays the number of pages that have been blocked by the URL filter engine. The summary grid is composed by two tabs: *Access report* and *Filter report*. The former shows the blocked URL divided by *Source IP Address*, *URL*, and *Users* that have been blocked, and the total count for every item, each in a table.

The latter tab shows in the first table, the following categories, that are those found in the *Web filter* (See *Menubar* ► *Proxy* ► *HTTP* ► *Web Filter*).

- General Use (Green).
- Parental Control (Yellow).
- Productivity (Blue).
- Security (Red).
- Uncategorized Sites (Dark Grey).

Like in the case of the *System* tab, a click on the small icon on the left-hand side of each event category causes the other categories to not be shown, while the current is further detailed and the pie chart is updated.

The other tables at the bottom show the counts of each the blocked objects: The *Source IP Addresses*, the *URLs*, and the *Users*.

8.1.5 Mail

The *Mail* tab displays all e-mails blocked as spam.

There is no summary grid in this tab, replaced by three tables, displaying counts for:

- From. The sender(s) of spam e-mails.
- To. The recipient(s) of spam e-mails.
- Source IP Address. The IP address from where spam e-mail have been sent.

8.1.6 Intrusion Attempts

The *Intrusion* attempts tab displays all tentative intrusions detected by the IPS (See *Menubar* ► *Services* ► *Intrusion Prevention*).

The tables at the bottom show counts of the following information:

- Intrusion attempts. The categories under which falls each attempt.
- Source IP Address. The IP address from where the attack originated.
- Destination IP Address. The IP Address to which the attack was launched.

8.1.7 Connections

The *Connections* tab displays the average number of connections started by the users of the appliance, grouped into:

- Local connections. Accesses via SSH or console.
- Hotspot users. Users accessing the Hotspot.
- IPsec users. Clients connected via IPsec.
- OpenVPN users. clients connected using VPN.

8.2 Traffic Monitoring

The ntopng software is the successor of the ntop network traffic analyser, which adds a more intuitive interface and more graphical representations of the traffic that flows through the appliance.

The management interface of ntopng provides now more usability and can be accessed easily accessed from any browser, and therefore has been integrated more tightly with the appliance interface than in previous versions.

In few words, the abilities of ntopng can be summarised as follows:

- Real time monitoring of every network interface of the appliance.
- Web-accessible management interface.
- Less resource needed compared to ntop.
- Integration of nDPI (Application firewall).
- Traffic analysis according to different parameters (protocol, source/destination).
- Export of reports in JSON format
- Storage of traffic statistics on disk.

The ntopng GUI is organised into four tabs: *Dashboard*, *Flows*, *Hosts*, and *Interfaces*. Moreover, there is also a search box to quickly display information about a given host.

In the footer of each tab, a couple of information are shown: Besides a copyright notice and a link to the ntop home page, there is a chart showing the network traffic over the last 20 seconds, updated in real time, and some numerical data about the current bandwidth used, the number of hosts and flows and the appliance's uptime.

8.2.1 Dashboard

The dashboard shows all connections that interest the appliance, that is, all established *Flows* in which the appliance is involved.

The page is divided into several diagrams, with the first one -a so-called Sankey diagram showing all flows moving on the appliance, updated in real time. The horizontal flows show the traffic between two hosts, while the vertical width of each flows is proportional to the bandwidth used by that flows, i.e., to the amount of data flowing. The connections -and therefore the direction of the data sent- are shown left to right: Hosts on the left hand-side of the diagram send data to hosts on the right-hand side and are identified by either their IP address or hostname. A click on one host leads to the *Overview* page in the *Hosts* tab, which shows several information about that host.

Below the Sankey diagram, four informative-only pie charts show in percentage the items that that generate the most traffic, divided into: Total by host (top left); application protocols (top right), ASNs (bottom left), and live flow senders (bottom right).

8.2.2 Flows

The active flows tab contains a big table with a number of information about the active flows:

- *Info*. A click on the icon opens a new page in which more detailed information about that flow is shown.
- *Application*. The application causing the flow. nDPI is used to recognise the application, therefore it might be necessary to wait for a couple of packets to see the correct application displayed: In this case, the *(Too Early)* message appears instead of the application name.
- *L4 Proto*. The network protocol used by the flow, which is usually TCP or UDP.
- *Client*. The hostname and port used by the flow on the client side. Clicking on either the hostname or port, more information will be shown in a new page about the network traffic flowing that host or port.
- *Server*. The hostname and port used by the flow on the server side. Like for the *Client* above, more information is shown when clicking on the hostname or port.

Hint: By clicking on the hostname or port, the table shows detailed information about it, opening a sub-tab in the *Hosts* tab.

- *Duration*. The length of the connection.
- *Breakdown*. The percentage of traffic generated by the client and by the server.
- *Throughput*. The amount of data currently exchanged between the client (on the left, in black) and server (on the right, in green).
- *Total Bytes*. The total data exchanged since the connection was first established

At the bottom of the table, on the left-hand side it is shown the total number of rows shown, while on the right-hand side it is possible to browse the various pages in which the table is split, when the number of rows is higher than the pagination.

A click on the *Info* icon will give detailed information about that particular flow. Besides those already described above, these additional data are displayed.

- *First Seen*. The timestamp when the connection was established, along with the time passed since.
- *Last Seen*. The timestamp in which the connection was last active and the time passed since that moment.
- *Client to Server Traffic*. The number of packets and bytes sent from the client to the server.
- *Server to Client Traffic*. The number of packets and bytes sent from the server to the client.
- *TCP Flags*. The TCP states of the current flow.

It is possible to go back to the list of flows by clicking on the **Flows** hyperlink on the left, right above the table.

8.2.3 Hosts

The *Hosts* tab allows to view several details about the involved parties of a flow: Host, port, application, flows and their duration, data exchanged, and so on.

Two representations are available: *Host List* and *Top Hosts (Local)*

The *Hosts List* representation shows information about all the hosts involved in some flow with the appliance and the following data about them:

- *IP Address*. The IP address or MAC Address of the host. The latter is shown if the DHCP lease for that host has expired.
- *Location*. Whether the host is in the local or in a remote network.
- *Symbolic Name*. If available, it is the hostname of the host.
- *Seen Since*. The timestamp of the first established connection.
- *ASN*.
- *Breakdown*. The trade-off between sent and received traffic.
- *Traffic*. The amount of data exchanged by the host.

A click on the IP address opens an overview of the host, showing several information about it, besides those listed above:

- *Last Seen*. The timestamp in which the connection was last active and the time passed since that moment.
- *Sent vs Received Traffic Breakdown*. The traffic generated or received by the host.
- *Traffic Sent*. The number of packets and bytes sent from the client to the server.
- *Traffic Received*. The number of packets and bytes sent from the server to the client.
- *JSON*. Download information about the host in JSON format.
- *Activity map*. How many flows have seen the host involved at a given timestamp. Each square shows a minute and the darker the colour, the more flows have taken place in that minute.

From here it is also possible to open additional informative tabs about that host. Each tab contains one or more pie charts (except for the *Contacts* and *Historical* tabs) above a textual summary of the data displayed.

- *Traffic*. The network protocol used by the host. (TCP, UDP and ICMP being the most common).
- *Packets*. The length in packets of each flow. (note: just my guess)
- *Protocols*. The application protocol used by the host.
- *Flows*. The table with all the network flows from the hosts.
- *Talkers*. The Sankey diagram of the connections, very similar to the one shown in the Dashboard, which however shows only the most active flows.
- *Contacts*. This tab is slightly different from the others. It shows on top an interaction map and on the bottom a list of connections that have the host as client or receiver.
- *Historical*. An interactive graph that shows the history of the traffic flow from and to the host in a given timespan (up to one year), that can be selected above the graph.

The *Top Hosts (Local)* representation shows a real-time graphic of the hosts that have active connections to the host. It displays the last 30 minutes.

8.2.4 Interfaces

The *Interfaces* tab allow to select the network interface, among the active ones, whose traffic should be displayed.

Note: It is currently not possible to select flows and/or hosts from different interfaces.

8.3 Live

When entering in the *Logs* section, or clicking on the **Live** entry on the sub-menu, the *Live log viewer* is shown, a box showing the list of all the log files available for real time viewing. Any number of logs to see can be chosen by ticking the corresponding checkboxes, that are displayed in a new window upon clicking on the **Show selected logs** button. To watch all the log files at once, simply tick the *Select all* checkbox right above the **Show selected logs** button and then click on the latter button. Otherwise, to view only one log file, simply click on the **Show this log only** link.

The window that opens contains two boxes, *Settings* at the top and *Live logs* at the bottom.

Warning: The list of log entries can become nearly unreadable if many logs are showed, due to the possible high number of log entries produced (especially by the firewall or proxy log, which can generate several log entries per second in case of heavy traffic). In this cases, the logs to be displayed can be configured in the *Settings* box.

Settings

This box allows to modify the settings of the log viewer, including which of the log files to show, their colour and options to highlight or find specific keywords.

On the right-hand side of the box appears the list of the logs that are currently displayed, and the colour with which they are highlighted, while on the left-hand side some additional control elements are shown, that help limit the output:

Filter

Only the log entries that contain the expression in this field are shown.

Additional filter

Like the filter above, but applied to the output of the first filter. In other words, only log entries containing both expressions are shown in the log.

Pause output

Clicking on this button will prevent new log entries from appearing on the live log. However, after clicking the button once more, all new entries will appear at once, quickly scrolling the old ones.

Highlight

All the log entries that contain this expression will be highlighted in the chosen colour. The difference with the filtering option is that all the content is still displayed and the log entries containing the expression will be highlighted with a coloured background.

Highlight color

Clicking on the coloured square gives the choice to select the colour that will be used for highlighting.

Autoscroll

This option is only available if the *Sort in reverse chronological order* option in the *Menubar ► Logs ► Settings* section is turned off. This causes all the new entries to be shown at the bottom of the page: If this option is enabled, the list is scrolled upwards to show the latest entries at the bottom of the page, otherwise only the older entries are show and the scrollbar on the right should be used to see the new ones.

To add or remove some log from the display, click on the **Show more** link right below the list of the log files on the top right. The controls will be replaced by a table from which the desired log files can be selected by ticking or unticking their respective checkboxes. To change the colour of a log file, click on the *colour palette* of that log type and then choose a new colour. To show the controls again, click on one of the *Close* links below the table or below the list of the displayed log files.

Live Logs

The logs chosen for viewing are shown in this box, which consists of a table divided in three columns.

Left Column

This column contains the log name, that is, the daemon or service producing the log entry.

Middle Column

The time stamp (date and time) of the event that has been recorded.

Right Column

The actual message generated by the service or daemon and recorded in the log files.

Note: Some log messages -especially Firewall entries- span more than one line, denoted by the expand button at the right of the message. To show the whole message, click on it or on the button.

Finally, there is also the chance to increase or decrease the window size by clicking on the **Increase height** or **Decrease height** buttons, respectively, which are situated on the heading of the box.

8.4 Common actions

The sub-menu entries *System*, *Service*, *Firewall*, and *Proxy* show log files for different services and daemons, grouped by similar characteristics. Several controls are available to search within the log, or view only some entries of the log, many of which are the same in all the services and daemons, with only the *System* menu item and the *HTTP report* tab under *Proxy* that have some additional control. These sub-menu entries have also a common structure of their pages, organised in two boxes: *Settings* at the top and *Log* at the bottom.

Filter

Only the lines that contain the entered expression are shown.

Jump to Date

Directly show log entries from this date.

Jump to Page

Directly show log entries from this page in the result set. The number of entries shown per page can be modified on the *Menubar* ► *Logs* ► *Settings* page.

Update

After changing any of the settings above, a click on this button refreshes the page content. The page is not refreshed automatically.

Export

When clicking on this button the log entries are exported to a text file.

Sign log

When clicking on this link, the current log is signed. This button is only available if Trusted Timestamping is enabled.

Older, Newer

These two buttons are present in the *Log* box and show up whenever the number of entries grows too much and are divided into two or more parts. They allow to browse older or newer entries of the search results by clicking on them.

Note: A message at the top of the page informs if on a given date there are no logs available: This can happen either if the daemon or service were not running, or if they did not produce any message.

In the remainder of this section, all the services and their peculiar settings are presented.

8.5 Summary

This page presents summaries for the logs produced by the appliance, separated by days and generated by the **logwatch** log monitoring software. Unlike the other parts of the log section, it has its own settings to control the level of details shown. The following control elements are available in the first box at the top of the page.

Month

Select from this drop-down menu the month in which the log messages were generated.

Day

The second drop-down menu allows to pick the day in which the log messages were generated.

<< >>

Browse the history, moving from one day (or part of it when too many messages have been generated) to another. The content of the page will be automatically refreshed.

Update

Immediately refresh the content of the page when the month/day combination has been changed.

Export

When clicking on this button, a text version of the summary is shown and can be saved on a local filesystem.

Below the *Settings* box, a variable number of boxes appears, depending on the running services that have log entries. The *Disk Space* box should at least be visible, showing the available disk space on the chosen date, while other boxes that can show up include *Postfix* (mail queue) and *Firewall* (accepted and dropped packets)

Note that the summaries are not available for the current day, as they are generated every night from the log files generated the day before.

8.6 System

In this section appears the log viewer for the various system log files. The upper box, *Settings*, defines the criteria to display the entries in the lower box. Besides the common actions, one additional control is available:

Section

The type of logs that should be displayed, either *All* or only those related to a given service or daemon. Among others, they include kernel messages, SSH access, NTP, and so on.

Following the choice of the section, click on the **Update** button to refresh the logs displayed in the *Log box* at the bottom of the page, in which the **Older** and **Newer** buttons allow to browse the pages.

8.7 Service

In this section appear the log entries for three of the most important services provided by the appliance: IDS, OpenVPN, and the anti-virus, each in its own tab. Only the common actions are available.

8.8 Firewall

The firewall log viewer contains the messages that record the firewall's activities. Only the common actions are available.

Information shown in the table are:

Time

The timestamp at which the message was generated.

Chain

The chain through which the packet has passed.

Iface

The interface through which the packet has passed.

Proto

The prototype of the packet.

Source, Src port

The IP address and port from which the packet has arrived.

MAC address

The MAC address of the source interface.

Destination, Dst port

The IP address and port to which the packet had to arrive.

8.9 Proxy

The proxy log viewer shows the logs for the four daemons that use the proxy. Each of them has its own tab: squid (*HTTP*), icap (*Content filter*), sarg (*HTTP report*), and smtpd (*SMTP*, email proxy).

8.9.1 HTTP and Content filter

In addition to the common actions, the log viewer for the HTTP proxy and content filter allow these values to be specified:

Source IP

Show only the log entries containing the selected source IP Address, chosen from a drop-down menu.

Ignore filter

A regular expression that filters out all the log entries that contain it.

Enable ignore filter

Tick this checkbox to temporarily disable the ignore filter.

Restore defaults

Clicking on this button will restore the default search parameters.

8.9.2 HTTP Report

The *HTTP report* tab has only one option: To enable or not the proxy analysis report generator, by ticking the *Enable* checkbox and clicking on the **Save** button afterwards. Once the report generator is activated, a click on the **Daily report**, **Weekly report**, and **Monthly report** links shows detailed HTTP reports.

8.9.3 SMTP

Only the common actions are available in the tab of the postfix daemon.

8.10 Settings

This page contains all the global configuration items for the appliance's logging facilities, organised into four boxes: *Log viewing options*, *Log summaries*, *Remote logging*, and *Firewall logging*.

Log Viewing Options

Number of lines to display

The *pagination* value, i.e., how many lines are displayed per log-page.

Sort in reverse chronological order

If this checkbox is ticked, then the newest log entries will be displayed first.

Log Summaries

Keep summaries for __ days

How long should the log summaries be stored on disk before deletion.

Detail level

The detail level for the log summary: the higher the level, the more log entries are saved and showed. The drop-down menu allows three levels of detail: Low, Medium, and High.

Remote Logging

Enabled (Remote Logging)

Ticking this box allows to enable remote logging. The next option allows to enter the hostname of the syslog server.

Syslog server

The hostname of the remote server, to which the logs will be sent. The server must support the latest IETF syslog protocol standards.

Firewall Logging

Log packets with BAD constellation of TCP flags

If this option is enabled the firewall will log packets with a bad constellation TCP flag (e.g., all flags are set).

Log NEW connections without SYN flag

With this option enabled, all new TCP connections without SYN flag will be logged.

Log accepted outgoing connections

To log all the accepted outgoing connections this checkbox must be ticked.

Log refused packets

All the refused packets will be logged by the firewall, if this option is enabled.

8.11 Trusted Timestamping

Trusted timestamping is a process that log files (but in general any document) undergo in order to track and certify their origin and compliance to the original. In other words, trusted timestamping allows to certify and verify that a log file has not been modified in any way by anyone, not even the original author. In the case of log files, trusted timestamping proves useful for example, to verify the accesses to the system or the connections from the VPN users, even in cases of independent audits.

Trusted timestamping is not enabled by default, but its activation only requires a click on the grey switch. When it turns green, some configuration options will show up.

Timestamp server URL

The URL of the timestamp server (also called TSA) is mandatory, since it will be this server that signs the log files.

Note: A valid URL of a valid TSA is needed to be able to use trusted timestamping. Several Companies can supply this kind of service.

HTTP authentication

If the timestamp server requires to authenticate, tick the box below the *HTTP authentication label*.

Username

The username used to authenticate on the timestamp server.

Password

The password used to authenticate on the timestamp server.

Public key of the timestamping server

To ease and to make the communication with the server more secure, the server's public key can be imported. the certificate file can be searched on the local computer by clicking on the **Browse...** button, and then uploaded to the appliance by clicking on the **Upload** button. After the certificate has been stored, next to the *Public key of the timestamping server* label, a **Download** link will appear, that can be clicked to retrieve the certificate, for example if it should be installed on another appliance.

After clicking on the **Save** button, the settings are stored and, on the next day, a new button will appear in the *Logs* section, on the right-hand side of the *Settings* box:

Verify log signature

When clicked it will show a message in a yellow callout to inform about the status of the log.