Connection of Extruded Subnets: A Solution Based on RSIP

Cédric de Launois, Aurélien Bonnet, and Marc Lobelle

Université catholique de Louvain, Département INGI
Place Ste-Barbe, 2, 1348 Louvain-la-Neuve, Belgium
Fax: +32 10 45 03 45
delaunoi@info.ucl.ac.be, tel: +32 10 47 24 04
ab@info.ucl.ac.be, tel: +32 10 47 87 18
ml@info.ucl.ac.be, tel: +32 10 47 32 74

Abstract. Many remote computers need to be securely connected to their organization main network through a public IP network (e.g. Internet). Our purpose is to integrate as seamlessly as possible remote networks in the organization network, i.e. to put these in exactly the same situation as if they were located inside the organization. After summarizing the state of the art, the paper presents a solution based on RSIP, to dynamically allocate an IP address of the organization to a host of the remote network requesting an external access. Security is provided by IPSec. We compare this solution with a former proposal based on DHCP and show that the two solutions are very close but that RSIP brings us closer to an ideal situation but at an extra cost.

Introduction

Many organizations are faced to the problem of securely connecting remote computers to their network to accommodate nomadic users, teleworkers (including students in the case of educational institutions), remote branches and facilities etc. These remote systems are often connected to the main network of the organization through a public IP network, which can be the Internet or a provider network used to implement private virtual networks. In most instances, the remote computer obtains a single dynamic IP address in the provider range and security is added by encrypting the traffic in the application (SSL [1]), in an application tunnel (SSH [2]) or at IP level (IPSEC [3]).

In some situations these already classical solutions are inadequate.

One approach is to solve each individual problem when it appears. Another is to try to specify the ideal situation and to try to implement it. We choose this second approach. Our purpose is to integrate as seamlessly as possible remote machines in the organization network, i.e. to put these in exactly the same situation as if they were located inside the organization.

Since remote infrastructures often involve several computers (remote offices, student flats with several rooms, etc), we focus on remote subnetworks, rather

E. Gregori et al. (Eds.): NETWORKING 2002, LNCS 2345, pp. 685–696, 2002.

than remote single computers. The latter case can always be considered as a particular case of the former.

We call subnets in this ideal situation "extruded subnets". They have the following properties (this is what we consider the ideal situation).

- The extruded subnet is connected to a gateway [1] that only needs a single dynamically allocated address in the provider range.
- Computers in the extruded subnet have addresses in the main network of the organization and are undistinguishable from computers located directly on the main network. In particular,
  - they can be used as clients as well as servers by any application;
  - no computer outside the company network and the extruded subnetwork can read, modify or detect traffic between a particular computer in the extruded subnetwork and the main network.
- Computers in the remote subnetwork have statically allocated permanent DNS names.
- Computers in the remote subnet use sparingly IP addresses, i.e. IP addresses are not allocated to computers that do not need it, e.g. that are either inexistent or stopped, etc.

We will first review the current techniques for connecting to a main network a subnetwork to which a temporary single IP address has been allocated. This review and the DHCP solution are based on [5] (where "extruded subnetworks" are called "remote bubbles"). Then the RSIP solution will be presented in detail. Finally the DHCP and RSIP based solutions will be compared and discussed.

1 Connection of a Subnetwork through a Single IP Address: State of the Art

This section presents the existing solutions to connect subnetworks through a single dynamically allocated IP address in ascending order of satisfaction of the requirements for extruded subnetworks.

1.1 Address Translation

NAT (Network Address Translation) is a solution where the gateway replaces the IP address in packets outgoing from the subnetwork by its own one, and the port number by one of its unused ones. The reverse substitution is performed on incoming packets. NAT allows client applications on computers of the subnetwork to invisibly access the Internet through the gateway. To other machines on the Internet, all this traffic will appear to be from or to the gateway (for more information see [6]). Not all "client" applications work with this scheme (e.g. FTP). It is therefore often complemented with specific application level proxies.

[1] We will use the generic name of "gateway" for the computer linking the subnetwork to the rest of the world

NAT is inadequate when the machines in the extruded subnet must be accessible from outside (e.g. for direct videoconference or for peer to peer applications).

1.2 Virtual Private Networks

VPNs have been introduced to let two networks communicate securely when the only connection between them is over a third network which they don't trust. VPNs use a gateway between each of the communicating networks and the untrusted network. Most current VPN packages use tunneling.

The gateway can encrypt packets entering the untrusted net and decrypt packets leaving it, in order to secure the tunnel.

Simple Tunneling Protocol. Most current operating systems can enable simple tunnels between two gateways (without any authentication or encryption : the tunnel is thus not secure). Gateways on each network encapsulate packets destined to the distant network in a packet destined to the remote gateway. Gateways identify each other using their static IP addresses. In our context, this gateway address is dynamically allocated by the provider. So it must be authenticated by other means than its IP address.

The IPSec Protocol and its Use in Virtual Private Networks. IPSec is a mechanism for adding security to IP. It can protect traffic between hosts, between network security gateways (routers, firewalls,. . .) and between hosts and security gateways. IPSec hosts and gateways are authenticated by cryptographic techniques independently of their IP addresses, which may be dynamically allocated. More informations can be found in [4], [3].

The VPN can be built by deploying IPSec gateways using IPSec in tunnel mode.

Current VPN solutions are inadequate in organizations that cannot afford to assign permanent IP addresses to machines in the remote subnets

1.3 Extruded Subnets

This first implementation is based on three different protocols : IPSec, DHCP (Dynamic Host Configuration Protocol), and NAT (Network Address Translation) or proxy ARP. More information is available in [5].

First step : Building Static IPSec VPNs. The first step is to set up an IPSec VPN between the gateway of the extruded subnet and a gateway in the main network like in the preceding solution.

This provides already the following features :

- The computer in the extruded subnet is logically neighbour of the main network.

- The external address of the gateway may be dynamically allocated.
- IPSec can provide security (authentication and confidentiality).

All the packets destined to the extruded subnets will be routed through the tunnels.

Second step : Adding Dynamic IP Address Allocation to the Computers in the Extruded Subnets. The Dynamic Host Configuration Protocol *(DHCP)* automates the process of configuring devices on IP networks. DHCP performs many of the functions a network administrator could carry out manually when connecting a new computer to a network (see [7]). With DHCP relay agents, remote machines can also be configured. A relay agent is used to forward DHCP messages between clients and server when the server and the client are not in the same network. The central DHCP server knows what set of IP addresses it must allocate to requests for each relay agent. The DHCP protocol with relay agents can be used to dynamically configure the computers of extruded subnetworks with the following advantages :

- the network configuration of the computer is easier (most of the parameters are transmitted by the protocol),
- addresses can be leased temporarily when needed, which, for instance, simplifies network administration of nomadic computers (laptops),
- subnetworks can be created without any administrative overhead for address allocation in the subnet,
- the DHCP protocol is available on many operating systems.

In this DHCP based implementation of extruded subnets, modified relay agents run on the gateways of the extruded subnets. The difference with standard relay agents is that addresses are assigned to the devices on the different subnets without regard to their localization. This solution is more economical in IP addresses, but routes must be explicitly configured for each individual device. When a device asks for a new DHCP configuration, the relay agent offers a dedicated IPSec tunnel opened between the gateway of the extruded subnet and one in the main network for this new IP address. The device has the illusion to be connected by a point-to-point link to the gateway in the main network.

Third Step. In the preceeding solution, all the packets sent by a device on an extruded subnet will be routed through the gateways even the packets destined to the subnet itself, as the different devices of our extruded subnet do not know they are on the same physical network.

Two different techniques may be used to give them that knownledge. In both, the machines on the extruded subnet are made to believe they are on a large network including all the remote subnets. In the first technique, instead of sending an address such as “a.b.c.d”, the DHCP server will send the private IP address “10.b.c.d” with the same three last bytes and a class A subnetwork mask instead of a point-to-point mask. This way, all devices in the extruded network can see each other. The gateway has to "NAT" (translate address) between 10.b.c.d and a.b.c.d for incoming and outgoing messages, with the aforementioned disadvantages of NAT.

In the other solution, the devices get a netmask covering the set of addresses allocatables to all the extruded subnets. Typically, when a machine on an extruded subnet wants to communicate with a machine on another extruded subnet, Proxy ARP on the gateway will answer so that all traffic to the remote machine will be sent to it. From there, it will be routed to the destination extruded subnet.

2 Using RSIP to Manage Address Allocation in Extruded Subnets

Another way to integrate computers on a remote subnet into the main network of an organization is using the new RSIP (*Realm Specific IP*) protocol [8], [9], [10]. RSIP has been designed as an alternative to NAT but with the additional requirement to preserve end-to-end packet integrity, a feature not provided by NAT. RSIP is based on the concept of granting a host (called RSIP host) from a network A a presence in another network, B, by allowing it to use resources (e.g. addresses and other routing parameters) from the network B. The gateway (called RSIP gateway) between networks A and B owns a pool of such resources, that it can allocate to RSIP hosts. For connecting a private network to a public one, a gateway on the boundary between these networks owns a pool of public IP addresses that it can allocate to hosts of its private network. See figure 1.

The problem of connecting an extruded subnet to a remote main network is similar since the previously described gateway may be split in two parts, connected via a tunnel. We will deal with the problem of the distance between these two networks in section 3. We may thus first focus on the simple problem of dynamically allocating IP addresses to hosts of a private network connected to a public one.

RSIP has been defined in two basic flavors : RSA-IP and RSAP-IP. When using RSAP-IP, the RSIP gateway maintains a pool of IP addresses as well as pools of port numbers per address. The gateway allocates each IP address with one or more port numbers. A host may only use the tuples address/port that have been assigned to it. When using RSA-IP, a RSIP gateway only maintains a pool of IP addresses to be leased by RSIP hosts. Upon request, the gateway allocates an address to the host, that may use it with any TCP or UDP port. This method is particularly interesting in our case and will be discussed below.

2.1 Using RSA-IP in Extruded Subnets

When a new computer is started in an extruded subnet based on RSIP :

- the computer boots with a private IP address;
- it registers with its RSIP gateway.

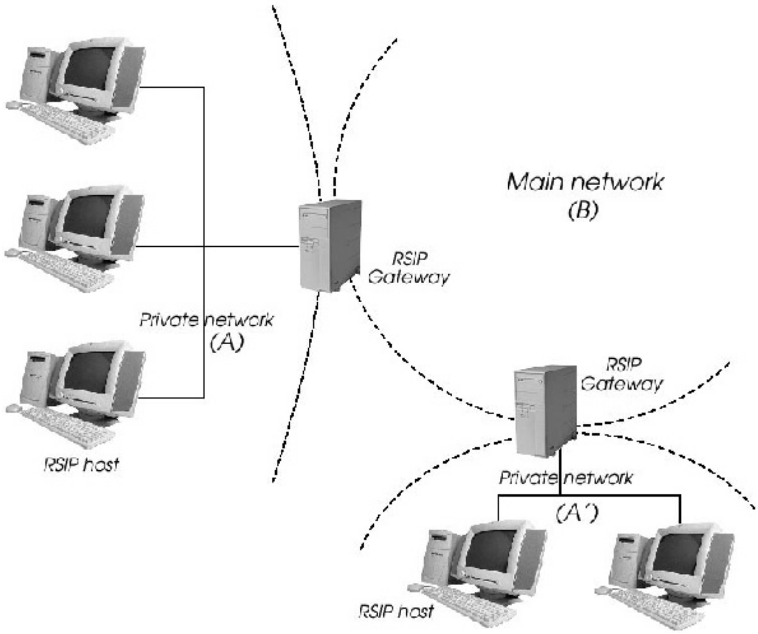

Fig. 1. Two private networks connected to the main network through RSIP gateways

- when it needs access to an external network, it requests an IP address from the gateway;
- the gateway delivers an address to the host, with an expiration time;
- the host uses this leased address for external accesses but still uses its private address to communicate with other hosts in the extruded subnet;
- when the lease time is about to expire, the host asks for a lease extension. If granted, the host may continue to use the address, otherwise it must release it.

2.2 The Routing Problem in RSIP-Based Extruded Subnets

Two main cases must be considered.

Communication between a Host X of the Extruded Subnet and a Host Y of an External Network (that may be another extruded subnet).

In this first case, the packets destined to hosts with public addresses must first be sent through the interface (often the interface of a tunnel to the gateway) corresponding to the public leased IP address. The gateway routes the packets it receives from X like regular packets. In the other direction, the packets originating from the public network can only be routed properly if the gateway is aware of the presence of a host with a public leased address inside the extruded subnet, and if a route or a tunnel to this host is available.

The RSIP gateway must establish this route each time it allocates an IP address to a host.

Communication between two hosts of two distinct extruded subnets is a particular case.

Communication between two Hosts X and Y in the Same Extruded Subnet. In this second case, when a host X wants to communicate with a host Y from the same extruded subnet, the routing will depend on whether host X contacts hosts Y using the leased public address of Y or not. If not, then host X will send its packets using its own private address (whether X possesses a leased address or not). No further routing is needed, Y will respond using its private address. If X uses the public address address of Y, the packets originating from X will reach the gateway first since it is the default gateway for IP packets destined to public hosts. Then, they will be routed to Y thanks to the special route established for packets coming from outside.

2.3 Using NAT/PAT in Coexistence with RSA-IP

The use of the RSA-IP protocol does not forbid to keep the NAT/PAT mechanism for hosts that only want to surf on the Internet or to use services for which a proxy exists.

The gateway must known which packets must be NAT'ed and which must not. The rule is to apply NAT only for packets with a private address in their header and destined to a public host (e.g. thanks to the iproute2 utility [11]).

3 Extension of the RSA-IP Protocol

The main drawback of the preceeding RSA-IP solution is, for our problem, its lack of scalability when there are several extruded subnets. RSA-IP, as described in [8], requires one pool of addresses per extruded subnet (the pool is kept on the gateway of each private network, see figure 1). Those addresses can thus only be allocated to hosts from that extruded subnet. This may lead to a waste of IP addresses if the range of addresses allocated to each remote subnetwork is statically allocated. It is much more efficient to maintain the pool in a unique, centralized, server. Moreover, the use of a central server makes maintenance and control a lot easier.

A way to obtain IP addresses from a central server is to use RSIP recursively : the RSIP gateways are themselves clients of a second level central RSIP server.

Another way, presented in this paper, is to extend the RSIP protocol (which is still an experimental) to make it support more possibilities.

A new agent is introduced in the system, the RSA-IP server, and is used to maintain, in a centralized way, the pool of addresses to be leased. The RSA-IP gateways are still located in the extruded subnets but don't own public addresses anymore : they are downgraded to proxies. See figure 2. A tunnel is established between the RSA-IP gateway and the RSA-IP server. Because of the use of this tunnel, the main network need not to be close to the extruded subnet. The tunnel may obviously be an IPSec tunnel.

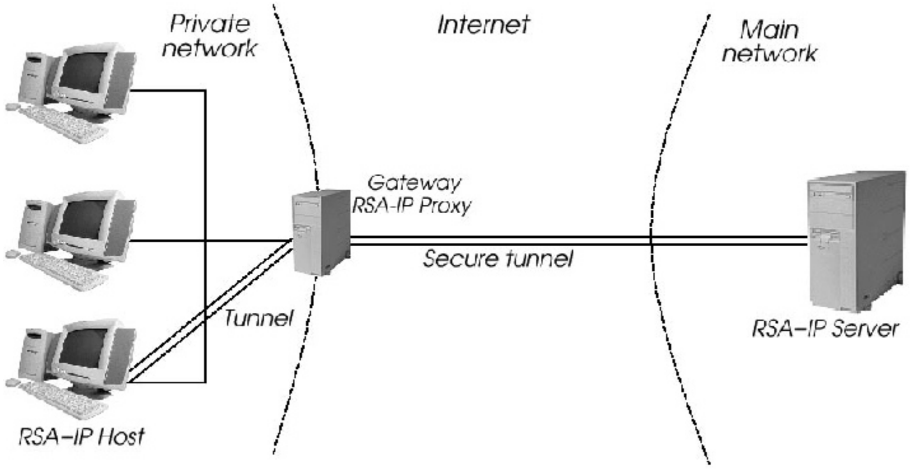

Fig. 2. A remote extruded subnet connected to the main network through a RSA-IP gateway and a RSA-IP server

A RSA-IP gateway just forwards requests from hosts of the extruded subnets to the server, which in turn replies just as if the requests were coming from a regular host. The specifications of this extension are beyond the scope of this paper. A prototype has been implemented [12].

For the purpose of dynamically allocating IP addresses to extruded subnets, the extension is equivalent to the recursive use of RSIP. However, the extension offers more possibilities, particularly for the dynamic binding of hosts to permanent domain names.

With this extension to RSIP, all the traffic to and from the extruded subnet travels through the tunnel. The Internet must thus route this traffic to the RSA-IP server and not directly to the gateway.

3.1 Binding Hosts in the Extruded Subnets to Permanent Domain Names

We want addresses in extruded subnets to be allocated dynamically but to be bound to permanent domain names. A partial solution is to let the server dynamically update the tables of the DNS server each time a resource is allocated to a host (e.g. using DNS Update protocol [13], [14]). However, this is only possible if the domain name of the host has been transmitted to the central RSIP server. This is not supported by the classical RSIP.

Moreover, at the time someone on the Internet tries to contact a server located in an extruded subnet, the latter may not yet have a dynamically leased public address. Those reasons led us to extend the classical RSIP solution.

The extension presented here and detailed in [12], proposes that a host sends its domain name when it registers with the central RSA-IP server. This way, when the host receives a leased IP address from the server, it becomes reachable by anyone using its domain name.

In addition, the extension also offers the possibility for a host to be warned (messages 2 and 3 in the figure 3) when it is contacted (message 1) by a public host, even if the contacted host has not yet requested a public IP address. In this case, the RSA-IP host may then request a public IP address from its gateway (messages 4 and 5) in order to be reachable by its correspondent. The solution is based on a two-way communication between the central RSA-IP server and the dynamical DNS server (messages 2 and 6).

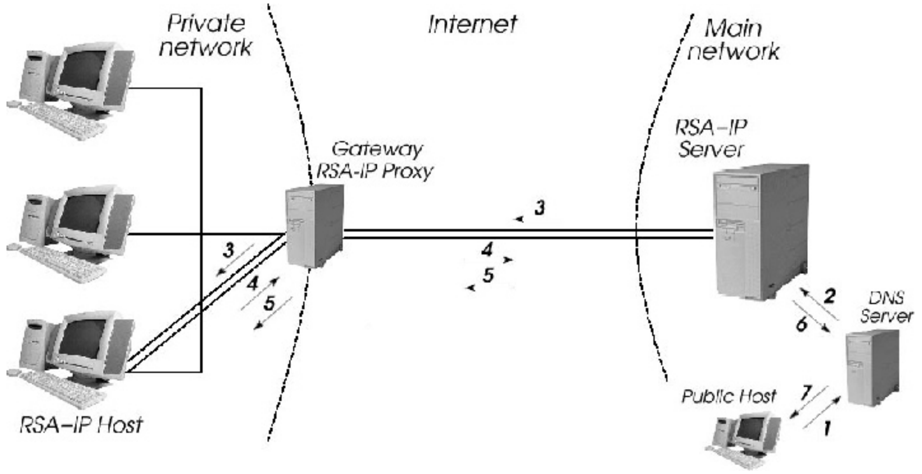

Fig. 3. Messages exchanged when a public host contacts a RSA-IP host which is not yet leasing an IP address

Thanks to the extension described above, DNS requests related to RSA-IP hosts can be handled through cooperation between RSIP and DNS servers.

This problem will not be further discussed in this paper.

4 Comparison between DHCP-Based and RSIP-Based Implementations of Extruded Subnets

We will compare in this section the use of the RSIP protocol instead of the DHCP protocol for the dynamic allocation of the public IP addresses to the hosts of the remote extruded subnets.

The services offered by RSA-IP and by DHCP are very close. However, RSA-IP provides somewhat different functionalities. For example, a RSIP gateway may be a policy enforcement point. In other words, it may have the ability to explicitly control which local addresses and ports are used to communicate with remote addresses and ports.

Both RSIP and DHCP have functionalities that can be used to spare IP addresses : RSA-IP as described in [8] and DHCP allow to specify expiration times for each allocated IP address. When this time expires, the RSA-IP or DHCP client may ask to extend its lease time. The RSA-IP gateway or the DHCP server may accept this extension or not. Thanks to these functionalities, we can dynamically allocate public IP addresses to hosts for the time they really need. This mechanism allows the saving of IP addresses.

DHCP and RSIP solutions are similar from the point of view of centralized address allocation. Both have been designed for on demand temporary allocation of IP addresses to hosts, but not to hosts in extruded subnets. However, both protocols can be used for this purpose.

A significant difference between DHCP and RSIP is how a host communicates with the server that allocates the addresses (DHCP server or RSIP gateway) when this server is on another physical network. A DHCP host communicates with the DHCP server through a DHCP relay agent located on its network because it has no IP address when the DHCP transaction is started. A RSIP host has already a local IP address when it starts the transaction and obtains a second (public) one through a RSIP transaction with the RSIP gateway. For this address to be useable, special routes must be set up (theoretically in the whole world) towards this address. Instead, a tunnel is usually set up between the host and the gateway and all external routes to RSIP hosts beyond a RSIP gateway point to this gateway. It must be noted that, before requesting an IP address, a RSIP host has to register with the RSIP gateway.

Thanks to relay agents, DHCP is very scalable regarding to the number of subnets. Besides, relay agents require no management. However, if only one relay agent is used in each subnet (i.e. the gateway), only simple networks (e.g. one ethernet) can be supported in the subnet. On the other hand, "standard RSIP" has been designed to allocate external addresses in large networks with any number of routers etc. So the scalability of RSIP is excellent regarding the size of the subnets, but RSIP has not been designed for handling several subnets. This means that each gateway must be managed by hand. This problem is solved either by using RSIP recursively or by the extension proposed to RSIP : the central RSIP server. With this, RSIP gets the same scalability as DHCP regarding the number of subnetworks.

DHCP has one significant advantage over RSIP : it is a well known and widely used protocol available on many operating systems. No special software must be added on the hosts of the extruded subnets. This is not the case with RSIP : software must be added on the RSIP hosts with current operating systems releases.

RSIP is designed to allocate a supplementary (external) IP address to a machine that has already an internal one. So local traffic can use the local address and external traffic can use the external address. Note that if RSIP is used to obtain the external address, DHCP can be used to obtain the local address.

DHCP is designed to allocate a first IP address to a machine. This single address serves two purposes : local and external. This brings a problem when addresses are allocated randomly in a set of subnets : when one cannot distinguish a machine on one's own subnet and on a remote subnet which makes routing between subnets impossible, proxy ARP creates the illusion that all subnets of a main network constitute a single network and remove the need of routing between them. The effect is the same as with the double address of RSIP.

Both the DHCP and RSIP based solution use IPSec tunnels to satisfy the security requirements of extruded subnets.

From a performance point of view, both solution are equivalent. They only differ in the address allocation to the hosts of the remote subnet, which is a relatively unfrequent operation without performance impact. During regular operation, the only overhead is that induced by IPSec. This is similar to what happens with any VPN. This overhead is negligible for ADSL and cable modem connections.

5 Conclusion

The problem exposed in this paper is to build extruded subnets, which are remote subnetworks virtually imported in another network, with IP addresses belonging to this network allocated on demand only to hosts that need it when they need it.

The solution proposed is based on the new RSIP protocol, modified to extend its functionality in order to manage IP adresses exported in the extruded subnets dynamically and centrally from the main network. Computers in the extruded subnets and others in the main appear to be in the same network.

The RSIP based solution has several advantages over the previous one based on DHCP. It has a better scalability regarding the size of the subnets and offers extended possibilities concerning the binding of RSA-IP host to permanent domain names : a RSA-IP host may be assigned an address at the time it is contacted by an external client. The extra cost of the RSIP solution is the necessity to add a software agent on each host in the extruded subnet. This agent must be designed for each operating system.

When these advantages are not useful, the DHCP solution is to be preferred since it does not require extra software on the client machines.

Acknowledgements. This work has been partially supported by grant FIRST 981/3749 by MRW-DGTRE ("Travaux réalisés avec l'aide du Ministère de la Région Wallonne, en application du programme de formation et d'impulsion à la recherche scientifique et technologique") and by a grant from FRIA (Fond pour la Formation à la Recherche dans l'Industrie et dans l'Agriculture, rue d'Egmont 5 - 1000 Bruxelles).

The definition and implementation of the extension to the RSA-IP is a common work of C. de Launois, G. Fauveaux and J. Honlet.

The authors express their thanks to the anonymous reviewers for their useful comments.

References

1. Freier, A.O.; Karlton, P.: Kocher P.C.; "The SSL Protocol Version 3.0", Internet Draft <draft-freier-ssl-version3-02.txt>, 1996.
2. Ylonen, T.; Kivinen, T.; Saarinen, M.; Rinne, T.: Lehtinen, S.: "SSH Protocol Architecture", Internet Draft <draft-ietf-secsh-architecture-09.txt>, work in progress, July 2001.
3. Kent, S.; Atkinson, R.: "Security Architecture for the Internet Protocol. Network Working Group", RFC 2401, 1998.
4. Doraswamy, N.; Harkins, D.: "IPSec : The New Security Standard for the Internet, Intranets and Virtual Private Network", Prentice Hall PTR, 1999.
5. Bonnet, A.; Lobelle, M.: "Extending a Campus Network with remote Bubbles using IPSec", I-NetSec'01, Leuven, 2001.
6. Ranch, D.: "Linux IP Masquerading HOWTO. Technical Report", 2000.
7. Droms, R.; Lemon, T.: "The DHCP Handbook", Macmillan Technical Publishing, 1999.
8. Borella, M.; Grabelsky, D.; Lo, J.; Taniguchi, K.: "Realm Specific IP : Protocol Specification", RFC 3103, October 2001.
9. Borella, M.; Lo, J.; Grabelsky, D.; Montenegro G.: "Realm Specific IP : Framework", RFC 3102, October 2001.
10. Montenegro G.; Borella, M.: "RSIP Support for End-to-end IPSEC", RFC 3104, October 2001.
11. Kuznetsov A. N.: "IP Command Reference", Institute for Nuclear Research, Moscow, April 1999.
12. de Launois C.; Fauveaux G.; Honlet J.: "Routeur d'accès à adresse dynamique", Université catholique de Louvain, Louvain-la-Neuve, 2001. http://openresources.info.ucl.ac.be/rsip/
13. Eastlake, G.: "Secure Domain Name System Dynamic Update", RFC 2137, April 1997.
14. Wellington, B.: "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, November 2000.