

Security Measures to couple mixed IPv4/IPv6 Networks over a pure IPv6 Infrastructure by making Use of NAT-PT

Thorsten Brikey

© SANS Institute 2003, Author retains full rights

1 Abstract

The next generation of the Internet Protocol (IPv6) was developed to improve the within the Internet widespread deployed Internet Protocol (IPv4). Among other things it enlarges the available addressing space and improves security.

Due to lack of unique IPv4 address space one strategy to couple existing IPv4 networks that uses private IPv4 addresses is to define a unique IPv6 network on top of the coupled IPv4 networks. At the border routers the use of NAT-PT (Network Address Translation – Protocol Translation) ensures that the IPv4 hosts have assigned virtual unique IPv6-addresses. IPv6 is used between the border routers. This ensures a smooth migration from a pure IPv4 to a pure IPv6 environment. The scope of this paper is to present a European test installation where NAT-PT is used to couple national networks over an IPv6 backbone. The description focuses on one national test installation with respect to security. The security impact on each communication layer will be discussed to ensure an acceptable level of security within the NAT-PT test installation.

2 Introduction

In the early days of the Internet each machine using the internet protocol (IPv4) was assigned its own globally unique IP-address (official IP-address) irrespective of the need to communicate with the Internet itself. During the last few years the number of applications using the internet protocol has enlarged in a dramatic way so that the official available addresses were in shortage. In order to extend the life of the IPv4 address space, address registries are requiring more justification than ever before, making it harder for organizations to acquire additional address space [1]. The foreseen lack of IP-address space was one of the major reasons to develop the next generation of IP (IPv6) [5]. Although this protocol was defined in the late nineties the number of manufacturers supporting IPv6 has increased only during the last two years.

Another strategy to reduce the demand for globally unique IP-addresses was the definition of private usable IP-address-ranges every enterprise can use within their own environment without registration [3]. The address space can thus be used by many enterprises. Addresses within this private address space will only be unique within the organization. Since these addresses were not allowed to be routed inside the internet the deployment of private IP-addresses improves the enterprises IT-security because all hosts using such addresses were not directly reachable over the Internet. Due to lack of official IPv4-addresses and the poor support for IPv6 many enterprises decided to use the private IPv4-address ranges.

In some cases the need to couple one or more IP-networks using private addresses may occur. Within the air traffic control (ATC) environment at the moment data like radar data or flight plans were exchanged between different national air traffic service providers (ATSP) via X.25. Since X.25 reaches its end of lifecycle at the end of this decade it has to be replaced by some other protocol like IP. As a consequence multiple existing networks making use of private IP-addresses have to be coupled. To investigate possible technical solutions for a future European ATC-network a Task Force was founded by Eurocontrol [9].

During the last two years this Eurocontrol Task Force developed recommendations for a future European ATC network. To evaluate this proposed network concept each

European ATSP was invited to build up a national test installation including ATC services. These national test networks are expected to be coupled at the end of May 2003 so that a real European ATC network on basis of IPv6 could be simulated. Pre-tests already have shown that the described mechanisms will work in principle.

3 NAP-PT within ATC network environment

To be open for future developments in telecommunications and to avoid multiple network address translation (NAT) processes at the national boundaries the Eurocontrol Task Force decided to define a unique European IPv6 network. In this concept the purely IPv4 based hosts have been assigned a fixed virtual IPv6 address. At the national boundaries network address translation – protocol translation (NAT-PT) [8] ensures the definite translation between real -- private -- IPv4 address and virtual unique IPv6 address (see figure (1)). As it is expected that the transition phase from a pure IPv4- to a pure IPv6-environment will last for a couple of years a further advantage of this concept is that it allows the coexistence of IPv4 and IPv6 networks on the same infrastructure.

The considered test infrastructure in principle is shown in figure (1). Several national ATSP networks are coupled over an international IPv6-backbone. The links between the national networks can be realized via different media like leased lines or the Internet. The NAT-PT routers are located at the national borders. The next generation of the Border Gateway Protocol BGP4+ [4] is used as routing protocol. The security policy demands that each national network manager is only responsible for traffic that enters his network.

3.1 Network Infrastructure

For reaching the IPv6 networks the border routers have to perform a simple routing process. To ensure that the IPv4 hosts are reachable from the outer world NAT-PT has to be taken out. Like the classical Network Address Translation (NAT) [2] NAT-PT could be configured in a static (fixed map between IPv4- and IPv6-address) as well as in a dynamic way (virtual IPv6 address will be assigned out of an address pool). Since the global network is logical based on pure IPv6 the DNS server is located in the IPv6 network. Because commercial IPv6- firewalls are not available on the market today only the IPv4-networks are protected by this kind of equipment. It has to be investigated which additional mechanisms have the ability to secure the pure IPv6 networks. Following NAT-PT restrictions will be investigated and their impact on security issues discussed. This should not affect IPv6 to IPv6 communication.




Figure 1 Network Infrastructure in Principle

3.2 NAT-PT Limitations

The NAT-PT translation method has some limitations that are similar to the classical NAT. For instance it is mandatory that all requests and responses pertaining to a session has to be routed via the same NAT-PT router. One way to guarantee this would be to have NAT-PT based on a border router that is unique to a stub domain, where all IP packets are either originated from the domain or destined to the domain [6]. This point is reflected in the chosen topology where NAT-PT is performed on the border routers. In the future when IPv6 will be widespread deployed it is expected to have NAT-PT routers only at the borders of IPv4 islands.

These border routers are an attractive target for any kind of attacker because on these machines they can get information of IPv4 hosts that have to communicate with international partners. Furthermore international communication processes can be disturbed when these border routers were attacked. For these reasons the border routers need a good protection which will be discussed in section 4.3.

3.2.1 Redundancy

As mentioned above, a session in- and outbound communications must traverse the same routing device due to the state information contained within the packet i.e. no asymmetrical routing can occur. This will affect the resilience of the network, as a new session will need to be recreated through an alternative path if the physical routing device fails and dynamic NAT-PT is used. The workaround to this is to use static NAT-PT tables within the routers. A further advantage of this action is that the network manager knows exactly the IP addresses of each host. While using dynamic NAT-PT it would be easier for an attacker to make unnoticed use of legal IP addresses out of the NAT-PT address pool so that he can get access to one of the national ATC networks.

3.2.2 Applications

Since NAT-PT performs address translation, applications that carry the IP address in the higher layers like DNS will not work. In this case Application Layer Gateways (ALG) need to be incorporated to provide support for those applications [6]. Because ATC applications are not standard applications it is necessary to define application driven system identification and authentication methods that are independent of the particular IP address (section 4.5).

3.2.3 End-to-End Security

One of the most important limitations of the NAT-PT is the fact that end-to-end network layer security is not possible. Also transport and application layer security may not be possible for applications that carry IP addresses to the application layer. This is an inherent limitation of the Network Address Translation function [7]. Independent of NAT-PT, end-to-end IPsec security is not possible across different address realms. The two end-nodes that seek IPsec network level security must both support one of IPv4 or IPv6. The impact of this fact is discussed in section 4.3.3. Concerning the international ATC network infrastructure real end-to-end IPsec is not possible between IPv4 domains.

3.2.4 IP-Field Translation

A number of IPv4 fields have changed meaning in IPv6 and translation is not straightforward. For example, the option headers semantics and syntax have changed significantly in IPv6 [7]. IP header conversion is complex and information may be lost across boundaries. Although an attacker could make use of an inconsistency within any kind of communication process in general there is no well-known attack yet that actually makes use of this NAT-PT behaviour. But the developments on this field have to be observed!

3.2.5 DNS

Since the European network is designed for IPv6 each IPv4 host within this network has assigned a virtual IPv6 address that over the entire network there is a unique IPv6 address scheme. For that reason the DNS server is located inside the IPv6 network and that would obviate the need for DNS-ALG intervention. It is clear that this scheme can not be deployed in combination with secure DNS. I.e., an authoritative DNS name server in the IPv6 domain cannot sign replies to queries that originate from the IPv4 world. As a result, an IPv4 end-node that demands DNS replies to be signed will reject replies that have been tampered with by NAT-PT. Therefore only servers in IPv6 domain that need to be accessible from the IPv4 world pay the price for the above limitation, as **IPv4 end-nodes may not access IPv6 servers due to DNS replies not being signed.**

4 Security Aspects

In the previous section NAT-PT and its limitations were discussed with respect of possibly security lacks. In this section measures will be presented which improve the security of coupled networks using NAT-PT. The major underlying policy is that **each company is responsible for the incoming data traffic**.

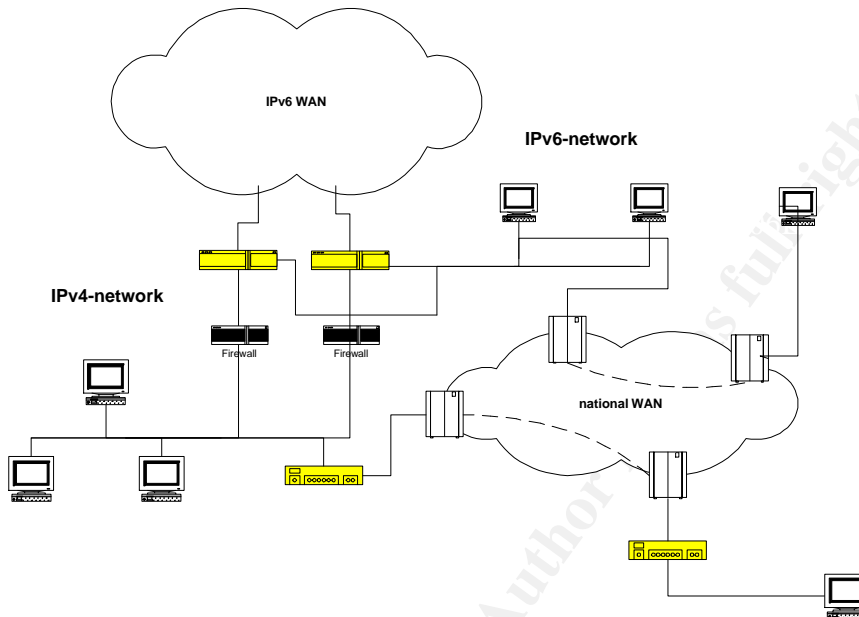


Figure 2 National Network Architecture

The discussion is focused on one national network of this future European ATC network which is shown in principle in figure (2). The assumption of the chosen national topology is that it represents the transition phase in which only new systems are capable in communicating via pure IPv6. The old systems are completely located in the IPv4 network and it is not expected to migrate them to IPv6 or to make them dual stack hosts. Therefore the description is restricted to two separated network clouds -- one for IPv4 and the other for IPv6 only. In this particular case NAT-PT is performed by two Cisco routers at the national border.

4.1 IPv4 Network

Since IPv4 is the default protocol used within the Internet a great variety of security equipment is available at the moment. Therefore the entry point to the IPv4 network is protected by two chained standard firewall systems produced by different manufacturers. This is senseful because if a new vulnerability is announced for one of these firewall systems the network behind the two firewalls is still protected by the stable system.

For the particular case of ATC applications the firewall systems can only act as stateful inspection packet-filter because ATC applications in general make use of propriety TCP/UDP-ports so that proxies are not applicable. As an additional factor to ensure that only expected communication between well known hosts occur an

Intrusion Detection System (IDS) is installed within the IPv4 cloud. But the precondition of using IDS is the deployment of static NAT-PT. In so far the network protection of the IPv4 network is based on standard methods.

4.2 IPv6 Network

As mentioned above commercial IPv6 firewalls that provide full firewall features are not available at the moment. Actually Check Point announced that firewall-1 NG is capable in dealing with a few IPv6-feature at the moment [12]. Therefore the only way to protect the IPv6 network on basis of the lower communication layers is to configure access lists on the border routers (section

4.3.1) that allows only the needed services to enter the IPv6 network.

To achieve an acceptable level of security for the higher layers it is important to disable all non needed TCP/UDP services on the hosts on which the applications run. Since within the ATC environment propriety TCP/UDP ports are used and knowledge about the applications is not widespread, the probability for a successful attack decreases.

4.3 NAT-PT Routers and WAN Links

The NAT-PT border routers are not only the first devices to be attacked from the outside they also ensure the communication to other ATSPs as well. Failed devices and misled traffic between European ATSPs due to an attack can cause serious delays for European flights because air traffic controllers have to rely on radar data or flight plans. Therefore the border routers have to be protected in particular.

4.3.1 Router Configuration

Within the ATC environment it is a wise decision to follow the recommendations made by the American National Security Agency (NSA) concerning the configuration of Cisco boxes [11]. In this paper five points to ensure a secure configuration of a Cisco Router were highlighted:

- Creation and maintenance of a router policy. The policy should identify who is allowed to what on the router. Especially remote access to the router should only be allowed from the locale network with a secure service like Secure Shell (SSH).
- Router configurations files should turned out offline.
- Access lists should ensure that only the needed protocols and services are allowed to pass the router.
- Deployment of latest official IOS version.
- The security of the router should be checked after **every** configuration change.

As mentioned in section (\ref{redundant}) a further step in securing the IPv4 hosts is the deployment of static NAT-PT. In this case an attacker can not reach unnoticed the IPv4 network cloud by using an IPv4 address out of the address pool provided by the NAT-PT router.

4.3.2 Routing

Since not all ATSP are connected to each other transit traffic from one national network to another through a third ATSP network may occur. As an example the German network has no direct link to Spain. If a German end-system has to communicate with its Spanish counterpart the data has to pass through France. Therefore the Border Gateway Protocol BGP4+ [4] is the most suitable routing protocol to fulfil these requirements. Nevertheless each country is free in choosing another routing protocol for its own purposes. But each ATSP has to ensure that the routing protocol used within its country has to be redistributed into BGP4+ at the border routers.

Regardless of the chosen routing protocol a manipulation of routing databases can easily be carried out. The easiest way to do this is to connect a router that distributes manipulated routing information to one or more border routers. In this particular case an attacker can interrupt data connections or manipulate ATC data like flight plans. To prevent a potential attacker from manipulating any kind of routing database routing authentication is enabled on each interface of at least the border routers. This ensures that any additional router illegally connected to the ATC WAN will not have the ability to manipulate routing information if the authentication passwords were kept secret by the national ATSP.

4.3.3 WAN Links

Due to NAT-PT end-to-end security is not possible for IPv4-to-IPv4 host communication (section 3.2.3) within the European ATC network. Therefore the traffic over the WAN links has to be protected in another way. This is valid in case of the situation when WAN links are established via the Internet. For this reason and in view of the fact that the links between the different national border routers are established via IPv6 the NAT-PT routers have to be regarded as some kind of end-system for the IPv6 networks. To make this clear one has to imagine the situation in which the national ATC networks are all IPv4 based. Then only the border routers are capable in communicating via IPv6. In view of the border routers end-to-end security is achievable simply in deploying IPv6sec for WAN links.

When the first pure IPv6 end-system wants to communicate with an IPv4 host behind a NAT-PT router then the IPsec-tunnel will be established between IPv6 host and NAT-PT router. Due to the logical separation of IPv4- and IPv6-networks this works without limitations.

During the test phase the encryption keys are exchanged manually between the national ATSP. In future a European ATC Certification Authority (CA) is planned to be set up to ensure the key exchange on a regular basis.

4.4 DNS Security

Due to the fact that all international communications are IPv6 based the DNS server is located in the IPv6 world. As discussed above (section 3.2.5) DNSsec is not applicable if DNS certificates have to cross the NAT-PT routers. Since the NAT-PT routers have to use DNS-ALG the NAT-PT router itself can be regarded as an end-system that queries the DNS. Therefore it is possible to use DNSsec between the NAT-PT routers and the DNS-server. Pre-tests have shown that Cisco routers fulfil

this requirement. All IPv6 hosts are not affected by this mechanism and can query the DNS-server via DNSsec directly.

Furthermore the specific ATC environment produces restrictions that cause a gain of security as well. One of those restrictions is that communication has only to be allowed between clear defined partners; the Flight Data Processing System (FDPS) at site A must exchange data with FDPS at site B but must not communicate with FDPS at site C. The result of this requirement is that NAT-PT has to make use of static mapping. When using dynamic NAT-PT where an IPv6 host got assigned an IPv4 address out of a pool of IPv4-addresses the situation can occur that there is a shortage of IPv4 address space. So a legal connection between an IPv6- and an IPv4-host can not be established if an attacker tries to block all IPv4 addresses out of the pool simply by requesting more IPv4 addresses than there are available. This kind of denial of service attack is not possible to be carried out when using static NAT-PT.

4.5 Application Security

All over the world each country has its own Air Traffic Control Service. Therefore a flight that crosses national boundaries is controlled by several ATC service providers. Having this in mind it is a factual issue that standards have to be defined as well for technical systems as for flight procedures. In this document flight procedures are out of scope so that the focus lays on the technical standards. During the last two decades X.25 was a defacto standard for the exchange of ATC data like flight plans or radar data. Since industry diminishes support for the X.25 protocol it was decided for the European environment to replace this protocol suite with the Internet Protocol (IP). This forced the European regulator Eurocontrol to redefine the mandatory standards for the European ATC providers for flight plan and radar data exchange. So a Task Force has to develop standards for ATC applications.

4.5.1 Application Handshake Process

The created Task Force started to develop an Interface Control Document (ICD) for flight plan exchange on IP basis [10]. The lower X.25 communication layers were replaced with the standard IP layers up to TCP/UDP. In addition the applications shall be independent of the IP version. It has to be mentioned that in contrast to X.25 where the both machines which exchange data are equal within the IP-world a client server architecture has to be used.

In view of data transfer UDP is an unreliable protocol it was decided to use TCP only. This ensures that each data packet sent by a host will reach the receiving host. To follow the TCP philosophy it was decided to agree on a fixed TCP port number for each ATC application, i.e. TCP port X has to be used for flight plan exchange, while TCP port Y is used for radar and so on. Therefore the TCP ports used by air traffic applications are in a special meaning propriety.

Since many applications of the same type are running on the same host the IP address is insufficient to identify the addressed host. Suppose the flight data processing system for two different airports are running on the same machine then the addressed airport can not be identified by the IP address of the host. Therefore a

new communication layer between layer four – TCP – and the application layer was created.

The design of this so called 4-b layer is similar to TCP itself. After the TCP connection has been established the calling host – the client – signals that it wants to communicate with a server by sending an identification message to the server which contains client **and** server identification string (ID). This process is similar to the TCP SYNC process. If the client has not sent a valid client ID the server sends a REJECT message. In case that the client is allowed to contact the server the server responds with a message that contains server and client (ID) as well (comparable to SYNC/ACK). To confirm that the server is still a valid communication partner for the client it sends an OK messages to the server. Otherwise the server got back a REJECT message. In case the server receives the clients OK message the application itself can start the data exchange. This message exchange process on top of the TCP layer is called the Tree Way Application Handshake (TWAH). The definition of the TWAH is intentional left open. For example it is allowed to be turned out more than once so that authentication messages could be exchanged during the second, third, and so on TWAH. In which exact way the TWAH has to be turned out has to agreed on a bilateral basis between the two organizations of the client and server. This is valid for the ID messages as well as for some authentication keys.

Although the Tree Way Application Handshake originally was designed for the flight plan exchange the definition of TWAH is quite general. To increase applications security it can be used for other applications as well and is one part of IT defense strategy.

5 Conclusion

The deployment of NAT-PT makes it possible to couple IPv4 networks which use private IPv4 addresses while defining a unique IPv6 address structure. It ensures a smooth transition from a mixed IPv4/IPv6 network infrastructure to a pure IPv6 world as well. Since during the last few years many manufacturers provide IPv6 capable products which have reached maturity a reliable network infrastructure that uses NAT-PT can be built up. This was demonstrated by the ATC test infrastructure. Although NAT-PT has some security limitations the use of suitable measures on all communication layers ensures a trustworthy level of security. This is all the more true as IPv6 firewalls that able to provide the same features as already for IPv4 are expected to be sold in nearest future. Therefore the use of NAT-PT is a suitable and secure way to migrate existing IPv4 infrastructure to a pure IPv6 network. To prove this statement during the year 2003 further investigations including hacking trials will be turned out to find security holes within the test-installation to evaluate the archived level of security.

References

- [1] E. Gerich, *Guidelines for Management of IP Address Space*, RFC 1466, May 1993
- [2] K. Egevang, P. Francis, *The IP Network Address Translator (NAT)*, RFC 1631, May 1994
- [3] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, E. Lear, *Address Allocation for Private Internets*, RFC 1918, February 1996
- [4] T. Bates, R. Chandra, D. Katz, Y. Rekhter, *Multiprotocol Extensions for BGP-4*, RFC 2283, February 1998
- [5] S. Deering, R. Hinden, *Internet Protocol Version 6 – Specification (IPv6)*, RFC 2460, December 1998
- [6] P. Srisuresh, M. Holdrege, *IP Network Address Translator (NAT) Terminology and Considerations*, RFC 2663, August 1999
- [7] E. Nordmark, *Stateless IP/ICMP Translator*, RFC 2765, February 2000
- [8] G. Tsirtsis, P. Srisuresh, *Network Address Translation – Protocol Translation (NAT-PT)*, RFC 2766, February 2000
- [9] Eurocontrol: Internet Protocol for Aeronautical eXchange Task Force (iPAX-TF), <http://www.eurocontrol.int/ipax>
- [10] Eurocontrol iPAX-TF: *Eurocontrol Standard Document for Flight Data Exchange, Interface Control Document, Part2, TCP/IP*, April 2003
<http://www.eurocontrol.int/ipax/docs>
- [11] NSA/SNAC *Router Security Configuration Guide*,
<http://nsa1.www.conxion.com/cisco/guides/cis-1.pdf>
- [12] CheckPoint
<http://nsa1.www.checkpoint.com/press/2002/ipv6\ 081402.html>

© SANS Institute 2003, Author retains full rights