

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
9 August 2007 (09.08.2007)

PCT

(10) International Publication Number
WO 2007/088174 A1

- (51) International Patent Classification:
H04L 29/06 (2006.01) H04L 29/12 (2006.01)
- (21) International Application Number:
PCT/EP2007/050931
- (22) International Filing Date: 31 January 2007 (31.01.2007)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
0601913.7 31 January 2006 (31.01.2006) GB
- (71) Applicant (for all designated States except US): TELEFONAKTIEBOLAGET LM ERICSSON (PUBL) [SE/SE]; S-164 83 Stockholm (SE).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): ARKKO, Jari [FI/FI]; Kauppalantie 25 A 7, FI-02700 Kauniainen (FI).
- (74) Agent: SUCKLING, Andrew, Michael; 4220 Nash Court, Oxford Business Park South, Oxford Oxfordshire OX4 2RU (GB).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,


AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— with international search report
— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: PACKET REDIRECTION IN A COMMUNICATION NETWORK



(57) Abstract: A method of controlling the re-direction of IP packets to an IP host having two or more different IP addresses comprises generating a first of said IP addresses as a one-way function of the second IP address. The method further comprises accepting a request to re-direct a packet destined to said first IP address to another IP address only if the other IP address is the second IP address.

WO 2007/088174 A1

PACKET REDIRECTION IN A COMMUNICATION NETWORK

Field of the Invention

The present invention relates to a method of controlling packet re-direction in a communication network. The method is particularly applicable to controlling re-direction of packets destined for a host that has two or more addresses, for example an address according to a first address realm and another address according to a second address realm.

Background of the Invention

Most Internet applications either suffer or have suffered from problems relating to Internet Protocol (IP) address realm transitions. Such problems are particularly prevalent since the format of IP addresses is currently changing with the introduction of IP version 6 (IPv6) address realm to replace the IP version 4 (IPv4) address realm (as the number of addresses available under the IPv4 realm is no longer sufficient to meet demand). Examples of these problems include the inability to place servers behind Network Address Translators (NATs), difficulty in communicating from the IPv4 Internet to the IPv6 Internet and vice versa, limitations in multimedia applications that do not employ global addresses, and so on. A Name Address Translator, as an example, translates between an internal IP address of a device and the global IP address of the device – where the internal IP address is allotted to a device by local network, which may not want the internal address to be sent outside the network's firewall.

A number of solutions have been devised to deal with these problems. For instance, many applications and protocols have built-in "NAT traversal" mechanisms and a plethora of IPv4-to-IPv6 transition mechanisms exists. These solutions have, however, created a number of problems of their own.

One problem is that the inherent ambiguity of local address space leads to trial-and-error approach to finding a working path to a peer. A peer found through a server might

be in the local network, or might be in a different local address space somewhere else. Moreover, the security of this approach relies entirely on probing and is susceptible to attackers on the path. For instance, an attacker in a local network would be capable of hijacking communications intended to someone else, even if the SIP server exchange that delivered the peer's address were performed securely.

Another problem is that configuration is needed in many of the solutions, and hence it becomes impossible for new devices to work "out of the box" in an automatic fashion. For instance, the Microsoft® Teredo IPv4-IPv6 transition service requires the configuration of the server's address, and typically requires the configuration of a shared key between hosts and servers.

In a communications network, a particular device may be identified by two or more addresses. The addresses may belong to different address realms from one another, or may belong to the same realm. In the communication network shown in figure 1, a host 1 in a first network (network 1) is communicating with a peer host 2 in another network (network 2) over, for example, an IP backbone 3. If host 1 has two different IP addresses, such as an IPv4 address and an IPv6 address, packets sent from host 2 to host 1 may identify host 1 as the destination by either the IPv4 address of host 1 or the IPv6 address of host 1. The choice of which address to use may for example depend on network availability or on other requirements – for example, some Internet technology is fundamentally limited to using IPv6 addresses. Thus, if host 2 is directing packets to the IPv6 address of host 1, host 1 might wish at some stage to instruct host 2 to switch to use the IPv4 address for host 1, or vice versa.

Assume that packets are being sent from a device in network 2, for example device 21, to host 1 of network 1, for example, in response to a request made by device in network 1. If the original request specified that the IPv6 address of host 1 should be used, the packets arriving at host 2 from device 21 will specify, as their destination address, the IPv6 address of host 1. If host 1 wishes to switch to using its IPv4 addresses, host 1 would send a re-direction message to host 2; the re-direction message would contain the IPv4 address for host 1, and would instruct host 2 to re-direct packets destined for host 1

to this IPv4 address rather than to the IPv6 destination address specified in the packet. Host 2 would then re-direct packets intended for host 1 to the IPv4 address of host 1 rather than to the IPv6 address.

One problem is that host 2 needs to be able to authenticate a received message that instructs it to re-direct packets to a new address. A malicious host, shown as host 4 in figure 1, could wish to intercept the packets passing from host 2 to host 1. The malicious host might therefore send a message to host 2 that contains an IP address for host 4, or for a device controlled by host 4, and instruct host 2 to re-direct packets destined to host 1 to this new address rather than to the IPv6 address for host 1 specified in the packets. If host 2 were to obey this instruction, host 2 would re-direct packets to the malicious host 4 rather than to the intended host 1. It is therefore important that host 2 should be able to distinguish between a genuine re-direction message sent by host 1 and a malicious re-direction message sent by host 4.

An IPv6 address typically has the form of a prefix and a suffix. The prefix of the address is a "routing prefix" and is allotted to, for example, a host such as the host 1 in figure 1. The suffix is an "interface identifier", and identifies a particular device. Whenever a device in the network controlled by the host 1, such as device 11 in figure 1, wishes to connect to the Internet, an IPv6 address is generated for that device by choosing a suffix and combining the suffix with the routing prefix allotted to host 1. It is then necessary to check that the generated IPv6 address is not already in use by another device in the network, ie to check that the particular interface identifier suffix chosen for device 11 is not already in use by another device in the network 1. This done by broadcasting a message that includes the IPv6 address generated for device 11 and asking other devices in the network, such as the devices 12,13 in figure 1, to indicate if they are already using the address. This process is, however, also vulnerable to attack by a malicious party. For example, if malicious host 4 were to read the broadcast message including the IPv6 address generated for device 11 and respond by claiming that the particular IPv6 address chosen for device 11 was already in use, the device 11 would not be allotted that IPv6 address. It would be necessary to generate another IPv6 address for device 11, and to broadcast a message including that address to check

whether it was in use. If the malicious host 4 were to respond to every such broadcast message by claiming that the IPv6 address included in the message was already in use, the malicious host 4 would be able to prevent an IPv6 address being allotted to device 11.

Summary of the invention

A first aspect of the present invention provides a method of controlling the re-direction of IP packets to an IP host having two or more different IP addresses, the method comprising the steps of:

generating a first of said IP addresses as a one-way function of the second IP address; and

at a peer host or network routing node, accepting a request to re-route a packet destined to said first IP address to another IP address only if the other IP address is the second IP address.

A “one-way function” is a function such that, for a given object and a given one-way function, it is straightforward to calculate the result of applying the one-way function to the object – but, for a given result, it is very difficult to determine the original object to which the one-way function was applied to obtain the known result, even if the one-way function is known. “One-way” functions of this type are well-known. A cryptographic hash function is one example of a “one way” function.

In the method of the invention, therefore, the peer host or routing node can easily tell whether the re-direction request is genuine, provided that they know the one-way function used to generate the first IP address from the second IP address. They simply need to apply the one-way function to the IP address contained in the re-direction request and see whether the result is the first IP address: if the result of applying the one-way function is the first IP address, then the request may be taken to be genuine - but if the result is not the first IP address then the re-direction request should be regarded as malicious.

In a preferred embodiment of the invention, the method further comprises including the first IP address in a packet destined for the second IP address, and, at a peer host or network routing node, accepting a request to re-route a packet destined to said second IP address to another IP address only if the other IP address is the first IP address included in the packet.

Generating the first IP address as a one-way function of the second IP address solves the problem of authenticating a request to re-direct a packet destined to the first IP address. However, it does not solve the problem of authenticating a request to re-direct a packet destined for the second IP address. The second IP address is not an output of the one-way function, and a malicious user could generate a spurious re-direction address that they control, but that is different to the first IP address. However, by including the first IP address within a packet destined to the second IP address, it is then possible to prevent packets being wrongly re-directed to a malicious user. When a request is received to re-direct a packet destined to the first IP address to another IP address, it can easily be checked whether the other IP address contained in the re-direction request is the same as the first IP address included in the packet. If the other IP address contained in the re-direction request is the same as the first IP address included in the packet then the packet is re-directed to that address, otherwise the re-direction request is ignored.

The first IP address may be encapsulated within the packet, for example as an inner packet.

A second aspect of the invention provides a method of controlling the re-direction of IP packets to an IP host having two or more different IP addresses, the method comprising the steps of: generating a first of said IP addresses as a one-way function of the second IP address; and transmitting the one-way function to a peer host or to a network routing node.

This method includes the steps that would be carried out at the host at which the first IP address is generated.

A third aspect of the invention provides a method of controlling the re-direction of IP packets to an IP host having two or more different IP addresses, the method comprising the steps of: receiving, at the peer host or network routing node, a request to re-route a packet destined to a first IP address to another IP address; applying a one-way function to the another IP address; and accepting the re-routing request only if the result of applying the one-way function to the another IP address is the first IP address.

This method includes the steps that would be carried out at the peer host or routing node. The one-way function applied is, of course, the same one-way function as used to generate the first IP address from the second IP address.

A fourth aspect of the invention provides a storage medium containing a program for controlling a processor to carry out a method of the first, second or third aspect.

Other aspects of the invention provide a communications network corresponding to the first aspect of the invention and apparatus corresponding to the second or third aspects of the invention.

Brief description of the drawings

Preferred embodiments of the present invention will now be described by way of illustrative example with reference to the accompanying figures in which:

Figure 1 is a schematic illustration of a communication network;

Figure 2 is a schematic illustration of a communication network employing the method of the invention.

Detailed description of the preferred embodiments

Figure 2 is a schematic illustration of a communications network employing the method of the present invention. The network corresponds to the network shown in figure 1, and contains a host 1 in a first network (network 1) that is communicating with a peer

host 2 in another network (network 2) over, for example, an IP backbone 3. Network 1 contains devices controlled by the host 1, and these are represented as devices 11,12, 13 in figure 1, and Network 2 contains devices controlled by the host 2, and these are represented as devices 21,22, 23 in figure 1. Host 1, and the devices 11,12,13 in network 1 each have two or more IP addresses. In the foregoing description it will be assumed that host 1 and the devices 11,12,13 in network 1 each have exactly two IP addresses, a first IP address and a second IP address, but the invention is not limited to this. It will also be assumed that the first IP address is an IPv6 address (“addr6”) and that the second IP address is an IPv4 address (“addr4”), but the invention is again not limited to this.

According to the present invention, the first IP address and the second IP address are bound to one another by a one-way function. It will be assumed, as an example, that the one-way function is a cryptographically secure hash function (or “cryptographic hash function”). That is, the first IP address ($IPad_1$) is determined from the second IP address ($IPad_2$) according to the general relation: $IPad_1 = \text{hash}(IPad_2)$. However, the one-way function is not limited to a cryptographic hash function and any cryptographically secure one-way function may be used.

The function $\text{hash}()$ is defined in such a way that it is very difficult to reverse the operation, i.e. obtain the value that the function has operated upon. The function $\text{hash}()$ can be public knowledge; its one-way functionality means that it does not need to be secret.

One form of an IPv6 address is a Hash-Based Address (HBA). HBAs are known, for example from [draft-ietf-shim6-hba] Marcelo Bagnulo, “Hash Based Addresses”. Internet Draft draft-ietf-shim6-hba-01.txt. Work In Progress, IETF, 2005. Simplifying a little, HBA forms a set of addresses bound to each other in a cryptographic manner:

$$\begin{aligned}
 \text{addr1} &= \text{prefix1} \mid \text{hash}(\text{modifier} \mid \text{prefix1} \dots \text{prefixn}) \\
 &\dots \\
 \text{addrn} &= \text{prefixn} \mid \text{hash}(\text{modifier} \mid \text{prefix1} \dots \text{prefixn})
 \end{aligned}
 \tag{1}$$

These address have the form stated above, with the “prefix1” being the routing prefix of the address and the “hash(modifier | prefix1 ... prefixn)” being the interface identifier suffix. The “prefix1”...“prefixn” terms in the hash term are the routing prefixes of the addresses. The modifier is provided in the hash term, as is known, to increase security by varying the term on which the hash function operates. The modifier may be, for example, a random number.

Addresses within this set are interchangeable and attackers cannot force other addresses to be used.

Another form of IPv6 address is a Cryptographically Generated Address (CGA). CGAs are described by Tuomas Aura in “Cryptographically Generated Addresses”. RFC 3972, IETF, March 2005. A CGA can be used to form an address from a public key generated in the node, and has the following form:

$$\text{Addr} = \text{prefix} | \text{hash}(\text{public key} | \text{modifier}) \quad (2)$$

These addresses have the form stated above, with the “prefix” being the routing prefix of the address and the “hash(public key | modifier)” being the interface identifier suffix.

A signature of the associated private key can secure any operation on this address.

However, a current limitation of these HBA or CGA IPv6 addresses is the lack of ability to employ them with IPv4 addresses. The invention broadens the application of cryptographically generated addresses.

Assume that packets are being sent from a device in network 2, for example device 21, to a host 1 of network 1, for example in response to a request made by a device in network 1. If the original request specified the first IP address of host 1 as the destination address for the packets, the packets arriving at host 2 from device 21 will

specify, as their destination address, the first IP address of host 1 (namely the address “addr6” in the example where the first IP address is an IPv6 address).

If host 1 wishes to switch to using the second IP address (namely to the address “addr4” in the example where the second IP address is an IPv4 address), it sends a re-redirectation request to the host 2 of the second network, specifying that packets destined for the first IP address addr6 should be re-directed to another IP address, namely to the second IP address addr4. As explained above, the first IP address addr6 was generated by applying a one-way function to the second IP address addr4.

When the host 2 of the second network receives a re-direction request which specifies that packets destined for the first IP address addr6 should be re-directed to another IP address, the host 2 determines whether the another IP address in the re-direction request is the same as the second IP address. If it is the same, the host 2 determines that the re-direction request has come from the first host and so re-directs packets in accordance with the request (so that packets arriving at the host 2 of the second network that are destined for the first IP address addr6 are re-directed to the second IP address addr4). If, however, the host 2 of the second network determines that the another IP address in the re-direction request is not the same as the second IP address, it determines that the re-direction request has not come from the first host and so ignores the re-direction request.

In order to check whether the another IP address contained in the re-direction request is the second IP address of the host 1 of the first network, the host 2 of the second network employs the hash function (or other one-way function) used to generate the first IP address addr6 of the host 1 from the second IP address addr4 of the host 1; the host 2 of the second network is aware of the hash function since the host 1 of the first network may have transmitted the hash function direct to the host 2 of the second network, or the host 1 of the first network may have made the hash function public knowledge. If the another IP address in the re-direction request is the second IP address addr4 of the host 1, the result of applying the hash function to the IP address specified in the re-direction request will be the first IP address addr6 of the host 1 (which is of course known to the

host 2). However, if the another IP address in the re-direction request is not the second IP address $addr_4$ of the host 1, but is a fake address $addr_4'$, the result of applying the hash function will not be the first IP address $addr_6$ of the host 1. Thus, if the result of applying the hash function to the other IP address contained in the re-direction request is the first IP address $addr_6$ of the host, this shows that the redirection request has indeed come from the first host. The host 2 of the second network is thus able to determine whether the re-direction request that it has received is genuine (ie, is from the host 1 of the first network). The second host ignores any re-direction request that it determines was not made by the first host.

In their simplest form, an IPv4 address (“ $addr_4$ ”) and an IPv6 address (“ $addr_6$ ”) can be bound to each other in the following form, which is similar to equation (1) above:

$$addr_6 = prefix | hash(modifier | prefix | addr_4) \quad (3)$$

As stated above, an IPv6 address consists generally of a prefix part and an interface identifier suffix. In equation (3), the interface identifier suffix of the IPv6 address is generated by a hash function on the IPv4 address. The “prefix” in the hash term is the routing prefix of the IPv6 address. The modifier is provided, as is known, to increase security by varying the term on which the hash function operates. The modifier may be, for example, a random number.

When the first address (in this example an IPv6 address) has been constructed in this manner, a multiaddressing protocol can be constructed as explained above so that it is not possible to redirect a packet destined to the first address ($addr_6$) to any other address except the second address ($addr_4$) used to generate the first address, because $addr_4$ is the only address which can be input to the hash to obtain $addr_6$ as the output.

However, forming the first IP address by a hash function on the second IP address does not solve the problem of authenticating a request to re-direct packets that are destined for the second IP address, since the second IP address ($addr_4$ in this example) is not an output from equation (3). An attacker can take any IPv4 address and any IPv6 prefix

and construct a fake IPv6 address (addr6') that the attacker controls but which is different from the true IPv6 address (addr6) of the host 1.

In a preferred embodiment of the invention, therefore, an additional requirement is placed on the multiaddressing protocol. This requirement is that redirection of packets destined to the second IP address (eg addr4) of the host 1 can only take place for a limited subset of traffic to and from the second IP address of the host. Specifically, only traffic that employs packets directed to or from the first IP address (addr6) of the host may be re-directed in this embodiment.

It is expected that in future, most if not all communications will employ IPv6 addresses, even if a communication is sent via nodes that are only connected through an IPv4 network. In this embodiment, even though an IPv4 address is used for communications, a packet will contain the IPv6 address of the destination. Thus, a packet 7 has an outer packet 5 with the second IP address (in this example an IPv4 address) and an inner, encapsulated packet 6 containing the first IP address (in this example an IPv6 address) and packet data fields. If a re-direction request is received at host 2 of the second network asking that packets destined to the IPv4 address addr4 of the host 1 are re-directed to an IPv6 address, the node 2 of the second network compares the IPv6 address in the re-direction request with the IPv6 address addr6 in the inner packet 6. If the IPv6 address in the re-direction request is the same as the IPv6 address in the inner packet 6, the packet 7 is re-directed from the IPv4 address in the outer packet 5 to the IPv6 address in the inner packet 6. If, however, the IPv6 address in the re-direction request is not the same as the IPv6 address in the inner packet 6, the host 2 ignores the re-direction request. Thus, in this embodiment, a request to redirect packets from the first IP address to the second IP address and a request to redirect packets from the second IP address to the first IP address can both be authenticated. The provision of the first IP address in the inner packet means that it is no longer possible for an attacker to redirect packets originating in the second network 2 and destined to the second IP address of host 1 to an IP address controlled by the attacker. Only packets that are destined for the IP address controlled by the attacker can be re-directed to the IP address controlled by the attacker.

The provision of the first IP address in the inner packet also makes it possible for TCP sessions to survive address changes; TCP requires the same IP addresses to be used.

The invention is not limited to including the first IP address in a packet by encapsulating the first IP address in an inner packet in the manner described above. The first IP address may be included in the packet 7 in any suitable manner.

This basic scheme with two addresses can be extended in various ways. In the following we discuss these extensions.

For NAT traversal, the owner of the first IP address (addr6) can employ both internal and external IP addresses as well as port numbers for UDP encapsulation. This makes it possible to securely bind the “real” IPv6 addresses used for the actual payload packets and the internal and external addresses. This overcomes one of the problems of Teredo, the need to rely on probing to determine address equivalence.

In the embodiment described above the IPv6 address and the IPv4 address are linked by equation (3), so that the IPv6 address is generally of an HBA form. For dynamic environments, CGA instead of HBA can be employed, so that the IPv6 address (or, in general terms, the first IP address) is generally of a CGA form. This allows multiple IPv6 addresses to be created and secured through signatures rather than the static HBA relationship. Similarly, IPv4 addresses can either appear as static input to the hash or in a signed statement from the address owner.

Generalizing, we can define Cryptographically Generated General Addresses (CGGAs) as follows:

A CGGA is an address set that satisfies equation (4):

$$\begin{aligned}
 \text{addr6_1} &= \text{prefix1} \mid \text{hash}(\text{material} \mid \text{prefix1} \dots \text{prefixn} \mid \text{id1} \dots \text{idn}) \\
 &\dots \\
 \text{addr6_n} &= \text{prefixn} \mid \text{hash}(\text{material} \mid \text{prefix1} \dots \text{prefixn} \mid \text{id1} \dots \text{idn})
 \end{aligned}
 \tag{4}$$

In equation (4), $\text{addr6_1} \dots \text{addr6_n}$ are the *core set* of IPv6 addresses (or, more generally the core set of first IP addresses), $\text{prefix1} \dots \text{prefixn}$ are the set of prefixes in that set, and *material* represents public keys (possibly several), random modifiers and other parameters. $\text{id1} \dots \text{idn}$ are the other identifiers bound to the core set. Examples of other identifiers include, for example, global and local IPv4 addresses, fixed IPv6 addresses that the creator of the set is unable to change (e.g. due to allocation from DHCPv6), or address and port number of a NAT device. Thus, in the example of figure 2, the IPv6 address of the host 1, would be generated by an equation of the form of equation (4) with the IPv4 address of the host 1 (addr4) used as one of the $\text{id1} \dots \text{idn}$ terms in the hash.

CGGAs provide assurances of address ownership, membership, and relationship. These assurances come in several variants:

- a) Given the input parameters, anyone can verify that the core set is correct. This can be used in multiaddressing operations that switch communications within the core set.
- b) Given the other identity parameters, anyone can verify that the core set is related to the other identities. This can be used in multiaddressing operations that switch communications from the core set to the other identifiers or switch communications from the core set to the other identifiers and the core set (e.g. UDP NAT traversal via a specific NAT device, carrying IPv6 traffic inside).

Similarly, communications may employ both a core set identifier and another identifier, such as when IPv6 packets are encapsulated inside IPv4 packets. Such communications can be switched from the other identifier back to using solely a core set identifier.

c) Given one or more signatures from the private keys associated with the input public keys, the private key holders can show that they are owners (or one of them) of the core set. This can be used in multiaddressing operations that move communications from the core set to previously unknown addresses.

The method of the invention has been described above as being performed in an IP host 2. The invention may alternatively be performed in a network routing node.

The method of the present invention may be implemented by a suitably programmed processor. A program for controlling a processor to perform the method of the invention may be stored on any suitable storage medium such as, for example, a computer disk, a magnetic disk or an optical disk.

CLAIMS:

1. A method of controlling the re-direction of IP packets to an IP host having two or more different IP addresses, the method comprising the steps of:
 - generating a first of said IP addresses as a one-way function of the second IP address; and
 - at a peer host or network routing node, accepting a request to re-direct a packet destined to said first IP address to another IP address only if the other IP address is the second IP address.
2. A method as claimed in claim 1 and comprising the further step of, at the peer host or network routing node, applying the one-way function to the another address contained in the re-direction request.
3. A method as claimed 1 and further comprising the steps including the first IP address in a packet destined for the second IP address, and, at a peer host or network routing node, accepting a request to re-route a packet destined to said second IP address to another IP address only if the other IP address is the first IP address included in the packet.
4. A method as claimed in claim 1, 2 or 3 wherein the first address is an IPv6 address.
5. A method as claimed in claim 1, 2, 3 or 4 wherein the second address is an IPv4 address.
6. A method as claimed in any preceding claim wherein the one-way function is a hash function.
7. A method as claimed in claim 6 when dependent on claim 4 wherein the step of generating the first of said IP addresses comprises generating the first of said IP addresses according to:

$$\text{addr6} = \text{prefix} \mid \text{hash}(\text{modifier} \mid \text{prefix} \mid \text{addr4})$$

where addr6 is an IPv6 address and addr is an IPv4 address.

8. A method as claimed in claim 6, wherein the step of generating the first of said IP addresses comprises generating the first of said IP addresses (IPad₁) according to:

$$\text{IPad}_1 = \text{prefix1} \mid \text{hash}(\text{material} \mid \text{prefix1} \dots \text{prefixn} \mid \text{id1} \dots \text{idn}),$$

where one of id1...idn is the second address.

9. A method of controlling the re-direction of IP packets to an IP host having two or more different IP addresses, the method comprising the steps of:

generating a first of said IP addresses as a one-way function of the second IP address; and

transmitting the one-way function to a peer host or to a network routing node.

10. A method of controlling the re-direction of IP packets to an IP host having two or more different IP addresses, the method comprising the steps of:

receiving, at the peer host or network routing node, a request to re-route a packet destined to a first IP address to another IP address;

applying a one-way function to the another IP address; and

accepting the re-routing request only if the result of applying the one-way function to the another IP address is the first IP address.

11. A storage medium containing a program for controlling a processor to carry out a method as defined in any one of claims 1 to 10.

12. A communication network comprising:

an IP host having two or more different IP addresses, a first of said IP addresses having been generated as a one-way function of the second IP address; and

means for accepting, at a peer host or network routing node, a request to re-direct a packet destined to said first IP address to another IP address only if the other IP address is the second IP address.

13. An apparatus for controlling the re-direction of IP packets to an IP host having two or more different IP addresses, the apparatus comprising:

means for generating a first of said IP addresses as a one-way function of the second IP address; and

means for transmitting the one-way function to a peer host or to a network routing node.

14. An apparatus for controlling the re-direction of IP packets to an IP host having two or more different IP addresses, the apparatus comprising:

means for receiving a request to re-route a packet destined to a first IP address to another IP address;

means for applying a one-way function to the another IP address; and

means for accepting the re-routing request only if the result of applying the one-way function to the another IP address is the first IP address.

15. An apparatus as claimed in claim 14 and further comprising:

means for receiving a request to re-route a packet destined to said second IP address to another IP address; and

means for determining whether the another IP address is included in the packet;




Figure 1




Figure 2

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2007/050931

A. CLASSIFICATION OF SUBJECT MATTER

INV. H04L29/06
ADD. H04L29/12

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, COMPENDEX, INSPEC, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	BAGNULO UC3M M: "Hash Based Addresses (HBA)" IETF STANDARD-WORKING-DRAFT, INTERNET ENGINEERING TASK FORCE, IETF, CH, vol. shim6, no. 1, 23 October 2005 (2005-10-23), XP015042675 ISSN: 0000-0004 cited in the application abstract page 3, line 1 - page 5, last line page 9, line 11 - page 13, line 28 -----	1-17
A	GB 2 367 986 A (ERICSSON TELEFON AB L M [SE]) 17 April 2002 (2002-04-17) abstract page 4, line 16 - page 5, last line page 9, line 19 - page 12, line 22; claim 1 -----	1-17

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *&* document member of the same patent family

Date of the actual completion of the international search

8 June 2007

Date of mailing of the international search report

18/06/2007

Name and mailing address of the ISA/
European Patent Office, P.B. 5618 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Lebas, Yves

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2007/050931

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
GB 2367986	A	17-04-2002	AU 2002232348 A1	03-10-2002
			DE 10296445 T5	29-04-2004
			JP 2004533741 T	04-11-2004
			WO 02076060 A2	26-09-2002
			US 2002133607 A1	19-09-2002
