SITE MULTIHOMING AND PROVIDER-INDEPENDENT ADDRESSING USING IPV6

Dirk Henrici, David Prantl, Paul Müller
University of Kaiserslautern, Integrated Communication Systems
67663 Kaiserslautern
Germany
{henrici, prantl, pmueller}@informatik.uni-kl.de

ABSTRACT

Using IPv6, multihoming and Internet service provider migration are still not satisfactorily solved problems. This leads to delay in the adaptation of the new protocol version. This contribution aims to address both of the two stated problems while retaining the advantages of strictly hierarchical addressing and routing.

The solution presented in this paper consists of two building blocks: So called "Unique Local Addresses" that are intended to be used instead of the deprecated IPv6 site local addresses can be employed as globally valid, provider-independent identifiers. Using address mapping at site exit routers, a feature-rich multihoming solution can be created without breaking the end-to-end model.

The proposed solution has many advantages: It is simple and compatible to current Internet standards. No changes at all are required at hosts, and the solution is designed to keep network management easy.

KEY WORDS

Internet protocols, communication protocols, multi-homing, Internet architecture, communication systems

1. Current Situation and Motivation

Although the current IPv6 standard [1] was published in 1998, practical adaptation still proceeds slowly excepting in Asia where IP address space is scarce and thus the demand for IPv6 introduction is high. Besides the "never touch a running system" maxim, important reasons for the slow adaptation are two still unsolved problems: Network renumbering is required when an Internet service provider (ISP) is changed, and a working multihoming solution is still missing in IPv6.

Both problems result from hierarchical routing that became mandatory for the IPv6 [2] to solve scalability problems. Using hierarchical routing, address assignment takes place strictly hierarchically: Tier-1 providers are assigned a certain address space and assign parts of it to Tier-2 providers that are their customers. The Tier-2 providers assign parts of their address space to their respective customers and so on. At the last level, companies or private persons get address space from their respective ISP. In contrast to that, in IPv4, the address assignment is provider-independent which led to dramatic growth of the routing tables in the core Internet, the so-called "default free zone". Using hierarchical routing, the experienced scalability problem can be solved: Routes can be aggregated efficiently thus keeping routing tables small [3]. This is the reason why hierarchical routing was demanded for in IPv6.

But hierarchical routing also has negative issues: A company setting up a corporate network must use the address space assigned by its provider. This binds a company to its ISP because if the provider changes, the company needs to renumber its whole network [4]. Such a renumbering is usually a complicated task although it is facilitated by auto-configuration features of IPv6 [5, 6]. The reason is that network addresses are often set explicitly in configuration files of servers, in access control lists of network equipment, in firewall rules, and in many more places. Because of the necessary renumbering, changing an ISP can become extremely difficult and costly since renumbering is a tedious and error-prone process [7]. Because nobody wants to be dependent on a particular ISP, companies are forced to stick to IPv4 with which they can have provider-independent address space.

Nowadays, companies of all sizes rely on Internet connectivity for doing their business. Because of that they might have more than one ISP to get a redundant connection to the Internet. This procedure is called "multihoming" or in this context more precisely named "site multihoming" in contrast to other forms of multihoming. As long as companies have provider-independent addresses they can route their data over the networks of both ISPs as need be. The ISP can be selected accounting for cost, shortest route, fastest route or other

criteria. But using IPv6 where no provider-independent addresses exist because of the strictly hierarchical routing, such a multihoming scenario is no longer that simple.

Using IPv6, each host would be assigned more than one IP-address: one from each ISP. Switching between ISPs is then no longer transparent for hosts: They need to check the usability of a certain address by themselves. How this can be done is not yet standardized. Also, switching between ISPs is no longer possible without dropping active TCP connections. Network administration also becomes more sophisticated because different sets of addresses need to be set in access control lists, firewall rules etc. In consequence, the current multihoming possibilities in IPv6 are all but satisfactory.

In this contribution, we present a viable solution for the IPv6 scenarios. The main goal is to render renumbering on ISP change unnecessary and to create a site multihoming solution with a feature set of the multihoming solutions offered by provider-independent addresses like in IPv4 [8]. This has to take place without affecting hierarchical routing since it is needed for ensuring future scalability of the Internet. The presented solution is compatible with existing Internet standards and no changes in hosts are required at all. The requirements stated in RFC 3582 [9] and in the more comprehensive multi6 working group Internet draft [10] are taken account of. The solution can be introduced without affecting the end-to-end model of networking which is relevant for P2P-applications and VoIP-protocols like SIP.

The paper is organized as follows: In chapter 2, relevant other works that have similar objectives are referenced and associated with our proposal. Chapter 3 presents our solution in two steps: Firstly, the introduction of an identifier address space, and secondly, using site exit routers for switching between identifiers and locators. In chapter 4, the effects of the presented solution are shown and relevant topics like security issues are discussed. The paper concludes with a summary in chapter 5.

2. Related Work

A variety of proposals have been made to solve the multihoming issues in IPv6, see for instance [11]. The problem was also discussed in the IETF working group "multi6" [12]. Within this group a document describing the goals of site-multihoming architectures was worked out, see [9]. Also a document analyzing the security threats that multihoming solutions face was created [13]. As the number of proposals is very high, the working group released an attempt to categorize solutions [14] to get a better overview. Currently, the IETF is working on an approach called "L3 shim" [15].

The solution that is presented in this paper takes up the goals stated by the IETF. It contains elements of the proposals "16+16" or "GSE/8+8", respectively, [16] and MHAP [17]. Also, similarities exist with the on-demand-tunneling idea of van Beijnum [18]. Note that the latter is a host based approach, not a site multihoming solution like the proposal presented in this paper.

3. Solution Proposal

The solution which will be presented in the following consists of two parts. At first, a new addressing space that is already proposed as IETF draft [19] is used to solve the renumbering problem. In a second step, this new address space is mapped to the address space of globally routable addresses at site exit routers.

3.1 First Step: "Unique Local Addresses"

To avoid the necessity to renumber the corporate network in case of an ISP change, provider-independent addresses are required. Since there is no equivalent to the IPv4 non hierarchical addresses, in IPv6 another solution needs to be implemented. Besides having provider-independent addresses, nowadays it is usual to number a local network with private addresses [20] and to map these to one or several public addresses by the site exit router(s) using network address translation (NAT). The form of NAT used for this purpose most frequently is "port address translation" (PAT), also known as IP-masquerading. PAT is deprecated in IPv6 because it affects end-to-end reachability negatively which is relevant for applications like IP-telephony or P2P-applications.

The idea is to use an equivalent for IPv4 private addresses for numbering local corporate networks and thereby avoiding the problems caused by PAT [21]. An equivalent for IPv4 private addresses was standardized with IPv6 in RFC 3513 [22]. These private addresses were deprecated in RFC 3879 [23], but a proposal for a successor is already available as Internet draft [19]. This draft introduces the so-called "Unique Local Addresses", in the following abbreviated "UL-addresses". These UL-addresses are globally unique but not routable in the Internet.

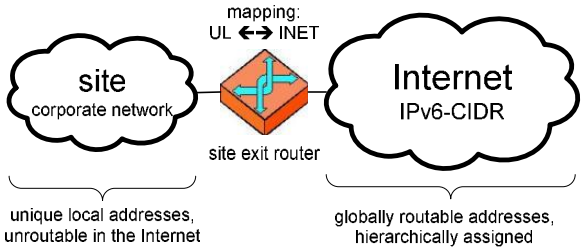

Figure 1: Use of UL-addresses in the corporate network

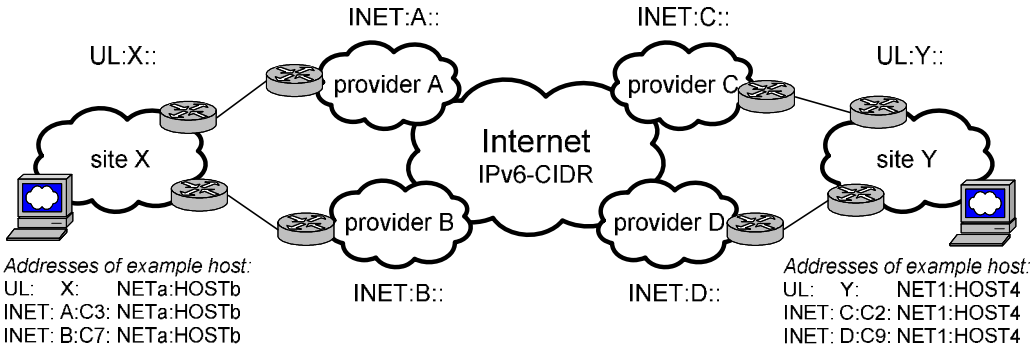

Figure 2: Example multihoming configuration

The UL-addresses can thus be used to number a corporate network globally uniquely. In the draft, UL-addresses are meant to coexist with globally routable addresses in the local networks, but we propose to use solely UL-addresses in the corporate network. The latter has the advantage that firewall rules and access control lists only need to be maintained for a single address range and thus become simpler and less error prone. Additionally, they do not need to be changed on ISP change any more.

The only locally valid UL-addresses are then mapped into standard globally routable addresses (in the following abbreviated "INET-addresses") at site exit routers. The network prefix of the addresses are simply exchanged, thus a simple 1:1-mapping is done. Such an operation is called "network mapping" or, in Cisco jargon, "over-lapping". Beneficially, it does not affect end-to-end-reachability in contrast to more complex forms of network address translation like PAT. But note that the proposed mapping alone would cause problems in protocols like SIP that include addresses in their payload. Because of that the following second step is designed in such a way so that this issue gets solved.

3.2 Second Step: "Identifiers vs. Locators"

IP-addresses are used for many purposes in the Internet [24]: When considering multihoming, they act as network interface identifiers, i.e. identify a particular destination, and also as locators, i.e. provide information on how to reach an interface in the network using routing. IP-addresses doing these two different things also lead to problems for other applications, for instance IP mobility. A split of identifiers and locators is therefore being discussed for multihoming by the IETF multi6 working group as well [24].

Seen from (and only from) the scope of the Internet, we propose to use the introduced UL-addresses as identifiers and the standard globally routable INET-addresses as locators. This approach solves the multihoming issues because hosts can use the static identifiers (i.e. the UL-addresses) for identification of a destination peer whereas the data can be routed in the Internet over one or several ISPs using the locators (i.e. the INET-addresses). In contrast to this identifier-locator-spit, within a site the UL-addresses are used as identifiers and locators in the same way INET-addresses are used nowadays.

One prerequisite for using UL-addresses as identifiers is to announce them using the domain name system (DNS). Normal IPv6-INET-addresses are stored as "AAAA" resource records in the DNS. The UL-addresses can simply be added as additional AAAA resource records. This has no negative implications for hosts that do not use or know UL-addresses because hosts will always use the right kind of address as long as the hosts follow the standards stated in [25]: It is specified that the address to be used is selected from the ones returned from a DNS query using a longest prefix match with the address of the host's network interface if no additional rules apply.

In a local network that adapts the solution proposed in this paper, only UL-addresses are used (see previous section 3.1). Thus, hosts within the local network use UL-addresses to communicate with each other. As explained above, since the interface addresses of the hosts are UL-addresses the host automatically use the UL-address-records from the DNS and ignore the INET-address-records for local communication.

If a host wants to communicate with a host at another site we need to distinguish between two cases: whether the remote site uses the proposed solution, too, or not. In case of the latter, a DNS query for the address of the remote hosts will only return INET-addresses and communication takes place as usual with an INET-address as destination address but with an UL-address as source address which is mapped to an INET-address at the site exit router. In case that the remote site also uses the proposed solution, the host in the local network will select the UL-address of the remote host for communication. The host in the local network will thus use his own UL-address as source address and the UL-address of the remote host as destination address. As the UL-addresses are not routable in the Internet they are mapped into corresponding routable INET-addresses at the site exit router. When the data packets reach the site exit router of the remote site, the mapping is undone and thus the UL-addresses are

restored. If (in case of multihoming) multiple INET-addresses belong to a single UL-address the site exit router of the local network can select which INET-address of the remote host and thus which ISPs to use. If a failure is detected the ISPs used can be changed by the site exit router without affecting existing TCP connections. Obviously, explicit host support is not required for this to work because rehoming is transparent to the local network since it is done solely by the site exit router(s).

3.3 Example

In figure 2 an example is given in which local and remote site both employ the proposed solution. The local network of both sites is then numbered with UL-addresses: The UL-addresses begin with a prefix indicating that the address is an UL-address. The second part ("X" or "Y", respectively, in the example) is assigned to a site by the number registry. The last part of the address is used for addressing sub networks and hosts within the site.

In the Internet, the UL-addresses are not routable. Only normal INET-addresses are valid there. Each site gets address space by its ISP in hierarchical manner. The INET-addresses start with a prefix indicating that the address is an INET-address. The second part is the address space prefix of the ISP from which a part is assigned to the particular site ("Cx" in the figure). The rest of the address is equal to the suffix of the corresponding UL-addresses.

All of the addresses are stored in the DNS as AAAA resource records. The site exit routers map between the UL-addresses and the INET-addresses. This way, the Internet with its INET-addresses works as a kind of tunnel between the sites in which UL-addresses are deployed.

4. Discussion of the Proposed Solution

In this section the implications of the proposed solution shall be discussed and important points to be considered in the implementation highlighted. Note that we will focus on some of the most important topics here: The compatibility to the existing Internet and existing Internet standards, security considerations, and performance issues.

4.1 Compatibility and Transition

A big advantage of the proposed solution is that it only requires UL-addresses to be announced using the existing domain name system and extensions in site exit routers. It neither requires a change in any existing Internet standard nor does it affect routing nor does it require any modifications at hosts. These three features constitute a big advantage compared to most other proposals since complete compatibility to the existing Internet is given.

The proposed solution is backward compatible in the sense that communication with sites not using the approach is still possible: If no UL-addresses exist, the INET-addresses are used. Since in that case no double network mapping takes place, the mapping between INET-addresses and UL-addresses is not completely transparent for hosts. This causes problems for protocols like SIP that encode addresses in their protocol messages. But that problem is much easier to solve than the NAT-issues that are common in today's IPv4 because the network mapping in our approach is a simple 1:1 mapping and not a one-to-many mapping. Besides a solution in form of the usual application level gateways that are used with today's NATs, the protocol stack at the host or alternatively a mapping gateway can solve the issue by providing the second mapping.

In the case that both communication peers employ the proposed scheme, no mapping (in case of local or VPN communication) or a double mapping of addresses (communication between two sites; one mapping at each site exit router) takes place. Thus, if a mapping occurs it is transparent to hosts and transport layer protocols. Thus no steps like application level gateways and so on are required in this scenario.

Because of the backward compatibility of our solution a gradual transition is possible and thus there is no need for a "flag day". This makes transition an easy task. The double mapping at the site exit routers makes the mapping transparent for hosts and protocols so that the end-to-end model is not broken and protocols like IPsec or SIP work without any changes.

4.2 Security

Implementation of an IPv6 multihoming solution must not introduce new threats to the Internet. Solutions using address rewriting like the one proposed need to pay particular attention to so-called "redirection attacks" [14]. Another important issue are denial-of-service (DoS) attacks in case reservation of resources or considerable processing takes place.

The site exit routers need to obtain the mapping between the used INET-addresses and the corresponding UL-addresses. If the mapping is obtained from the DNS on packet reception, which would be the straightforward implementation, particularly the site exit router of the destination host would be vulnerable to DoS-attacks: For each packet reaching the site exit router with yet unknown mapping between INET-addresses and corresponding UL-address it would need to issue a DNS-query. This would be an attack that is as simple but more effective than TCP-SYN-flooding. To carry all the required information about the mappings in the packets themselves using IPv6-extension-headers is no solution because an attacker could provide a faked mapping to pass access control devices

(e.g. firewalls) that rely on the validity of the UL-address. A viable solution is ensuring that the sender's address is not faked and forcing the party sending the initial packet to have more effort than the destination site exit router for getting the needed mapping. This can be achieved by implementing a four-way-handshake between the two involved site exit routers similar like the one done in SCTP [26]. The idea is to counteract attacks like TCP-SYN-flooding by not allocating resources until the sender has been proven valid. Similarly, we would not issue the DNS-query before the sender has been proven valid. Additionally, the destination site exit router can give the sending one a challenge [27] that requires more effort than getting the needed trusted address mapping from the DNS before starting the DNS-query. Thus, an additional handshake sequence is required between the site exit routers prevents redirection and spoofing attacks in case the needed address mappings are yet unknown.

In consequence, careful consideration is required to make the approach secure but implementation is possible without introducing new security threats to the Internet. Techniques like ingress filtering [28] that are considered useful in today's Internet are further on applicable and make still sense.

4.3 Performance

After obtaining the needed address mappings by the site exit routers communication is very efficient. There exists no additional overhead like in tunnelling solutions [18]. The involved network mapping is an efficient operation since only network prefixes need to be exchanged. This is much less effort than is required in the today widespread forms of NAT (PAT or IP-masquerading, respectively).

If the needed address mappings are not available at the involved site exit routers the required mappings need to be obtained from the DNS and the described handshake-sequence needs to take place. This leads to a short delay before the actual data packets can start to flow. But this is not too much of a problem because the delay only occurs at the first time of communication with a given destination and the delay is shorter than the still tolerable ones occurring using on-demand dial-in, for instance DSL-routers establishing a connection on demand.

Additionally, the data flow can already start together with the third message of the handshake-sequence, so that the delay introduced by the handshake sequence is reduced to a bit more than a single round-trip-time between the involved site exit routers. If the sending host and the site exit router on the sending site are supplied by the same DNS-server the network mapping needed by the site exit router will be still in the DNS-cache of the DNS-server so that the delay for obtaining this mapping will be very short (approx. one roundtrip time to the local DNS-server).

In consequence, the delay introduced for obtaining the needed mapping between network prefixes is not high enough to be a problem. After obtaining the mappings, communication is more efficient than is usual today since network mapping is much simpler than PAT.

The presented solution is scalable because it does not affect the hierarchical routing scheme that became mandatory in IPv6 to provide scalability. The solution does not presume extension of the existing Internet infrastructure because it only relies on the existing DNS. Nevertheless, the solution is that scalable to fit the demand of single dial-in-computers as far as the demand of large companies.

4.4 Extensions

There are some extensions conceivable that further improve the proposed solution. If host support for the solution would be introduced the required network mapping for the site exit router on the sending side can be provided by the host (e.g. by an extension header in the first packet) because it anyway needs to perform a DNS query to resolve the hostname. This further reduces initial communication setup time a bit.

More interestingly, host support would offer the possibility for host triggered multihoming. In contrast to the stated scenario in which the site exit router chooses the optimal ISP in case several ones are available, there are applications for which it would be advantageous if a host could specify the ISP to be used. For instance, a user of a VoIP-application could specify to always use the ISP with better delay times although it is more expensive. This can be realized using IPv6-extension headers in which a host/ an application can specify which ISP to prefer. Of course, if such a feature is not wanted (untrusted hosts etc.) an administrative policy in the site exit router could prevent host-based provider selection.

But note that host support is not required at all for the proposed solution to work. It would just lead to additional features and a shorter time for connection establishment.

5. Summary

There are two integral problems still impeding the widespread use of the IPv6: site renumbering and site multihoming. In this paper a viable solution for these two problems was presented. It is based on the use of Unique-Local-Addresses as replacement for the deprecated site local addresses [23]. The UL-addresses are not only used locally but as identifier for remote hosts as well. "Classical", hierarchically assigned and globally routable IPv6-addresses (INET-addresses) work as locators. Mapping between the two kinds of addresses is done at the site exit routers.

After motivation of the work and introducing the ideas the solution is based on, important related topics were discussed. It was clarified that the presented solution is fully backward compatible to the existing Internet and existing Internet standards and that it can be implemented without creating new security problems. As major advantage of the approach it was shown, that introduction of the scheme only required insertion of UL-addresses into the DNS-system and modification of the site exit routers. No changes at all at hosts are required, and routing in the Internet is not affected. The principle of assigning INET-addresses strictly hierarchically is not touched. As the DNS is employed, no additional infrastructure needs to be built. Interestingly, although a mapping solution, the scheme does not break the end-to-end model of networking.

With the proposed solution the need to renumber local networks on ISP change is eliminated, and network management of IPv6 networks becomes easier since only UL-addresses need to be considered. For building VPNs to other companies, the provider-independent UL-addresses are also well suited. The possibility for multihoming with the ability to switch between ISPs without affecting existing connections is a major advantage of the solution. Host extensions are not required but can be introduced to enable even host-based ISP-selection.

References:

[1] Deering, S. et al.: RFC 2460 - Internet Protocol, Version 6 (IPv6) Specification, *Network Working Group, 1998*

[2] Hinden, R. et al.: RFC 2374 - An IPv6 Aggregatable Global Unicast Address Format, *Network Working Group, 1998*

[3] Rekhter, Y.; Li, T. (eds.): RFC 1518, An Architecture for IP Address Allocation with CIDR, *Network Working Group, 1993*

[4] Ferguson, P.; Berkowitz H.: RFC 2071, Network Renumbering Overview: Why would I want it and what is it anyway?, *Network Working Group, 1997*

[5] Thomson, S. et al: RFC 2462, IPv6 Stateless Address Autoconfiguration, *Network Working Group, 1998*

[6] Droms, R. (ed.): RFC 3315, Dynamic Host Configuration Protocol for IPv6 (DHCPv6), *Network Working Group, 2003*

[7] Carpenter, B.; Rekhter, Y.: RFC 1900, Renumbering Needs Work, *Network Working Group, 1996*

[8] Abley, J. et al.: IPv4 Multihoming Practices and Limitations, work-in-progress: draft-ietf-multi6-v4-multihoming-03, *IETF multi6 working group, 2005*

[9] Abley, J. et al.: RFC 3582 - Goals for IPv6 Site-Multihoming Architectures, *Network Working Group, 2003*

[10] Lear, E.: Things MULTI6 Developers should think about, work-in-progress: draft-ietf-multi6-things-to-think-about-01, *IETF multi6 working group, 2005*

[11] Dumore, M. (ed.): Report on IETF Multihoming Solutions,v2, *6net-project (IST-2001-32603), 2003*

[12] IETF multi6-Working Group: Site Multihoming in IPv6 (multi6), Charter, *see http://www.ietf.org/html.charters/multi6-charter.html, 2005*

[13] Nordmark, E.; Li, T.: Threats relating to IPv6 multihoming solutions, work-in-progress: draft-ietf-multi6-multihoming-threats-03.txt, *IETF multi6 working group, 2005*

[14] Huston, G.: Architectural Approaches to Multi-Homing for IPv6, work-in-progress: draft-ietf-multi6-architecture-04.txt, *IETF multi6 working group, 2005*

[15] Nordmark, E.; Bagnulo, M.: Multihoming L3 Shim Approach, work-in-progress: draft-ietf-multi6-l3shim-00.txt, *IETF multi6 working group, 2005*

[16] O'Dell, M.: GSE - An Alternate Addressing Architecture for IPv6, *see website: http://arneill-py.sacramento.ca.us/ipv6mh/draft-ipng-gseaddr-00.txt, 1997*

[17] Py, M.: Multi Homing Aliasing Protocol (MHAP) intro, *see website: http://arneill-py.sacramento.ca.us/ipv6mh/draft-py-mhap-intro-00.txt, 2003*

[18] van Beijnum, I.: On Demand Tunneling For Multihoming, *see website: http://www.muada.com/drafts/draft-van-beijnum-multi6-odt-00.txt, 2004*

[19] Hinden, R.; Haberman, B.: Unique Local IPv6 Unicast Addresses, work-in-progress: draft-ietf-ipv6-unique-local-addr-09.txt, *IETF IPv6 working group, 2005*

[20] Rekhter, Y. et al: RFC 1918, Address Allocation for Private Internets, *Network Working Group, 1996*

[21] Hain, T: RFC 2993, Architectural Implications of NAT, *Network Working Group, 2003*

[22] Hinden, R.; Deering, S.: RFC 3513 - Internet Protocol Version 6 (IPv6) Addressing Architecture, *Network Working Group, 2003*

[23] Huitema, C.; Carpenter, B.: RFC 3879 - Deprecating Site Local Addresses, *Network Working Group, 2004*

[24] Nordmark, E.: Multi6 Application Referral Issues, work-in-progress: draft-ietf-multi6-app-refer-00.txt, *IETF multi6 working group, 2005*

[25] Draves, R.: RFC 3484, Default Address Selection for Internet Protocol version 6 (IPv6), *Network Working Group, 2003*

[26] Stewart, R. et al.: RFC 2960 - Stream Control Transmission Protocol, *Network Working Group, 2000*

[27] Aura, T. et al.: DOS-resistant Authentication with Client Puzzles; *Lecture Notes In Computer Science, vol. 2133, revised Papers from the 8th International Workshop on Security Protocols, pp. 170-177, 2000*

[28] Ferguson, P; Senie, D.: RFC 2827, Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing, *Network Working Group, 2000*