

# Security and Privacy with IPv6

Latif Ladid  
*President, IPv6 Forum*  
*Chair, European IPv6 Task Force*  
*Trustee Emeritus, Internet Society*  
*Luxembourg*  
*latif.ladid@ipv6forum.com*

Jimmy McGibney, John Ronan and Mícheál Ó Foghlú  
*Telecommunications Software & Systems Group*  
*Waterford Institute of Technology*  
*Cork Road, Waterford, Ireland*  
*{jmcgibney, jronan, mofoghlu}@tssg.org*

## 1 Introduction

The Internet today provides generic communication infrastructure for packet-based communications. Several edge networks that carry both business and non-business oriented traffic communicate with each other via this public infrastructure. This infrastructure is based on an agreed suite of protocols, generally denoted TCP/IP [RFC 791, RFC 793] in reference to the two most significant of these protocols.

Currently, version 4 of the Internet Protocol (known as IPv4) is the *de facto* standard for Internet connectivity. Internet standards emerged from research work in the United States, sponsored by DARPA, a US defense research agency. TCP/IP became popular in universities and was incorporated into the UNIX family of operating systems used in university computer science faculties. Over time TCP/IP became a world-wide network linking universities and research bodies, and became the foundation for the web.

With IPv4, there have been a variety of exploits on end and intermediate systems due to protocol design as well as implementation problems resulting in substantial loss of revenues. IPv4 has therefore been supplemented by the IPsec protocols to provide for security needs at the network layer. The new version of the Internet Protocol, IPv6 [RFC 2460], by contrast, has mandated support for IPsec as part of its basic design.

IPsec helps serve the data privacy and integrity needs of the data in transit across the Internet in addition to providing authenticity of the data's origin. Traditionally, the term security addresses requirements of privacy, authentication, integrity and availability of data and services. At the network layer, IPsec provides for the first three of these needs. Consequently, IPv6 provides for these requirements too.

This paper is structured as follows. Section 2 presents an argument for the deployment of IPv6 as the key enabler for restoration of the end-to-end model and how it would, most likely, increase demand for new, innovative applications and services. Section 3 then provides a brief overview of the current position in the area of network security, so that IPv6 security issues can be understood in this context. The following section describes IPsec along with an overview of the current state of the IP protocol and the advantages IPv6 has to offer over IPv4, from a security perspective. Finally, we present a brief discussion of security aspects of IPv4-IPv6 coexistence, conclude in section 6.

## 2 IPv6, Mobility and the End-to-End Argument

At its core, the case for IPv6 is based on the available range of addresses for Internet devices in IPv4 (32 bit) versus IPv6 (128 bit). It is true that, using a NAT-enabled [RFC 2663, RFC 3022] infrastructure, IPv4 can continue to be deployed in Western Europe and the developed world for the near term. Here there are enough IPv4 addresses for Internet Service Providers to assign limited addresses to publicly accessible machines and hide the complexity of the real diversity of internal machines (as they are given a private [RFC 1918] Class A address range such as 10.x.y.z). However, it is already the case that IPv6 has become the Internet protocol of choice for Japan, Korea, China, India, and the Far East in general (and to a lesser extent for South America and Africa as well). Population pressures, and the political residue of the differential treatment of these regions, have prompted strong governmental and industrial support for IPv6. In the Far East the IPv6 address allocations are perceived as being more culturally neutral and essential to overcome the immediate addressing concerns today. This is because of the control the USA had over the allocation of IPv4 addresses, and the perception that effectively this allocation favored the western economies.

The argument so far has made no mention of mobile devices. When these are introduced to the equation, timescales are shortened further and a large address range is even more important, as there is the potential for multiple devices per person (rather than just one or two, a desktop and/or laptop). Recognition of this was a major factor in the selection of IPv6 as the protocol of choice by the 3GPP for the deployment of a world-wide UMTS 3G mobile telephone network. So there is no need to look further than this one simple argument: in a world wide mobile network, based on IP protocols, it makes sense to use IPv6 rather than IPv4.

The counter-argument to IPv6 is relatively simple: what does IPv6 give that forces users and providers into a migration to IPv6? The answer for many is nothing. Whilst there are more addresses, many can survive happily with existing IPv4 address ranges, especially in Europe and North America, and use NAT to create privately addressed subnets where needed. These counter arguments are as much about the financial justification for upgrading corporate infrastructures as they are about IPv6. Meanwhile, the more populous countries, where the issues that IPv6 solves are more pressing, may push ahead of Europe and North America in IPv6. If this prediction is correct, this means there is a business argument for creating products for deployment in developing countries, which are by definition large markets, then a push for IPv6 makes sense more generally.

There are other reasons why IPv6 is the obvious choice. With IPv4 it requires a considerable infrastructure to allow the assignment of addresses on the network: a router and a DHCP server are needed. In IPv6 this auto-configuration [RFC 2462] is part of the basic infrastructure, supported by routers themselves. It is close to zero management in that once the router knows the network prefix, the addresses for devices are assigned automatically. Furthermore these are predictable as they are derived from the MAC address of the Ethernet device. These are called Global Addresses [RFC 3587]. Even if no router is available, the device configures itself with a Link Local address, and may be able to tunnel over an IPv4 network. While such addresses may be useful for local traffic, effectively for clients, these addresses are not useful as a generic global identifier for devices, thus missing out on one of the key benefits of IPv6 (the use of the address itself as a unique global identifier of an endpoint). The trend has been in many areas of distributed computing for every node to be a server as well as a client (especially so in peer-to-peer networking) and this then raises the issue of uniquely identifying the server, potentially in a global context.

As mentioned, one positive potential for the use of IPv6 is the restoration of an end-to-end Internet where every endpoint has a unique address, thus allowing every machine to potentially offer a services to every other machine without the “problem” (from a service

access point of view) of firewalls and NAT. Throughout the history of the Internet there has been a debate about the potential benefits of end-to-end communication [Saltzer, 1981], [Reed, 1998], [Clark, 2002]. However, re-establishing this paradigm would raise huge security concerns. Currently many companies and ISPs use NAT as a way of hiding the real identity of endpoints inside their domains (changing all outgoing packets to a single publicly visible IPv4 address). NAT is used not only to solve the issue of the shortage of IPv4 addresses, but as a primitive security tool which it was not designed to be. Furthermore, it doesn't actually give the user their perceived level of security in many cases. Of course, many applications have now been engineered to tunnel over the protocols that are allowed through NAT gateways and firewalls (primarily web protocol http on port 80 of the remote server). The normal response of the IPv6 community is that more secure networking is possible with IPv6 as IPsec is mandated, thus providing the potential for a network-level securely encrypted session. Of course, this may not address the requirements of central control in companies and ISPs as to what kind of traffic is and is not allowed into and out of their networks, but this requirement can easily be dealt with by addition of suitable filters to border (and indeed internal) routers. This debate is important and as new solutions emerge, innovative applications and services need to be able to integrate with current practices. There may be an IPv6 world where it is equally difficult to get through corporate choke points of various kinds.

A very promising facility offered by IPv6 is the use of Mobile IPv6. When Mobile IPv6 is deployed, this framework can allow for the transfer of a session from one IPv6 endpoint to another relatively seamlessly. In contrast Mobile IPv4 is routed via the original home node and so can be very inefficient. Clearly it is important for this type of macro-mobility to allow users to move from one place to another (potentially administered by different authorities) and to continue to use a service without interruption. Of course the issues of coverage and of organizational relationships to allow such roaming are larger than the issue of being able to negotiate a new IPv6 address and continue with a service originally accessed from a different IPv6 address. As Internet applications have traditionally not included large elements of network management, there is much work to be done in establishing such interrelationships. It is more likely to be tackled in the 3G world (where potentially revenue can fund such research) rather than in the open access WiFi world where there is less incentive to creating roaming agreements, and less agreement as to who the operators are and what their responsibilities are.

### **3 Security Challenges**

As computing and computer-based communications gain an ever-increasing foothold in our lives, the need for security is paramount. The field of security incorporates concepts of authentication (including identification and trust), confidentiality, integrity, availability, access control and non-repudiation. The International Telecommunications Union [ITU, 2003] defines security services along these lines. In general, such security requirements are critical for enterprises that use the Internet or Internet-like infrastructure for their day-to-day business.

It is fair to say that a constant war is being waged between those who own and manage systems and those who wish to attack them. Access to the Internet is relatively easy and cheap, with users enjoying a high level of anonymity if desired. In many ways, those who wish to breach security have all the aces:

- The standards used for basic Internet protocols are public. This means that attackers know much more about how the Internet works than they would if a closed network were used.
- Even though the number and diversity of systems is increasing, as is their complexity, the level of technical sophistication required to carry out attacks is falling [Manikopoulos, 2002].

- Modern systems in general are very complex, operating at several layers. Complexity usually makes for bad security. The sheer number of ways in which modern systems can be used makes comprehensive testing an extremely difficult problem, and production systems almost inevitably have flaws.
- The speed of development of the Internet has been huge. This means that much of the software used was developed with its main functionality in mind, with less thought being given to the security aspects. The most secure systems are those that were designed with security in mind from the start. IPv4, for example, was not designed with security as a priority.
- The trend toward code mobility in the past decade has provided all sorts of opportunities for the development of viruses and worms, which are in effect autonomous agents that, once released, can reproduce. The resulting combinatorial explosion allows the attacker's wishes to be carried out on a very large scale.

These attacks manifest themselves as identity impersonation (referred to as spoofing), loss of privacy, loss of data integrity (e.g. credit card transaction details being modified in transit), communications monitoring, and denial-of-service. Such attacks are the result of discovering exploits that emerge from the flaws in the basic protocol design (e.g. WEP in IEEE 802.11b, also known as WiFi) or from the incorrect implementation of protocols, applications, and operating systems (e.g. not enforcing the use of strong encryption in a WiFi network, or the selection of a weak cipher on an encrypted communications link). Exploiting such discoveries will remain as long as protocol implementations do.

Defenses available to those charged with managing computer systems include the strict enforcement of a comprehensive security policy (the importance of having a policy and enforcing it cannot be overstated), avoidance of insecure technologies and protocols where possible, use of the best available and most secure technologies, and keeping up to date with events in the world of security, especially in order to patch systems when a new exploit is discovered. Just as attackers quickly share information about new flaws using the Internet, system administrators can be warned almost immediately and patches disseminated quickly.

## 4 Network layer security

### 4.1 IPsec overview

The term IPsec refers to a suite of protocols from the IETF providing network layer encryption and authentication for IP-based networks. The objective of IPsec is to authenticate and (optionally) encrypt *all traffic* at the IP level. As it operates at the IP level, it is independent of applications and transport.

IPsec first arose from a workshop held by the Internet Architecture Board (IAB) in 1994 on security in the Internet architecture, from which recommendations were published in RFC 1636. The key IPsec specifications are provided in:

- RFC 2401 (Security architecture)
- RFC 2402 (Authentication)
- RFC 2406 (Encryption)
- RFC 2408 (Security Associations & Key management)

The main present-day use of IPsec is in establishing virtual private networks for connecting remote offices and users to the enterprise using the public Internet, for low-cost remote access for teleworkers (via local call to ISP) and for extranet connectivity (secure communication with partners, suppliers and so on).

IPsec is implemented by means of one of two alternative IP header extensions. The first, Authentication Header (AH) [RFC 2402], provides authentication but not privacy. The alternative, Encapsulating Security Payload (ESP), provides packet encryption and, optionally, authentication. AH adds an additional header field to the traditional IP packet, normally based on a message authentication code (key-based hash of the packet data). With ESP, the content is encrypted and encapsulated between header and trailer fields.

To use IPsec, a pair of hosts must first negotiate a Security Association (SA). This acts as a virtual connection, for which various attributes are set such as the type of protection, keys, and cryptographic algorithms to be used. An SA specifies a one-way relationship, so two SAs are required for a duplex connection. An SA caters for AH or ESP, but not both.

IPsec can be deployed in either transport mode or tunnel mode. Transport mode is typically used for end-to-end communication and protects the IP packet payload only – a consequence of this is that the traffic pattern between hosts is not protected, but there is no need for involvement of intermediate devices (that may not be trusted). Tunnel mode, by contrast, is typically used for connecting secure gateways (e.g. firewalls or routers) and protects the entire IP packet, including the header. A significant advantage of tunnel mode is that hosts do not need to be IPsec-enabled, which will often be the case in mixed IPv4-IPv6 environments.

Cryptographic key management is a significant issue with IPsec – i.e. how to generate and distribute secret keys? Within a small organization, this can be done manually by the system administrator but it does not scale well. A number of automated approaches exist, most notably Internet Security Association and Key Management Protocol (ISAKMP) [RFC 2408] and Internet Key Exchange (IKE) [RFC 2409].

While the objective of introducing security mechanisms like IPsec is to ensure data privacy and authenticity, the mere usage of such mechanisms may not render the security at other layers (application, transport, etc.) redundant and ensure end-to-end communications are fully secure, forever. The framework provided by IPsec is generic enough to allow additional complementing security mechanisms (like PGP, S/MIME, etc.).

In conclusion, it can be said that IPsec provides a level of security for all applications and hosts and allows the deployment of new applications and the addition of new on-site hosts without needing any extra configuration. It also readily supports the secure addition of off-site users and partners. A further benefit of IPsec is that the architecture is independent of specific cryptographic methods and new, stronger, algorithms can be used as they become available. Note though that IPsec cannot strictly provide user level authentication, but rather packet source (i.e. host) authentication. This is not a concern if the user is working, say, at a Windows desktop, but would be an issue with a multi-user OS.

#### *Performance issues*

Implementing security with IPv6 of course has a performance overhead, as does any deployment of cryptography [Ferguson, 2003] – this is particularly significant where protection is applied to *all* traffic, though only a small portion may be security-sensitive. Most IP traffic in organizations does not require IPsec protection. For example, there is nothing to be gained from protecting a user's connection from home into the corporate network with an IPsec VPN when all they are doing is general web browsing or downloads.

IPsec can have a significant impact on throughput and processing power required in devices. What may be insignificant in terms of processing power for a standard user's desktop, can very quickly become the proverbial 'straw' when IPsec is deployed on a VPN server with several hundred users or on a low powered mobile or portable device [Welcher, 2004] [Ronan, 2004].

## 4.2 *Security with IPv4*

IPv4 does not mandate IPsec and hence does not have the security mechanism inbuilt at the network layer. However, most major vendors today support IPsec in their products. IPsec is realized in various forms, the most common forms being bump-in-the-stack (BIS), bump-in-the-code (BIC) and bump-in-the-wire (BIW).

One desirable attribute on an Internet-like infrastructure is to have an authenticated source. Such authentication relies on the global addressability of hosts and devices. Hosts and devices connected behind a NAT have private addresses that are mapped onto globally routable addresses, behind the translation device. Hence only the translating host can be authenticated at the destination and not the real source that sits behind the translating device. The source will have to be authenticated separately at the translating device. A similar scheme has to be adopted at the destination host, should such a host be behind a translating device. Thus it is evident that a direct end-to-end authentication is not possible with translating devices in between. The IETF is actively trying to find a solution for this problem.

For true end-to-end authentication, the basic need therefore is a distinct untranslated IP address. On the global Internet today, the lack of IPv4 address space thwarts the use of end-to-end authentication.

## 4.3 *Security with IPv6*

IPsec is mandated in the protocol. Every implementation claiming support for IPv6 is expected to provide IPsec as part of the protocol.

To effectively use IPsec, there is a need for a key management framework (ISAKMP and IKE) to make an end-to-end secure communication truly happen. Such key management mechanisms, though used extensively with IPsec, are independent and are not a part of IPv6. Therefore, Public Key Infrastructures (PKIs) are required for wide scale deployment. PKIs function as authoritative sources for certified keys of hosts and services on the Internet and are somewhat similar operationally to the Domain Name Service (DNS). There is no accepted standard for PKI, yet. It is also very unlikely that there will be a single PKI for the entire Internet; it is neither acceptable operationally nor does it go with the Internet philosophy. Present day implementations use static key allocations and often do a manual exchange of keys.

An alternative for the provision of public-key authentication for IP addresses without relying on any trusted third parties, PKI, or other global infrastructure, is the use of cryptographically generated addresses (CGAs). CGAs are currently under investigation in the IETF, and provide an intermediate level of security below strong public-key authentication and above routing-based methods. The idea is to form the last 64 bits of an IPv6 address, the interface identifier, by computing a 64-bit one-way hash of the node's public signature key. The node signs its data with the corresponding private key and sends the public key along with the signed data. The recipient hashes the public key and compares the hash to the interface identifier of the source IP address before verifying the signature on the transmitted data. This prevents anyone except the node itself from sending data for its address. As only IPv6 addresses have a 64-bit interface identifier, CGAs consequently can only be used with IPv6.

In many instances, a constant 64-bit interface identifier is used to form a global IPv6 address (stateless address auto configuration). In the event that secure transfers are not using tunnel mode, the IPv6 source and destination addresses are visible rendering the fact that the occurrence of the session itself can be noticed by an intermediate snooper. In cases where the devices move between networks, it then becomes possible to track the movement of the

device and hence the sessions it participates in. This is considered a serious threat to privacy, especially for mobile and wireless users. RFC 3041 is proposed as a solution to this. The solution involves the use of a pseudorandom number as an interface identifier that changes over time, to generate an IPv6 address. This makes it difficult for an eavesdropper to correlate activity based on an address and hence makes it extremely difficult (if not unfeasible) for an eavesdropper to detect or track a given device (and thus potentially a user).

## **5 IPv4-IPv6 transition**

The transition from IPv4 to IPv6 will not of course happen overnight. Organizations that are adopting IPv6 are generally doing so piecemeal, largely due to the need to support legacy systems and applications. As with any new technology adoption, it is prudent to proceed with caution and first select pilot portions of the network for migration. The effect is that IPv4 and IPv6 will need to coexist for a considerable period of time. This means a dual-stack approach for systems, as well as extensive use of tunneling to deliver IPv6 packets over IPv4 networks (and vice versa). This coexistence phase presents several security challenges.

One of the greatest enemies of security is complexity. In general, the more complex a system is, the greater the risk of human error and the more opportunities exist for attack. For example, dual IPv4-IPv6 routers need more configuration than IPv4-only or IPv6-only routers. [Convery, 2004] report a 50% increase in the number of lines of a typical firewall configuration when IPv6 is added. The bigger it is, the greater the chance of misconfiguration. In addition, there are now two distinct protocols that can be attacked rather than just one.

Existing IPv4 systems have deployed security technologies that are well understood. Problems have been ironed out over time based on experience. Very strict change control is required as IPv6 is rolled out as experimental deployment can provide ways for long-established safeguards to be bypassed. In fact it is difficult to secure new IPv6 deployments on existing networks as some network protection mechanisms like intrusion detection systems may not yet support IPv6.

## **6 Conclusion**

All end-to-end security models today inherently imply security above the transport layer. PGP S/MIME and SSL secure higher layer objects and hand them down to the lower layers. Additionally, link layer security mechanisms ensure privacy on the physical communications link, hop-by-hop. IPsec in IPv6 implies security at the network layer. It complements the security mechanisms at the other layers and does not eliminate the need for them.

Users are becoming increasingly mobile and are demanding increasing flexibility, making perimeter security (firewalls, etc) less effective for organizations. In a world where applications are increasingly developed as Web services and port tunneling techniques are well advanced, firewalls, once seen as critical to system security, are increasingly perceived as having limitations [Singer, 2003]

Business applications will benefit by taking advantage of the IPv6 security infrastructure. There is an implicit need here for confidentiality as well as authentication. While security mechanisms today provide for confidentiality of objects, data in transit (transport payload) as well as link layer encryption, there is no specific security mechanism at the network layer. The most important benefits for such a specific community are twofold. All sources of data can be authenticated and data confidentiality can be provided with the use of IPsec. Given that such a community most likely already has a specific PKI developed for its own use, deployment of IPsec with this PKI becomes simply a matter of integration of the two.

Network management data is collected to analyze and monitor the traffic across the network. This information is strategic to decision makers in the corporate entity in order to provision a network for future growth. This data could be perceived as commercially sensitive and hence there is a need to secure such data. From the service provider perspective, the capability to collect accurate billing data is critical. This data needs to be secure and authentic; otherwise it could result in inappropriate, inaccurate or non-existent billing with consequent revenue losses (this is perceived to be especially important in the 3G wireless context).

To ensure that every end-to-end session is private in the real sense of the word, a large support infrastructure to support security is required. A public key infrastructure (PKI) is required with the objective of providing certified public keys for every potential IPv6 host. An IPv6 host that intends to communicate securely with a remote host will require to have the latter's public key to begin secure communication. Alternatively, with the use of CGAs, IPv6 can provide a level of security even without these supporting infrastructures.

In the current Internet scenario, conservative address allocation policies as well as asymmetric user traffic characteristics have resulted in the widespread use of network address translation devices. NATs have broken the peer-to-peer model of the Internet. With its huge address space, IPv6 is expected to re-enable transparent end-to-end applications and services over the Internet. Security will play a vital role in sustaining this attribute.

It is worth concluding with two final remarks. Firstly, despite the promise of a secure world of mobile IPv6-enabled devices communicating end-to-end, facilitated by IPsec, it is important to recognize that no security technology is a panacea – rather, security technologies are only useful in the context of a good, frequently-updated, security policy that is adhered to. For example, having a highly secure IPsec connection will not protect a host against Internet worms if they have penetrated the corporate network. The worm will quite happily make its way down the ‘secure’ tunnel and attempt to infect the host. Secondly, it should also be recognized that there is no such thing as perfect security – there will always be a need to balance security requirements with business and user demands for flexibility and freedom to get on with what they are doing. The key to effective security is to understand where to strike this balance.

## References

- [Clark, 2002] David D. Clark, John Wroclawski, Karen R. Sollins, Robert Braden “Tussle in cyberspace: Defining tomorrow's Internet” *SIGCOMM'02*, (August 2002) pages 19-23, Pittsburgh, Pennsylvania, USA.
- [Convery, 2004] S. Convery & D. Miller, “IPv6 and IPv4 Threat Comparison and Best-Practice Evaluation”, v1.0, Cisco Systems Technical Report, March 2004
- [Ferguson, 2004] Niels Ferguson & Bruce Schneier, “Practical Cryptography”, John Wiley & Sons, 2003
- [ITU, 2003] ITU-T, *Security in Telecommunications and Information Technology*, International Telecommunication Union, Geneva, 2003.
- [Manikopoulos, 2002] C. Manikopoulos, S. Papavassiliou, “Network Intrusion and Fault Detection: A Statistical Anomaly Approach”, *IEEE Communications Magazine*, October 2002
- [Reed, 1998] David P. Reed, Jerome H. Saltzer, and David D. Clark. “Comment on Active Networking and End-to-End Arguments.” *IEEE Network* 12, 3 (May/June 1998) pages 69–71.

- [RFC 791] J. Postel, "Internet Protocol", *RFC 791*, September 1981
- [RFC 793] J. Postel, "Transmission Control Protocol", *RFC 793*, September 1981
- [RFC 1636] R. Braden et al, "Report of IAB Workshop on Security in the Internet Architecture", *RFC 1636*, February 1994
- [RFC 1918] Y. Rekhter et al, "Address Allocation for Private Internets", *RFC 1918*, February 1996
- [RFC 2401] S. Kent & R. Atkinson, "Security Architecture for the Internet Protocol", *RFC 2401*, November 1998
- [RFC 2402] S. Kent & R. Atkinson, "IP Authentication Header", *RFC 2402*, November 1998
- [RFC 2406] S. Kent & R. Atkinson, "IP Encapsulating Security Payload (ESP)", *RFC 2406*, November 1998
- [RFC 2408] D. Maughan et al, "Internet Security Association and Key Management Protocol (ISAKMP)", *RFC 2408*, November 1998
- [RFC 2409] D. Harkins & D. Carrel, "The Internet Key Exchange (IKE)", *RFC 2409*, November 1998
- [RFC 2460] S. Deering & R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", *RFC 2460*, December 1998
- [RFC 2462] S. Thomson & T. Narten, "IPv6 Stateless Address Autoconfiguration", *RFC 2462*, December 1998
- [RFC 2663] P. Srisuresh & M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", *RFC 2663*, August 1999
- [RFC 3022] P. Srisuresh & K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", *RFC 3022*, January 2001
- [RFC 3041] T. Narten & R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", *RFC 3041*, January 2001
- [RFC 3587] R. Hinden, S. Deering & E. Nordmark, "IPv6 Global Unicast Address Format", *RFC 3587*, August 2003
- [Ronan, 2004] J. Ronan et al, "Performance Implications of IPsec Deployment", IPS2004 Interdomain Performance and Simulation (Intermon Workshop), Budapest Hungary, March 2004
- [Saltzer, 1981] Jerome H. Saltzer, David P. Reed, and David D. Clark, *ACM Transactions on Computer Systems* 2, 4 (November 1984) pages 277-288. An earlier version appeared in the *Second International Conference on Distributed Computing Systems* (April, 1981) pages 509-512.
- [Singer, 2003] A. Singer, "Life without firewalls", *USENIX ;login:* magazine, December 2003
- [Welcher, 2004] P. J. Welcher & I. Engle, "Case Study: IPsec VPN Performance", <http://www.netcraftsmen.net/welcher/papers/ipsec-perf-01.html>