

## Reflections on Internet Transparency

### Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The IETF Trust (2007).

### Abstract

This document provides a review of previous IAB statements on Internet transparency, as well a discussion of new transparency issues. Far from having lessened in relevance, technical implications of intentionally or inadvertently impeding network transparency play a critical role in the Internet's ability to support innovation and global communication. This document provides some specific illustrations of those potential impacts.

### Table of Contents

1. Introduction .....	2
2. Additional Transparency Issues .....	4
2.1. Application Restriction .....	4
2.2. Quality of Service (QoS) .....	6
2.3. Application Layer Gateways (ALGs) .....	7
2.4. IPv6 Address Restrictions .....	8
2.4.1. Allocation of IPv6 Addresses by Providers .....	8
2.4.2. IKEv2 .....	8
2.5. DNS Issues .....	9
2.5.1. Unique Root .....	9
2.5.2. Namespace Mangling .....	9
2.6. Load Balancing and Redirection .....	10
3. Security Considerations .....	11
4. References .....	11
4.1. Informative References .....	11
Acknowledgments .....	13
Appendix A - IAB Members at the Time of Approval .....	14

## 1. Introduction

In the past, the IAB has published a number of documents relating to Internet transparency and the end-to-end principle, and other IETF documents have also touched on these issues as well. These documents articulate the general principles on which the Internet architecture is based, as well as the core values that the Internet community seeks to protect going forward. This document reaffirms those principles, describes the concept of "oblivious transport" as developed in the DARPA NewArch project [NewArch], and addresses a number of new transparency issues.

A network that does not filter or transform the data that it carries may be said to be "transparent" or "oblivious" to the content of packets. Networks that provide oblivious transport enable the deployment of new services without requiring changes to the core. It is this flexibility that is perhaps both the Internet's most essential characteristic as well as one of the most important contributors to its success.

"Architectural Principles of the Internet" [RFC1958], Section 2 describes the core tenets of the Internet architecture:

However, in very general terms, the community believes that the goal is connectivity, the tool is the Internet Protocol, and the intelligence is end to end rather than hidden in the network.

The current exponential growth of the network seems to show that connectivity is its own reward, and is more valuable than any individual application such as mail or the World-Wide Web. This connectivity requires technical cooperation between service providers, and flourishes in the increasingly liberal and competitive commercial telecommunications environment.

"The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture" [RFC3724], Section 4.1.1 describes some of the desirable consequences of this approach:

One desirable consequence of the end-to-end principle is protection of innovation. Requiring modification in the network in order to deploy new services is still typically more difficult than modifying end nodes. The counterargument - that many end nodes are now essentially closed boxes which are not updatable and that most users don't want to update them anyway - does not apply to all nodes and all users. Many end nodes are still user configurable and a sizable percentage of users are "early adopters," who are willing to put up with a certain amount of technological grief in order to try out a new idea. And, even for

the closed boxes and uninvolved users, downloadable code that abides by the end-to-end principle can provide fast service innovation. Requiring someone with a new idea for a service to convince a bunch of ISPs or corporate network administrators to modify their networks is much more difficult than simply putting up a Web page with some downloadable software implementing the service.

Yet, even while the Internet has greatly expanded both in size and in application diversity, the degree of transparency has diminished. "Internet Transparency" [RFC2775] notes some of the causes for the loss of Internet transparency and analyzes their impact. This includes discussion of Network Address Translators (NATs), firewalls, application level gateways (ALGs), relays, proxies, caches, split Domain Name Service (DNS), load balancers, etc. [RFC2775] also analyzes potential future directions that could lead to the restoration of transparency. Section 6 summarizes the conclusions:

Although the pure IPv6 scenario is the cleanest and simplest, it is not straightforward to reach it. The various scenarios without use of IPv6 are all messy and ultimately seem to lead to dead ends of one kind or another. Partial deployment of IPv6, which is a required step on the road to full deployment, is also messy but avoids the dead ends.

While full restoration of Internet transparency through the deployment of IPv6 remains a goal, the Internet's growing role in society, the increasing diversity of applications, and the continued growth in security threats has altered the balance between transparency and security, and the disparate goals of interested parties make these tradeoffs inherently complex.

While transparency provides great flexibility, it also makes it easier to deliver unwanted as well as wanted traffic. Unwanted traffic is increasingly cited as a justification for limiting transparency. If taken to its logical conclusion, this argument will lead to the development of ever more complex transparency barriers to counter increasingly sophisticated security threats. Transparency, once lost, is hard to regain, so that such an approach, if unsuccessful, would lead to an Internet that is both insecure and lacking in transparency. The alternative is to develop increasingly sophisticated host-based security mechanisms; while such an approach may also fail to keep up with increasingly sophisticated security threats, it is less likely to sacrifice transparency in the process.

Since many of the fundamental forces that have led to a reduction in the transparency of the IPv4 Internet also may play a role in the IPv6 Internet, the transparency of the IPv6 Internet is not pre-

ordained, but rather represents an ideal whose maintenance will require significant ongoing effort.

As noted in [NewArch], the technical cooperation that once characterized the development of the Internet has increasingly given way to a tussle between the interests of subscribers, vendors, providers, and society at large. Oblivious transport may be desired by developers seeking to deploy new services; providers may desire to block unwanted traffic in the core before it impacts subscribers; vendors and providers may wish to enable delivery of "value added" services in the network that enable them to differentiate their offerings; subscribers may be sympathetic to either point of view, depending on their interests; society at large may wish to block "offensive" material and monitor traffic that shows malicious intent.

While there is no architectural "fix" that can restore oblivious transport while satisfying the interests of all parties, it is possible for providers to provide subscribers with information about the nature of the services being provided. Subscribers need to be aware of whether they are receiving oblivious transport, and if not, how the service affects their traffic.

Since the publication of the previously cited IAB statements, new technologies have been developed, and views on existing technology have changed. In some cases, these new technologies impact oblivious transport, and subscribers need to be aware of the implications for their service.

## 2. Additional Transparency Issues

### 2.1. Application Restriction

Since one of the virtues of the Internet architecture is the ease with which new applications can be deployed, practices that restrict the ability to deploy new applications have the potential to reduce innovation.

One such practice is filtering designed to block or restrict application usage, implemented without customer consent. This includes Internet, Transport, and Application layer filtering designed to block or restrict traffic associated with one or more applications.

While provider filtering may be useful to address security issues such as attacks on provider infrastructure or denial of service attacks, greater flexibility is provided by allowing filtering to be determined by the customer. Typically, this would be implemented at the edges, such as within provider access routers (e.g., outsourced

firewall services), customer premise equipment (e.g., access firewalls), or on hosts (e.g., host firewalls). Deployment of filtering at the edges provides customers with the flexibility to choose which applications they wish to block or restrict, whereas filtering in the core may not permit hosts to communicate, even when the communication would conform to the appropriate use policies of the administrative domains to which those hosts belong.

In practice, filtering intended to block or restrict application usage is difficult to successfully implement without customer consent, since over time developers will tend to re-engineer filtered protocols so as to avoid the filters. Thus over time, filtering is likely to result in interoperability issues or unnecessary complexity. These costs come without the benefit of effective filtering since many application protocols began to use HTTP as a transport protocol after application developers observed that firewalls allow HTTP traffic while dropping packets for unknown protocols.

In addition to architectural concerns, filtering to block or restrict application usage also raises issues of disclosure and end-user consent. As pointed out in "Terminology for Describing Internet Connectivity" [RFC4084], services advertised as providing "Internet connectivity" differ considerably in their capabilities, leading to confusion. The document defines terminology relating to Internet connectivity, including "Web connectivity", "Client connectivity only, without a public address", "Client only, public address", "Firewalled Internet Connectivity", and "Full Internet Connectivity". With respect to "Full Internet Connectivity" [RFC4084], Section 2 notes:

Filtering Web proxies, interception proxies, NAT, and other provider-imposed restrictions on inbound or outbound ports and traffic are incompatible with this type of service. Servers ... are typically considered normal. The only compatible restrictions are bandwidth limitations and prohibitions against network abuse or illegal activities.

[RFC4084], Section 4 describes disclosure obligations that apply to all forms of service limitation, whether applied on outbound or inbound traffic:

More generally, the provider should identify any actions of the service to block, restrict, or alter the destination of, the outbound use (i.e., the use of services not supplied by the provider or on the provider's network) of applications services.

In essence, [RFC4084] calls for providers to declare the ways in which the service provided departs from oblivious transport. Since the lack of oblivious transport within transit networks will also affect transparency, this also applies to providers over whose network the subscriber's traffic may travel.

## 2.2. Quality of Service (QoS)

While [RFC4084] notes that bandwidth limitations are compatible with "Full Internet Connectivity", in some cases QoS restrictions may go beyond simple average or peak bandwidth limitations. When used to restrict the ability to deploy new applications, QoS mechanisms are incompatible with "Full Internet Connectivity" as defined in [RFC4084]. The disclosure and consent obligations referred to in [RFC4084], Section 4 also apply to QoS mechanisms.

Deployment of QoS technology has potential implications for Internet transparency, since the QoS experienced by a flow can make the Internet more or less oblivious to that flow. While QoS support is highly desirable in order for real-time services to coexist with elastic services, it is not without impact on packet delivery.

Specifically, QoS classes such as "default" [RFC2474] or "lower effort" [RFC3662] may experience higher random-loss rates than others such as "assured forwarding" [RFC2597]. Conversely, bandwidth-limited QoS classes such as "expedited forwarding" [RFC3246] may experience systematic packet loss if they exceed their assigned bandwidth. Other QoS mechanisms such as load balancing may have side-effects such as re-ordering of packets, which may have a serious impact on perceived performance.

QoS implementations that reduce the ability to deploy new applications on the Internet are similar in effect to other transparency barriers. Since arbitrary or severe bandwidth limitations can make an application unusable, the introduction of application-specific bandwidth limitations is equivalent to application blocking or restriction from a user's standpoint.

Using QoS mechanisms to discriminate against traffic not matching a set of services or addresses has a similar effect to deployment of a highly restrictive firewall. Requiring an authenticated RSVP reservation [RFC2747][RFC3182] for a flow to avoid severe packet loss has a similar effect to deployment of authenticated firewall traversal.

As with filtering, there may be valid uses for customer-imposed QoS restrictions. For example, a customer may wish to limit the bandwidth consumed by peer-to-peer file sharing services, so as to limit the impact on mission-critical applications.

### 2.3. Application Layer Gateways (ALGs)

The IAB has devoted considerable attention to Network Address Translation (NAT), so that there is little need to repeat that discussion here. However, with the passage of time, it has become apparent that there are problems inherent in the deployment of Application Layer Gateways (ALGs) (frequently embedded within firewalls and devices implementing NAT).

[RFC2775], Section 3.5 states:

If the full range of Internet applications is to be used, NATs have to be coupled with application level gateways (ALGs) or proxies. Furthermore, the ALG or proxy must be updated whenever a new address-dependent application comes along. In practice, NAT functionality is built into many firewall products, and all useful NATs have associated ALGs, so it is difficult to disentangle their various impacts.

With the passage of time and development of NAT traversal technologies such as IKE NAT-T [RFC3947], Teredo [RFC4380], and STUN [RFC3489], it has become apparent that ALGs represent an additional barrier to transparency. In addition to posing barriers to the deployment of new applications not yet supported by ALGs, ALGs may create difficulties in the deployment of existing applications as well as updated versions. For example, in the development of IKE NAT-T, additional difficulties were presented by "IPsec Helper" ALGs embedded within NATs.

It should be stressed that these difficulties are inherent in the architecture of ALGs, rather than merely an artifact of poor implementations. No matter how well an ALG is implemented, barriers to transparency will emerge over time, so that the notion of a "transparent ALG" is a contradiction in terms.

In particular, DNS ALGs present a host of issues, including incompatibilities with DNSSEC that prevent deployment of a secure naming infrastructure even if all the endpoints are upgraded. For details, see "Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status" [RFC4966], Section 3.

## 2.4. IPv6 Address Restrictions

[RFC2775], Section 5.1 states:

Note that it is a basic assumption of IPv6 that no artificial constraints will be placed on the supply of addresses, given that there are so many of them. Current practices by which some ISPs strongly limit the number of IPv4 addresses per client will have no reason to exist for IPv6.

Constraints on the supply of IPv6 addresses provide an incentive for the deployment of NAT with IPv6. The introduction of NAT for IPv6 would represent a barrier to transparency, and therefore is to be avoided if at all possible.

### 2.4.1. Allocation of IPv6 Addresses by Providers

In order to encourage deployments of IPv6 to provide oblivious transport, it is important that IPv6 networks of all sizes be supplied with a prefix sufficient to enable allocation of addresses and sub-networks for all the hosts and links within their network. Initial address allocation policy suggested allocating a /48 prefix to "small" sites, which should handle typical requirements. Any changes to allocation policy should take into account the transparency reduction that will result from further restriction. For example, provider provisioning of a single /64 without support for prefix delegation or (worse still) a longer prefix (prohibited by [RFC4291], Section 2.5.4 for non-000/3 unicast prefixes) would represent a restriction on the availability of IPv6 addresses that could represent a barrier to transparency.

### 2.4.2. IKEv2

Issues with IPv6 address assignment mechanisms in IKEv2 [RFC4306] are described in [RFC4718]:

IKEv2 also defines configuration payloads for IPv6. However, they are based on the corresponding IPv4 payloads, and do not fully follow the "normal IPv6 way of doing things"... In particular, IPv6 stateless autoconfiguration or router advertisement messages are not used; neither is neighbor discovery.

IKEv2 provides for the assignment of a single IPv6 address, using the INTERNAL\_IP6\_ADDRESS attribute. If this is the only attribute supported for IPv6 address assignment, then only a single IPv6 address will be available. The INTERNAL\_IP6\_SUBNET attribute enables the host to determine the sub-networks accessible directly through the secure tunnel created; it could potentially be used to assign one



or more prefixes to the IKEv2 initiator that could be used for address creation.

However, this does not enable the host to obtain prefixes that can be delegated. The INTERNAL\_IP6\_DHCP attribute provides the address of a DHCPv6 server, potentially enabling use of DHCPv6 prefix delegation [RFC3633] to obtain additional prefixes. However, in order for implementers to utilize these options in an interoperable way, clarifications to the IKEv2 specification appear to be needed.

## 2.5. DNS Issues

### 2.5.1. Unique Root

In "IAB Technical Comment on the Unique DNS Root" [RFC2826], the technical arguments for a unique root were presented.

One of the premises in [RFC2826] is that a common namespace and common semantics applied to these names is needed for effective communication between two parties. The document argues that this principle can only be met when one unique root is being used and when the domains are maintained by single owners or maintainers.

Because [RFC4084] targets only IP service terms and does not talk about namespace issues, it does not refer to [RFC2826]. We stress that the use of a unique root for the DNS namespace is essential for proper IP service.

### 2.5.2. Namespace Mangling

Since the publication of [RFC2826], there have been reports of providers implementing recursive nameservers and/or DNS forwarders that replace answers that indicate that a name does not exist in the DNS hierarchy with a name and an address record that hosts a Web service that is supposed to be useful for end-users.

The effect of this modification is similar to placement of a wildcard in top-level domains. Although wildcard labels in top-level domains lead to problems that are described elsewhere (such as "The Role of Wildcards in the Domain Name System" [RFC4592]), they do not strictly violate the DNS protocol. This is not the case where modification of answers takes place in the middle of the path between authoritative servers and the stub resolvers that provide the answers to applications.

[RFC2826] section 1.3 states:

Both the design and implementations of the DNS protocol are heavily based on the assumption that there is a single owner or maintainer for every domain, and that any set of resources records associated with a domain is modified in a single-copy serializable fashion.

In particular, the DNSSEC protocol described in "Protocol Modifications for the DNS Security Extensions" [RFC4035] has been designed to verify that DNS information has not been modified between the moment they have been published on an authoritative server and the moment the validation takes place. Since that verification can take place at the application level, any modification by a recursive forwarder or other intermediary will cause validation failures, disabling the improved security that DNSSEC is intended to provide.

## 2.6. Load Balancing and Redirection

In order to provide information that is adapted to the locale from which a request is made or to provide a speedier service, techniques have been deployed that result in packets being redirected or taking a different path depending on where the request originates. For example, requests may be distributed among servers using "reverse NAT" (which modifies the destination rather than the source address); responses to DNS requests may be altered; HTTP "gets" may be re-directed; or specific packets may be diverted onto overlay networks.

Provided that these services are well-implemented, they can provide value; however, transparency reduction or service disruption can also result:

- [1] The use of "reverse NAT" to balance load among servers supporting IPv6 would adversely affect the transparency of the IPv6 Internet.
- [2] DNS re-direction is typically based on the source address of the query, which may not provide information on the location of the host originating the query. As a result, a host configured with the address of a distant DNS server could find itself pointed to a server near the DNS server, rather than a server near the host. HTTP re-direction does not encounter this issue.
- [3] If the packet filters that divert packets onto overlay networks are misconfigured, this can lead to packets being misdirected onto the overlay and delayed or lost if the far end cannot return them to the global Internet.

- [4] The use of anycast needs to be carefully thought out so that service can be maintained in the face of routing changes.

### 3. Security Considerations

Several transparency issues discussed in this document (NATs, transparent proxies, DNS namespace mangling) weaken existing end-to-end security guarantees and interfere with the deployment of protocols that would strengthen end-to-end security.

[RFC2775], Section 7 states:

The loss of transparency at the Intranet/Internet boundary may be considered a security feature, since it provides a well defined point at which to apply restrictions. This form of security is subject to the "crunchy outside, soft inside" risk, whereby any successful penetration of the boundary exposes the entire Intranet to trivial attack. The lack of end-to-end security applied within the Intranet also ignores insider threats.

Today, malware has evolved to increasingly take advantage of the application-layer as a rich and financially attractive source of security vulnerabilities, as well as a mechanism for penetration of the Intranet/Internet boundary. This has lessened the security value of existing transparency barriers and made it increasingly difficult to prevent the propagation of malware without imposing restrictions on application behavior. However, as with other approaches to application restriction (see Section 2.1), these limitations are most flexibly imposed at the edge.

### 4. References

#### 4.1. Informative References

- [NewArch] Clark, D. et al., "New Arch: Future Generation Internet Architecture",  
<http://www.isi.edu/newarch/iDOCS/final.finalreport.pdf>
- [RFC1958] Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, June 1996.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black,  
"Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.
- [RFC2597] Heinanen, J., Baker, F., Weiss, W., and J. Wroclawski,  
"Assured Forwarding PHB Group", RFC 2597, June 1999.

- [RFC2747] Baker, F., Lindell, B., and M. Talwar, "RSVP Cryptographic Authentication", RFC 2747, January 2000.
- [RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, February 2000.
- [RFC2826] Internet Architecture Board, "IAB Technical Comment on the Unique DNS Root", RFC 2826, May 2000.
- [RFC3182] Yadav, S., Yavatkar, R., Pabbati, R., Ford, P., Moore, T., Herzog, S., and R. Hess, "Identity Representation for RSVP", RFC 3182, October 2001.
- [RFC3246] Davie, B., Charny, A., Bennet, J., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)", RFC 3246, March 2002.
- [RFC3489] Rosenberg, J., Weinberger, J., Huitema, C., and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", RFC 3489, March 2003.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC3662] Bless, R., Nichols, K., and K. Wehrle, "A Lower Effort Per-Domain Behavior (PDB) for Differentiated Services", RFC 3662, December 2003.
- [RFC3724] Kempf, J., Ed., Austein, R., Ed., and IAB, "The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture", RFC 3724, March 2004.
- [RFC3947] Kivinen, T., Swander, B., Huttunen, A., and V. Volpe, "Negotiation of NAT-Traversal in the IKE", RFC 3947, January 2005.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.
- [RFC4084] Klensin, J., "Terminology for Describing Internet Connectivity", BCP 104, RFC 4084, May 2005.

- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4306] Kaufman, C., Ed., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.
- [RFC4592] Lewis, E., "The Role of Wildcards in the Domain Name System", RFC 4592, July 2006.
- [RFC4718] Eronen, P. and P. Hoffman, "IKEv2 Clarifications and Implementation Guidelines", RFC 4718, October 2006.
- [RFC4966] Aoun, C. and E. Davies, "Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status", RFC 4966, July 2007.

#### Acknowledgments

The authors would like to acknowledge Jari Arkko, Stephane Bortzmeyer, Brian Carpenter, Spencer Dawkins, Stephen Kent, Carl Malamud, Danny McPherson, Phil Roberts and Pekka Savola for contributions to this document.

## Appendix A - IAB Members at the Time of Approval

Bernard Aboba  
Loa Andersson  
Brian Carpenter  
Leslie Daigle  
Elwyn Davies  
Kevin Fall  
Olaf Kolkman  
Kurtis Lindqvist  
David Meyer  
David Oran  
Eric Rescorla  
Dave Thaler  
Lixia Zhang

## Authors' Addresses

Bernard Aboba  
Microsoft Corporation  
One Microsoft Way  
Redmond, WA 98052

EMail: [bernarda@microsoft.com](mailto:bernarda@microsoft.com)  
Phone: +1 425 706 6605  
Fax: +1 425 936 7329

Elwyn B. Davies  
Consultant  
Soham, Cambs  
UK

Phone: +44 7889 488 335  
EMail: [elwynd@dial.pipex.com](mailto:elwynd@dial.pipex.com)

## Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

## Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.