

Resolution of FQDN Conflicts among DHCP Clients
<[draft-ietf-dhc-ddns-resolution-10.txt](#)>

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on March 6, 2006.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

DHCP provides a mechanism for host configuration that includes dynamic assignment of IP addresses and fully qualified domain names. To maintain accurate name to IP address and IP address to name mappings in the DNS, these dynamically assigned addresses and fully qualified domain names require updates to the DNS. This document identifies situations in which conflicts in the use of fully qualified domain names may arise among DHCP clients and servers, and describes a strategy for the use of the DHCID DNS resource record in

resolving those conflicts.

Table of Contents

1.	Terminology	3
2.	Introduction	3
3.	Issues with DNS Update in DHCP Environments	3
3.1.	Client Misconfiguration	4
3.2.	Multiple DHCP Servers	5
4.	Use of the DHCID RR	5
5.	DNS RR TTLs	6
6.	Procedures for Performing DNS Updates	7
6.1.	Error Return Codes	7
6.2.	Dual IPv4/IPv6 Client Considerations	7
6.3.	Adding A and/or AAAA RRs to DNS	7
6.3.1.	Initial DHCID RR Query	8
6.3.2.	DNS UPDATE When FQDN in Use	8
6.3.3.	FQDN in Use by another Client	9
6.4.	Adding PTR RR Entries to DNS	9
6.5.	Removing Entries from DNS	9
6.6.	Updating Other RRs	10
7.	Security Considerations	11
8.	Acknowledgements	12
9.	References	12
9.1.	Normative References	12
9.2.	Informative References	12
	Authors' Addresses	14
	Intellectual Property and Copyright Statements	15

1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [1].

This document assumes familiarity with DNS terminology defined in [RFC 1035](#) [4] and DHCP terminology defined in [RFC 2131](#) [5] and [RFC 3315](#) [6].

FQDN, or Fully Qualified Domain Name, is the full name of a system, rather than just its hostname. For example, "venera" is a hostname and "venera.isi.edu" is an FQDN. See [RFC 1983](#) [7].

DOCSIS, or Data-Over-Cable Service Interface Specifications, is defined by CableLabs (www.cablelabs.com).

2. Introduction

"The Client FQDN Option" [8] includes a description of the operation of DHCPv4 [5] clients and servers that use the DHCPv4 client FQDN option. And, "The DHCPv6 Client FQDN Option" [9] includes a description of the operation of DHCPv6 [6] clients and servers that use the DHCPv6 client FQDN option. Through the use of the client FQDN option, DHCP clients and servers can negotiate the client's FQDN and the allocation of responsibility for updating the DHCP client's A and/or AAAA RRs. This document identifies situations in which conflicts in the use of FQDNs may arise among DHCP clients and servers, and describes a strategy for the use of the DHCID DNS resource record [2] in resolving those conflicts.

In any case, whether a site permits all, some, or no DHCP servers and clients to perform DNS updates ([RFC 2136](#) [3], [RFC 3007](#) [10]) into the zones that it controls is entirely a matter of local administrative policy. This document does not require any specific administrative policy, and does not propose one. The range of possible policies is very broad, from sites where only the DHCP servers have been given credentials that the DNS servers will accept, to sites where each individual DHCP client has been configured with credentials that allow the client to modify its own FQDN. Compliant implementations MAY support some or all of these possibilities. Furthermore, this specification applies only to DHCP client and server processes; it does not apply to other processes that initiate DNS updates.

3. Issues with DNS Update in DHCP Environments

There are two DNS update situations that require special consideration in DHCP environments: cases where more than one DHCP client has been configured with the same FQDN and cases where more than one DHCP server has been given authority to perform DNS updates in a zone. In these cases, it is possible for DNS records to be modified in inconsistent ways unless the updaters have a mechanism that allows them to detect anomalous situations. If DNS updaters can detect these situations, site administrators can configure the updaters' behavior so that the site's policies can be enforced. This specification describes a mechanism designed to allow updaters to detect these situations, and suggests that DHCP implementations use this mechanism by default.

3.1. Client Misconfiguration

Administrators may wish to maintain a one-to-one relationship between active DHCP clients and FQDNs, and to maintain consistency between a client's A, AAAA, and PTR RRs. Clients that are not represented in the DNS, or clients that inadvertently share an FQDN with another client may encounter inconsistent behavior or may not be able to obtain access to network resources. Whether each DHCP client is configured with a FQDN by its administrator or whether the DHCP server is configured to distribute the clients' FQDN, the consistency of the DNS data is entirely dependent on the accuracy of the configuration procedure. Sites that deploy Secure DNS [\[11\]](#) may configure credentials for each client and its assigned FQDN in a way that is more error-resistant, as both the FQDN and credentials must match.

Consider an example in which two DHCP clients in the "org.nil" network are both configured with the hostname "foo". The clients are permitted to perform their own DNS updates. The first client, client A, is configured via DHCP. It adds an A RR to "foo.org.nil", and its DHCP server adds a PTR RR corresponding to its assigned IP address. When the second client, client B, boots, it is also configured via DHCP, and it also begins to update "foo.org.nil".

At this point, the "org.nil" administrators may wish to establish some policy about DHCP clients' FQDNs. If the policy is that each client that boots should replace any existing A RR that matches its FQDN, Client B can proceed, though Client A may encounter problems. In this example, Client B replaces the A RR associated with "foo.org.nil". Client A must have some way to recognize that the RR associated with "foo.org.nil" now contains information for Client B, so that it can avoid modifying the RR. When Client A's assigned IP address expires, for example, it should not remove a RR that reflects Client B's DHCP assigned IP address.

If the policy is that the first DHCP client with a given FQDN should be the only client associated with that FQDN, Client B needs to be able to determine if it is not the client associated with "foo.org.nil". It could be that Client A booted first, and that Client B should choose another FQDN. Or it could be that B has booted on a new subnet, and received a new IP address assignment, in which case B should update the DNS with its new IP address. It must either retain persistent state about the last IP address it was assigned (in addition to its current IP address) or it must have some other way to detect that it was the last updater of "foo.org.nil" in order to implement the site's policy.

3.2. Multiple DHCP Servers

It is possible to arrange for DHCP servers to perform A and/or AAAA RR updates on behalf of their clients. If a single DHCP server manages all of the DHCP clients at a site, it can maintain a database of the FQDNs in use, and can check that database before assigning a FQDN to a client. Such a database is necessarily proprietary, however, and the approach does not work once more than one DHCP server is deployed.

When multiple DHCP servers are deployed, the servers require a way to coordinate the identities of DHCP clients. Consider an example in which DHCP Client A boots, obtains an IP address from Server S1, presenting the hostname "foo" in a Client FQDN option [8] in its DHCPREQUEST message. Server S1 updates the FQDN "foo.org.nil", adding an A RR containing the IP address assigned to A. The client then moves to another subnet, served by Server S2. When Client A boots on the new subnet, Server S2 will assign it a new IP address, and will attempt to add an A RR containing the newly assigned IP address to the FQDN "foo.org.nil". At this point, without some communication mechanism which S2 can use to ask S1 (and every other DHCP server that updates the zone) about the client, S2 has no way to know whether Client A is currently associated with the FQDN, or whether A is a different client configured with the same FQDN. If the servers cannot distinguish between these situations, they cannot enforce the site's naming policies.

4. Use of the DHCID RR

A solution to both of these problems is for the updater (a DHCP client or DHCP server) to be able to determine which DHCP client has been associated with a FQDN, in order to offer administrators the opportunity to configure updater behavior.

For this purpose, a DHCID RR, specified in [2], is used to associate

client identification information with a FQDN and the A, AAAA, and PTR RRs associated with that FQDN. When either a client or server adds A, AAAA, or PTR RRs for a client, it also adds a DHCID RR that specifies a unique client identity, based on data from the client's DHCPREQUEST message. In this model, only one client is associated with a given FQDN at a time.

By associating this ownership information with each FQDN, cooperating DNS updaters may determine whether their client is currently associated with a particular FQDN and implement the appropriately configured administrative policy. In addition, DHCP clients which currently have FQDNs may move from one DHCP server to another without losing their FQDNs.

The specific algorithm utilizing the DHCID RR to signal client ownership is explained below. The algorithm only works in the case where the updating entities all cooperate -- this approach is advisory only and is not a substitute for DNS security, nor is it replaced by DNS security.

5. DNS RR TTLs

RRs associated with DHCP clients may be more volatile than statically configured RRs. DHCP clients and servers that perform dynamic updates should attempt to specify resource record TTLs which reflect this volatility, in order to minimize the possibility that answers to DNS queries will return records that refer to DHCP IP address assignments that have expired or been released.

The coupling among primary, secondary, and caching DNS servers is 'loose'; that is a fundamental part of the design of the DNS. This looseness makes it impossible to prevent all possible situations in which a resolver may return a record reflecting a DHCP assigned IP address that has expired or been released. In deployment, this rarely, if ever, represents a significant problem. Most DHCP-managed clients are infrequently looked-up by name in the DNS, and the deployment of IXFR ([RFC 1995](#) [12]) and NOTIFY ([RFC 1996](#) [13]) can reduce the latency between updates and their visibility at secondary servers.

We suggest these basic guidelines for implementers. In general, the TTLs for RRs added as a result of DHCP IP address assignment activity SHOULD be less than the initial lease time or lifetime. The RR TTL on a DNS record added SHOULD NOT exceed 1/3 of the lease time or lifetime, and SHOULD be at least 10 minutes. We recognize that individual administrators will have varying requirements: DHCP servers and clients SHOULD allow administrators to configure TTLs and

upper and lower bounds on the TTL values, either as an absolute time interval or as a percentage of the lease time or lifetime.

While clients and servers MAY update the TTL of the records as the lease or lifetime is about to expire, there is no requirement that they do so as this puts additional load on the DNS system with likely little benefit.

6. Procedures for Performing DNS Updates

6.1. Error Return Codes

Certain RCODEs defined in [RFC 2136](#) [3] indicate that the destination DNS server cannot perform an update: FORMERR, SERVFAIL, REFUSED, NOTIMP. If one of these RCODEs is returned, the updater MUST terminate its update attempt. Because these errors may indicate a misconfiguration of the updater or of the DNS server, the updater MAY attempt to signal to its administrator that an error has occurred, e.g. through a log message.

6.2. Dual IPv4/IPv6 Client Considerations

At the time of publication of this document, a small minority of DHCP clients support both IPv4 and IPv6. We anticipate, however, that a transition will take place over a period of time, and more sites will have dual-stack clients present. IPv6 clients require updates of AAAA RRs; IPv4 client require updates of A RRs. The administrators of mixed deployments will likely wish to permit a single FQDN to contain A and AAAA RRs from the same client.

Sites that wish to permit a single FQDN to contain both A and AAAA RRs MUST make use of DHCPv4 clients and servers that support using the DHCP Unique Identifier (DUID) for DHCPv4 client identifiers such that this DUID is used in computing the RDATA of the DHCID RR by both DHCPv4 and DHCPv6 for the client, see Node-Specific Client Identifiers for DHCPv4 [14]. Otherwise, a dual-stack client that uses older-style DHCPv4 client identifiers (see [5] and [15]) will only be able to have either its A or AAAA records in DNS under a single FQDN because of the DHCID RR conflicts that result.

6.3. Adding A and/or AAAA RRs to DNS

When a DHCP client or server intends to update A and/or AAAA RRs, it starts with the update query in [Section 6.3.1](#).

As the update sequence below can result in loops, implementers SHOULD limit the total number of attempts for a single transaction.

6.3.1. Initial DHCID RR Query

The updater prepares a DNS UPDATE query that includes as a prerequisite the assertion that the FQDN does not exist. The update section of the query attempts to add the new FQDN and its IP address mapping (A and/or AAAA RRs) and the DHCID RR with its unique client identity.

If the update operation succeeds, the A and/or AAAA RR update is now complete (and a client updater is finished, while a server would then proceed to perform a PTR RR update).

If the update returns YXDOMAIN, the updater can now conclude that the intended FQDN is in use and proceeds to [Section 6.3.2](#).

If any other status is returned, the updater SHOULD NOT attempt an update (see [Section 6.1](#)).

6.3.2. DNS UPDATE When FQDN in Use

The updater next attempts to confirm that the FQDN is not being used by some other client by preparing an UPDATE query in which there are two prerequisites. The first prerequisite is that the FQDN exists. The second is that the desired FQDN has attached to it a DHCID RR whose contents match the client identity. The update section of the UPDATE query contains:

1. A delete of any existing A RRs on the FQDN if this is an A update or an AAAA update and the updater does not desire A records on the FQDN.
2. A delete of the existing AAAA RRs on the FQDN if the updater does not desire AAAA records on the FQDN or this update is adding an AAAA and the updater only desires a single IP address on the FQDN.
3. An add of the A RR that matches the DHCP binding if this is an A update.
4. Adds of the AAAA RRs that match the DHCP bindings if this is an AAAA update.

If the update succeeds, the updater can conclude that the current client was the last client associated with the FQDN, and that the FQDN now contains the updated A and/or AAAA RRs. The update is now complete (and a client updater is finished, while a server would then proceed to perform a PTR RR update).

If the update returns NXDOMAIN, the FQDN is no longer in use and the updater proceeds back to [Section 6.3.1](#).

If the update returns NXRRSET, there are two possibilities - there

are no DHCID RRs for the FQDN or the DHCID RR does not match. In either case, the updater proceeds to [Section 6.3.3](#).

[6.3.3](#). FQDN in Use by another Client

As the FQDN appears to be in use by another client or is not associated with any client, the updater can decide (based on some administrative configuration outside of the scope of this document) whether to let the existing owner of the FQDN keep that FQDN, and to (possibly) perform some FQDN disambiguation operation on behalf of the current client, or to replace the RRs on the FQDN with RRs that represent the current client. If the configured policy allows replacement of existing records either if another DHCID RR is present or no DHCID RR is present, the updater submits a query that deletes all RRs for the FQDN (with a prerequisite that a DHCID RR exists or does not exist) and adds the A and/or AAAA and DHCID RRs that represent the IP address and client identity of the new client.

Techniques that may be considered to disambiguate FQDNs include adding some suffix or prefix to the hostname portion of the FQDN or randomly generating a hostname.

DISCUSSION:

The updating entity may be configured to allow the existing DNS records on the FQDN to remain unchanged, and to perform disambiguation on the FQDN of the current client in order to attempt to generate a similar but unique FQDN for the current client. In this case, once another candidate FQDN has been generated, the updater should restart the process of adding A and/or AAAA RRs as specified in this section.

[6.4](#). Adding PTR RR Entries to DNS

The DHCP server submits a DNS query that deletes all of the PTR RRs associated with the client's assigned IP address, and adds a PTR RR whose data is the client's (possibly disambiguated) FQDN. The server MAY also add a DHCID RR as specified in [Section 4](#), in which case it would include a delete of all of the DHCID RRs associated with the client's assigned IP address, and adds a DHCID RR for the client.

There is no need to validate the DHCID RR for PTR updates as the DHCP server (or servers) only assigns an address to a single client at a time.

[6.5](#). Removing Entries from DNS

The most important consideration in removing DNS entries is be sure

that an entity removing a DNS entry is only removing an entry that it added, or for which an administrator has explicitly assigned it responsibility.

When an address' lease time or valid lifetime expires or a DHCP client issues a DHCPRELEASE [5] or Release [6] request, the DHCP server SHOULD delete the PTR RR that matches the DHCP binding, if one was successfully added. The server's update query SHOULD assert that the domain name (PTRDNAME field) in the PTR record matches the FQDN of the client whose address has expired or been released and should delete all RRs for the FQDN.

The entity chosen to handle the A or AAAA records for this client (either the client or the server) SHOULD delete the A or AAAA records that was added when the address was assigned to the client. However, the updater should only remove the DHCID RR if there are no A or AAAA RRs remaining for the client.

In order to perform this A or AAAA RR delete, the updater prepares an UPDATE query that contains a prerequisite that asserts that the DHCID RR exists whose data is the client identity described in [Section 4](#) and contains an update section that deletes the client's specific A or AAAA RR.

If the query succeeds, the updater prepares a second UPDATE query that contains three prerequisites and contains an update section that deletes all RRs for the FQDN. The first prerequisite asserts that the DHCID RR exists whose data is the client identity described in [Section 4](#). The second prerequisite asserts that there are no A RRs. The third prerequisite asserts that there are no AAAA RRs.

If either query fails, the updater MUST NOT delete the FQDN. It may be that the client whose address has expired has moved to another network and obtained an address from a different server, which has caused the client's A or AAAA RR to be replaced. It may also be that some other client has been configured with a FQDN that matches the FQDN of the DHCP client, and the policy was that the last client to specify the FQDN would get the FQDN. In these cases, the DHCID RR will no longer match the updater's notion of the client identity of the client pointed to by the FQDN.

[6.6.](#) Updating Other RRs

The procedures described in this document only cover updates to the A, AAAA, PTR, and DHCID RRs. Updating other types of RRs is outside the scope of this document.

7. Security Considerations

Administrators should be wary of permitting unsecured DNS updates to zones, where or not they are exposed to the global Internet. Both DHCP clients and servers SHOULD use some form of update request authentication (e.g., TSIG [16]) when performing DNS updates.

Whether a DHCP client may be responsible for updating an FQDN to IP address mapping, or whether this is the responsibility of the DHCP server is a site-local matter. The choice between the two alternatives may be based on the security model that is used with the Dynamic DNS Update protocol (e.g., only a client may have sufficient credentials to perform updates to the FQDN to IP address mapping for its FQDN).

Whether a DHCP server is always responsible for updating the FQDN to IP address mapping (in addition to updating the IP to FQDN mapping), regardless of the wishes of an individual DHCP client, is also a site-local matter. The choice between the two alternatives may be based on the security model that is being used with dynamic DNS updates. In cases where a DHCP server is performing DNS updates on behalf of a client, the DHCP server should be sure of the FQDN to use for the client, and of the identity of the client.

Currently, it is difficult for DHCP servers to develop much confidence in the identities of their clients, given the absence of entity authentication from the DHCP protocol itself. There are many ways for a DHCP server to develop a FQDN to use for a client, but only in certain relatively rare circumstances will the DHCP server know for certain the identity of the client. If DHCP Authentication [17] becomes widely deployed this may become more customary.

One example of a situation that offers some extra assurances is when the DHCP client is connected to a network through a DOCSIS cable modem, and the Cable Modem Termination System (head-end) of the cable modem ensures that MAC address spoofing simply does not occur. Another example of a configuration that might be trusted is when clients obtain network access via a network access server using PPP. The Network Access Server (NAS) itself might be obtaining IP addresses via DHCP, encoding client identification into the DHCP client-id option. In this case, the NAS as well as the DHCP server might be operating within a trusted environment, in which case the DHCP server could be configured to trust that the user authentication and authorization processing of the NAS was sufficient, and would therefore trust the client identification encoded within the DHCP client-id.

8. Acknowledgements

Many thanks to Mark Beyer, Jim Bound, Ralph Droms, Robert Elz, Peter Ford, Olafur Gudmundsson, Edie Gunter, Andreas Gustafsson, David W. Hankins, R. Barr Hibbs, Kim Kinnear, Stuart Kwan, Ted Lemon, Ed Lewis, Michael Lewis, Josh Littlefield, Michael Patton, and Glenn Stump for their review and comments.

9. References

9.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Stapp, M., Gustafsson, A., and T. Lemon, "A DNS RR for Encoding DHCP Information ([draft-ietf-dnsext-dhcid-rr](#)-*)", February 2005.
- [3] Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", [RFC 2136](#), April 1997.

9.2. Informative References

- [4] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [5] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.
- [6] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.
- [7] Malkin, G., "Internet Users' Glossary", [RFC 1983](#), August 1996.
- [8] Stapp, M. and Y. Rekhter, "The DHCP Client FQDN Option ([draft-ietf-dhc-fqdn-option](#)-.txt)", February 2005.
- [9] Volz, B., "The DHCPv6 Client FQDN Option ([draft-ietf-dhc-dhcpv6-fqdn](#)-.txt)", February 2005.
- [10] Wellington, B., "Secure Domain Name System (DNS) Dynamic Update", [RFC 3007](#), November 2000.
- [11] Eastlake, D., "Domain Name System Security Extensions", [RFC 2535](#), March 1999.

- [12] Ohta, M., "Incremental Zone Transfer in DNS", [RFC 1995](#), August 1996.
- [13] Vixie, P., "A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)", [RFC 1996](#), August 1996.
- [14] Lemon, T. and B. Sommerfeld, "Node-Specific Client Identifiers for DHCPv4 ([draft-ietf-dhc-3315id-for-v4](#)-*txt)", June 2005.
- [15] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", [RFC 2132](#), March 1997.
- [16] Vixie, P., Gudmundsson, O., Eastlake, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", [RFC 2845](#), May 2000.
- [17] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", [RFC 3118](#), June 2001.

Authors' Addresses

Mark Stapp
Cisco Systems, Inc.
1414 Massachusetts Ave.
Boxborough, MA 01719
USA

Phone: 978.936.1535
Email: mjs@cisco.com

Bernie Volz
Cisco Systems, Inc.
1414 Massachusetts Ave.
Boxborough, MA 01719
USA

Phone: 978.936.0382
Email: volz@cisco.com

Intellectual Property Statement

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.

Disclaimer of Validity

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 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.

Copyright Statement

Copyright (C) The Internet Society (2005). 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.

Acknowledgment

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

