

Network Working Group
Request for Comments: 4703
Category: Standards Track

M. Stapp
B. Volz
Cisco Systems, Inc.
October 2006

Resolution of Fully Qualified Domain Name (FQDN) Conflicts
among Dynamic Host Configuration Protocol (DHCP) Clients

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

The Dynamic Host Configuration Protocol (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 (FQDNs) 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 it describes a strategy for the use of the DHCID DNS resource record (RR) in resolving those conflicts.

Table of Contents

1. Introduction	3
2. Terminology	3
3. Issues with DNS Update in DHCP Environments	4
3.1. Client Misconfiguration	4
3.2. Multiple DHCP Servers	5
4. Use of the DHCID RR	5
5. Procedures for Performing DNS Updates	6
5.1. Error Return Codes	6
5.2. Dual IPv4/IPv6 Client Considerations	6
5.3. Adding A and/or AAAA RRs to DNS	7
5.3.1. Initial DHCID RR Request	7
5.3.2. DNS UPDATE When FQDN in Use	7
5.3.3. FQDN in Use by Another Client	8
5.4. Adding PTR RR Entries to DNS	8
5.5. Removing Entries from DNS	9
5.6. Updating Other RRs	10
6. Security Considerations	10
7. Acknowledgements	11
8. References	11
8.1. Normative References	11
8.2. Informative References	11

1. Introduction

"The Client FQDN Option" [8] includes a description of the operation of [4] clients and servers that use the DHCPv4 client FQDN option. "The DHCPv6 Client FQDN Option" [9] includes a description of the operation of [5] 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 it 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 ([3], [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.

2. 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 [1].

This document assumes familiarity with DNS terminology defined in [6] and DHCP terminology defined in [4] and [5].

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 [7].

DOCSIS, or Data-Over-Cable Service Interface Specifications, is defined by CableLabs.

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 an 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 [10] 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 "example.com" 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.example.com", 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.example.com".

At this point, the "example.com" 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.example.com". Client A must have some way to recognize that the RR associated with "foo.example.com" 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 an 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.example.com". 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.example.com" 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 an 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 DHCPv4 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.example.com", 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.example.com". At this point, without some communication mechanism that 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 an 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 an 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 DHCP 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 that 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. Procedures for Performing DNS Updates

5.1. Error Return Codes

Certain RCODEs defined in [3] indicate that the destination DNS server cannot perform an update, i.e., FORMERR, SERVFAIL, REFUSED, NOTIMP. If one of these RCODEs is returned, the updater MUST terminate its update attempt. Other RCODEs [13] may indicate that there are problems with the key being used and may mean to try a different key, if available, or to terminate the operation. Because some errors may indicate a misconfiguration of the updater or the DNS server, the updater MAY attempt to signal to its administrator that an error has occurred, e.g., through a log message.

5.2. 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 clients 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 [11]. Otherwise, a dual-stack client that uses older-style DHCPv4 client identifiers (see [4] and [12]) 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.

5.3. 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 request in Section 5.3.1.

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

5.3.1. Initial DHCID RR Request

The updater prepares a DNS UPDATE request that includes as a prerequisite the assertion that the FQDN does not exist. The update section of the request 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 request 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 response to the UPDATE returns YXDOMAIN, the updater can now conclude that the intended FQDN is in use and proceeds to Section 5.3.2.

If any other status is returned, the updater SHOULD NOT attempt an update (see Section 5.1).

5.3.2. 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 request 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 request 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, or if this update is adding an A and the updater only desires a single IP address 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 if this update is adding an AAAA and the updater only desires a single IP address on the FQDN.
3. An add (or adds) 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.

Whether A or AAAA RRs are deleted depends on the updater or updater's policy. For example, if the updater is the client or configured as the only DHCP server for the link on which the client is located, the updater may find it beneficial to delete all A and/or AAAA RRs and then add the current set of A and/or AAAA RRs, if any, for the client.

If the UPDATE request 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 response to the UPDATE request returns NXDOMAIN, the FQDN is no longer in use, and the updater proceeds back to Section 5.3.1.

If the response to the UPDATE request 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 5.3.3.

5.3.3. 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 SHOULD either choose another FQDN and restart the update process with this new FQDN or terminate the update with a failure.

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.

5.4. Adding PTR RR Entries to DNS

The DHCP server submits a DNS UPDATE request 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 would add 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.

5.5. Removing Entries from DNS

The most important consideration in removing DNS entries is to 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 [4] or Release [5] request, the DHCP server SHOULD delete the PTR RR that matches the DHCP binding, if one was successfully added. The server's UPDATE request 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 were 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 request 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 UPDATE request succeeds, the updater prepares a second UPDATE request that contains three prerequisites and 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 request 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. Or, the DNS data may have been removed or altered by an administrator.

5.6. 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.

6. Security Considerations

Administrators should be wary of permitting unsecured DNS updates to zones, whether 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 [13]) 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 an 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 [14] 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.

7. 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, Pekka Savola, and Glenn Stump for their review and comments.

8. References

8.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Stapp, M., Lemon, T., and A. Gustafsson, "A DNS Resource Record (RR) for Encoding Dynamic Host Configuration Protocol (DHCP) Information (DHCID RR)", RFC 4701, October 2006.
- [3] Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.
- [4] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [5] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

8.2. Informative References

- [6] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [7] Malkin, G., "Internet Users' Glossary", FYI 18, RFC 1983, August 1996.
- [8] Stapp, M., Volz, B., and Y. Rekhter, "The Dynamic Host Configuration Protocol (DHCP) Client Fully Qualified Domain Name (FQDN) Option", RFC 4702, October 2006.

- [9] Volz, B., "The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Client Fully Qualified Domain Name (FQDN) Option", RFC 4704, October 2006.
- [10] Wellington, B., "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, November 2000.
- [11] Lemon, T. and B. Sommerfeld, "Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)", RFC 4361, February 2006.
- [12] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.
- [13] Vixie, P., Gudmundsson, O., Eastlake, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", RFC 2845, May 2000.
- [14] 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

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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 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.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

