

DHC Working Group
Internet-Draft
Expires: September 2000

M. Stapp
Y. Rekhter
Cisco Systems, Inc.
March 10, 2000

Interaction between DHCP and DNS
[<draft-ietf-dhc-dhcp-dns-12.txt>](#)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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."

To view the entire list of Internet-Draft Shadow Directories, see <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on September 2000.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

DHCP provides a powerful mechanism for IP host configuration. However, the configuration capability provided by DHCP does not include updating DNS, and specifically updating the name to address and address to name mappings maintained in the DNS.

This document specifies how DHCP clients and servers should use the Dynamic DNS Updates mechanism in [RFC2136\[5\]](#) to update the DNS name to address and address to name mappings so that the mappings for DHCP clients will be consistent with the IP addresses that the clients acquire via DHCP.

Table of Contents

- [1.](#) Terminology [3](#)
- [2.](#) Introduction [3](#)
- [3.](#) Models of Operation [3](#)
- [4.](#) Issues with DDNS in DHCP Environments [4](#)
 - [4.1](#) Name Collisions [5](#)
 - [4.2](#) Multiple DHCP servers [6](#)
 - [4.3](#) Use of the DHCID RR [6](#)
 - [4.3.1](#) Format of the DHCID RRDATA [6](#)
 - [4.4](#) DNS RR TTLs [8](#)
- [5.](#) Client FQDN Option [8](#)
 - [5.1](#) The Flags Field [9](#)
 - [5.2](#) The RCODE Fields [10](#)
 - [5.3](#) The Domain Name Field [10](#)
- [6.](#) DHCP Client behavior [10](#)
- [7.](#) DHCP Server behavior [12](#)
- [8.](#) Procedures for performing DNS updates [14](#)
 - [8.1](#) Adding A RRs to DNS [14](#)
 - [8.2](#) Adding PTR RR Entries to DNS [15](#)
 - [8.3](#) Removing Entries from DNS [15](#)
 - [8.4](#) Updating other RRs [16](#)
- [9.](#) Security Considerations [16](#)
- [10.](#) Acknowledgements [17](#)
 - References [17](#)
 - Authors' Addresses [18](#)
 - Full Copyright Statement [19](#)

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](#)[6].

2. Introduction

DNS ([RFC1034](#)[1], [RFC1035](#)[2]) maintains (among other things) the information about mapping between hosts' Fully Qualified Domain Names (FQDNs) [RFC1594](#)[4] and IP addresses assigned to the hosts. The information is maintained in two types of Resource Records (RRs): A and PTR. The A RR contains mapping from a FQDN to an IP address; the PTR RR contains mapping from an IP address to a FQDN. The Dynamic DNS Updates specification ([RFC2136](#)[5]) describes a mechanism that enables DNS information to be updated over a network.

DHCP [RFC2131](#)[3] provides a mechanism by which a host (a DHCP client) can acquire certain configuration information, along with its IP address(es). However, DHCP does not provide any mechanisms to update the DNS RRs that contain the information about mapping between the host's FQDN and its IP address(es) (A and PTR RRs). Thus the information maintained by DNS for a DHCP client may be incorrect - a host (the client) could acquire its address by using DHCP, but the A RR for the host's FQDN wouldn't reflect the address that the host acquired, and the PTR RR for the acquired address wouldn't reflect the host's FQDN.

The Dynamic DNS Update protocol can be used to maintain consistency between the information stored in the A and PTR RRs and the actual address assignment done via DHCP. When a host with a particular FQDN acquires its IP address via DHCP, the A RR associated with the host's FQDN would be updated (by using the Dynamic DNS Updates protocol) to reflect the new address. Likewise, when an IP address is assigned to a host with a particular FQDN, the PTR RR associated with this address would be updated (using the Dynamic DNS Updates protocol) to reflect the new FQDN.

Although this document refers to the A and PTR DNS record types and to DHCP assignment of IPv4 addresses, the same procedures and requirements apply for updates to the analogous RR types that are used when clients are assigned IPv6 addresses via DHCPv6.

3. Models of Operation

When a DHCP client acquires a new address, a site's administrator may desire that one or both of the A RR for the client's FQDN and the PTR RR for the acquired address be updated. Therefore, two separate Dynamic DNS Update transactions occur. Acquiring an address

via DHCP involves two entities: a DHCP client and a DHCP server. In principle each of these entities could perform none, one, or both of the transactions. However, in practice not all permutations make sense. This document covers these possible design permutations:

1. DHCP client updates the A RR, DHCP server updates the PTR RR
2. DHCP server updates both the A and the PTR RRs

The only difference between these two cases is whether the FQDN to IP address mapping is updated by a DHCP client or by a DHCP server. The IP address to FQDN mapping is updated by a DHCP server in both cases.

The reason these two are important, while others are unlikely, has to do with authority over the respective DNS domain names. A DHCP client may be given authority over mapping its own A RRs, or that authority may be restricted to a server to prevent the client from listing arbitrary addresses or associating its address with arbitrary domain names. In all cases, the only reasonable place for the authority over the PTR RRs associated with the address is in the DHCP server that allocates the address.

In any case, whether a site permits all, some, or no DHCP servers and clients to perform DNS updates into the zones which 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 which allow the client to modify its own domain name. 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 which initiate dynamic DNS updates.

This document describes a new DHCP option which a client can use to convey all or part of its domain name to a DHCP server.

Site-specific policy determines whether DHCP servers use the names that clients offer or not, and what DHCP servers may do in cases where clients do not supply domain names.

4. Issues with DDNS 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 anomolous situations. If DNS updaters can

Stapp & Rekhter

Expires September 2000

[Page 4]

detect these situations, site administrators can configure the updaters' behavior so that the site's policies can be enforced. We use the term "Name Collisions" to refer to cases where more than one DHCP client has been associated with a single FQDN. This specification describes a mechanism designed to allow updaters to detect these situations, and requires that DHCP implementations use this mechanism by default.

4.1 Name Collisions

How can the entity updating an A RR (either the DHCP client or DHCP server) detect that a domain name has an A RR which is already in use by a different DHCP client? Similarly, should a DHCP client or server update a domain name which has an A RR that has been configured by an administrator? In either of these cases, the domain name in question would either have an additional A RR, or would have its original A RR replaced by the new record. Either of these effects may be considered undesirable by some sites. Different authority and credential models have different levels of exposure to name collisions.

1. Client updates A RR, uses Secure DNS Update with credentials that are associated with the client's FQDN, and exclusive to the client. Name collisions in this scenario are unlikely (though not impossible), since the client has received credentials specific to the name it desires to use. This implies that the name has already been allocated (through some implementation- or organization-specific procedure) to that client.
2. Client updates A RR, uses Secure DNS Update with credentials that are valid for any name in the zone. Name collisions in this scenario are possible, since the credentials necessary for the client to update DNS are not necessarily name-specific. Thus, for the client to be attempting to update a unique name requires the existence of some administrative procedure to ensure client configuration with unique names.
3. Server updates the A RR, uses a name for the client which is known to the server. Name collisions in this scenario are likely unless prevented by the server's name configuration procedures. See [Section 9](#) for security issues with this form of deployment.
4. Server updates the A RR, uses a name supplied by the client. Name collisions in this scenario are highly likely, even with administrative procedures designed to prevent them. (This scenario is a popular one in real-world deployments in many types of organizations.) See [Section 9](#) for security issues with this type of deployment.

Scenarios 2, 3, and 4 rely on administrative procedures to ensure name uniqueness for DNS updates, and these procedures may break down. Experience has shown that, in fact, these procedures will break down at least occasionally. The question is what to do when

these procedures break down or, for example in scenario #4, may not even exist.

In all cases of name collisions, the desire is to offer two modes of operation to the administrator of the combined DHCP-DNS capability: first-update-wins (i.e., the first updating entity gets the name) or most-recent-update-wins (i.e., the last updating entity for a name gets the name).

4.2 Multiple DHCP servers

If multiple DHCP servers are able to update the same DNS zones, or if DHCP servers are performing A RR updates on behalf of DHCP clients, and more than one DHCP server may be able to serve addresses to the same DHCP clients, the DHCP servers should be able to provide reasonable and consistent DNS name update behavior for DHCP clients.

4.3 Use of the DHCID RR

A solution to both of these problems is for the updating entities (both DHCP clients and DHCP servers) to be able to detect that another entity has been associated with a DNS name, and to offer administrators the opportunity to configure update behavior.

Specifically, a DHCID RR, described in DHCID RR[12] is used to associate client identification information with a DNS name and the A RR associated with that name. When either a client or server adds an A RR for a client, it also adds a DHCID RR which specifies a unique client identity (based on a "client specifier" created from the client's client-id or MAC address). In this model, only one A RR is associated with a given DNS name at a time.

By associating this ownership information with each A RR, cooperating DNS updating entities may determine whether their client is the first or last updater of the name (and implement the appropriately configured administrative policy), and DHCP clients which currently have a host name may move from one DHCP server to another without losing their DNS name.

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

4.3.1 Format of the DHCID RRDATA

The DHCID RR used to hold the DHCP client's identity is formatted as

follows:

The name of the DHCID RR is the name of the A or PTR RR which refers to the DHCP client.

The RDATA section of a DHCID RR in transmission contains RDLENGTH bytes of binary data. From the perspective of DHCP clients and servers, the DHC resource record consists of a 16-bit identifier type, followed by one or more bytes representing the actual identifier. There are two possible forms for a DHCID RR - one that is used when the client's link-layer address is being used to identify it, and one that is used when some DHCP option that the DHCP client has sent is being used to identify it.

DISCUSSION:

Implementors should note that the actual identifying data is never placed into the DNS directly. Instead, the client-identity data is used as the input into a one-way hash algorithm, and the output of that hash is then used as DNS RRDATA. This has been specified in order to avoid placing data about DHCP clients that some sites might consider sensitive into the DNS.

When the updater is using the client's link-layer address, the first two bytes of the DHCID RRDATA MUST be zero. To generate the rest of the resource record, the updater MUST compute a one-way hash using the MD5[13] algorithm across a buffer containing the client's network hardware type and link-layer address. Specifically, the first byte of the buffer contains the network hardware type as it appears in the DHCP htype field of the client's DHCPREQUEST message. All of the significant bytes of the chaddr field in the client's DHCPREQUEST message follow, in the same order in which the bytes appear in the DHCPREQUEST message. The number of significant bytes in the chaddr field is specified in the hlen field of the DHCPREQUEST message.

When the updater is using a DHCP option sent by the client in its DHCPREQUEST message, the first two bytes of the DHCID RR MUST be the option code of that option, in network byte order. For example, if the DHCP client identifier option is being used, the first byte of the DHCID RR should be zero, and the second byte should be 61 decimal. The rest of the DHCID RR MUST contain the results of computing a one-way hash across the payload of the option being used, using the MD5 algorithm. The payload of a DHCP option consists of the bytes of the option following the option code and length.

In order for independent DHCP implementations to be able to use the DHCID RR as a prerequisite in dynamic DNS updates, each updater must

be able to reliably choose the same identifier that any other would choose. To make this possible, we specify a prioritization which

will ensure that for any given DHCP client request, any updaters will select the same client-identity data. All updaters MUST use this order of prioritization by default, but all implementations SHOULD be configurable to use a different prioritization if so desired by the site administrators. Because of the possibility of future changes in the DHCP protocol, implementors SHOULD check for updated versions of this draft when implementing new DHCP clients and servers which can perform DDNS updates, and also when releasing new versions of existing clients and servers.

DHCP clients and servers should use the following forms of client identification, starting with the most preferable, and finishing with the least preferable. If the client does not send any of these forms of identification, the DHCP/DDNS interaction is not defined by this specification. The most preferable form of identification is the Globally Unique Identifier Option [TBD]. Next is the DHCP Client Identifier option. Last is the client's link-layer address, as conveyed in its DHCPREQUEST message. Implementors should note that the link-layer address cannot be used if there are no significant bytes in the chaddr field of the DHCP client's request, because this does not constitute a unique identifier.

4.4 DNS RR TTLs

RRs associated with DHCP clients may be more volatile than statically configured RRs. DHCP clients and servers which perform dynamic updates should attempt to specify resource record TTLs which reflect this volatility, in order to minimize the possibility that there will be stale records in resolvers' caches. A reasonable basis for RR TTLs is the lease duration itself: TTLs of 1/2 or 1/3 the expected lease duration might be reasonable defaults. Because configured DHCP lease times vary widely from site to site, it may also be desirable to establish a fixed TTL ceiling. DHCP clients and servers MAY allow administrators to configure the TTLs they will supply, possibly as a fraction of the actual lease time, or as a fixed value.

5. Client FQDN Option

To update the IP address to FQDN mapping a DHCP server needs to know the FQDN of the client to which the server leases the address. To allow the client to convey its FQDN to the server this document defines a new DHCP option, called "Client FQDN". The FQDN Option also contains Flags and RCode fields which DHCP servers can use to convey information about DNS updates to clients.

Clients MAY send the FQDN option, setting appropriate Flags values, in both their DISCOVER and REQUEST messages. If a client sends the

FQDN option in its DISCOVER message, it MUST send the option in

subsequent REQUEST messages.

The code for this option is 81. Its minimum length is 4.

Code	Len	Flags	RCODE1	RCODE2	Domain Name
81	n				...

5.1 The Flags Field

0	1	2	3	4	5	6	7
+--+--+--+--+--+--+							
MBZ		E O S					
+--+--+--+--+--+--+							

When a DHCP client sends the FQDN option in its DHCPDISCOVER and/or DHCPREQUEST messages, it sets the right-most bit (labelled "S") to indicate that it will not perform any Dynamic DNS updates, and that it expects the DHCP server to perform any FQDN-to-IP (the A RR) DNS update on its behalf. If this bit is clear, the client indicates that it intends to maintain its own FQDN-to-IP mapping update.

If a DHCP server intends to take responsibility for the A RR update whether or not the client sending the FQDN option has set the "S" bit, it sets both the "O" bit and the "S" bit, and sends the FQDN option in its DHCPOFFER and/or DHCPACK messages.

The data in the Domain Name field may appear in one of two formats: ASCII, or DNS-style binary encoding (without compression, of course), as described in [RFC1035\[2\]](#). A client which sends the FQDN option MUST set the "E" bit to indicate that the data in the Domain Name field is DNS binary encoded. If a server receives an FQDN option from a client, and intends to include an FQDN option in its reply, it MUST use the same encoding that the client used. The DNS encoding is recommended. The use of ASCII-encoded domain-names is fragile, and the use of ASCII encoding in this option should be considered deprecated.

The remaining bits in the Flags field are reserved for future assignment. DHCP clients and servers which send the FQDN option MUST set the MBZ bits to 0, and they MUST ignore values in the part of the field labelled "MBZ".

5.2 The RCODE Fields

The RCODE1 and RCODE2 fields are used by a DHCP server to indicate to a DHCP client the Response Code from any A or PTR RR Dynamic DNS Updates it has performed. The server may also use these fields to indicate whether it has attempted such an update before sending the DHCPACK message. Each of these fields is one byte long.

Implementors should note that EDNS0 describes a mechanism for extending the length of a DNS RCODE to 12 bits. EDNS0 is specified in [RFC2671](#)[8]. Only the least-significant 8 bits of the RCODE from a Dynamic DNS Update will be carried in the Client FQDN DHCP Option. This provides enough number space to accomodate the RCODEs defined in the Dynamic DNS Update specification.

5.3 The Domain Name Field

The Domain Name part of the option carries all or part of the FQDN of a DHCP client. A client may be configured with a fully-qualified domain name, or with a partial name that is not fully-qualified. If a client knows only part of its name, it MAY send a single label, indicating that it knows part of the name but does not necessarily know the zone in which the name is to be embedded. The data in the Domain Name field may appear in one of two formats: ASCII (with no terminating NULL), or DNS encoding as specified in [RFC1035](#)[2]. If the DHCP client wishes to use DNS encoding, it MUST set the third-from-rightmost bit in the Flags field (the "E" bit); if it uses ASCII encoding, it MUST clear the "E" bit.

A DHCP client that can only send a single label using ASCII encoding includes a series of ASCII characters in the Domain Name field, excluding the "." (dot) character. The client SHOULD follow the character-set recommendations of [RFC1034](#)[1] and [RFC1035](#)[2]. A client using DNS binary encoding which wants to suggest part of its FQDN MAY send a non-terminal sequence of labels in the Domain Name part of the option.

6. DHCP Client behavior

The following describes the behavior of a DHCP client that implements the Client FQDN option.

If a client that owns/maintains its own FQDN wants to be responsible for updating the FQDN to IP address mapping for the FQDN and address(es) used by the client, then the client MUST include the Client FQDN option in the DHCPREQUEST message originated by the client. A DHCP client MAY choose to include the Client FQDN option in its DISCOVER messages as well as its REQUEST messages. The rightmost ("S") bit in the Flags field in the option MUST be set to

0. Once the client's DHCP configuration is completed (the client receives a DHCPACK message, and successfully completes a final check on the parameters passed in the message), the client MAY originate an update for the A RR (associated with the client's FQDN). The update MUST be originated following the procedures described in [RFC2136\[5\]](#) and [Section 8](#). If the DHCP server from which the client is requesting a lease includes the FQDN option in its ACK message, and if the server sets both the "S" and the "O" bits (the two rightmost bits) in the option's flags field, the DHCP client MUST NOT initiate an update for the name in the Domain Name field.

A client can choose to delegate the responsibility for updating the FQDN to IP address mapping for the FQDN and address(es) used by the client to the server. In order to inform the server of this choice, the client SHOULD include the Client FQDN option in its DHCPREQUEST message. The rightmost (or "S") bit in the Flags field in the option MUST be set to 1. A client which delegates this responsibility MUST NOT attempt to perform a Dynamic DNS update for the name in the Domain Name field of the FQDN option. The client MAY supply an FQDN in the Client FQDN option, or it MAY supply a single label (the most-specific label), or it MAY leave that field empty as a signal to the server to generate an FQDN for the client in any manner the server chooses.

Since there is a possibility that the DHCP server may be configured to complete or replace a domain name that the client was configured to send, the client might find it useful to send the FQDN option in its DISCOVER messages. If the DHCP server returns different Domain Name data in its OFFER message, the client could use that data in performing its own eventual A RR update, or in forming the FQDN option that it sends in its REQUEST message. There is no requirement that the client send identical FQDN option data in its DISCOVER and REQUEST messages. In particular, if a client has sent the FQDN option to its server, and the configuration of the client changes so that its notion of its domain name changes, it MAY send the new name data in an FQDN option when it communicates with the server again. This may allow the DHCP server to update the name associated with the PTR record, and, if the server updated the A record representing the client, to delete that record and attempt an update for the client's current domain name.

A client that delegates the responsibility for updating the FQDN to IP address mapping to a server might not receive any indication (either positive or negative) from the server whether the server was able to perform the update. In this case the client MAY use a DNS query to check whether the mapping is updated.

A client MUST set the RCODE1 and RCODE2 fields in the Client FQDN

option to 0 when sending the option.

Stapp & Rekhter

Expires September 2000

[Page 11]

If a client releases its lease prior to the lease expiration time and the client is responsible for updating its A RR, the client SHOULD delete the A RR (following the procedures described in [Section 8](#)) associated with the leased address before sending a DHCP RELEASE message. Similarly, if a client was responsible for updating its A RR, but is unable to renew its lease, the client SHOULD attempt to delete the A RR before its lease expires. A DHCP client which has not been able to delete an A RR which it added (because it has lost the use of its DHCP IP address) should attempt to notify its administrator.

7. DHCP Server behavior

When a server receives a DHCPREQUEST message from a client, if the message contains the Client FQDN option, and the server replies to the message with a DHCPACK message, the server may be configured to originate an update for the PTR RR (associated with the address leased to the client). Any such update MUST be originated following the procedures described in [Section 8](#). The server MAY complete the update before the server sends the DHCPACK message to the client. In this case the RCODE from the update MUST be carried to the client in the RCODE1 field of the Client FQDN option in the DHCPACK message. Alternatively, the server MAY send the DHCPACK message to the client without waiting for the update to be completed. In this case the RCODE1 field of the Client FQDN option in the DHCPACK message MUST be set to 255. The choice between the two alternatives is entirely determined by the configuration of the DHCP server. Servers SHOULD support both configuration options.

When a server receives a DHCPREQUEST message containing the Client FQDN option, the server MUST ignore the values carried in the RCODE1 and RCODE2 fields of the option.

In addition, if the Client FQDN option carried in the DHCPREQUEST message has the "S" bit in its Flags field set, then the server MAY originate an update for the A RR (associated with the FQDN carried in the option) if it is configured to do so by the site's administrator, and if it has the necessary credentials. The server MAY be configured to use the name supplied in the client's FQDN option, or it MAY be configured to modify the supplied name, or substitute a different name.

Any such update MUST be originated following the procedures described in [Section 8](#). The server MAY originate the update before the server sends the DHCPACK message to the client. In this case the RCODE from the update [[RFC2136](#)] MUST be carried to the client in the RCODE2 field of the Client FQDN option in the DHCPACK message. Alternatively the server MAY send the DHCPACK message to the client

without waiting for the update to be completed. In this case the

Stapp & Rekhter

Expires September 2000

[Page 12]

RCODE2 field of the Client FQDN option in the DHCPACK message MUST be set to 255. The choice between the two alternatives is entirely up to the DHCP server. In either case, if the server intends to perform the DNS update and the client's REQUEST message included the FQDN option, the server SHOULD include the FQDN option in its ACK message, and MUST set the "S" bit in the option's Flags field.

Even if the Client FQDN option carried in the DHCPREQUEST message has the "S" bit in its Flags field clear (indicating that the client wants to update the A RR), the server MAY be configured by the local administrator to update the A RR on the client's behalf. A server which is configured to override the client's preference SHOULD include an FQDN option in its ACK message, and MUST set both the "O" and "S" bits in the FQDN option's Flags field. The update MUST be originated following the procedures described in [Section 8](#). The server MAY originate the update before the server sends the DHCPACK message to the client. In this case the RCODE from the update [[RFC2136](#)] MUST be carried to the client in the RCODE2 field of the Client FQDN option in the DHCPACK message. Alternatively, the server MAY send the DHCPACK message to the client without waiting for the update to be completed. In this case the RCODE2 field of the Client FQDN option in the DHCPACK message MUST be set to 255. Whether the DNS update occurs before or after the DHCPACK is sent is entirely up to the DHCP server's configuration.

When a DHCP server sends the Client FQDN option to a client in the DHCPACK message, the DHCP server SHOULD send its notion of the complete FQDN for the client in the Domain Name field. The server MAY simply copy the Domain Name field from the Client FQDN option that the client sent to the server in the DHCPREQUEST message. The DHCP server MAY be configured to complete or modify the domain name which a client sent, or it MAY be configured to substitute a different name. If the server initiates a DDNS update which is not complete until after the server has replied to the DHCP client, the server's The server MUST use the same encoding format (ASCII or DNS binary encoding) that the client used in the FQDN option in its DHCPREQUEST, and MUST set the "E" bit in the option's Flags field accordingly.

If a client's DHCPREQUEST message doesn't carry the Client FQDN option (e.g., the client doesn't implement the Client FQDN option), the server MAY be configured to update either or both of the A and PTR RRs. The updates MUST be originated following the procedures described in [Section 8](#).

If a server detects that a lease on an address that the server leases to a client has expired, the server SHOULD delete any PTR RR which it added via dynamic update. In addition, if the server added

an A RR on the client's behalf, the server SHOULD also delete the A

RR. The deletion MUST follow the procedures described in [Section 8](#).

If a server terminates a lease on an address prior to the lease's expiration time, for instance by sending a DHCPNAK to a client, the server SHOULD delete any PTR RR which it associated with the address via DNS Dynamic Update. In addition, if the server took responsibility for an A RR, the server SHOULD also delete that A RR. The deletion MUST follow the procedures described in [Section 8](#).

8. Procedures for performing DNS updates

8.1 Adding A RRs to DNS

When a DHCP client or server intends to update an A RR, it first prepares a DNS UPDATE query which includes as a prerequisite the assertion that the name does not exist. The update section of the query attempts to add the new name and its IP address mapping (an A RR), and the DHCID RR with its unique client-identity.

If this update operation succeeds, the updater can conclude that it has added a new name whose only RRs are the A and DHCID RR records. The A RR update is now complete (and a client updater is finished, while a server might proceed to perform a PTR RR update).

If the first update operation fails with YXDOMAIN, the updater can conclude that the intended name is in use. The updater then attempts to confirm that the DNS name is not being used by some other host. The updater prepares a second UPDATE query in which the prerequisite is that the desired name has attached to it a DHCID RR whose contents match the client identity. The update section of this query deletes the existing A records on the name, and adds the A record that matches the DHCP binding and the DHCID RR with the client identity.

If this query succeeds, the updater can conclude that the current client was the last client associated with the domain name, and that the name now contains the updated A RR. The A 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 second query fails with NXRRSET, the updater must conclude that the client's desired name is in use by another host. At this juncture, the updater can decide (based on some administrative configuration outside of the scope of this document) whether to let the existing owner of the name keep that name, and to (possibly) perform some name disambiguation operation on behalf of the current client, or to replace the RRs on the name with RRs that represent the current client. If the configured policy allows replacement of existing records, the updater submits a query that deletes the

existing A RR and the existing DHCID RR, adding A and DHCID RRs that represent the IP address and client-identity of the new client.

DISCUSSION:

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

8.2 Adding PTR RR Entries to DNS

The DHCP server submits a DNS query which deletes all of the PTR RRs associated with the lease IP address, and adds a PTR RR whose data is the client's (possibly disambiguated) host name. The server also adds a DHCID RR specified in [Section 4.3](#).

8.3 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 a lease expires or a DHCP client issues a DHCPRELEASE 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 name in the PTR record matches the name of the client whose lease has expired or been released.

The entity chosen to handle the A record for this client (either the client or the server) SHOULD delete the A record that was added when the lease was made to the client.

In order to perform this delete, the updater prepares an UPDATE query which contains two prerequisites. The first prerequisite asserts that the DHCID RR exists whose data is the client identity described in [Section 4.3](#). The second prerequisite asserts that the data in the A RR contains the IP address of the lease that has expired or been released.

If the query fails, the updater MUST NOT delete the DNS name. It may be that the host whose lease on the server has expired has moved to another network and obtained a lease from a different server, which has caused the client's A RR to be replaced. It may also be that some other client has been configured with a name that matches

the name of the DHCP client, and the policy was that the last client

Stapp & Rekhter

Expires September 2000

[Page 15]

to specify the name would get the name. In this case, the DHCID RR will no longer match the updater's notion of the client-identity of the host pointed to by the DNS name.

8.4 Updating other RRs

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

9. Security Considerations

Unauthenticated updates to the DNS can lead to tremendous confusion, through malicious attack or through inadvertent misconfiguration. Administrators should be wary of permitting unsecured DNS updates to zones which are exposed to the global Internet. Both DHCP clients and servers SHOULD use some form of update request origin authentication procedure (e.g., Simple Secure DNS Update[11]) 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 DNS name 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 its clients, given the absence of entity authentication from the DHCP protocol itself. There are many ways for a DHCP server to develop a DNS name to use for a client, but only in certain relatively unusual circumstances will the DHCP server know for certain the identity of the client. If DHCP Authentication[10] becomes widely deployed this may become more customary.

One example of a situation which offers some extra assurances is one where the DHCP client is connected to a network through an MCNS cable modem, and the CMTS (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 one where clients obtain network access via a network access server using PPP. The NAS itself might be obtaining IP addresses via DHCP, encoding a client identification into the DHCP client-id option. In this case, the network access server 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 procedure of the remote access server was sufficient, and would therefore trust the client identification encoded within the DHCP client-id.

10. Acknowledgements

Many thanks to Mark Beyer, Jim Bound, Ralph Droms, Robert Elz, Peter Ford, Edie Gunter, Andreas Gustafsson, 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.

References

- [1] Mockapetris, P., "Domain names - Concepts and Facilities", [RFC 1034](#), Nov 1987.
- [2] Mockapetris, P., "Domain names - Implementation and Specification", [RFC 1035](#), Nov 1987.
- [3] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.
- [4] Marine, A., Reynolds, J. and G. Malkin, "FYI on Questions and Answers to Commonly asked ``New Internet User'' Questions", [RFC 1594](#), March 1994.
- [5] Vixie, P., Thomson, S., Rekhter, Y. and J. Bound, "Dynamic Updates in the Domain Name System", [RFC 2136](#), April 1997.
- [6] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997.
- [7] Eastlake, D., "Domain Name System Security Extensions", [RFC 2535](#), March 1999.
- [8] Vixie, P., "Extension Mechanisms for DNS (EDNS0)", [RFC 2671](#), August 1999.
- [9] Vixie, P., Gudmundsson, O., Eastlake, D. and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG) ([draft-ietf-dnsext-tsig](#)-*)", July 1999.

- [10] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages ([draft-ietf-dhc-authentication-*](#))", June 1999.
- [11] Wellington, B., "Simple Secure DNS Dynamic Updates ([draft-ietf-dnsext-simple-secure-update-*](#))", June 1999.
- [12] Gustafsson, A., "A DNS RR for encoding DHCP client identity ([draft-ietf-dnsext-dhcid-rr-*](#))", October 1999.
- [13] Rivest, R., "The MD5 Message Digest Algorithm", [RFC 1321](#), April 1992.

Authors' Addresses

Mark Stapp
Cisco Systems, Inc.
250 Apollo Dr.
Chelmsford, MA 01824
US

Phone: 978.244.8498
EMail: mjs@cisco.com

Yakov Rekhter
Cisco Systems, Inc.
170 Tasman Dr.
San Jose, CA 95134
US

Phone: 914.235.2128
EMail: yakov@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

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

