Network Working Group
Request for Comments: 4339
Category: Informational

J. Jeong, Ed.
ETRI/University of Minnesota

February 2006

IPv6 Host Configuration of DNS Server Information Approaches

Status of This Memo

Copyright Notice

IESG Note

This document describes three different approaches for the configuration of DNS name resolution server information in IPv6 hosts.

There is not an IETF consensus on which approach is preferred. The analysis in this document was developed by the proponents for each approach and does not represent an IETF consensus.

The ’RA option’ and ’Well-known anycast’ approaches described in this document are not standardized. Consequently the analysis for these approaches might not be completely applicable to any specific proposal that might be proposed in the future.

Abstract

This document describes three approaches for IPv6 recursive DNS server address configuration. It details the operational attributes of three solutions: RA option, DHCPv6 option, and well-known anycast addresses for recursive DNS servers. Additionally, it suggests the deployment scenarios in four kinds of networks (ISP, enterprise, 3GPP, and unmanaged networks) considering multi-solution resolution.

Table of Contents

1. Introduction

Neighbor Discovery (ND) for IP Version 6 and IPv6 Stateless Address Autoconfiguration provide ways to configure either fixed or mobile nodes with one or more IPv6 addresses, default routes, and some other parameters [1][2]. To support the access to additional services in the Internet that are identified by a DNS name, such as a web server, the configuration of at least one recursive DNS server is also needed for DNS name resolution.

This document describes three approaches of recursive DNS server address configuration for IPv6 host: (a) RA option [6], (b) DHCPv6 option [3]-[5], and (c) well-known anycast addresses for recursive DNS servers [7]. Also, it suggests the applicable scenarios for four kinds of networks: (a) ISP network, (b) enterprise network, (c) 3GPP network, and (d) unmanaged network.

This document is just an analysis of each possible approach, and it does not recommend a particular approach or combination of approaches. Some approaches may even not be adopted at all as a result of further discussion.

Therefore, the objective of this document is to help the audience select the approaches suitable for IPv6 host configuration of recursive DNS servers.

2. Terminology

This document uses the terminology described in [1]-[7]. In addition, a new term is defined below:

- o Recursive DNS Server (RDNSS): Server which provides a recursive DNS resolution service.

3. IPv6 DNS Configuration Approaches

In this section, the operational attributes of the three solutions are described in detail.

3.1. RA Option

The RA approach defines a new ND option, called the RDNSS option, that contains a recursive DNS server address [6]. Existing ND transport mechanisms (i.e., advertisements and solicitations) are used. This works in the same way that nodes learn about routers and prefixes. An IPv6 host can configure the IPv6 addresses of one or more RDNSSes via RA message periodically sent by a router or solicited by a Router Solicitation (RS).

This approach needs RDNSS information to be configured in the routers doing the advertisements. The configuration of RDNSS addresses can be performed manually by an operator or in other ways, such as automatic configuration through a DHCPv6 client running on the router. An RA message with one RDNSS option can include as many RDNSS addresses as needed [6].

Through the ND protocol and RDNSS option, along with a prefix information option, an IPv6 host can perform network configuration of its IPv6 address and RDNSS simultaneously [1][2]. The RA option for RDNSS can be used on any network that supports the use of ND.

The RA approach is useful in some mobile environments where the addresses of the RDNSSes are changing because the RA option includes a lifetime field that allows client to use RDNSSes nearer to the client. This can be configured to a value that will require the client to time out the entry and switch over to another RDNSS address [6]. However, from the viewpoint of implementation, the lifetime field would seem to make matters a bit more complex. Instead of just writing to a DNS configuration file, such as resolv.conf for the list of RDNSS addresses, we have to have a daemon around (or a program that is called at the defined intervals) that keeps monitoring the lifetime of RDNSSes all the time.

The preference value of RDNSS, included in the RDNSS option, allows IPv6 hosts to select primary RDNSS among several RDNSSes [6]; this can be used for the load balancing of RDNSSes.

3.1.1. Advantages

The RA option for RDNSS has a number of advantages. These include:

1. The RA option is an extension of existing ND/Autoconfig mechanisms [1][2] and does not require a change in the base ND protocol.

2. This approach, like ND, works well on a variety of link types, including point-to-point links, point-to-multipoint, and multipoint-to-multipoint (i.e., Ethernet LANs). RFC 2461 [1] states, however, that there may be some link types on which ND is not feasible; on such links, some other mechanisms will be needed for DNS configuration.

3. All the information a host needs to run the basic Internet applications (such as the email, web, ftp, etc.) can be obtained with the addition of this option to ND and address autoconfiguration. The use of a single mechanism is more reliable and easier to provide than when the RDNSS information is

learned via another protocol mechanism. Debugging problems when multiple protocol mechanisms are being used is harder and much more complex.

4. This mechanism works over a broad range of scenarios and leverages IPv6 ND. This works well on links that are high performance (e.g., Ethernet LANs) and low performance (e.g., cellular networks). In the latter case, by combining the RDNSS information with the other information in the RA, the host can learn all the information needed to use most Internet applications, such as the web, in a single packet. This not only saves bandwidth, but also minimizes the delay needed to learn the RDNSS information.

5. The RA approach could be used as a model for similar types of configuration information. New RA options for other server addresses, such as NTP server address, that are common to all clients on a subnet would be easy to define.

3.1.2. Disadvantages

1. ND is mostly implemented in the kernel of the operating system. Therefore, if ND supports the configuration of some additional services, such as DNS servers, ND should be extended in the kernel and complemented by a user-land process. DHCPv6, however, has more flexibility for the extension of service discovery because it is an application layer protocol.

2. The current ND framework should be modified to facilitate the synchronization between another ND cache for RDNSSes in the kernel space and the DNS configuration file in the user space. Because it is unacceptable to write and rewrite to the DNS configuration file (e.g., resolv.conf) from the kernel, another approach is needed. One simple approach to solve this is to have a daemon listening to what the kernel conveys, and to have the daemon do these steps, but such a daemon is not needed with the current ND framework.

3. It is necessary to configure RDNSS addresses at least at one router on every link where this information needs to be configured via the RA option.

3.1.3. Observations

The proposed RDNSS RA option, along with the IPv6 ND and Autoconfiguration, allows a host to obtain all of the information it needs to access basic Internet services like the web, email, ftp, etc. This is preferable in the environments where hosts use RAs to autoconfigure their addresses and all the hosts on the subnet share the same router and server addresses. If the configuration information can be obtained from a single mechanism, it is preferable because it does not add additional delay, and because it uses a minimum of bandwidth. Environments like this include homes, public cellular networks, and enterprise environments where no per host configuration is needed.

DHCPv6 is preferable where it is being used for address configuration and if there is a need for host specific configuration [3]-[5]. Environments like this are most likely to be the enterprise environments where the local administration chooses to have per host configuration control.

3.2. DHCPv6 Option

DHCPv6 [3] includes the "DNS Recursive Name Server" option, through which a host can obtain a list of IP addresses of recursive DNS servers [5]. The DNS Recursive Name Server option carries a list of IPv6 addresses of RDNSSes to which the host may send DNS queries. The DNS servers are listed in the order of preference for use by the DNS resolver on the host.

The DNS Recursive Name Server option can be carried in any DHCPv6 Reply message, in response to either a Request or an Information request message. Thus, the DNS Recursive Name Server option can be used either when DHCPv6 is used for address assignment, or when DHCPv6 is used only for other configuration information as stateless DHCPv6 [4].

Stateless DHCPv6 can be deployed either by using DHCPv6 servers running on general-purpose computers, or on router hardware. Several router vendors currently implement stateless DHCPv6 servers. Deploying stateless DHCPv6 in routers has the advantage that no special hardware is required, and it should work well for networks where DHCPv6 is needed for very straightforward configuration of network devices.

However, routers can also act as DHCPv6 relay agents. In this case, the DHCPv6 server need not be on the router; it can be on a general purpose computer. This has the potential to give the operator of the DHCPv6 server more flexibility in how the DHCPv6 server responds to individual clients that can easily be given different configuration information based on their identity, or for any other reason. Nothing precludes adding this flexibility to a router, but generally, in current practice, DHCP servers running on general-purpose hosts tend to have more configuration options than those that are embedded in routers.

DHCPv6 currently provides a mechanism for reconfiguring DHCPv6 clients that use a stateful configuration assignment. To do this, the DHCPv6 server sends a Reconfigure message to the client. The client validates the Reconfigure message, and then contacts the DHCPv6 server to obtain updated configuration information. By using this mechanism, it is currently possible to propagate new configuration information to DHCPv6 clients as this information changes.

The DHC Working Group has standardized an additional mechanism through which configuration information, including the list of RDNSSes, can be updated. The lifetime option for DHCPv6 [8] assigns a lifetime to configuration information obtained through DHCPv6. At the expiration of the lifetime, the host contacts the DHCPv6 server to obtain updated configuration information, including the list of RDNSSes. This lifetime gives the network administrator another mechanism to configure hosts with new RDNSSes by controlling the time at which the host refreshes the list.

The DHC Working Group has also discussed the possibility of defining an extension to DHCPv6 that would allow the use of multicast to provide configuration information to multiple hosts with a single DHCPv6 message. Because of the lack of deployment experience, the WG has deferred consideration of multicast DHCPv6 configuration at this time. Experience with DHCPv4 has not identified a requirement for multicast message delivery, even in large service provider networks with tens of thousands of hosts that may initiate a DHCPv4 message exchange simultaneously.

3.2.1. Advantages

The DHCPv6 option for RDNSS has a number of advantages. These include:

1. DHCPv6 currently provides a general mechanism for conveying network configuration information to clients. Configuring DHCPv6 servers in this way allows the network administrator to configure RDNSSes, the addresses of other network services, and location-specific information, such as time zones.

2. As a consequence, when the network administrator goes to configure DHCPv6, all the configuration information can be managed through a single service, typically with a single user interface and a single configuration database.

3. DHCPv6 allows for the configuration of a host with information specific to that host, so that hosts on the same link can be configured with different RDNSSes and with other configuration information.

4. A mechanism exists for extending DHCPv6 to support the transmission of additional configuration that has not yet been anticipated.

5. Hosts that require other configuration information, such as the addresses of SIP servers and NTP servers, are likely to need DHCPv6 for other configuration information.

6. The specification for configuration of RDNSSes through DHCPv6 is available as an RFC. No new protocol extensions (such as new options) are necessary.

7. Interoperability among independent implementations has been demonstrated.

3.2.2. Disadvantages

The DHCPv6 option for RDNSS has a few disadvantages. These include:

1. Update currently requires a message from server (however, see [8]).

2. Because DNS information is not contained in RA messages, the host must receive two messages from the router and must transmit at least one message to the router. On networks where bandwidth is at a premium, this is a disadvantage, although on most networks it is not a practical concern.

3. There is an increased latency for initial configuration. In addition to waiting for an RA message, the client must now exchange packets with a DHCPv6 server. Even if it is locally installed on a router, this will slightly extend the time required to configure the client. For clients that are moving rapidly from one network to another, this will be a disadvantage.

3.2.3. Observations

In the general case, on general-purpose networks, stateless DHCPv6 provides significant advantages and no significant disadvantages. Even in the case where bandwidth is at a premium and low latency is desired, if hosts require other configuration information in addition to a list of RDNSSes or if hosts must be configured selectively, those hosts will use DHCPv6 and the use of the DHCPv6 DNS recursive name server option will be advantageous.

However, we are aware of some applications where it would be preferable to put the RDNSS information into an RA packet; for example, in a mobile phone network, where bandwidth is at a premium and extremely low latency is desired. The DNS configuration based on RA should be standardized so as to allow these special applications to be handled using DNS information in the RA packet.

3.3. Well-known Anycast Addresses

Anycast uses the same routing system as unicast [9]. However, administrative entities are local ones. The local entities may accept unicast routes (including default routes) to anycast servers from adjacent entities. The administrative entities should not advertise their peer routes to their internal anycast servers, if they want to prohibit external access from some peers to the servers. If some advertisement is inevitable (such as the case with default routes), the packets to the servers should be blocked at the boundary of the entities. Thus, for this anycast, not only unicast routing but also unicast ND protocols can be used as is.

First of all, the well-known anycast addresses approach is much different from that discussed by the IPv6 Working Group in the past [7]. Note that "anycast" in this memo is simpler than that of RFC 1546 [9] and RFC 3513 [10], where it is assumed to be prohibited to have multiple servers on a single link sharing an anycast address. That is, on a link, an anycast address is assumed to be unique. DNS clients today already have redundancy by having multiple well-known anycast addresses configured as RDNSS addresses. There is no point in having multiple RDNSSes sharing an anycast address on a single link.

The approach with well-known anycast addresses is to set multiple well-known anycast addresses in clients’ resolver configuration files from the beginning as, say, factory default. Thus, there is no transport mechanism and no packet format [7].

An anycast address is an address shared by multiple servers (in this case, the servers are RDNSSes). A request from a client to the anycast address is routed to a server selected by the routing system. However, it is a bad idea to mandate "site" boundary on anycast addresses, because most users do not have their own servers and want to access their ISPs across their site boundaries. Larger sites may also depend on their ISPs or may have their own RDNSSes within "site" boundaries.

3.3.1. Advantages

The basic advantage of the well-known addresses approach is that it uses no transport mechanism. Thus, the following apply:

1. There is no delay to get the response and no further delay by packet losses.

2. The approach can be combined with any other configuration mechanisms, such as the RA-based approach and DHCP-based approach, as well as the factory default configuration.

3. The approach works over any environment where DNS works.

Another advantage is that this approach only needs configuration of the DNS servers as a router (or configuration of a proxy router). Considering that DNS servers do need configuration, the amount of overall configuration effort is proportional to the number of DNS servers and it scales linearly. Note that, in the simplest case, where a subscriber to an ISP does not have a DNS server, the subscriber naturally accesses DNS servers of the ISP, even though the subscriber and the ISP do nothing and there is no protocol to exchange DNS server information between the subscriber and the ISP.

3.3.2. Disadvantages

The well-known anycast addresses approach requires that DNS servers (or routers near to them as a proxy) act as routers to advertise their anycast addresses to the routing system, which requires some configuration (see the last paragraph of the previous section on the scalability of the effort). In addition, routers at the boundary of the "site" might need the configuration of route filters to prevent providing DNS services for parties outside the "site" and the possibility of denial of service attacks on the internal DNS infrastructure.

3.3.3. Observations

If other approaches are used in addition, the well-known anycast addresses should also be set in RA or DHCP configuration files to reduce the configuration effort of users.

The redundancy by multiple RDNSSes is better provided by multiple servers with different anycast addresses than by multiple servers sharing the same anycast address, because the former approach allows stale servers to generate routes to their anycast addresses. Thus, in a routing domain (or domains sharing DNS servers), there will be only one server with an anycast address unless the domain is so large that load distribution is necessary.

Small ISPs will operate one RDNSS at each anycast address that is shared by all the subscribers. Large ISPs may operate multiple RDNSSes at each anycast address to distribute and reduce load, where the boundary between RDNSSes may be fixed (redundancy is still provided by multiple addresses) or change dynamically. DNS packets with the well-known anycast addresses are not expected (though not prohibited) to cross ISP boundaries, as ISPs are expected to be able to take care of themselves.

Because "anycast" in this memo is simpler than that of RFC 1546 [9] and RFC 3513 [10], where it is assumed to be administratively prohibited to have multiple servers on a single link sharing an anycast address, anycast in this memo should be implemented as UNICAST of RFC 2461 [1] and RFC 3513 [10]. As a result, ND-related instability disappears. Thus, in the well-known anycast addresses approach, anycast can and should use the anycast address as a source unicast (according to RFC 3513 [10]) address of packets of UDP and TCP responses. With TCP, if a route flips and packets to an anycast address are routed to a new server, it is expected that the flip is detected by ICMP or sequence number inconsistency, and that the TCP connection is reset and retried.

4. Interworking among IPv6 DNS Configuration Approaches

Three approaches can work together for IPv6 host configuration of RDNSS. This section shows a consideration on how these approaches can interwork.

For ordering between RA and DHCP approaches, the O (Other stateful configuration) flag in the RA message can be used [6][28]. If no RDNSS option is included, an IPv6 host may perform DNS configuration through DHCPv6 [3]-[5] regardless of whether the O flag is set or not.

The well-known anycast addresses approach fully interworks with the other approaches. That is, the other approaches can remove the configuration effort on servers by using the well-known addresses as the default configuration. Moreover, the clients preconfigured with the well-known anycast addresses can be further configured to use other approaches to override the well-known addresses, if the configuration information from other approaches is available. Otherwise, all the clients need to have the well-known anycast addresses preconfigured. In order to use the anycast approach along with two other approaches, there are three choices as follows:

1. The first choice is that well-known addresses are used as last resort, when an IPv6 host cannot get RDNSS information through RA and DHCP. The well-known anycast addresses have to be preconfigured in all of IPv6 hosts' resolver configuration files.

2. The second is that an IPv6 host can configure well-known addresses as the most preferable in its configuration file even though either an RA option or DHCP option is available.

3. The last is that the well-known anycast addresses can be set in RA or DHCP configuration to reduce the configuration effort of users. According to either the RA or DHCP mechanism, the well-known addresses can be obtained by an IPv6 host. Because this approach is the most convenient for users, the last option is recommended.

Note: This section does not necessarily mean that this document suggests adopting all of these three approaches and making them interwork in the way described here. In fact, as a result of further discussion some approaches may not even be adopted at all.

5. Deployment Scenarios

Regarding the DNS configuration on the IPv6 host, several mechanisms are being considered by the DNSOP Working Group, such as RA option, DHCPv6 option, and well-known preconfigured anycast addresses as of today, and this document is a final result from the long thread. In this section, we suggest four applicable scenarios of three approaches for IPv6 DNS configuration.

Note: In the applicable scenarios, authors do not implicitly push any specific approaches into the restricted environments. No enforcement is in each scenario, and all mentioned scenarios are probable. The main objective of this work is to provide a useful guideline for IPv6 DNS configuration.

5.1. ISP Network

A characteristic of an ISP network is that multiple Customer Premises Equipment (CPE) devices are connected to IPv6 PE (Provider Edge) routers and that each PE connects multiple CPE devices to the backbone network infrastructure [11]. The CPEs may be hosts or routers.

If the CPE is a router, there is a customer network that is connected to the ISP backbone through the CPE. Typically, each customer network gets a different IPv6 prefix from an IPv6 PE router, but the same RDNSS configuration will be distributed.

This section discusses how the different approaches to distributing DNS information are compared in an ISP network.

5.1.1. RA Option Approach

When the CPE is a host, the RA option for RDNSS can be used to allow the CPE to get RDNSS information and /64 prefix information for stateless address autoconfiguration at the same time when the host is attached to a new subnet [6]. Because an IPv6 host must receive at least one RA message for stateless address autoconfiguration and router configuration, the host could receive RDNSS configuration information in the RA without the overhead of an additional message exchange.

When the CPE is a router, the CPE may accept the RDNSS information from the RA on the interface connected to the ISP and copy that information into the RAs advertised in the customer network.

This approach is more valuable in the mobile host scenario, in which the host must receive at least an RA message for detecting a new network, than in other scenarios generally, although the administrator should configure RDNSS information on the routers. Secure ND [12] can provide extended security when RA messages are used.

5.1.2. DHCPv6 Option Approach

DHCPv6 can be used for RDNSS configuration through the use of the DNS option, and can provide other configuration information in the same message with RDNSS configuration [3]-[5]. The DHCPv6 DNS option is already in place for DHCPv6, as RFC 3646 [5] and DHCPv6-lite or stateless DHCP [4] is not nearly as complex as a full DHCPv6 implementation. DHCP is a client-server model protocol, so ISPs can handle user identification on its network intentionally; also, authenticated DHCP [13] can be used for secure message exchange.

The expected model for deployment of IPv6 service by ISPs is to assign a prefix to each customer, which will be used by the customer gateway to assign a /64 prefix to each network in the customer’s network. Prefix delegation with DHCP (DHCPv6 PD) has already been adopted by ISPs for automating the assignment of the customer prefix to the customer gateway [15]. DNS configuration can be carried in the same DHCPv6 message exchange used for DHCPv6 to provide that information efficiently, along with any other configuration information needed by the customer gateway or customer network. This service model can be useful to Home or SOHO subscribers. The Home or SOHO gateway, which is a customer gateway for ISP, can then pass that RDNSS configuration information to the hosts in the customer network through DHCP.

5.1.3. Well-known Anycast Addresses Approach

The well-known anycast addresses approach is also a feasible and simple mechanism for ISP [7]. The use of well-known anycast addresses avoids some of the security risks in rogue messages sent through an external protocol such as RA or DHCPv6. The configuration of hosts for the use of well-known anycast addresses requires no protocol or manual configuration, but the configuration of routing for the anycast addresses requires intervention on the part of the network administrator. Also, the number of special addresses would be equal to the number of RDNSSes that could be made available to subscribers.

5.2. Enterprise Network

An enterprise network is defined as a network that has multiple internal links, one or more router connections to one or more providers, and is actively managed by a network operations entity [14]. An enterprise network can get network prefixes from an ISP by either manual configuration or prefix delegation [15]. In most cases, because an enterprise network manages its own DNS domains, it operates its own DNS servers for the domains. These DNS servers within enterprise networks process recursive DNS name resolution requests from IPv6 hosts as RDNSSes. The RDNSS configuration in the enterprise network can be performed as it is in Section 4, in which three approaches can be used together as follows:

1. An IPv6 host can decide which approach is or may be used in its subnet with the O flag in RA message [6][28]. As the first choice in Section 4, well-known anycast addresses can be used as a last resort when RDNSS information cannot be obtained through either an RA option or a DHCP option. This case needs IPv6 hosts to preconfigure the well-known anycast addresses in their DNS configuration files.

2. When the enterprise prefers the well-known anycast approach to others, IPv6 hosts should preconfigure the well-known anycast addresses as it is in the first choice.

3. The last choice, a more convenient and transparent way, does not need IPv6 hosts to preconfigure the well-known anycast addresses

because the addresses are delivered to IPv6 hosts via either the RA option or DHCPv6 option as if they were unicast addresses. This way is most recommended for the sake of the user’s convenience.

5.3. 3GPP Network

The IPv6 DNS configuration is a missing part of IPv6 autoconfiguration and an important part of the basic IPv6 functionality in the 3GPP User Equipment (UE). The higher-level description of the 3GPP architecture can be found in [16], and transition to IPv6 in 3GPP networks is analyzed in [17] and [18].

In the 3GPP architecture, there is a dedicated link between the UE and the GGSN called the Packet Data Protocol (PDP) Context. This link is created through the PDP Context activation procedure [19]. There is a separate PDP context type for IPv4 and IPv6 traffic. If a 3GPP UE user is communicating by using IPv6 (i.e., by having an active IPv6 PDP context), it cannot be assumed that the user simultaneously has an active IPv4 PDP context, and DNS queries could be done using IPv4. A 3GPP UE can thus be an IPv6 node, and somehow it needs to discover the address of the RDNSS. Before IP-based services (e.g., web browsing or e-mail) can be used, the IPv6 (and IPv4) RDNSS addresses need to be discovered in the 3GPP UE.

Section 5.3.1 briefly summarizes currently available mechanisms in 3GPP networks and recommendations. 5.3.2 analyzes the Router Advertisement-based solution, 5.3.3 analyzes the Stateless DHCPv6 mechanism, and 5.3.4 analyzes the well-known addresses approach. Section 5.3.5 summarizes the recommendations.

5.3.1. Currently Available Mechanisms and Recommendations

3GPP has defined a mechanism in which RDNSS addresses can be received in the PDP context activation (a control plane mechanism). That is called the Protocol Configuration Options Information Element (PCO-IE) mechanism [20]. The RDNSS addresses can also be received over the air (using text messages) or typed in manually in the UE. Note that the two last mechanisms are not very well scalable. The UE user most probably does not want to type IPv6 RDNSS addresses manually in the user’s UE. The use of well-known addresses is briefly discussed in section 5.3.4.

It is seen that the mechanisms above most probably are not sufficient for the 3GPP environment. IPv6 is intended to operate in a zero-configuration manner, no matter what the underlying network infrastructure is. Typically, the RDNSS address is needed to make an IPv6 node operational, and the DNS configuration should be as simple as the address autoconfiguration mechanism. Note that there will be additional IP interfaces in some near-future 3GPP UEs; e.g., 3GPP-specific DNS configuration mechanisms (such as PCO-IE [20]) do not work for those IP interfaces. In other words, a good IPv6 DNS configuration mechanism should also work in a multi-access network environment.

From a 3GPP point of view, the best IPv6 DNS configuration solution is feasible for a very large number of IPv6-capable UEs (even hundreds of millions in one operator’s network), is automatic, and thus requires no user action. It is suggested that a lightweight, stateless mechanism be standardized for use in all network environments. The solution could then be used for 3GPP, 3GPP2, and other access network technologies. Thus, not only is a light, stateless IPv6 DNS configuration mechanism needed in 3GPP networks, but also 3GPP networks and UEs would certainly benefit from the new mechanism.

5.3.2. RA Extension

Router Advertisement extension [6] is a lightweight IPv6 DNS configuration mechanism that requires minor changes in the 3GPP UE IPv6 stack and Gateway GPRS Support Node (GGSN, the default router in the 3GPP architecture) IPv6 stack. This solution can be specified in the IETF (no action is needed in the 3GPP) and taken in use in 3GPP UEs and GGSNs.

In this solution, an IPv6-capable UE configures DNS information via an RA message sent by its default router (GGSN); i.e., the RDNSS option for a recursive DNS server is included in the RA message. This solution is easily scalable for a very large number of UEs. The operator can configure the RDNSS addresses in the GGSN as a part of normal GGSN configuration. The IPv6 RDNSS address is received in the Router Advertisement, and an extra Round Trip Time (RTT) for asking RDNSS addresses can be avoided.

When one considers the cons, this mechanism still requires standardization effort in the IETF, and the end nodes and routers need to support this mechanism. The equipment software update should, however, be pretty straightforward, and new IPv6 equipment could support RA extension already from the beginning.

5.3.3. Stateless DHCPv6

A DHCPv6-based solution needs the implementation of Stateless DHCP [4] and DHCPv6 DNS options [5] in the UE, and a DHCPv6 server in the operator’s network. A possible configuration is such that the GGSN works as a DHCP relay.

The pros of a stateless DHCPv6-based solution are:

1. Stateless DHCPv6 is a standardized mechanism.
2. DHCPv6 can be used for receiving configuration information other than RDNSS addresses; e.g., SIP server addresses.
3. DHCPv6 works in different network environments.
4. When DHCPv6 service is deployed through a single, centralized server, the RDNSS configuration information can be updated by the network administrator at a single source.

Some issues with DHCPv6 in 3GPP networks are listed below:

1. DHCPv6 requires an additional server in the network unless the (Stateless) DHCPv6 functionality is integrated into an existing router. This means that there might be one additional server to be maintained.
2. DHCPv6 is not necessarily needed for 3GPP UE IPv6 addressing (3GPP Stateless Address Autoconfiguration is typically used) and is not automatically implemented in 3GPP IPv6 UEs.
3. Scalability and reliability of DHCPv6 in very large 3GPP networks (with tens or hundreds of millions of UEs) may be an issue; at least the redundancy needs to be taken care of. However, if the DHCPv6 service is integrated into the network elements, such as a router operating system, scalability and reliability is comparable with other DNS configuration approaches.
4. It is sub-optimal to utilize the radio resources in 3GPP networks for DHCPv6 messages if there is a simpler alternative is available.
   * The use of stateless DHCPv6 adds one round-trip delay to the case in which the UE can start transmitting data right after the Router Advertisement.
5. If the DNS information (suddenly) changes, Stateless DHCPv6 cannot automatically update the UE; see [21].

5.3.4. Well-known Addresses

Using well-known addresses is also a feasible and light mechanism for 3GPP UEs. Those well-known addresses can be preconfigured in the UE software and the operator can make the corresponding configuration on the network side. Thus, this is a very easy mechanism for the UE, but it requires some configuration work in the network. When using well-known addresses, UE forwards queries to any of the preconfigured addresses. In the current proposal [7], IPv6 anycast addresses are suggested.

Note: An IPv6 DNS configuration proposal, based on the use of well-known site-local addresses, was developed by the IPv6 Working Group; it was seen as a feasible mechanism for 3GPP UEs, although no IETF consensus was reached on this proposal. In the end, the deprecation of IPv6 site-local addresses made it impossible to standardize a mechanism that uses site-local addresses as well-known addresses. However, as of this writing, this mechanism is implemented in some operating systems and 3GPP UEs as a last resort of IPv6 DNS configuration.

5.3.5. Recommendations

It is suggested that a lightweight, stateless DNS configuration mechanism be specified as soon as possible. From a 3GPP UE and network point of view, the Router Advertisement-based mechanism looks most promising. The sooner a light, stateless mechanism is specified, the sooner we can stop using well-known site-local addresses for IPv6 DNS configuration.

5.4. Unmanaged Network

There are four deployment scenarios of interest in unmanaged networks [22]:

1. A gateway that does not provide IPv6 at all,

2. A dual-stack gateway connected to a dual-stack ISP,

3. A dual-stack gateway connected to an IPv4-only ISP, and

4. A gateway connected to an IPv6-only ISP.

5.4.1. Case A: Gateway Does Not Provide IPv6 at All

In this case, the gateway does not provide IPv6; the ISP may or may not provide IPv6. Automatic or Configured tunnels are the recommended transition mechanisms for this scenario.

The case where dual-stack hosts behind an NAT need access to an IPv6 RDNSS cannot be entirely ruled out. The DNS configuration mechanism has to work over the tunnel, and the underlying tunneling mechanism could implement NAT traversal. The tunnel server assumes the role of a relay (for both DHCP and well-known anycast addresses approaches).

The RA-based mechanism is relatively straightforward in its operation, assuming the tunnel server is also the IPv6 router emitting RAs. The well-known anycast addresses approach also seems simple in operation across the tunnel, but the deployment model using well-known anycast addresses in a tunneled environment is unclear or not well understood.

5.4.2. Case B: A Dual-stack Gateway Connected to a Dual-stack ISP

This is similar to a typical IPv4 home user scenario, where DNS configuration parameters are obtained using DHCP. The exception is that Stateless DHCPv6 is used, as opposed to the IPv4 scenario, where the DHCP server is stateful (it maintains the state for clients).

5.4.3. Case C: A Dual-stack Gateway Connected to an IPv4-only ISP

This is similar to Case B. If a gateway provides IPv6 connectivity by managing tunnels, then it is also supposed to provide access to an RDNSS. Like this, the tunnel for IPv6 connectivity originates from the dual-stack gateway instead of from the host.

5.4.4. Case D: A Gateway Connected to an IPv6-only ISP

This is similar to Case B.

6. Security Considerations

As security requirements depend solely on applications and differ from application to application, there can be no generic requirement defined at the IP or application layer for DNS.

However, note that cryptographic security requires configured secret information and that full autoconfiguration and cryptographic security are mutually exclusive. People insisting on secure, full autoconfiguration will get false security, false autoconfiguration, or both.

In some deployment scenarios [17], where cryptographic security is required for applications, the secret information for the cryptographic security is preconfigured, through which application-specific configuration data, including those for DNS, can be securely configured. Note that if applications requiring cryptographic security depend on DNS, the applications also require cryptographic security to DNS. Therefore, the full autoconfiguration of DNS is not acceptable.

However, with full autoconfiguration, weaker but still reasonable security is being widely accepted and will continue to be acceptable.

That is, with full autoconfiguration, which means there is no cryptographic security for the autoconfiguration, it is already assumed that the local environment is secure enough that the information from the local autoconfiguration server has acceptable security even without cryptographic security. Thus, the communication between the local DNS client and local DNS server has acceptable security.

In autoconfiguring recursive servers, DNSSEC may be overkill, because DNSSEC [23]-[25] needs the configuration and reconfiguration of clients at root key roll-over [26][27]. Even if additional keys for secure key roll-over are added at the initial configuration, they are as vulnerable as the original keys to some forms of attack, such as social hacking. Another problem of using DNSSEC and autoconfiguration together is that DNSSEC requires secure time, which means secure communication with autoconfigured time servers, which requires configured secret information. Therefore, in order that the autoconfiguration may be secure, configured secret information is required.

If DNSSEC [23]-[25] is used and the signatures are verified on the client host, the misconfiguration of a DNS server may simply be denial of service. Also, if local routing environment is not reliable, clients may be directed to a false resolver with the same IP address as the true one.

6.1. RA Option

The security of RA option for RDNSS is the same as the ND protocol security [1][6]. The RA option does not add any new vulnerability.

Note that the vulnerability of ND is not worse and is a subset of the attacks that any node attached to a LAN can do independently of ND. A malicious node on a LAN can promiscuously receive packets for any router’s MAC address and send packets with the router’s MAC address as the source MAC address in the L2 header. As a result, the L2 switches send packets addressed to the router to the malicious node. Also, this attack can send redirects that tell the hosts to send their traffic somewhere else. The malicious node can send unsolicited RA or NA replies, answer RS or NS requests, etc. All of this can be done independently of implementing ND. Therefore, the RA option for RDNSS does not add to the vulnerability.

Security issues regarding the ND protocol were discussed by the IETF SEND (Securing Neighbor Discovery) Working Group, and RFC 3971 for the ND security has been published [12].

6.2. DHCPv6 Option

The DNS Recursive Name Server option may be used by an intruder DHCP server to cause DHCP clients to send DNS queries to an intruder DNS recursive name server [5]. The results of these misdirected DNS queries may be used to spoof DNS names.

To avoid attacks through the DNS Recursive Name Server option, the DHCP client SHOULD require DHCP authentication (see "Authentication of DHCP messages" in RFC 3315 [3][13]) before installing a list of DNS recursive name servers obtained through authenticated DHCP.

6.3. Well-known Anycast Addresses

The well-known anycast addresses approach is not a protocol, thus there is no need to secure the protocol itself.

However, denial of service attacks on the DNS resolver system might be easier to achieve as the anycast addresses used are by definition well known.

7. Contributors

Ralph Droms
Cisco Systems, Inc.
1414 Massachusetts Ave.
Boxboro, MA 01719
US

Phone: +1 978 936 1674
EMail: rdroms@cisco.com

Robert M. Hinden
Nokia
313 Fairchild Drive
Mountain View, CA 94043
US

Phone: +1 650 625 2004
EMail: bob.hinden@nokia.com

Ted Lemon
Nominum, Inc.
950 Charter Street
Redwood City, CA 94043
US

EMail: Ted.Lemon@nominum.com

Masataka Ohta
Tokyo Institute of Technology
2-12-1, O-okayama, Meguro-ku
Tokyo 152-8552
Japan

Phone: +81 3 5734 3299
Fax: +81 3 5734 3299
EMail: mohta@necom830.hpcl.titech.ac.jp

Soohong Daniel Park
Mobile Platform Laboratory, SAMSUNG Electronics
416 Maetan-3dong, Yeongtong-Gu
Suwon, Gyeonggi-Do 443-742
Korea

Phone: +82 31 200 4508
EMail: soohong.park@samsung.com

Suresh Satapati
Cisco Systems, Inc.
San Jose, CA 95134
US

EMail: satapati@cisco.com

Juha Wiljakka
Nokia
Visiokatu 3
FIN-33720, TAMPERE
Finland

Phone: +358 7180 48372
EMail: juha.wiljakka@nokia.com

8. Acknowledgements

This document has greatly benefited from inputs by David Meyer, Rob Austein, Tatuya Jinmei, Pekka Savola, Tim Chown, Luc Beloeil, Christian Huitema, Thomas Narten, Pascal Thubert, and Greg Daley. Also, Tony Bonanno proofread this document. The authors appreciate their contribution.

9. References

9.1. Normative References

[1] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

[2] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

[3] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[4] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, April 2004.

[5] Droms, R., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, December 2003.

9.2. Informative References

[6] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Option for DNS Configuration", Work in Progress, September 2005.

[7] Ohta, M., "Preconfigured DNS Server Addresses", Work in Progress, February 2004.

[8] Venaas, S., Chown, T., and B. Volz, "Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 4242, November 2005.

[9] Partridge, C., Mendez, T., and W. Milliken, "Host Anycasting Service", RFC 1546, November 1993.

[10] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.

[11] Lind, M., Ksinant, V., Park, S., Baudot, A., and P. Savola, "Scenarios and Analysis for Introducing IPv6 into ISP Networks", RFC 4029, March 2005.

[12] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

[13] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 2001.

[14] Bound, J., "IPv6 Enterprise Network Scenarios", RFC 4057, June 2005.

[15] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

[16] Wasserman, M., "Recommendations for IPv6 in Third Generation Partnership Project (3GPP) Standards", RFC 3314, September 2002.

[17] Soininen, J., "Transition Scenarios for 3GPP Networks", RFC 3574, August 2003.

[18] Wiljakka, J., "Analysis on IPv6 Transition in Third Generation Partnership Project (3GPP) Networks", RFC 4215, October 2005.

[19] 3GPP TS 23.060 V5.4.0, "General Packet Radio Service (GPRS); Service description; Stage 2 (Release 5)", December 2002.

[20] 3GPP TS 24.008 V5.8.0, "Mobile radio interface Layer 3 specification; Core network protocols; Stage 3 (Release 5)", June 2003.

[21] Chown, T., Venaas, S., and A. Vijayabhaskar, "Renumbering Requirements for Stateless Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 4076, May 2005.

[22] Huitema, C., Austein, R., Satapati, S., and R. van der Pol, "Unmanaged Networks IPv6 Transition Scenarios", RFC 3750, April 2004.

[23] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

[24] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.

[25] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.

[26] Kolkman, O. and R. Gieben, "DNSSEC Operational Practices", Work in Progress, October 2005.

[27] Guette, G. and O. Courtay, "Requirements for Automated Key Rollover in DNSSEC", Work in Progress, January 2005.

[28] Park, S., Madanapalli, S., and T. Jinmei, "Considerations on M and O Flags of IPv6 Router Advertisement", Work in Progress, March 2005.

Author’s Address

Jaehoon Paul Jeong (editor)
ETRI/Department of Computer Science and Engineering
University of Minnesota
117 Pleasant Street SE
Minneapolis, MN 55455
US

Phone: +1 651 587 7774
Fax: +1 612 625 2002
EMail: jjeong@cs.umn.edu
URI: http://www.cs.umn.edu/˜jjeong/

Full Copyright Statement

Intellectual Property

Acknowledgement

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