

IPv6 Operations Working Group  
Internet-Draft  
Expires: June 12, 2008

M. Stenberg  
O. Troan  
Cisco Systems, Inc.  
December 10, 2007

**IPv6 Prefix Delegation routing state maintenance approaches  
draft-stenberg-v6ops-pd-route-maintenance-00**

Status of this Memo

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

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

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

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

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

This Internet-Draft will expire on June 12, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

The maintenance of Prefix Delegation (PD) routing state is an issue that people have discussed in the IETF DHC WG, and there have been drafts on the topic. However, as the pros and cons of the different routing state maintenance solutions have not been examined thoroughly, this text attempts to shed some light on both the actual problem and the various alternative solutions.

Table of Contents

- [1.](#) Introduction . . . . . [3](#)
- [2.](#) Different approaches for maintaining routing state . . . . . [5](#)
  - [2.1.](#) Backend provisioning system responsible for routing state . . . . . [5](#)
  - [2.2.](#) Delegating router responsible for routing state . . . . . [6](#)
    - [2.2.1.](#) Approach 1: On-demand lease query . . . . . [7](#)
    - [2.2.2.](#) Approach 2: Anticipatory lease query . . . . . [7](#)
    - [2.2.3.](#) Approach 3: Persistent storage . . . . . [8](#)
  - [2.3.](#) Requesting router responsible for routing state . . . . . [9](#)
    - [2.3.1.](#) Approach 1: Layer-2 detection of link state . . . . . [9](#)
    - [2.3.2.](#) Approach 2: Keepalive . . . . . [9](#)
    - [2.3.3.](#) Approach 3: Short lifetimes . . . . . [10](#)
    - [2.3.4.](#) Approach 4: Routing protocol to the requesting router . . . . . [10](#)
- [3.](#) Security Considerations . . . . . [12](#)
- [4.](#) IANA Considerations . . . . . [13](#)
- [5.](#) Summary . . . . . [14](#)
- [6.](#) Informative references . . . . . [15](#)
- [Appendix A.](#) Acknowledgements . . . . . [16](#)
- [Appendix B.](#) Revision history . . . . . [17](#)
- Authors' Addresses . . . . . [18](#)
- Intellectual Property and Copyright Statements . . . . . [19](#)

**1. Introduction**

A prefix delegation deployment consists of Requesting Routers (RR), Delegating Routers (DR) and possibly a backend provisioning system (see Figure 1). The delegated prefix has to be routed in the network. This document explores various alternatives for how the route for the delegated prefix can be injected in the network, and how the routing state can be maintained.

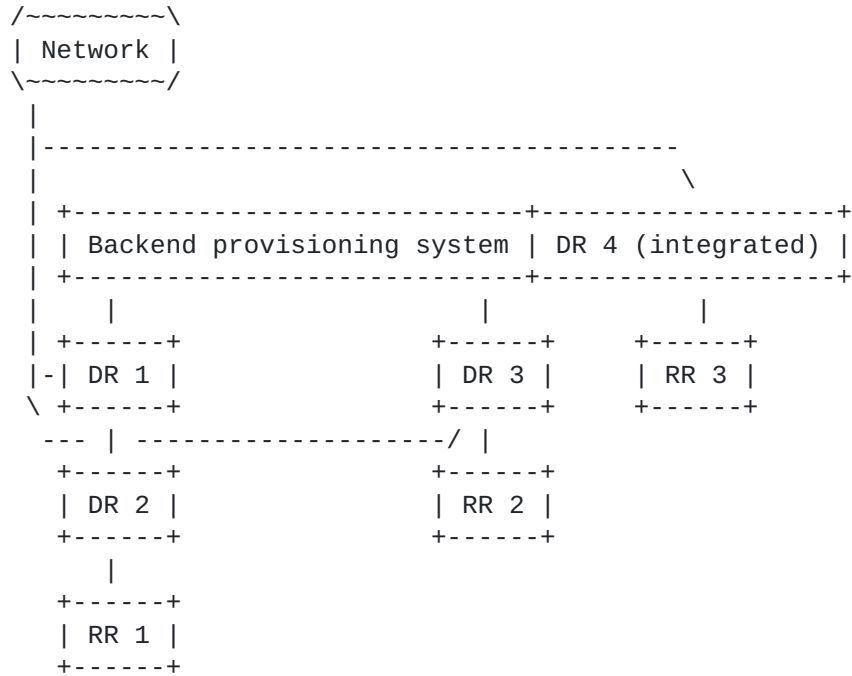


Figure 1: Possible prefix delegation deployment cases.

Prefix delegation is a stateful protocol. The RR needs to maintain state so that it can sub-delegate prefixes to downstream links. The RR maintains soft-state which can be recovered by redoing the prefix request (for example, using Dynamic Host Configuration Protocol for IPv6 (DHCPv6) [[RFC3315](#)] with the Prefix Delegation options defined in [[RFC3633](#)]).

If the DR should do route injection on behalf of the RR, it needs to maintain state. The backend provisioning system must maintain a list of prefixes delegated, as a prefix delegation is a long-lived entity (lifetime of a customer relationship, as in months or years). The backend system and DR might run on the same router. This document focuses on the case where the backend system and DR are separate and the DR has little or no persistent storage. Therefore the DR 4 case (in Figure 1) is not covered here, as it is trivial - the backend has

nonvolatile storage for the prefixes, and it can re-inject the routes when the integrated DR 4 restarts.

The DR's routing state that needs to be maintained can be divided into two distinct categories: local routing state (that is, a local RIB entry containing the next-hop and the interface the assigned prefix is connected to), and global (AS-wide) routing state which requires advertising the route via a routing protocol.

Advertising a route per delegation from the DR can be avoided if there is an aggregate prefix covering the delegation. This requires stringent address allocation procedures and prohibits an RR from moving to a different DR.

## **2. Different approaches for maintaining routing state**

As any router (or the backend system for that matter) can go offline and come back up later, it is necessary for the system to recover from these intermittent failures.

The problem is how to delegate responsibility for route maintenance to one (or more) of the three components of the system, and letting it take care of maintaining the required routing state in place for the RR's prefix.

### **2.1. Backend provisioning system responsible for routing state**

Considering the backend provisioning system is the only component in the system that actually requires significant amount of nonvolatile storage, from data system point of view it would be ideal to have the backend provisioning system responsible for maintaining the routing state as well.

It would mean that the backend provisioning system should, when DR restarts, (securely) re-inject the local or global routing state into the DR(s). Similarly, whenever first hop DR-RR link goes down, the backend provisioning system should remove the routing state for the RR.

In practise, this is infeasible as a general-purpose, standard-oriented solution - there is a significant amount of both information access and routing state configuration change methods that cannot be done with the currently available standards.

- o There is no standard way for the backend to detect when the DR is restarted, or when the first hop DR-RR link goes down without some sort of keepalive mechanism, and if the keepalive mechanism fails, there are multiple potential reasons for that which require different reactions from the backend system. This leads to implementation complexity.
- o None of the current routing protocols are suitable for altering remote router's local routing information, and therefore some protocol development would be in order for this approach to be usable. As alternative, some out of band configuration method might be used for this, such as say, XML over some secure transport, but those are also currently not defined. There are also security implications with this solution when passing administrative boundaries.
- o The backend may lack the information to identify the DR to the routing system. With multiple DRs, if the delegation protocol

does not contain everything needed to re-inject the route later on to the specific DR, the re-injection cannot be performed (without some statistic configuration, at least). For example, DHCPv6 does not uniquely identify the relays. And if interim DRs do not have backend provisioning system-addressible addresses, there is a problem. All DRs may not have global unicast addresses, and this is problematic especially in configurations spanning multiple administrative domains.

In addition to lacking the ways of retrieving/altering the DR state, there are also some issues with the backend system performing active role for the maintenance of the routing state:

- o Redundancy of DR, or the links between DR and backend system, makes it difficult for the backend system to judge the state of the DR accurately without significant extra configuration data about the deployed configuration.
- o Lack of scalability; the benefit of having 'backend' provisioning system disappears as it will need to take care of maintaining routes of every one of its DRs.

Having noted these disclaimers about "standard"-oriented ways of shifting the routing state maintenance to the backend system, it must be noted that given some proprietary solutions for the state retrieval from the DR, as well as method of pushing routing state, the backend system may be viable candidate for maintaining the routing state. This type of approach would consist of (for example) Bidirectional Forwarding Detection (BFD - [[I-D.ietf-bfd-base](#)]) or some layer-2 way for the backend provisioning system to make sure that the RR connectivity is available; if it is down for significant period of time, the route for the allocated prefix could be removed until either connectivity resumes (that is, either the RR reconnects or some intermediate network failure is repaired). If the first-hop DR requires some sort of updating whenever connectivity is restored, some way of accessing DR and querying and updating it's state would be needed.

Having the backend provisioning system responsible for the routing state of the whole system seems possible only in some walled-garden proprietary applications; for general-purpose approach, we look elsewhere, that is, delegating router/requesting router-based options.

## **2.2. Delegating router responsible for routing state**

As the DR is part of the local routing infrastructure, placing the responsibility for routing state in the DR seems sensible. With that

design decision, the next problem is when the routing information is updated after the DR restarts:

### **2.2.1. Approach 1: On-demand lease query**

In the on-demand lease query case as defined in [[I-D.ietf-dhc-dhcpv6-leasequery](#)], the routing state maintenance problem is assumed to be local, and therefore the DR will receive packets both from the network at large as well as the RR even after a loss of local state caused by a restart.

When traffic arrives to DR either from the RR, or from the network to the DR for a prefix without local routing information, the DR will perform lease query, acquire the allocated prefix, and update the routing information appropriately.

This approach, while simple to specify, has some major issues:

- o It depends on the aggregated prefix to cause the inbound traffic to wind up in the DR. This assumption may not be valid, depending on the address assignment policies of the organization. Geographical or network topological hierarchical address assignment at large seems to be a failure, and it is unclear if all deployments can really implement this.
- o It requires the incoming traffic both from the RR and the network, for which no route exists, to trigger the lease query. This has two negative side effects: it requires support from the fast path hardware in the DR, and potentially causes large amount of spurious requests to the backend provisioning system (up to the desired rate that is considered harmful to the system).
- o It requires simulated ordering of the unordered transaction stream, to ensure that the routing state is maintained correctly. The DR cannot be argued to be particularly stateless anymore.

### **2.2.2. Approach 2: Anticipatory lease query**

Anticipatory, or bulk lease query, solves the routing state problem by requesting ALL prefix information from the backend provisioning system at the DR restart time. There are two different ways: The first approach is asynchronous, that is, the old state is fetched while handling the delegation requests, requiring synchronization algorithm between the bulk data retrieved from the backend system, and the requests served during that. For synchronization, some sort of ordering of the transaction stream is needed. The second alternative is synchronous: the bulk query is performed first, and only then the RRs' requests are handled.

Bulk query has several advantages over the on-demand case:

- o No need for triggering based on either inbound or outbound traffic for the prefix.
- o If DR handles the query synchronously, we can avoid the ordering of the transaction stream and the associated complexity rising from it.
- o Given reasonable TCP transport scheme, the transfer of the state is more efficient than the on-demand case in terms of total number of packets.
- o Does not require changes to fast path hardware, as no new triggers are needed from the traffic. Instead, simple additional code in the system initialization is enough.

But, unfortunately it has also some disadvantages:

- o It causes more uneven load on the backend provisioning system than the on-demand case. If the prefix is not being actively used at the time, it will not cause traffic in the on-demand case, but it will in the bulk case.
- o Synchronization is non-trivial if the DR serves RR requests during the bulk retrieval of the data.
- o Doesn't work very well with virtual interfaces - it is hard to retrieve state at boot time if the interfaces themselves get up only at some point, and with their transient nature mapping a DUID to individual customers is difficult.

### **2.2.3. Approach 3: Persistent storage**

It is possible for the DR to store the route information to be injected either locally, or on some adjacent storage node. The clear advantage of this is the lack of traffic on the wire.

Unfortunately, it has also some problems - the data being possibly outdated due to lack of synchronization, and the management overhead when the customers for example move around would be significant.

However, in most deployment scenarios persistent storage at or near all routers is not desirable or possible in the first place, so this is listed simply for the sake of completeness.



### **2.3. Requesting router responsible for routing state**

The most interested party in the routing state of the given prefix is the RR itself; therefore, giving the responsibility for maintaining its routing state to it seemed to be idea worth considering.

Due to the operators wariness of the systems not under their direct control, even with the RR responsible for maintenance of the state, the real route injection should be handled by the DR.

The nice thing about some of the RR-oriented solutions is that they can be deployed without any changes to the rest of the infrastructure.

#### **2.3.1. Approach 1: Layer-2 detection of link state**

If the RR implementation gets notifications about the state of the link layer, it can actually detect the state of the network link going down and coming back up; performing reconfiguration to ensure that the routing state is still up seems like a trivial solution in this case.

This solution can be the best one when operating over connection-oriented media (PPPoE, L2TP) but it doesn't work on say, Ethernet without direct connection between the RR and DR.

#### **2.3.2. Approach 2: Keepalive**

If the RR doesn't have L2 way of detecting DR being restarted, it can maintain a keepalive mechanism using, for example, BFD to send self-addressed echo packets to the DR and waiting for their replies. The implementations SHOULD do this only if there is no traffic from the network within a desired period of time - see IPv6 Neighbor Unreachability Detection (NUD)'s definition of forward progress detection as a way to send keepalives only when truly necessary in [[RFC2461](#)].

Assuming sub-second round-trips (reasonable assumption in most modern network environments), the longest factor for the determining the keepalive timeout is the recovery speed of the DR (by orders of magnitude), as it can take from some seconds (hot standby) to minutes (non-HA restart, or cold standby with huge configurations). The initial keepalive timer should be some fraction of the highest delay in the system, that is, the DR recovery time. The subsequent retries if no reply is received within reasonable timeframe should be calculated based on the link delay, and jitter, to ensure that the reply is unlikely to be coming back by the time the keepalive message is re-sent.

As far as overhead is concerned, assuming the cold standby/restart taking minute(s), with a keepalive per 60 seconds for example, the QoS would remain roughly same as with faster intervals (as the DR going down would cause interruption in the routing in the order of minute(s) in any case). This value would cause overhead of 0.017pps per RR, and it is unlikely to be the straw that breaks camel's back for the DR, as the BFD handling on should on the receiver side should be simply a matter of re-forwarding the traffic within fast path, and therefore it should not require any routing plane resources to handle. With any traffic, even NUD packets should outnumber the keepalive traffic.

As far as resource utilization is concerned, this solution involves only routing plane of the RR, the data link between RR and DR, and fast path of DR which bounces the packets back. Therefore it can be argued to be fairly lightweight general-purpose solution.

### **2.3.3. Approach 3: Short lifetimes**

The current best practice for maintaining the routing state is to set short configuration lifetimes (DHCP T1/T2 values). It causes extra traffic and load on the whole DHCP infrastructure. That is because during every reconfiguration, even with the DR constantly up and running, the backend system is queried. The transaction involves all three components. Due to that, every RR will cause constant load on the backend system itself over the time, making the solution simply not scale well.

### **2.3.4. Approach 4: Routing protocol to the requesting router**

The final RR-based approach consists of the RR actually running a routing protocol; this way, the RR router can simply advertise the prefix as it receives it, and everything just works. Or not, as it may be.

The downside is the security, or complete lack of it. The DR accepting arbitrary RR-advertised prefixes (assuming no state at the DR) should be acceptable only if DR and RR are within the same administrative domain. For that case, this is probably the cleanest solution of all.

If administrative boundaries are crossed, the DR will not take prefix advertisement at face value. The DR will have extra overhead of checking the backend provisioning system for AAA purposes before actually doing anything with the prefix. This can imply look-up for validity using the prefix and the interface the advertisement came from, including the DUID or some other identifier within the route advertisement message, or using some real AAA mechanism if the

routing protocol supports one. If minimal changes to the routing protocol implementation are desired, it is also possible to ignore the advertisement itself, and just trigger on-demand lease query, thereby using the routing protocol just as an alternative keepalive mechanism the with most of the logic shoved in DR instead of RR.




Figure 2: Multihomed deployment.

There is also a case where the RR HAS to run its own routing protocol; in multihomed situation like Figure 2, with the same routable prefix advertised via two different DRs, there is no other practical way to get the system working. Of course, static route configuration is always an alternative but it is seldom desirable.

The routing-protocol-based solutions all require a significant level of trust between RR and DR; regrettably the current routing protocols are not designed with AAA (or security for that matter) in mind, and therefore when crossing administrative boundaries, the alternatives are either using them as-is as a hint that something needs to be done, or significantly extending the protocols in the AAA direction. Adding extra complexity to the DR's routing protocol implementation or configuration is not desirable in general.

Finally, the current prefix delegation solution (DHCPv6 PD) does not provide the information about which routing protocol to use, and there is no routing protocol auto-negotiation protocol. Therefore the auto-configuration of the RR with arbitrary routing protocol cannot be done currently.

### **3. Security Considerations**

The backend-oriented solution detailed in [Section 2.1](#) implies a significant level of trust between the DR and the backend provisioning system. The system's configuration is simpler if the backend provisioning system can inject arbitrary routes to the DR, but allowing injection of routes for only specific sub-prefixes of a specific prefix is considerably more secure solution. Unfortunately it requires advance configuration of the prefix(es) involved.

The delegating router-based solutions detailed in [Section 2.2](#) do not have any security issues, assuming the delegation protocol itself is secured, or can be assumed to be used only within a trusted network.

The requesting router-based solutions in [Section 2.3](#), even incorrectly implemented, at most just cause extra load to the DR. As noted in [Section 2.3.4](#), even when running routing protocol from the RR, ideally the DR should consider the advertisements only a hint at best if not part of the same administrative domain. This may not be ideal if the routing protocol information should be propagated as-is onward, as in the the multihoming cases. Unfortunately, those cases also most likely cross administrative boundaries (the requesting router being part of one domain, and connected to delegating routers in most likely more than one), the providers will not most likely trust the routing protocol to be used as-is at the delegating routers, and their complexity will increase due to the required AAA/policy checks. This is a potential security risk in a critical part of the network infrastructure.

#### **4. IANA Considerations**

As this document is informational in nature and only summarizes current best practices, it does not require action from IANA.

## 5. Summary

The backend provisioning system should not be assigned the responsibility for the maintenance of the route. As seen in [Section 2.1](#), that approach has significant obstacles without any clear benefits.

If the link layer state can be used to detect the (potential) restart of delegating router, the requesting router-based simple reconfiguration described in [Section 2.3.1](#) seems to be the best choice.

When link layer state is not available, there is no clear 'best' solution. The tradeoff seems to be between increasing the complexity of the delegation protocol and the delegating router/backend system (as described in the lease query cases in [Section 2.2.1](#) and [Section 2.2.2](#)), decreasing scalability of the system significantly by using low lifetimes for configuration (as described in [Section 2.3.3](#)), or small overhead of the keepalive (as described in [Section 2.3.2](#)). The keepalive is most likely the preferred solution among these, as it causes load mostly only to the requesting router itself.

Only in multihoming cases, given some extensions to the current prefix delegation protocol, should routing protocol on the requesting router be considered, as described in [Section 2.3.4](#). Multihoming solution itself is challenging to do securely, as noted in [Section 3](#), due to lack of AAA support in routing protocols.

## **6. Informative references**

- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", [RFC 3633](#), December 2003.
- [I-D.ietf-dhc-dhcpv6-leasequery]  
Brzozowski, J., "DHCPv6 Leasequery",  
[draft-ietf-dhc-dhcpv6-leasequery-01](#) (work in progress),  
December 2006.
- [I-D.ietf-bfd-base]  
Katz, D. and D. Ward, "Bidirectional Forwarding  
Detection", [draft-ietf-bfd-base-06](#) (work in progress),  
March 2007.
- [RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor  
Discovery for IP Version 6 (IPv6)", [RFC 2461](#),  
December 1998.

**Appendix A. Acknowledgements**

Thanks to Bernie Volz for feedback during writing of the document.  
Also thanks to the assorted people on the DHC mailing list, who  
submitted comments to the second iteration of the draft.



**Appendix B. Revision history**

October 2006: [draft-stenberg-pd-route-maintenance-00](#) - initial version (IETF67)

December 2007: [draft-stenberg-v6ops-pd-route-maintenance-00](#) - resubmitted at request of some people, with small editorial changes and clear indication of preference for the L2 or keepalive solutions.

Authors' Addresses

Markus Stenberg  
Cisco Systems, Inc.  
Shinjuku Mitsui Building, 2-1-1, Nishi-Shinjuku  
Shinjuku-Ku, Tokyo-to 1630409  
JP

Email: [mstenber@cisco.com](mailto:mstenber@cisco.com)

Ole Troan  
Cisco Systems, Inc.  
GPK02/2/250 Longwater Avenue  
Reading, England RG2 6GB  
UK

Email: [ot@cisco.com](mailto:ot@cisco.com)

## Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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

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

## Intellectual Property

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

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

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

## Acknowledgment

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