Internet Engineering Task Force
Internet-Draft
Expires: July 1, 2004

P. Savola
CSC/FUNET
Jan 2004

Simple IPv6-in-IPv4 Tunnel Establishment Procedure (STEP)
draft-savola-v6ops-conftun-setup-02.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."

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 July 1, 2004.

Abstract

This memo describes a set of operational procedures, a UDP encapsulation for configured tunnels, and one implementation mechanism to provide a very simple and straightforward way to easily manage IPv6-over-IPv4 configured tunnels between an ISP and a customer. The configured tunnels work even if the IPv4 addresses change dynamically, or are private addresses; the procedure provides at least a /64 prefix per customer and requires no administrative set-up. A simple form of NAT traversal is also supported.

Table of Contents

1. Introduction

A need for a simple mechanism to set up IPv6-over-IPv4 configured tunnels between a customer and the ISP seems to have been demonstrated in 3GPP analysis [17] as well as Unmanaged [18] and ISP analysis [19]. Most currently proposed mechanisms (like 6to4 [7] or ISATAP [8]) appear to be unnecessarily complex or otherwise problematic in these particular scenarios.

ISPs that already have access infrastructure (L2TP Access Concentrator (LAC), L2TP Network Servers (LNS), PPP Termination Aggregators (PTA). etc.), IPv6/PPP/L2TP/UDP/IP could be readily provided using L2TP [9] and IPv6 over PPP [10] as long as the customer operating systems also support these mechanisms. This approach, however, is not suitable for 3GPP and Enterprise environments. See Appendix A for a more detailed comparison.

This memo documents a set of operational procedures which require no additional protocol specification to provide a very simple and suitably elegant solution to these problems.

One observation made prior to designing the procedure was that a signalling protocol is not really needed if the existing mechanisms for e.g., optional prefix delegation are used, and the ISP can authenticate the user otherwise; this simplifies the procedure significantly.

The second section gives a brief problem statement which also describes the applicability of the solution. The third section explains the overview of the procedure. The fourth and the fifth sections describe the customer- and ISP-side procedures in more detail. The sixth section describes issues related to a simple form of NAT traversal, and specifies how to optionally encapsulate IPv6-over-IPv4 packets over UDP.

In appendix A, we compare the mechanism to several other proposed mechanisms and techniques: pure configured tunnels, the use of Layer 2 Tunneling Protocol (L2TP [9]), use of 6to4 [7], an instance of Tunnel Broker concept -- TSP [11], ISATAP [8], and Teredo [12].

2. Problem Statement

There are ISPs which are willing to provide IPv6 connectivity to their customers, but may not be able to do it natively due to a number of reasons. Such ISPs want to find a method to help in providing IPv6-over-IPv4 tunnels to the customers, with the following characteristics:

- The IPv4 address of the customer may be either static or dynamic, and may be a private address [6] as well, if the customer chooses to NAT the (public) IP address given by the ISP.

- The ISP may want to offer the tunnel service either requiring prior agreement with the user, or to every customer who wishes to try it.

- The customer may have one or more nodes which should obtain IPv6 connectivity.

- The configured tunnel may be set up either from the customer's gateway, or if the gateway does not support IPv6, from a node inside the customer's network, when NAT traversal is used. No more than one node behind a NAT'ed public IPv4 address needs to participate in the IPv4-in-IPv6 tunnel service (but many more can use the IPv6 service, of course).

- The solution should be as simple as possible, requiring no new protocols or substantial modifications to IPv6 or IPv4 implementations either at the ISP or customer side.

2.1 Non-problems

The problem statement explicitly excludes:

- Support for third party ISPs: the methods described here work to an extent with a lower amount of security even if the ISP providing the service is not the user's own ISP. Typically, the third party ISP would have to be able to authenticate the user somehow; this could be done using a static IPv4 address (rather insecure), IPsec Security Association, or an unspecified mechanism. However, third party ISPs are not considered an important scenario for the IPv6 deployment, and are considered out of scope.

- More complex forms of NAT traversal: the case where the tunnel endpoint is visible (from the ISP point of view) behind a public IPv4 address, and no other tunnel endpoints are using that address is in scope. However, the case where multiple nodes would want to initiate a tunnel from behind a "big" NAT, which maps them all to a single address, is defined out of scope. The customer which has multiple nodes can still use IPv6 behind such a NAT by selecting one of the nodes to provide IPv6 access through the tunnel, and have IPv6 connectivity routed or proxied as normal by the tunnel endpoint node. The case where the ISP has deployed a single "big" NAT affecting many customers can be addressed by the ISP deploying

the tunnel router inside the privately addressed infrastructure (remember that third party ISPs were out-of-scope as well), so that no NAT traversal is needed in the first place, as the connectivity to the ISP's tunnel router is native IPv6 or a configured tunnel with a static public IPv4 address.

- Short-cut paths between the users (e.g, like 6to4 [7] or ISATAP [8]): all the IPv6-over-IPv4 traffic flows through the tunnel router; short-cut mechanisms are believed to be non-essential in this environment of "short" tunnels, and add to complexity and security risks. If the load on the tunnel router rises too high, one could switch to offering native service instead, or deploying additional tunnel routers.

3. Overview of the Procedure

Throughout this memo, two major operational modes, "managed" and "ad-hoc" are described. It's expected that some ISPs would like to use one, and some the other, and both approaches are described.

The procedure can be summarized as follows:

1. If the ISP requires prior agreement ("managed mode"), the customer contacts the ISP off-band and registers as an IPv6 user.

2. The customer discovers (using one of a number of mechanisms) the IPv4 tunnel end-point address of the ISP, and creates a configured tunnel (encapsulating in either IP (protocol 41) tunnel [1] or UDP (Section 6)) to the address, and sends a normal Neighbor Discovery [2] (ND) Route Solicitation (RS) or a DHCPv6 [4] SOLICIT or prefix delegation request [5] message over the tunnel.

3. The ISP's tunnel router sets up a configured tunnel towards the customer's IPv4 address; the address may be obtained using a number of mechanisms, or created ad-hoc ("ad-hoc mode") when tunnel packets arrive. In the managed more, the tunnel interface is typically pre-configured prior to receiving any packets from the customer.

4. The ISP's tunnel router sends a normal ND Route Advertisement (RA) or a further DHCPv6 message over the tunnel to the customer; the prefix advertised is obtained using one of a number of mechanisms. The customer automatically configures the prefix and the addresses and uses them normally.

Note that the description includes DHCPv6, prefix delegation etc.

just for completeness. It is assumed than in most cases a simple ND RS/RA exchange will suffice. However, as the procedure is agnostic of the prefix assignment methods used, any other mechanism can be used as well.

No new protocols are needed. Both in the managed and ad-hoc modes, the customer can learn the tunnel address off-band.

In the managed mode, the ISP has to know the IPv4 address assigned to the customer, configure a new IPv6 tunnel interface for the customer, and reserve the IPv6 prefix that will be assigned; these have to be configured on the tunnel router using operator-specific management techniques (e.g, RADIUS).

In the ad-hoc mode, on the other hand, the tunnel router has to implement a simple mechanism to allocate a new configured tunnel, after successful validation (or authentication) procedures as discussed in Section 5.5, for tunnel packets received from different customers, and algorithmically derive an IPv6 prefix to be assigned to the customer.

4. Customer-side Procedures

4.1 Possible Prior Agreement with the ISP

The ISP may require prior agreement or notification before a customer is allowed to use their tunnel service. In that case, the customer must contact the ISP using off-band mechanisms. Even if not required, special requirements (e.g., a static IPv6 prefix when IPv4 address is dynamic, or a need for an IPv6 /48 prefix) may be easier to fulfill if the user has contacted the ISP beforehand and the ISP has made arrangements; only a /64 prefix (which will be dynamic if the IPv4 address is dynamic) will be available in ad-hoc mode.

4.2 Learning and Configuring the Tunnel Endpoint

To get started, the customer has to learn the IPv4 address of the ISP's tunnel router somehow. Possibilities include, for example:

- Using off-band mechanisms, e.g., from the ISP's web page.
- Using DNS to look up a service name by appending it to the DNS search path provided by DHCPv4 (e.g. "tunnel-service.example.com").
- Using a (yet unspecified) DHCPv4 option.
- Using a pre-configured or pre-determined IPv4 anycast address,

whether in the private or public space; however, note the considerations about embedding addresses in the nodes [14].

- Using other, unspecified methods.

This memo does not (at least yet) take a stance on the selection of the mechanism even though some are more problematic than others, but it is assumed that the first or the second option should be enough for everyone considering that the customer's own ISP is providing IPv6 service.

Once the IPv4 address has been learned, it is configured as the tunnel end-point for the configured IPv6-over-IPv4 tunnel. Unless the user has a private IPv4 address, implying being behind a NAT, IP encapsulation must be used; otherwise, the encapsulation can be selected as described in Section 6. Note that this configuration can even be done transparently to the user, with very little or no configuration.

4.3 Tunnel Activation

Next, IPv6 is activated over the tunnel as normal; this could be done either by a Neighbor Discovery RS, DHCPv6 Solicitation message, DHCPv6 Prefix Delegation request message, or by simple manual configuration (note: manual configuration does not work in the "ad-hoc" operation, because there is no trigger to bring up the ISP's interface).

The tunnel router responds to this query as normal by sending a Route Advertisement or continuing with DHCP message exchanges.

4.4 Providing Connectivity to Other Nodes

If the customer has multiple nodes, they can each obtain their own tunnel in the same manner as long as the nodes are not behind a NAT. However, this is unoptimal especially if such nodes have internal communications.

Instead, the customer may want to set up one node to as a Neighbor Discovery proxy [15] for the /64 route advertisement received, or if a less specific prefix (e.g., a /48) is being used, as a router for the internal network(s). This does not need to be any more complicated than just setting up the tunnel on one node.

5. ISP-side Procedures

5.1 Possible Prior Agreements with the Customers

The ISP may operate in either or both "managed" and "ad-hoc" modes. In only the managed mode, a prior agreement with the customer is needed to allow the customer to use IPv6 using this procedure. In only the ad-hoc mode, no agreements with the customers are needed. In both modes, "basic service" (a /64 prefix which will be dynamic if the customer's IPv4 address is dynamic) can be offered, but more advanced services (e.g., prefix delegation or a static prefix) are offered to those with a prior agreement.

ISPs which operate in the managed mode must configure (e.g., manually or using a script or configuration tool) the configured tunnels on the tunnel router; also, they may want to create a link between the stable customer identification and their IPv6 properties (e.g., a prefix) especially if the IPv4 address is dynamic, to maintain the stability of IPv6 properties even when the IPv4 address may change.

5.2 Learning the Customers' Tunnel Endpoint Addresses

The ISP must somehow obtain the tunnel endpoint address to be configured for a configured tunnel. Every active customer has its own configured tunnel interface on the tunnel router.

When operating in the managed mode, this could be done from e.g. DHCPv4 leases, RADIUS or Diameter databases, other databases or some other means. This information will be used to update the tunnel end-point address on the configured tunnel interface when changing or as appropriate; the updates can be done e.g. using management tools, scripting, etc. -- because a change of IPv4 address must be reflected without delay to the tunnel end-point address, this configuration update should be immediately triggered by changes in the used database or lease.

When operating in the ad-hoc mode, the tunnel server should create a new configured tunnel interface for each IPv6-over-IPv4 tunnel with a different IPv4 source address. The ISP should be aware of a potential for a resource exhaustion if the number of customers rises too high, but actual DoS attacks are not possible if the ISP has secured its network as described in Section 5.5. Automatic creation of configured tunnel interfaces requires only rather trivial implementation [XXX: does this need elaboration?]. Performing "garbage collection" on such tunnels, e.g. in a Least-Recently-Used (LRU) manner may be called for if the number of tunnels rises too high. However, this should only be done after sufficiently long period has passed, as not to disturb the existing (but maybe dormant) IPv6 connections over the tunnels.

If the ISP wants to support NAT traversal when protocol 41 forwarding [16] is not implemented in all the NAT boxes used by the customers, the ISP must also provide the support for UDP decapsulation at UDP port TBD. The users should default to use protocol 41, but if the initiating packet is encapsulated in UDP, the configured tunnel type may be changed if supported. NAT traversal is further discussed in Section 6.

5.3 Prefix Advertisement or Delegation

Each customer should be provided with at least a /64 prefix; this is both practical (because /64 is required by Stateless Address Autoconfiguration [3]), and architecturally correct (providing the possibility to connect more than one node without an IPv6 NAT).

In the managed mode, the ISP may advertise a static or dynamic IPv6 /64 prefix using RAs, provide a prefix delegation, or something else (e.g., manual configuration if the IPv4 address is static).

In the ad-hoc mode, the ISP must ensure that a sufficiently large pool of /64 prefixes are available. The prefixes can be allocated either in a sequential fashion and advertised in RA's, or automatically calculated, with some assumptions, from the used IPv4 addresses. For example, if the ISP uses IPv4 network 10.0.0.0/8 for its customers, it needs 24 bits to uniquely identify each customer -- this calls for assigning an IPv6 /40 prefix to be used for advertising /64's; in this example, a customer with address "10.1.2.3" might get advertised an IPv6 prefix "2001:db8:FF01:0203::/64", where "01:0203" corresponds to the client address and "2001:db8:FF00::" the /40 allocated to the ad-hoc tunneling operations by the ISP. Mapping the most interesting bits (for the ISP) of an IPv4 address to the IPv6 prefix allows even large ISPs to easily give each user an algorithmically derived IPv6 prefix.

5.4 Tunnel Activation and Maintenance

When the router receives e.g. ND RS, DHCPv6 SOLICIT or prefix delegation request from the configured tunnel, it responds normally, as on any other interface. (When in ad-hoc mode, setting up the tunnel from the received IPv6-over-IPv4 packet may take a while, but the processing continues when set up.)

The ISP should avoid sending periodic messages (e.g., unsolicited route advertisements) to the tunnel, or decrease the interval used for sending them: if the customer disconnects for some time, and someone else gets the same address, it might be disturbing to the new, potentially non-IPv6 aware customer to receive "weird" protocol 41 or UDP packets meant to the previous customer. The similar effect

occurs if someone in the Internet is trying to communicate with an IPv6 user, but the user has changed its address in the meantime, and packets may go to someone else's IPv4 address. However, this is no different to the situation with IPv4 today, except that the packets may be discarded by the operating system and never even be noticed; but if they are noticed, e.g., by a personal firewall, they may not be recognized and may cause more alarm.

5.5 Secure Operations for Tunnel Service

The ISP should perform IPv4 ingress filtering at its borders towards peers and upstreams, by disallowing packets with the source addresses belonging to its own site or its customers. In particular, the ISP must block the tunnel router's address from being used as a source address from the outside; blocking the use of the customer prefixes would be preferred as well.

The ISP must perform IPv4 ingress filtering towards the customers, in particular those that use the tunnel service, so that they will not be able to forge the IPv4 source address of the packets. In particular, they must not be able to spoof the address of the tunnel router to the other customers.

Both of these are very simple operations especially in the minimal case of blocking only the abuse of the tunnel router address.

Naturally, the ISP should perform IPv6 ingress filtering as well, but that is orthogonal to the security of this procedure.

In addition, the ISP must ensure, especially if in ad-hoc mode, that only a selected subset of source addresses is able to communicate with the tunnel router's designated tunnel address. For example, creating dynamic interfaces with packets from outside of the ISP's network could easily be used in a resource exhaustion attack. In addition, to curtail internal resource exhaustion attacks, it makes very much sense to ingress filter all the customers which are allowed to use the ad-hoc tunnel service. With these precautions, resources may only be exhausted by a real resource starvation, not through an attack; on the other hand, if the ISP does not bother to add such checks, it only harms itself for being susceptible to various forms of attacks!

5.6 Sufficient Tunnel Service Provisioning

The ISP must naturally ensure that the tunnel router is capable to handle the amount of users and the traffic that goes through it. It should also be noted that all the traffic between the users of the ISP go through the same router; "shortcuts" routes are not deemed

necessary. The increases in the latency are not significant as the tunnel router is deployed close to the IPv4 access router (or even co-located with it) topology-wise.

Typically, these are not believed to be problems. If the number of users or the amount of traffic generated increases, starting deploying native IPv6 access instead eliminates the problem, or the ISP could deploy more tunnel routers in a load-balancing configuration -- depending on the mechanism used to find the tunnel service, this could be e.g., through DNS load-balancing, anycasting the tunnel service address, etc.

6. NAT Traversal

NAT traversal may be desirable especially in the case when the customer gets one IPv4 address from the ISP, which is assigned on the IPv4 gateway, and NAT'ed access is provided to the customer's nodes. If the gateway cannot be IPv6-enabled, the customer may want to obtain the access from an internal node to bypass the gateway and the NAT.

There are two ways to do this: (1) ensuring that the NAT forwards protocol 41 packets [16], or (2) providing a minimal UDP encapsulation to the tunnel packets.

Forwarding protocol 41 packets is simple if implemented by the NAT gateway; this requires no administrative set-up. This works (basically) for one node behind the NAT at the time -- however, multiple nodes behind a single public IPv4 address was considered out of scope.

If protocol 41 packets are not forwarded, a minimal UDP encapsulation may be needed: instead of using protocol 41, UDP is used with the minimal headers. This adds 8 bytes to the packets, and should be taken into consideration with configured tunnel MTU calculations [1]. The routers should use a UDP port TBD by default to de-multiplex UDP configured tunnels. The tunnels are identified in SNMP by udp(8) IANAtunnelType [20].

The customer using private addresses behind a NAT must select which method to use. It is recommended to try protocol 41 first; if no response is received, UDP encapsulation may be tried instead. Note that this choice of the encapsulation may be completely transparent to the user as well.

With UDP encapsulation, the algorithmically derived prefix assignment is kept simple with the assumption that every tunnel service user can be identified with a public IPv4 address; whether the tunnel

originates inside a NAT does not matter. That way, for UDP-encapsulated configured tunnels, the tunnel router only additionally needs to keep a record of the source UDP port each customer uses.

However, NAT mappings must be maintained whether UDP or protocol 41 is being used. It is recommended to do this by the customer activating ND Neighbor Unreachability Detection (NUD) on the configured tunnel; the default value for DELAY_FIRST_PROBE_TIME is 5 seconds [2]; this is enough, and could even be increased to e.g., 60 seconds for this link type. Increasing it further may have adverse effects as the NAT UDP/proto-41 mapping lifetimes typically vary from 60..200 seconds. No protocol is proposed to discover and use the most optimal lifetimes for the particular NAT; this is not believed to be worth the robustness losses.

7. Acknowledgements

This procedure was inspired by a need to severely simplify ISATAP [8]. Suresh Satapati coined up the name, and provided useful feedback. Gert Doering, Marc Blanchet and Janos Mohacsi participated in the discussion clarifying the applicability.

8. IANA Considerations

This memo requests an allocation of a "privileged" UDP port (TBD).

9. Security Considerations

The requirements for reasonably secure operations within an ISP are described in Section 5.5; with these in place, it is difficult to imagine a case where stronger mechanisms such as IPsec for IPv6-over-IPv4 tunnels would be needed.

A particular case occurs when an IPv4 address of the user changes, and the user's IPv6 prefix changes as well; this may be allocated to a different IPv6 user. However, this is no different than IPv4 address re-use threats. [XXX: can be considered more if really needed.]

When the ISP operates in the ad-hoc mode, and there is an event where all the IPv4 addresses change simultaneously, there may be a large number of simultaneous updates to update the tunnel point addresses in the tunnel router. This situation should be taken into consideration e.g. if renumbering.

The case where a third party ISP provides the service was decreed out of scope, because it is impractical and economically unfeasible, and

has a number of security problems as well. Similarly, multiple users behind a single NAT'ted public IPv4 address seems to be only relevant in the third party case and is equally out of scope; this would have security implications as well, as it would be relatively easy to hijack someone else's IPv6 prefix.

Normative References

[1] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", draft-ietf-v6ops-mech-v2-01 (work in progress), October 2003.

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

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

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

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

Informative References

[6] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G. and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[7] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.

[8] Templin, F., Gleeson, T., Talwar, M. and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", draft-ietf-ngtrans-isatap-17 (work in progress), January 2004.

[9] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G. and B. Palter, "Layer Two Tunneling Protocol "L2TP"", RFC 2661, August 1999.

[10] Haskin, D. and E. Allen, "IP Version 6 over PPP", RFC 2472, December 1998.

[11] Blanchet, M., "Tunnel Setup Protocol (TSP)A Control Protocol to Setup IPv6 or IPv4 Tunnels", draft-vg-ngtrans-tsp-01 (work in

progress), July 2002.

[12] Huitema, C., "Teredo: Tunneling IPv6 over UDP through NATs", draft-huitema-v6ops-teredo-00 (work in progress), June 2003.

[13] Massar, J., "SixXS Heartbeat Protocol", draft-massar-v6ops-heartbeat-00 (work in progress), January 2004.

[14] Plonka, D., "Embedding Globally Routable Internet Addresses Considered Harmful", draft-ietf-grow-embed-addr-00 (work in progress), December 2003.

[15] Thaler, D. and M. Talwar, "Bridge-like Neighbor Discovery Proxies (ND Proxy)", draft-thaler-ipv6-ndproxy-01 (work in progress), October 2003.

[16] Palet, J., "Forwarding Protocol 41 in NAT Boxes", draft-palet-v6ops-proto41-nat-03 (work in progress), October 2003.

[17] Wiljakka, J., "Analysis on IPv6 Transition in 3GPP Networks", draft-ietf-v6ops-3gpp-analysis-07 (work in progress), October 2003.

[18] Huitema, C., "Evaluation of Transition Mechanisms for Unmanaged Networks", draft-ietf-v6ops-unmaneval-00 (work in progress), June 2003.

[19] Lind, M., "Scenarios and Analysis for Introducing IPv6 into ISP Networks", draft-ietf-v6ops-isp-scenarios-analysis-00 (work in progress), December 2003.

[20] Thaler, D., "IP Tunnel MIB", draft-thaler-inet-tunnel-mib-00 (work in progress), October 2003.

Author's Address

Pekka Savola
CSC/FUNET

Espoo
Finland

EMail: psavola@funet.fi

Appendix A. Comparison to Other Mechanisms and Procedures

This mechanism can be compared to several other proposed mechanisms and proposals: pure configured tunnels, the use of Layer 2 Tunneling Protocol (L2TP [9]), use of 6to4 [7], an instance of Tunnel Broker concept -- TSP [11], ISATAP [8], and Teredo [12].

Since obtaining IPv6 connectivity without the support of your own ISP is out-of-scope, we exclude 6to4 and Teredo from the comparison. Now's let's take a look at the rest.

A.1 Configured Tunnels

Configured tunnels are preferable in every case where they can be used. However, it's difficult to manage them especially in the cases where dynamic (but public) IPv4 addresses are being used, when the user needs IPv6 connectivity to nodes behind the user's own NAT gateway (which doesn't implement protocol-41 forwarding), or when the amount of configuration must be kept to the minimum.

A.2 L2TP

L2TP could be leveraged by using UDP (passes NATs) to encapsulate PPP frames toward the customers. The customers would have to have an L2TP client and IPv6-capable PPP, and the ISP would have to have an L2TP server, a management system for the IPv6 attributes (e.g., RADIUS), and a configured address pool. Additionally, as IPV6CP PPP negotiation does not allow prefix delegation, DNS resolver configuration, etc., one might have to run (especially if more than one address is required) an additional protocol, e.g. DHCPv6 for prefix delegation, on the link.

The ISPs which already have L2TP, PPP and RADIUS infrastructures (e.g., for dial-up IPv6 users, or certain classes of xDSL users), the additional set-up complexity would not be high; for those which do not, this would be a rather complicated set of operations. Naturally, the customer operating systems would have to support L2TP and PPP as well.

L2TP clearly has its strenghts, but some ISPs might see it as too complicated to set up. Also, if the ISP wishes to offer an "ad-hoc" operations (as seems to be the case in 3GPP at least), the amount of infrastructure required might be too high.

A.3 Tunnel Broker Solutions

A large number of custom tunnel brokering solutions have been used. For example, many (open-source) VPN products offer the capabilities

of providing IPv6 service through the VPN. Multiple tunnel brokers have also been deployed, e.g. by SixXS [13], Viagenie, and others. We'll look at Tunnel Setup Protocol (TSP) as an instance of this model.

TSP provides a similar set of functions as STEP. However, TSP has an overhead of a signalling protocol, which is not needed in STEP. TSP offers a custom way of prefix delegation, while STEP relies on standard mechanisms like DHCPv6 Prefix Delegation, or the use of ND proxying. TSP works also for tunnel configuration across ISPs, which was out of scope for STEP. STEP is transparent to the user, but TSP requires a client software and some form of set-up.

A.4 ISATAP

ISATAP provides many features, like automatic tunneling between ISATAP nodes in the same ISATAP site, which was decreed insecure and out of scope for STEP. The insecurity rises from the applying ISATAP in scenarios where the bounds of an ISATAP site are larger than the bounds of an administrative domain, leading to e.g., issues with the trust of the pseudo-interface when a packet with Hop Limit=255 and a link-local address is received. STEP has no such assumptions, and it's security properties are about the same as using bidirectional configured tunnels.

Appendix B. Multiple Users Behind a NATted IPv4 Address

The scenario where multiple users are behind a single NAT'ed IPv4 address (e.g., when using a third party ISP) was decreed out of scope. However, the possibility to achieve that is shown, even though it is not considered to be useful.

Providing service to multiple users would require nothing more than a change in the algorithm used to derive the customer's prefix. Of course, at first glance this appears to be problematic, as mapping 16 additional bits to the IPv6 address may seem like a challenge. However, this is not the case; such support is only needed for (out-of-scope) third party ISP case, which must operate in the managed mode, and the prefix assignment cannot be done based on the address anyway, so there is no real problem with this approach.

Intellectual Property Statement

Full Copyright Statement

Acknowledgment

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