Network Working Group
Internet-Draft
Expires: April 1, 2005

F. Parent
Hexago
A. Durand
SUN Microsystems,inc.
A. Baudot
France Telecom R&D

Oct 2004

Goals for Registered Assisted Tunneling
draft-ietf-v6ops-assisted-tunneling-requirements-01

Status of this Memo

Copyright Notice

Abstract

This document defines requirements for a tunnel set-up protocol that could be used by an ISP to jumpstart its IPv6 offering to its customers by providing them IPv6 connectivity through tunneling.

Table of Contents

1. Goal and Scope of the Document

The v6ops working group has worked on requirements and scenarios for IPv6 deployment by soliciting input from network operators. This work has identified a need for an "assisted tunneling" mechanism. For example, an ISP starting its IPv6 offering to its customers without upgrading its access network to support IPv6 could use a "tunnel brokering solution" (section 5.1 [I-D.ietf-v6ops-isp-scenarios-analysis]) a la [RFC3053]. What has been identified as missing from that document is a tunnel set-up protocol.

In an ISP network, getting IPv6 connectivity to the customers involves upgrading the access network to support IPv6, which can take a long time and/or be costly. A tunneled infrastructure can be used as a low cost migration path (section 5.1 [I-D.ietf-v6ops-isp-scenarios-analysis]).

With such an infrastructure, the ISP can connect its customers to its IPv6 network using its production IPv6 address space, thus facilitating migration towards native IPv6 deployment. The IPv6 deployment roadmap for connecting customers may become:

- assisted tunneling infrastructure to early adopters,
- native IPv6 to customers where economically justified,
- native IPv6 to all customers.

Contrary to automatic tunneling mechanism where the IPv4 address is embedded inside the IPv6 address, no special format are imposed on the IPv6 address used in assisted tunneling. Prefix delegation is also possible. As the addressing space used during the transition to native remains the same, the customer routing, filtering, accounting stay the same, and there is no need to maintain any kind of relay.

"Assisted tunneling" is used in this document to describe a transition mechanism where the parameters to configure a bi-directional tunnel between an end-node (or leaf network) and a router in the core of an ISP are exchanged through a tunnel set-up

protocol. Although this exchange can be automated, this remains different from transition mechanisms like 6to4, Teredo or ISATAP. In particular, assisted tunneling enables explicit access control to the tunneled IPv6 connectivity, where the other transion mechanisms have to rely on other kinds of control (e.g., access control based on the IPv4 address). Also, assisted tunneling protocol negotiate the tunnel parameters and does not depend on having the IPv4 address inside the IPv6 address, for example.

This document analyzes the requirements for such a tunnel set-up protocol. The v6ops WG scenario and evaluation documents for deploying IPv6 within common network environments are used as input to this document.

2. Applicability

Assisted tunneling is applicable in different IPv6 transition scenarios. The focus of this document is to define the requirements to apply this mechanism in the IPv4 ISP context making the following assumptions:

- ISP is offering IPv6 connectivity to its customers initially using controlled tunneling infrastructure [I-D.ietf-v6ops-isp-scenarios-analysis], section 5.1 "Steps in Transitioning Customer Connection Networks"

- Provider network where deployment of IPv6 is done in a more controlled manner or when the provider cannot rely on IPv4 related authentication (e.g. roaming customers, users not connected to ISP access network). In such networks, ease of debugging, traceability, filtering and so on are important features.

- The customer configuration may be diverse, and not necessarily predictable by the ISP. The protocol must be able to adapt to the following cases, for example by choosing the most optimal tunnel encapsulation depending on the presence of a NAT.

  * a single node,

  * a leaf network,

  * using a globally routable IPv4 address,

  * behind a NAT (customer or ISP owned),

  * using dynamic IPv4 address (internally or externally to the NAT)

There are actually two cases where the IPv4 address of the customer tunnel end point can be dynamic, and both must be supported:

- The device used as tunnel end point is using a dynamic IPv4 address provided by the ISP.

- The device used as tunnel end point is located behind a customer owned NAT box that is also acting as a local DHCP server. In that case, the device IPv4 address may change after a reboot.

Althought the main focus of this document is the ISP scenario, assisted tunneling is applicable in other scenarios:

In unmanaged networks [RFC3904], assisted tunneling is applicable in the case A where a gateway does not provide IPv6 at all (section 3), and case C where a dual-stack gateway is connected to an IPv4-only ISP (section 5).

In the enterprise scenario [I-D.ietf-v6ops-ent-analysis], assisted tunneling can be used to support remote users connecting to the enterprise network (section 7.5.2).

3. Requirements for Simplicity

The tunnel set-up protocol must be simple to implement and easy to deploy. In particular, it should not depend on any complex, yet to be designed, protocols or infrastructure pieces.

This protocol is a transition mechanism, thus does not need to be perfect. As a matter of fact, making it perfect would be counter productive, at it would first delay its definition, then make its deployment more cumbersome and, last but not least, diminish the incentives to deploy native IPv6.

4. Protocol Requirements

Assisted tunneling is aimed at production deployment which requires the user to be authenticated (such as using a shared secret mechanism Section 4.3). This can be to offer the tunneling service to roaming users (users that are not directly connected to the ISP access network), and/or restrict the service to specific users.

From a user point of view, an initial registration process may be required before using the service. If the service provider uses an existing authentication database (Section 4.2), this step may not be needed.

From an ISP point of view, this makes a clear link between a tunnel and a user (account). It provides some means to :

- meet requirements such as tracing (section 5.4 [I-D.ietf-v6ops-isp-scenarios-analysis])

- control service provision to valid and/or identified or even selected users (section 5.2 [I-D.ietf-v6ops-isp-scenarios-analysis])

- less prone to denial of service attacks: Since every tunnel

request are authenticated, it is more difficult to request multiple tunnels to saturate the service.

- stay in touch with users

4.1 Address and Prefix Delegation

The protocol must support delegation of an IPv6 prefix. The length of the IPv6 prefix delegated must be configurable on the server. The protocol must be able to offer stable IPv6 prefixes to the authenticated customers.

Assignment of an IPv6 address (/64) to the end-node must be supported.

Since no special address format is imposed, the ISP's address space can be used in the delegation (section 5.1 [I-D.ietf-v6ops-isp-scenarios-analysis])

4.2 Registration

The registration of credentials is external to the protocol. The user may require registration prior to using this service (through some web based service or other means). Or service provider may use an existing authentication database to pre-register its users.

In order to allow a service provider to use its existing authentication database, an implementation may provide hooks to facilitate integration with the ISP management infrastructure (e.g. RADIUS for AAA, billing) ([I-D.ietf-v6ops-isp-scenarios-analysis], section 5.2).

The protocol may send information about registration procedure when a non-registered client requests registered mode (ex: URL to provider registration web page).

4.3 Authentication

Authentication can be used to control user has access to the IPv6 services (section 5.2 [I-D.ietf-v6ops-isp-scenarios-analysis])

The authentication mechanism supported should be compatible with standardized methods that are generally deployed. In order to assure interoperability, at least one common authentication method must be supported. Other authentication may be supported and should be negotiated between the client and server (e.g., SASL [RFC2222]).

4.4 Confidentiality

Assisted tunneling can be used across networks which are not under the service provider control (e.g., roaming users). The tunnel set-up protocol should allow protection of the authentication data Section 4.3. This can be achieve by selecting an authentication mechanism that protects the credentials (e.g., digest-md5).

Protecting the tunneled data (IPv6 in this case) should be possible. A possible usage scenario is when an enterprise's users is working off-site and tunneling to the enterprise network (7.5.2 [I-D.ietf-v6ops-ent-analysis]). Mechanisms do exist to make this possible, such as using IPsec over IPv6 [RFC2401]. [I-D.tschofenig-v6ops-secure-tunnels] may be applicable here but is not analysed further.

4.5 Service Discovery

In order to facilitate deployment, the implementation should allow a mechanism to discover the address of the server that will provide the tunnel connectivity.

This discovery should be automatic when the protocol is used within an ISP network. There is no service discovery requirements when used outside the provider network (roaming users, 3rd party ISP).

Tunnel end-point discovery mechanism work ([I-D.palet-v6ops-tun-auto-disc] may be applicable here.

4.6 NAT Traversal

Tunneling through IPv4 NAT must be supported. The protocol should detect if an IPv4 NAT is in the path during the set-up phase (Section 5.2). If a NAT is present, an extra level of encapsulation is necessary to tunnel IPv6 across the NAT. If no NAT is detected, IPv6-over-IPv4 tunneling (IP protocol 41) is usually enough (see also Section 4.7).

NAT traversal is identified as a requirement in ISP scenarios (section 5.1 [I-D.ietf-v6ops-isp-scenarios-analysis]) and unmanaged

scanarios (section 7, Recommendation 1 [RFC3904])

4.7 Firewall Traversal

Even if no NAT is in the tunnel path, there may be a firewall which prohibits IP protocol 41. In such case, the tunnel encapsulation selection based on NAT detection (Section 5.2) will select a tunnel that will not work.

In some cases, when the firewall is managed by the ISP or customer, it can be configured to allow IP protocol 41. In such cases this may not be an issue (section 5.1 [I-D.ietf-v6ops-isp-scenarios-analysis])

But in order to be functional in any situation (e.g., firewall lacking feature), the assisted tunneling implementation must allow a user to explicitly specify the desired tunnel encapsulation, regardless of the NAT detection process.

4.8 Accounting

The assisted tunneling should include tools for managing and monitoring the provided service. Such information can be used to plan service capacity (traffic load) or billing information.

Some useful accounting data are (not exhaustive list):

- Tunnel counters (traffic in/out)
- User utilization (tunnel uptime)
- System logging (authentication failures, resource exhaustion, etc.)

The interface used to provide such information can be through SNMP or an AAA protocol (e.g., RADIUS accounting).

5. General Requirements

5.1 Scalability

The tunnel set-up protocol must be scalable. Typically, this protocol should be deployable in an ISP or enterprise network.

A scalability requirement which is not related to the protocol itself is to be able to deploy multiple servers inside the ISP network. To do so, the server implementation would possibly need some load

balancing feature and an IPv6 IGP.

5.2 NAT Considerations

The assisted tunnel established by the protocol to be designed must work with the existing infrastructure, in particular it must be compatible with the various customer premise equipments available today. This means that, in particular, the tunnels must be able to traverse one or many NAT boxes of different kinds. There is no requirement for any particular NAT traversal technology. However, as NAT traversal usually requires an extra layer of encapsulation, the

tunnel set-up protocol should be able to detect automatically the presence of one or more NAT boxes in the path.

The implementationÂ must provide an option to turn on extra encapsulation manually. In order to assure interoperability, at least one common tunnel encapsulation type must be supported.

5.3 Keep-alive

When a tunnel has to cross a NAT box, the mapping established by the NAT must be preserved as long as the tunnel is in use. This is usually achieved by sending keep alive messages across the tunnel. Also, the same keep alive messages can enable the ISP tunnel end point to perform garbage collection of its resources when tunnels are not in use anymore. To enable those two functionalities, the tunnel set-up protocol must include the transmission of keep-alive messages. A client may choose not to send those messages (for example on ISDN type links). In this case, the client should be able to handle a tunnel disconnect event and be able to restart the set-up phase to re-establish the tunnel.

5.4 Security

The tunnel set-up protocol must not introduce any new vulnerability to the network. See security considerations in Section 7.

5.5 Traceability

In some production environment, traceability is an important consideration ([I-D.ietf-v6ops-isp-scenarios-analysis], section 5.4). The tunnel set-up protocol must be instrumentable to enable the collection of usage data that can be used, for example, for capacity planning.

5.6 Phase Out

This assisted tunneling mode is only a transition mechanism to enable ISP to jump start IPv6 service without requiring an immediate global upgrade of access networks and instead enabling a progressive roll out. Once IPv6 is available natively in the access network

connecting a customer, there is no reason to keep using tunnels. So the implementation should have a provision to enable the ISP to signal the user to use native IPv6 instead.

5.7 Extensibility

The protocol must be extensible to support tunnel encapsulation other than IPv6 over IPv4 and IPv6 over transport over IPv4. In

particular, encapsulation of IPv4 over IPv6 (section 7 [I-D.ietf-v6ops-isp-scenarios-analysis], section 7 [RFC3904], section 6 [I-D.ietf-v6ops-ent-analysis]) or IPv6 over IPv6 could be defined.

6. Compatibility with other Transition Mechanisms

The tunnel set-up protocol is not required to be compatible with any existing transition mechanism. Although, a great deal of experience can be drawn from the operation of tunnel brokers currently using the TSP protocol [I-D.blanchet-v6ops-tunnelbroker-tsp].

7. Security Considerations

The establishment of a tunnel can be compared to Mobile IP technology, where traffic can be redirected to go from one place to another one. So similar threats exists. In particular, when a customer is asking for the set-up of a tunnel ending at IP address X, the ISP should check:

- the customer is allowed to set-up this tunnel, i.e. he "owns" the IPv6 prefix.

- the customer is allowed to terminate the tunnel where he said he would, i.e. he "owns" the IPv4 tunnel endpoint.

The first check is an authentication issue. The second may be more complex. The protocol must make sure that the tunnel is established to the legitimate and authenticated destination. IPv4 return routability checks could help this validation. Also, when the user is within the ISP access network, strict ingress filtering can help prevent IPv4 address spoofing.

8. Acknowlegements

This draft has greatly benefited from substantial inputs by Pekka Savola.

The following people provided feedback on this work (in no particular

order): Carl Williams, Brian Carpenter, Christian Huitema, Jordi Palet, Jeroen Massar, Erik Nordmark, Soohong Daniel Park, Suresh Satapati, Fred Tremplin, Karim El-Malki, Tim Chown, Gert Doering, Soliman Hesham.

Thanks to Mark Prior and Bernard Tuy for providing input from an ISP perspective to validate many requirements.

9. References

9.1 Normative References

[I-D.ietf-v6ops-ent-analysis]
Bound, J., "IPv6 Enterprise Network Analysis", draft-ietf-v6ops-ent-analysis-00 (work in progress), September 2004.

[I-D.ietf-v6ops-isp-scenarios-analysis]
Lind, M., Ksinant, V., Park, S., Baudot, A. and P. Savola, "Scenarios and Analysis for Introducing IPv6 into ISP Networks", draft-ietf-v6ops-isp-scenarios-analysis-03 (work in progress), June 2004.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3053] Durand, A., Fasano, P., Guardini, I. and D. Lento, "IPv6 Tunnel Broker", RFC 3053, January 2001.

[RFC3904] Huitema, C., Austein, R., Satapati, S. and R. van der Pol, "Evaluation of IPv6 Transition Mechanisms for Unmanaged Networks", RFC 3904, September 2004.

9.2 Informative References

[I-D.blanchet-v6ops-tunnelbroker-tsp]
Parent, F. and M. Blanchet, "IPv6 Tunnel Broker with the Tunnel Setup Protocol(TSP)", draft-blanchet-v6ops-tunnelbroker-tsp-01 (work in progress), June 2004.

[I-D.huitema-v6ops-teredo]
Huitema, C., "Teredo: Tunneling IPv6 over UDP through NATs", draft-huitema-v6ops-teredo-02 (work in progress), June 2004.

[I-D.ietf-ngtrans-isatap]

Templin, F., Gleeson, T., Talwar, M. and D. Thaler,
"Intra-Site Automatic Tunnel Addressing Protocol
(ISATAP)", draft-ietf-ngtrans-isatap-22 (work in
progress), May 2004.

[I-D.palet-v6ops-tun-auto-disc]
Palet, J. and M. Diaz, "Evaluation of v6ops Auto-discovery
for Tunneling Mechanisms",
draft-palet-v6ops-tun-auto-disc-01 (work in progress),

June 2004.

[I-D.tschofenig-v6ops-secure-tunnels]
Tschofenig, H., "Using IPsec to Secure IPv6-over-IPv4 Tunnels", draft-tschofenig-v6ops-secure-tunnels-01 (work in progress), July 2004.

[RFC2222] Myers, J., "Simple Authentication and Security Layer (SASL)", RFC 2222, October 1997.

[RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[RFC2831] Leach, P. and C. Newman, "Using Digest Authentication as a SASL Mechanism", RFC 2831, May 2000.

Authors' Addresses

Florent Parent
Hexago
2875 boul. Laurier, suite 300
Sainte-Foy, QC G1V 2M2
Canada

EMail: Florent.Parent@hexago.com

Alain Durand
SUN Microsystems,inc.
17 Network circle UMPK17-202
Menlo Park, CA 94025
USA

EMail: Alain.Durand@sun.com

Alain Baudot
France Telecom R&D

42, rue des coutures
14066 Caen,
France

EMail: alain.baudot@rd.francetelecom.com

Appendix A. Changes from version 00

- Non-registered mode removed (now covered in generic zero-conf tunneling draft). Text throughout the document changed to reflect this.

-

- Renamed title from "requirements" to "goals"

- Changed imperatives to lowercase

- /128 endpoints replaced by /64

- Removed DNS considerations

- Added many references to other v6ops scenario documents

- Removed the appendix on protocol requirements summary

- Removed references to 3GPP scenario

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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