Internet Engineering Task Force
Internet-Draft
Expires: August 18, 2005

J. Palet
Consulintel
K. Nielsen
Ericsson
F. Parent
Hexago
A. Durand
Sun Microsystems, inc.
R. Suryanarayanan
Samsung India Software Operations
P. Savola
CSC/FUNET

February 14, 2005

Goals for Tunneling Configuration
draft-palet-v6tc-goals-tunneling-00.txt

Status of this Memo

Copyright Notice

Abstract

This memo describes the set of goals for a tunneling setup protocol that could be used in several network cases to jumpstart its IPv6 offering to its customers by providing them IPv6 connectivity through a simplistic tunneling mechanism.

The basic network cases, which may have different sets of goals, are also introduced, including 3GPP and other Service Providers. Two cases are analyzed in the Service Provider scenario, one which apply to simplistic mechanism where the user is already authenticated by other network existing means, and another which also takes care of the user authentication.

Table of Contents

1. Introduction 4
2. Terminology 5
3. Assumptions and Prerequisites 6
4. Goals of the Tunneling Configuration Protocol 7
4.1 General 7
4.1.1 Simplicity 7
4.1.2 Easy to deploy and Easy to Phase-out 7
4.2 Tunnel Set-up 8
4.2.1 Tunnel End-Point Auto-Discovery and tunnel establishment 8
4.2.2 Tunnel End-Point Reachability Detection 8
4.2.3 Scalability and Load-Balancing 9
4.2.4 Latency in Set-up Phases 9
4.2.5 Tunnel Link Sustainability 9
4.2.6 NAT Traversal 10
4.2.7 Firewall Traversal 10
4.2.8 Use Native Connectivity when Available 10
4.3 IPv6 Configuration 10
4.3.1 IPv6 Address Assignment 10
4.3.2 IPv6 Address Stability 11
4.3.3 IPv6 Prefix Delegation 11
4.3.4 IPv6 DNS 11
4.4 Implementation Considerations 11
4.4.1 Private and Public IPv4 Addresses 11
4.4.2 Extensibility 11
4.4.3 Stateful or Stateless 12
4.5 Management and Security 12
4.5.1 Security 12
4.5.2 Traceability 12
4.5.3 Registration 12
4.5.4 Authentication 13
4.5.5 Confidentiality 13

4.5.6 Accounting 13
5. Applicability of the Tunneling Configuration to Different Network Cases 14
5.1 3GPP Access Networks 14
5.1.1 Simplicity 15
5.1.2 Automated IPv6-in-IPv4 tunnel establishment 15
5.1.3 IPv6 Address Assignment and Prefix Delegation 15
5.1.4 15
5.1.5 15
5.1.6 16
5.1.7 16
5.2 Narrowband Access Networks 16
5.3 Broadband Access Networks 16
5.4 Unmanaged Networks 16
5.5 Enterprise Networks 16
6. Conclusions 17
7. Security Considerations 17
8. Acknowledgements 17
9. References 17
9.1 Normative References 17
9.2 Informative References 17
Authors' Addresses 18
Intellectual Property and Copyright Statements 21

1. Introduction

Regardless of which is the network that is involved in the transition from IPv4 to IPv6, generally this could involve several phases, often in different networks parts (i.e., core, access).

The transition of the core network usually can be done in a much more easy way than the access network. This is the case even if the core network is only connected via tunnels to other IPv6 networks. The setup of those tunnels involves a small effort and is not a big trouble, even in the case of a manual configuration.

However, this is not the case for the access network, which may involve different types of layer two technologies. In all the cases, in order to facilitate the transition of those access networks, which will be impossible to do manually and efficiently, there is a need for an automatic IPv6-in-IPv4 tunneling mechanism. The goal is to provide bidirectional IPv6-in-IPv4 tunneled connectivity between dual-stack end-nodes located at an IPv4-only access network and dual-stack tunnel servers located at IPv6/IPv4 network boundaries.

This should be applicable to all kind of ISPs and access networks. They could be regular ISPs, providing service through DSL, PSTN, ISDN, cable, PLC or any other technologies, but could be also a wireless ISP, or even an enterprise with its own service provider infrastructure for the employees at remote locations.

In order to simplify the text, "customers" is used in the rest of this document to refer to both Customers of Service Providers (3GPP, other ISPs) and users (Enterprise, others).

In this document, the refereces to "Service Provider" is a general one, meaning whatever network/technology is used for providing access to IP connectivity.

In the case of an ISP starting its IPv6 offering to its customers, without initially upgrading its access network to support IPv6, as indicated in section 5.1 of [3], could use a "tunnel brokering solution", as described in [5]. However the tunnel set-up protocol has been identified as a missing piece.

Similarly, in an 3GPP, ISP or enterprise network, the provision of the native IPv6 connectivity to the customers/users, can take a long time and may be costly, while a tunneled infrastructure can be used as a low cost transition path, which can be deployed easily and in a short time, enabling progressive native IPv6 deployment when/where justified.

Such tunneling infrastructure can connect the customers/users to the IPv6 network using available production IPv6 address space, thus facilitating the transition towards native IPv6 deployment, so the roadmap may become:

- Tunneling infrastructure for early adopters.
- Native IPv6 to some customers/user groups once economically justified.
- Native IPv6 to all customers/users.

"Tunneling Configuration" (TC) is used in this document to describe a protocol which takes care of the setup and maintenance of the bidirectional tunnel between a dual-stack end-node (or leaf network) and a dual-stack tunnel server. The exchange of parameters needed for the setup and maintenance of the tunnel (such as address, prefix, routing, encapsulation, filtering, authentication, accounting, ...), should be automated to avoid manual user/operator intervention.

The tunneling configuration protocol is envisaged to be deployed as an initial and temporary mechanism to provide basic IPv6 connectivity services only. This basic IPv6 connectivity may be limited, in the sense than may be not 100% comparable to a native IPv6 service. However this basic IPv6 connectivity should be enough while it allows the communication through IPv6 during the transition phase, until the native IPv6 service is available, and consequently is expected that the tunneling will be phased out as soon as native IPv6 access service is available.

This document analyzes the goals for a such tunnel setup protocol, taking in consideration the different possible common network cases for deploying IPv6.

2. Terminology

Tunneling-Configured Site (TCS): A logical network over which IPv6 connectivity is provided by means of Tunneling Configuration. At least one dual-stack node is required in this logical network.

Tunnel End-Point (TEP): A dual-stack node performing IPv6-in-IPv4 tunnel encapsulation/decapsulation in accordance with Tunneling Configuration. There will be always two TEPs in order to establish the communication, the local one at the customer site (TCS) and the remote one at the ISP site.

Tunnel Server (TS): A dual-stack server node with IPv6 connectivity and which provides IPv6 connectivity to client nodes by performing

IPv6-in-IPv4 tunnel encapsulation/decapsulation to/from client nodes in accordance with Tunneling Configuration. A Tunnel Server is likely to be a dual-stack router, but could be also a node behaving as a router. The TS is often the ISP TEP.

Tunnel Client (TCL): A dual-stack node that obtains IPv6 connectivity by means of Tunneling Configuration. A tunnel client relies on IPv6-in-IPv4 tunnel encapsulation/decapsulation to/from Tunnel Servers for IPv6 communications to native IPv6 nodes. This is often the customer TEP.

Direct Tunneling (DT): Direct tunnelling here refer to the case where end-hosts located within different Tunneling-Configured Sites, in the same ISP network, may circumvent the Tunnel Server and communicate directly using the tunnel protocol.

CPE: Customer Premises Equipment.

3. Assumptions and Prerequisites

Tunneling Configuration may be applicable in different IPv6 transition network cases. The focus of this document is to define the goals to apply this mechanism in the Service Provider context making the following assumptions and prerequisites:

- The customer configuration may be diverse and not necessarily predictable. Consequently the tunneling configuration protocol must be able to adapt to different cases or combinations of:
  - The TCS is a single node or leaf network.
  - The TCL at the TCS has a global IPv4 address or is behind one or more NATs.
  - The TCL at the TCS has a static or dynamic IPv4 address.
  - In case of NAT, the external IPv4 address is a static or dynamic.
  - In case of NAT, it can be customer or ISP owned.
- IPv4 multicast is not widely available, so the tunnel configuration protocol should work in IPv4 network environments where IPv4 multicast is not provided.
- The tunnel configuration protocol should 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 tunneling configuration protocol is provided within a restrictive timescale, in the sense that it should be phased out as soon as native access can be provided.

- The tunneling configuration is a protocol to be used in the transition phase, thus does not need to be perfect. As a matter of fact, making it perfect would be counter productive, as it would first delay its definition, then make its deployment more cumbersome and, last but not least, diminish the incentives to deploy native IPv6. Furhermore, should not rely in a complex set of devices, which may not be readily available, and could even mean a more expensive cost than the support of native IPv6 itself.

4. Goals of the Tunneling Configuration Protocol

As introduced above, there are different ISP types and different access networks. This means that that there are different goals related to different network cases and situations. For instance, factors as presence of NAT or not, low/high bandwidth, expensive/cheap, strong internal access control or not, etc.

Different access media or network cases brings up different sets of goals. Obviously, once choice will be to create a protocol for each specific case, but this is not optimal. Instead, the motivation of this document is to combine all the goals and look for a common solution, which can fit as much as possible and in the most optimal way, in all the cases. Some of those goals could be conflictual and that need to be resolved as well.

This section groups these different goals.

4.1 General

4.1.1 Simplicity

The Tunneling Configuration protocol should be easy to implement, implying a lightweight protocol. The protocol should provide a reasonable, even if limited, set of basic IPv6 connectivity features.

4.1.2 Easy to deploy and Easy to Phase-out

The Tunneling Configuration protocol should be easy to deploy into the existing IPv4 and IPv6 network infrastructure.

The Tunneling Configuration protocol should have no major impact on protocols and infrastructure nodes deployed in existing infrastructures providing IPv4 and native IPv6 connectivity.

The Tunneling Configuration protocol should coexist and work seamlessly together with any native IPv6 infrastructure that gradually may be implemented in the network. The Tunneling Configuration protocol should have no negative implications on how such infrastructure is implemented.

The Tunneling Configuration protocol should be easy to take out of service once native IPv6 is available.

4.2 Tunnel Set-up

4.2.1 Tunnel End-Point Auto-Discovery and tunnel establishment

The tunnel protocol should provide a mechanism for the automated discovery of the Tunnel End-Point, by the virtue of which end-hosts automatically and at run-time can determine the IPv4 addresses of available Tunnel Servers.

The discovery mechanism should rely on intrinsic services, read services already universally deployed, to the particular network environment. It should not require the addition of additional IP network infrastructure elements for this function only.

The mechanism should be fully dynamic in the sense that it must not require IP address information such as the IPv4 address of a Tunnel Server and/or the IPv6 address(es) to use for IPv6 connectivity to be configured on the Tunnel Clients beforehand.

The analysis done in [11] may apply.

The Tunneling Configuration protocol should provide for the set of IPv6-in-IPv4 tunnels, based on IPv6-in-IPv4 encapsulation as defined in [10], from dual-stack nodes, attached to IPv4-only networks, to Tunnel Servers.

The IPv6-in-IPv4 tunnels and the IPv6 connectivity should be established in an automated manner, i.e., without requiring manual intervention at any of the tunnel end-points at tunnel establishment time. We can typically describe it as a "plug and play" protocol, which can be triggered through the execution of a simple program.

4.2.2 Tunnel End-Point Reachability Detection

The Tunneling Configuration protocol should allow for means for one tunnel end-point to verify the reachability of other tunnel end-points towards which it intends to send packets in a method similar to IPv6 NUD.

The unicast neighbor reachability discovery functions provided by IPv6 Neighbor Discovery ([6]), i.e., unicast NS/NA exchanges, should be supported on the tunnel link.

It is preferable that a Tunnel Server monitors the reachability of the tunnel client towards which it is sending packets. Full emulation of IPv6 NUD mechanism is however not an explicit goal.

4.2.3 Scalability and Load-Balancing

In order to ensure the scalability of the tunnel service, in terms of not limiting the number of simultaneous connections to the service and consequently limiting possible service denial situations, it should be possible for a Service Provider to load-balance those connections among several available Tunnel Servers.

Load balancing should be planned already during the early phases of deployment. Given adequate planning it should be possible for an ISP to seamlessly deploy additional Tunnel Servers in order to support an increased amount of Tunnel Clients.

This may be achieved, for example, by using the load balancing functions provided by the Tunnel Server End-point discovery mechanism as detailed in [12].

4.2.4 Latency in Set-up Phases

In certain type of networks, keeping tunnels active all the time is not possible. In such environments, the protocol must be able to set-up tunnels on demand when the IPv6 connectivity either natively or through tunneling is unavailable. The tunnel will be set-up only once though for the tunnel client and not per session.

The Tunneling Configuration protocol must then have a low enough latency to enable quasi-instant configuration. Latency is usually a function of the number of packet exchanges required, so minimizing this parameter is important.

4.2.5 Tunnel Link Sustainability

The tunnel link established in between a host deploying Tunneling Configuration and an associated Tunnel Server should be expected to remain in administrative active state for the lifetime of the IPv6 address provided to the host.

The Tunneling Configuration protocol should not mandate keep-alive messages to be transmitted by the host simply in order to sustain tunnel link connectivity. However, this may be required when a

tunnel has to cross a NAT box. In this case, 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 3GPP or 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.

4.2.6 NAT Traversal

The Tunneling Configuration protocol should be able to detect the presence of one or more NATs in its path.

The tunnel should be able to traverse NAT, if present, so it may be necessary to choose among several tunnel encapsulation protocols for the most optimal one.

4.2.7 Firewall Traversal

Even if no NAT is in the tunnel path, there may be a firewall which prohibits proto-41. In such case, the tunnel encapsulation selection based on NAT detection could select a tunnel that will not work.

To cope with this situation, the Tunnel Configuration protocol implementation may allow a user to explicitly specify the desired tunnel encapsulation, regardless of the NAT detection process.

4.2.8 Use Native Connectivity when Available

The node should not use the Tunneling Configuration protocol when native IPv6 connectivity is available.

The fact that a node should not initiate the Tunneling Configuration protocol when native IPv6 connectivity is available is not considered to be a functional goal on the tunnel protocol per se. For example, rather it is related to the activation and deactivation of the protocol.

4.3 IPv6 Configuration

4.3.1 IPv6 Address Assignment

Assignment of at least one globally routable IPv6 unicast address

(/128) to the end-node should be supported.

No goals are defined as to how address configuration should be performed. This may be done based on stateless or stateful IPv6 address configuration mechanisms or by some altogether different mechanism particular to the Tunneling Configuration protocol.

4.3.2 IPv6 Address Stability

The IPv6 address is "transient" and may change, but the protocol should offer a mechanism to provide IPv6 address stability (for example, a cookie mechanism). The implementation of this mechanism should allow this feature to be turned off.

It is preferable that the address assignment provides a stable address, that is, an address that can be used for IPv6 connectivity for a certain amount of time rather than solely one address per higher layer session initiation.

4.3.3 IPv6 Prefix Delegation

Prefix Delegation support may depend on the different deployment cases. It is not however required that a Tunneling Configuration protocol supporting only basic requirements provides support for prefix delegation.

4.3.4 IPv6 DNS

Dual-stack nodes could use both IPv4 and IPv6 DNS discovery mechanisms and both, IPv4 and IPv6 transport for DNS services.

Consequently IPv6 DNS discovery and IPv6 transport for DNS services should not be a goal of the Tunneling Configuration protocol.

4.4 Implementation Considerations

4.4.1 Private and Public IPv4 Addresses

The Tunneling Configuration protocol should work over IPv4 sites deploying both private and public IPv4 addresses.

Furthermore, the Tunneling Configuration protocol should work with both dynamic and static IPv4 address allocation.

4.4.2 Extensibility

The Tunneling Configuration protocol should be extensible to support tunnel encapsulation other than IPv6 in IPv4 and IPv6 in transport in

IPv4. In particular, encapsulation of IPv4 in IPv6 or IPv6 in IPv6 could be defined.

4.4.3 Stateful or Stateless

By a stateful mechanism we mean a mechanism that require the Tunnel Server to maintain tunnel state per client it serves.

Tunnel state here is considered to be any parameter kept by the server per client and without which the server is unable to serve the client (receive packets from/send packets to).

Tunnel state must be distinguished from state used to optimize the packet delivery function of the tunnel server and which is kept in a fixed or upper limited amount of memory space, such as, e.g., reachability information.

It should be emphasized that this document makes no deliberate assumptions on whether a Tunneling Configuration protocol should be based on a stateful or stateless Tunnel Server mechanism. Indeed it is anticipated that the goals of Tunneling Configuration as put forward here could be served both by a stateless as well as by a stateful mechanism.

4.5 Management and Security

4.5.1 Security

The Tunneling Configuration protocol should not impose any new vulnerability to the existing network infrastructure.

The Tunneling Configuration protocol should not impose any new vulnerability to the nodes implementing it than what is already present in existing multi-access IPv6 networks, where multiple hosts are served by the same router or possibly multiple routers.

4.5.2 Traceability

In some environments, traceability is an important consideration. The Tunneling Configuration protocol should be instrumentable to enable the collection of usage data which can be used, for example, for capacity planning.

4.5.3 Registration

The registration of credentials should be external to the Tunneling Configuration protocol. The user may require registration prior to using this service (through some web based service or other means).

Alternatively, the service provider may use an existing authentication database to pre-register its users, or even not require registration at all, depending on the network configuration.

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).

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

4.5.4 Authentication

Authentication can be used to control that the user has access to the IPv6 services.

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 should be supported. Other authentication may be supported and should be negotiated between the client and server (e.g., SASL [13]).

4.5.5 Confidentiality

Tunneling Configuration can be used across networks which are not under the service provider control (e.g., roaming users). The Tunneling Configuration protocol should allow protection of the authentication data. 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, for instance by means of using IPsec tunneling.

4.5.6 Accounting

The Tunneling Configuration 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. Applicability of the Tunneling Configuration to Different Network Cases

Note: Section to be completed in a new version.

The goals enumerated in the precedent section are not all necessarily applicable and not in the same degree to different network cases. However seems feasible to build a common ground to these different network cases in order to define a single Tunneling Configuration protocol, which can accommodate to the different combinations of those goals and network cases.

The v6ops working group has identified and analyzed several deployment scenarios for IPv4/IPv6 transition mechanisms in various stages of the IPv6 deployment and its coexistence with IPv4.

This work has been carried out for a number of different network environments each with their particular characteristics including 3GPP [1], Unmanaged [2], ISP [3] and Enterprise networks [4].

The following sub-sections basically look into those different network environments to provide an analysis of the common and uncommon goals.

5.1 3GPP Access Networks

IPv6-in-IPv4 tunneling is envisaged to be deployed in 3GPP networks as an initial and temporary mechanism to provide limited and basic IPv6 connectivity services only. The IPv6-in-IPv4 tunneling mechanism demanded by the 3GPP environment falls within the realm of Tunneling Configuration.

Native IPv6-like 3GPP connectivity services, e.g. services including flexible charging and quality of service on demand, will be feasible, in 3GPP environments by virtue of true native IPv6 only. This is due to the interrelation between the native IPv6 3GPP service and various 3GPP signaling interfaces. The latter which is not envisaged upgraded to support the IPv6-in-IPv4 tunneling situation.

It is important to note that the IPv6 connectivity provided by 3GPP Tunneling Configuration IPv6-in-IPv4 tunneling does not compare with the native IPv6 3GPP connectivity in terms of the services offered.

This differentiates the 3GPP IPv6-in-IPv4 tunneling transition case somewhat from some of the other transition scenarios considered in the IETF v6ops WG and unlike some of these scenarios, the 3GPP IPv6-in-IPv4 tunneling deployment case is not a case of progressive and gradual roll out of native IPv6-like services. Rather, Tunneling Configuration will in the 3GPP environment be deployed for the following purposes:

- To provide temporary provisioning of basic IPv6 services, which users may deploy for the simplest IPv6 services only.

- To allow an Operator, possibly a native IPv6 enabled Operator, to provide basic IPv6 services to users roaming into foreign networks which supports IPv4 bearer connectivity only.

The scope of Tunneling Configuration in the 3GPP network environment is restricted to an absolute minimal set of functions required to provide automatic IPv6 connectivity establishment to dual stack nodes by means of IPv6-in-IPv4 encapsulation as defined in [10] to Tunnel Servers.

Tunneling Configuration in the 3GPP network environment does not attempt to provide emulation of the full set of native IPv6 connectivity functions as defined by [14], [15] and [16].

With this in mind the following goals are applicable to the 3GPP Access Networks:

5.1.1 Simplicity

.

5.1.2 Automated IPv6-in-IPv4 tunnel establishment

.

5.1.3 IPv6 Address Assignment and Prefix Delegation

It is only an explicit goal to have a /128 address allocated for global connectivity on the tunnel link.

5.1.4

.

5.1.5

.

5.1.6

.

5.1.7

.

5.2 Narrowband Access Networks

Somehow this type of networks are very similar to the 3GPP case, because the main constrain is the low bandwidth and the high cost of the usage of the access network. Examples of this are PSTN and ISDN access networks.

.

5.3 Broadband Access Networks

This is typically the case when an ISP is offering IPv6 connectivity to its customers, initially using controlled tunneling infrastructure, as described in section 5.1 "Steps in Transitioning Customer Connection Networks" of [3].

.

5.4 Unmanaged Networks

In unmanaged networks [2], Tunneling Configuration 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 case the link is not IPv4 capable, Tunneling Configuration is applicable by means of IPv4 in IPv6 tunneling.

This will actually fall into the same set of goals as already described for narrow or broadband access networks, depending on the media.

5.5 Enterprise Networks

In the enterprise scenario [4], Tunneling Configuration can be used to support both, remote users connecting to the enterprise network (section 7.5.2) and internal users if the native infrastructure is not yet available.

6. Conclusions

TBD.

7. Security Considerations

TBD.

8. Acknowledgements

This memo was written starting from previous existing work at v6ops, such as "Goals for Zero-Configuration Tunneling in 3GPP" [7], "Zero-Configuration Tunneling Requirements" [8] and "Goals for Registered Assisted Tunneling" [9]. The authors would also like to acknowledge inputs from Tim Chown and the European Commission support in the co-funding of the Euro6IX project, where this work is being developed.

9. References

9.1 Normative References

[1] Wiljakka, J., "Analysis on IPv6 Transition in 3GPP Networks", Internet-Draft draft-ietf-v6ops-3gpp-analysis-11, October 2004.

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

[3] Lind, M., Ksinant, V., Park, S., Baudot, A. and P. Savola, "Scenarios and Analysis for Introducing IPv6 into ISP Networks", Internet-Draft draft-ietf-v6ops-isp-scenarios-analysis-03, June 2004.

[4] Bound, J., "IPv6 Enterprise Network Scenarios", Internet-Draft draft-ietf-v6ops-ent-scenarios-05, July 2004.

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

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

9.2 Informative References

[7] Nielsen, k., "Goals for Zero-Configuration Tunneling in 3GPP", Internet-Draft draft-nielsen-v6ops-3GPP-zeroconf-goals-00, October 2004.

[8] Suryanarayanan, R., "Zero-Configuration Tunneling Requirements", Internet-Draft draft-suryanarayanan-v6ops-zeroconf-reqs-00, October 2004.

[9] Parent, F., "Goals for Registered Assisted Tunneling", Internet-Draft draft-ietf-v6ops-assisted-tunneling-requirements-01 , October 2004.

[10] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", Internet-Draft draft-ietf-v6ops-mech-v2-06, September 2004.

[11] Palet, J. and M. Diaz, "Analysis of IPv6 Tunnel End-point Discovery Mechanisms", Internet-Draft draft-palet-v6ops-tun-auto-disc-03, January 2005.

[12] Palet, J., "IPv6 Tunnel End-point Automatic Discovery Mechanism", Internet-Draft draft-palet-v6ops-solution-tun-auto-disc-01, October 2004.

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

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

[15] Loughney, J., "IPv6 Node Requirements", Internet-Draft draft-ietf-ipv6-node-requirements-11, August 2004.

[16] IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites", RFC 3177, September 2001.

Authors' Addresses

Jordi Palet Martinez
Consulintel
San Jose Artesano, 1
Alcobendas - Madrid
E-28108 - Spain

Phone: +34 91 151 81 99
Fax: +34 91 151 81 98
Email: jordi.palet@consulintel.es

Karen Egede Nielsen
Ericsson
Skanderborgvej 232
8260 Viby J
zip - Denmark

Phone: +45 89 38 51 00
Fax:
Email: karen.e.nielsen@ericsson.com

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

Phone:
Fax:
Email: florent.parent@hexago.com

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

Phone:
Fax:
Email: alain.durand@sun.com

Radhakrishnan Suryanarayanan
Samsung India Software Operations
No. 3/1 Millers Road
Bangalore
India

Phone: +91 80 51197777
Fax:
Email: rkrishnan.s@samsung.com

Pekka Savola
CSC/FUNET
Espoo
Finland

Phone:
Fax:
Email: psavola@funet.fi

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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