

Network Working Group
INTERNET-DRAFT

Expires: March 2, 2005

K. Nielsen
Ericsson
M. Morelli
Telecom Italia Lab
J. Palet
Consulintel
J. Soininen
Nokia
J. Wiljakka
Nokia

September 2, 2004

Goals for Zero-Configuration Tunneling
<[draft-nielsen-v6ops-zeroconf-goals-01.txt](#)>

Status of this memo

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#) ([BCP 79](#)).

By submitting this Internet-Draft, I accept the provisions of Section [#3] of [RFC 3667](#) ([BCP 78](#)).

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 cite them other than as "work in progress".

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

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

Abstract

This document describes the set of goals to be fulfilled by a Zero-Configuration Tunneling protocol.

Zero-Configuration Tunneling here denotes a minimalistic IPv6-in-IPv4 automatic tunneling mechanism that could be used by a Service

Provider to offer IPv6 connectivity to its customers in early phases of IPv4 to IPv6 transition.

Table of Contents

- [1. Introduction.....2](#)
- [2. Terminology.....4](#)
- [3. Scope and Limitations.....4](#)
 - [3.1. IPv6 address allocation, Scope and Limitations:.....4](#)
 - [3.2. IPv6 tunnel link characteristics, Scope and Limitations:....5](#)
- [4. Assumptions and Prerequisites.....5](#)
 - [4.1. Applicability Assumptions.....5](#)
 - [4.2. 3GPP Compliance Prerequisite.....6](#)
- [5. Timing.....6](#)
- [6. Goals.....7](#)
 - [6.1. Simplicity.....7](#)
 - [6.2. Automated IPv6-in-IPv4 tunnel establishment.....7](#)
 - [6.3. Use native when available.....7](#)
 - [6.4. Easy to deploy and Easy to Phase-out with no modifications on existing equipment.....8](#)
 - [6.5. Tunnel Server End-point Discovery.....8](#)
 - [6.6. Address Assignment.....8](#)
 - [6.7. Tunnel Link Sustainability.....9](#)
 - [6.8. Tunnel End-Point Reachability Detection.....9](#)
 - [6.9. Private and public IPv4 addresses.....9](#)
 - [6.10. Scalability, Load Balancing.....9](#)
 - [6.11. Security.....10](#)
- [7. Non Goals.....10](#)
 - [7.1. NAT and Firewall Traversal.....10](#)
 - [7.2. IPv6 DNS.....10](#)
 - [7.3. Extensibility.....11](#)
 - [7.4. Registration burden.....11](#)
- [8. Stateful or Stateless.....11](#)
- [9. Security Considerations.....11](#)
 - [9.1. Threats to existing network infrastructures.....12](#)
 - [9.2. Threats to nodes implementing Zero-Configuration Tunneling.12](#)
 - [9.2.1. Threats to end-hosts.....13](#)
 - [9.2.2. Threats to Tunnel Servers.....13](#)
 - [9.2.2.1. Tunnel State related risks.....13](#)
 - [9.2.2.2. Traffic related risks.....14](#)
 - [9.2.2.3. Packet Delivery related threats.....14](#)
 - [9.3. Implications of Direct Tunneling.....14](#)
- [10. Acknowledgments.....15](#)
- [11. Authors Addresses.....15](#)
- [12. Changes from \[draft-nielsen-v6ops-zeroconf-goals-00.txt\]\(#\).....16](#)
- [13. Informative References.....17](#)
- [Appendix A Out of Scope.....18](#)

[Appendix B](#) Open Issues.....[18](#)

[1. Introduction](#)

Nielsen, et. al.

[Page 2]

The IETF v6ops Working Group has identified and analyzed deployment scenarios for IPv4/IPv6 transition mechanisms in various stages of IPv6 deployment and IPv6 and IPv4 coexistence.

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

The work has identified a need for automatic IPv6-in-IPv4 tunneling mechanisms that 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 within the Service Providers network.

The term Zero-Configuration Tunneling is used in this document to denote an IPv6-in-IPv4 tunneling mechanism that fulfills the goals as put forward here.

A Zero-Configuration Tunneling mechanism provides a basic set of tunneling features only, and intentionally so. The scope of Zero-Configuration Tunneling is not to provide emulation of the full suite of native IPv6 connectivity functions as defined by [7], [8] and [9]; rather the focus is to provide a minimal set of features required for automatic establishment of IPv6 connectivity.

The starting point for the definition of the set of goals to be fulfilled by a Zero-Configuration Tunneling mechanism has been the set of features required for the IPv6-in-IPv4 tunneling mechanism envisaged to be needed during the early phases of IPv6 transition in 3GPP environments as described in [4].

The applicability of Zero-Configuration Tunneling may widen to other transition scenarios.

For scenarios demanding advanced tunneling features, for example full emulation of native (though tunneled) IPv6 connectivity, a more full-fledged tunneling mechanism is envisaged to be deployed, see [5]. With respect to the latter, an analysis of appropriate mechanisms for automatic discovery of the tunnel endpoint is being done in [6].

It should be emphasized that unless otherwise specified then in this document the reference, IPv6-in-IPv4 encapsulation as defined in [1], refers to the aspects of Protocol-41 encapsulation related to IPv4 header construction (except for source and destination address determination), MTU and Fragmentation, Hop Limits and ICMP handling as detailed in Section 3.1-3.6 of [1]. The particular aspects of Configured IPv6-In-IPv4 Tunneling in the areas of IPv4 source and destination address determination, tunnel link characteristics and

IPv6 Neighbor Discovery operation are not intended referred to by the above reference.

Nielsen, et. al.

[Page 3]

2. Terminology

Zero-Configuration Tunneling site:

A logical IPv4 network over which IPv6 connectivity is provided to dual-stack nodes by means of Zero-Configuration Tunneling.

Tunnel End-point:

A dual-stack node performing IPv6-in-IPv4 tunnel encapsulation/decapsulation in accordance with Zero-Configuration Tunneling.

Tunnel Server:

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 Zero-Configuration Tunneling.

A Tunnel Server is likely to be a dual-stack router.

Tunnel Client:

A dual-stack node that obtains IPv6 connectivity by means of Zero-Configuration Tunneling. A tunnel client relies on IPv6-in-IPv4 tunnel encapsulation/decapsulation to/from Tunnel Servers for IPv6 communications to native IPv6 nodes.

Direct Tunneling:

Direct tunnelling here refer to the case where end-hosts located within the same Zero-Configuration Tunnelling site may circumvent the Tunnel Server and communicate directly using the tunnel protocol.

3. Scope and Limitations

The scope of Zero-Configuration Tunneling 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 [1] to tunnel servers under the assumptions and prerequisites described in [Section 4](#).

Zero-Configuration Tunneling does not attempt to provide emulation of the full set of native IPv6 connectivity functions as defined by [7], [8] and [9].

3.1. IPv6 address allocation, Scope and Limitations:

The primary goal of Zero-Configuration Tunneling is to provide IPv6 connectivity to nodes on an individual basis. By this it is meant that it is only an explicit goal to have a /128 address allocated for

global connectivity on the tunnel link. As such optimal IPv6

Nielsen, et. al.

[Page 4]

connectivity provisioning in Personal Area Network (PAN) scenarios is not explicitly within the scope of Zero-Configuration Tunneling.

It is not explicitly within the scope of Zero-Configuration Tunneling to support usage of privacy IPv6 extensions as defined in [12].

It is not explicitly within the scope to support usage of IPv6 multicast.

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

3.2. IPv6 tunnel link characteristics, Scope and Limitations:

Direct tunneling is neither an explicit goal nor explicitly excluded in Zero-Configuration Tunneling.

It is not an explicit requirement for the zero-configuration tunnel link to support IPv6 link-local multicast.

The tunnel protocol should allow for the formation of a link-local address on the tunnel link. Though no particular usage of such an address is explicitly demanded by the goals set forward here.

It is an explicit goal that nodes attached to a tunnel link must be able to ascertain the reachability of neighbors with which it is communicating (or wish to start communicate). This may be achieved using IPv6 Neighbor Discovery mechanism ([13]) based on unicast link-local packet exchanges (or link-local multicast if such is supported) but it may also be achieved by altogether different mechanisms.

4. Assumptions and Prerequisites

4.1. Applicability Assumptions

Zero-Configuration Tunneling is a tunneling mechanism by the virtue of which dual-stacks hosts, attached to IPv4-only networks links, can use IPv6-in-IPv4 encapsulation as defined in [1] to tunnel servers for global IPv6 connectivity.

The aim of the document is to define the set of goals to be fulfilled by zero-configured tunneling when the following assumptions are made on the deployment environment:

- IPv4 source addresses spoofing within the Zero-Configuration Tunneling site is prevented.
- The Zero-Configuration Tunneling site is protected from proto-41

encapsulated packets arriving from external IPv4 networks.

Nielsen, et. al.

[Page 5]

- At least one authoritative DNS server is IPv4-enabled and at least one recursive DNS server supports IPv4. Further IPv4 DNS Server discovery is provided by already existing means/means outside the scope of the tunnel protocol.
- There are no NATs in between the tunnel endpoints in the Zero-Configuration Tunneling site.
- The Zero-Configuration Tunneling network is fully penetrable for intra-site IPv6-in-IPv4 Protocol 41 traffic.
- The user is being authenticated to the network by means external to the tunneling protocol and other than that no additional authentication/registration mechanisms are required.

The above assumptions are believed to be readily applicable to the 3GPP tunneling transition scenario described in [4], [section 3.1](#).

It is a prerequisite that the tunnel protocol must work in IPv4 network environments where IPv4 multicast is not provided.

[4.2. 3GPP Compliance Prerequisite](#)

It is a prerequisite that Zero-Configuration Tunneling should be applicable in 3GPP wireless networks. When considering the characteristics of 3GPP network links and mobile terminals / User Equipment (UE), the following points need to be taken into account:

- Link bandwidth (tunnel overhead / usage cost)
- Link latency
- UE battery power and derived implications on memory and processing power

It is thus an explicit requirement for Zero-Configuration Tunneling to comply well with the constrained conditions put on the above parameters by the 3GPP environments. The latter which commonly is recognized as translating into requirements for the protocol to operate with a limited number of message exchanges, small packet sizes and simple message processing.

Here we shall refer to a protocol as being lightweight when its demands on message exchanges, packet sizes and message processing complexity are sufficiently light for it to be readily applicable in environments characterized by the constrained conditions of 3GPP networks (as described above).

As a mean to ensure that the protocol be lightweight it is considered preferable for the protocol to provide a simple set of functions only, even if it provides only a basic IPv6 service compared to the native one. It is although acknowledged that additional functionality doesn't necessarily automatically add complexity to the demands on the aforementioned parameters.

5. Timing

Nielsen, et. al.

[Page 6]

For the purpose of 3GPP deployment it is a prerequisite that this tunneling protocol is provided within a very restrictive timescale.

Zero-Configuration Tunneling is envisaged to be deployed in 3GPP networks as an initial and temporary mechanism to provide limited IPv6 connectivity services. Native IPv6 like connectivity is in 3GPP environments envisaged to be feasible by virtue of true native IPv6 only.

Trial deployments, in which zero-configuration type of IPv6 connectivity is provided in 3GPP environments, are starting up using experimental protocols at the time of writing this document.

6. Goals

The goals to be achieved by Zero-Configuration Tunneling are detailed in the following subsections.

6.1. Simplicity

By simplicity, we understand a tunnel protocol that is easy to implement in the targeted environment. Additionally, the protocol should provide a reasonable, limited set of basic IPv6 connectivity features.

Further by simplicity we imply that the protocol must be lightweight.

6.2. Automated IPv6-in-IPv4 tunnel establishment

The protocol should provide for the set up of IPv6-in-IPv4 tunnels, based on IPv6-in-IPv4 encapsulation as defined in [1], from dual-stack nodes, attached to IPv4-only networks, to Tunnel Servers.

The IPv6-in-IPv4 tunnels and the IPv6 connectivity must be established in an automated manner, i.e. without requiring manual intervention at any of the tunnel end-points at tunnel establishment time.

The mechanism must 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.

6.3. Use native when available

The tunnel protocol should allow the usage of native IPv6 connectivity whenever such is available.

The protocol must in no way restrict the native IPv6 capabilities of

the client node.

Nielsen, et. al.

[Page 7]

The node should not use Zero-Configuration Tunneling when native IPv6 connectivity is available.

Comment: The fact that a node should not use Zero-Configuration Tunneling when native IPv6 connectivity is available is not considered to be a functional requirement on the tunnel protocol. Rather it is considered related to when, and how, the node (implementation) uses the Zero-Configuration Tunneling function.

6.4. Easy to deploy and Easy to Phase-out with no modifications on existing equipment

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

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

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

The tunnel protocol must be easy to take out of service once native IPv6 is available.

6.5. Tunnel Server End-Point Auto-Discovery

The tunnel protocol must provide a mechanism for automated end-point discovery 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.

Comment: The analysis done in [6] may apply.

6.6. Address Assignment

The tunnel protocol must allow for the assignment of at least one globally routable (/128) IPv6 unicast address to use for tunneled IPv6 connectivity over the link provided by the Zero-Configuration Tunneling mechanism.

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.

Nielsen, et. al.

[Page 8]

6.7. Tunnel Link Sustainability

The tunnel link established in between a host deploying Zero-Configuration Tunneling 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 tunnel protocol must not mandate keep-alive messages to be transmitted by the host simply in order to sustain tunnel link connectivity.

Motivation: The fact that a 3GPP terminal, for the single purpose of transmitting keep-alive messages, could have to wake up the radio periodically, send a packet over the radio and possibly wait for response is undesirable for the following reasons:

- The terminal cannot, as otherwise anticipated, be in dormant mode all the time it is idle. This has severe implications for the battery consumption of the device.
- Radio resources are costly and sparse and consequently not to be used for what is considered to be unnecessary traffic.

6.8. Tunnel End-Point Reachability Detection

The tunnel protocol must allow for means for one tunnel end-point to verify the reachability of other tunnel end-points towards which it intends to send packets.

The unicast neighbor reachability discovery functions provided by IPv6 Neighbor Discovery ([13]), i.e., unicast NS/NA exchanges, should be supported on the tunnel link.

6.9. Private and public IPv4 addresses

The tunnel protocol must work over IPv4 sites deploying both private and public IPv4 addresses.

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

Motivation: Private IPv4 addresses are widely used in current 3GPP networks.

6.10. Scalability, 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.

Nielsen, et. al.

[Page 9]

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

Comment: This may be achieved using load balancing functions provided by the Tunnel Server End-point Discovery mechanism as detailed in [14].

6.11. Security

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

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

7. Non Goals

Non-goals of zero-configured tunneling are detailed in the following subsections.

With the term Non-goals we refer to features that generally are believed to be applicable to tunneling, but which are not among the minimal set of required features of Zero-Configuration Tunneling. The latter primarily because of the prerequisites for Zero-Configuration Tunneling and/or because of the assumptions made on the applicability environments for Zero-Configuration Tunneling, e.g., see [Section 4](#).

7.1. NAT and Firewall Traversal

NAT and Firewall traversal is not required due to the assumptions on the applicability environment.

Moreover to minimize the tunneling overhead applied to the packets as well as in order to minimize the number of tunnel set-up signaling messages exchanged on the link, it is preferable that the protocol does not deploy the UDP encapsulation techniques, on which mechanisms able to traverse NATs and Firewalls normally rely.

7.2. IPv6 DNS

By virtue of the assumptions on the applicability environments, the dual stack end-hosts can use IPv4 DNS discovery mechanisms and IPv4 transport for DNS services.

Given that IPv4 based DNS services are already available, it is not

considered a requirement that the end-host should be able to deploy

IPv6 based DNS services. Consequently, the tunnel protocol does not need to provide IPv6 DNS discovery mechanisms.

7.3. Extensibility

As a minimal tunneling mechanism Zero-Configuration Tunneling targets IPv6 connectivity provisioning only. The protocol does not need to be readily extendable to other encapsulation mechanisms, e.g., IPv4-in-IPv6.

7.4. Registration burden

Tunnel service registration is not required due to the assumptions on the applicability environment.

In order to keep the simplicity and minimize the tunnel overhead it is desirable that the tunnel protocol not in itself (e.g., in order to meet the goals put forward in this document) mandates authenticated registration of the user.

8. 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 Zero-Configuration Tunneling solution should be based on a stateful or stateless Tunnel Server mechanism. Indeed it is anticipated that the goals of zero-configuration as put forward here could be served both by a stateless as well as by a stateful mechanism.

9. Security Considerations

It is assumed that the following assumptions of [Section 4](#) are valid in the particular network environment:

- IPv4 source addresses spoofing within the Zero-Configuration Tunneling site is prevented.
- The Zero-Configuration Tunneling site is protected from proto-41

encapsulated packets arriving from external IPv4 networks.

It is worthwhile to note that together these assumptions imply that the IPv4 source of all protocol-41 tunneled packets is legitimate.

9.1. Threats to existing network infrastructures

As stated in [Section 6.11](#) the tunnel protocol should not impose any new vulnerability to the existing network infrastructure.

The following have been identified as potential threats opened up for by the deployment of Zero-Configuration Tunneling:

- As the tunnel service is un-authenticated (not registered) it may be possible to use a tunnel server to reflect tunneled packets into the network, similar to the 6to4-reflection attacks identified in [10].
- The Zero-configuration site must be kept fully penetrable for intra-site IPv6-in-IPv4 protocol-41 encapsulated packets. This may open up for threats to end-hosts that rely on the network infrastructure to filter out bogus protocol-41 encapsulated packets.
- Zero-configuration tunneling may open up for threats to other mechanisms in the network that rely on Protocol-41 encapsulation.

Detailed analysis of the validity of these threats will have to depend on the particular zero-configuration solution. In general it could be noted that attacks based on the above threats largely should be preventable if the end-hosts in the network implement appropriate drop policies, either simple drop all protocol-41 policies or more differentiated policies based, e.g., on source addresses.

9.2. Threats to nodes implementing Zero-Configuration Tunneling

The following considerations apply to the situation where Zero-Configuration Tunneling is deployed in between tunnel servers and end-hosts only.

Special security considerations for the usage of Zero-Configuration Tunneling for direct tunneling in between end-hosts is given in [Section 9.3](#).

As stated in [Section 6.11](#) the tunnel protocol should not impose any new vulnerability to the nodes implementing the tunnel protocol than what is already present in existing multi-access IPv6 networks where multiple hosts are served by the same router or possible multiple routers.

Here it is implicitly assumed that the tunnel server(s) take the role of default routers and the end-nodes using Zero-Configuration

Tunneling for IPv6 connectivity the role of hosts in multi-access environments.

Nielsen, et. al.

[Page 12]

9.2.1. Threats to end-hosts

Given that all IPv4 sources of protocol-41 tunneled packets can be assumed to be legitimate, threats stemming from encapsulated packets sourced by nodes (addresses) other than nodes (addresses) which the end-hosts recognize as tunnel servers (identified by addresses) can, if not already an intrinsic part of the zero-configuration protocol, easily be mitigated by the implementation of appropriate differentiated (source addresses) drop policies in the end-hosts, i.e., accept only if source is tunnel server.

In current multi-access IPv6 networks hosts need to trust on the benevolence of their default routers as well as hosts must trust that anyone impersonating as a router is indeed one, see, e.g., the trust models and threats described in [11].

Future multi-access IPv6 networks may rely on SEND mechanisms, i.e., mechanisms developed in the SEND WG in order to mitigate the threats described in [11], to establish a trust relationship in between host and routers.

Given that IPv4 source address spoofing is not possible in Zero-Configuration Tunneling sites, then

- for an end-host to trust that packets it perceives as stemming from tunnel servers do actually stem from such - as well as
- for an end-host to trust on the benevolence of its tunnel servers,

it suffices that a trustworthy tunnel server end-point discovery mechanism, read discovery of benevolent tunnel servers IPv4 address(es), is implemented.

In open environments, such as, e.g., the 3GPP environment, it is assumed a prerequisite that a trustworthy zero-configuration tunnel server end-point discovery mechanism is implemented.

9.2.2. Threats to Tunnel Servers

Zero-Configuration Tunneling may be deployed over very large IPv4 sites with low density of active tunnel clients but with a very high number of dormant, but potential tunnel clients. Therefore Denial of Service prevention by strict over provisioning of Tunnel Server capacity is unlikely to be performed.

9.2.2.1. Tunnel State related risks

If the Tunnel Server relies on state to be kept per tunnel client that it serves, the server risks resource exhaustion.

In this situation it is a security prerequisite that no node, whether located within or outside the Zero-Configuration Tunneling site, can

initiate initialization of tunnel state for other entities than itself.

Given this prerequisite, then for tunnel server resource exhaustion by tunnel state creation to be categorized as a security threat, rather than a case of under provisioning, requires a large number of tunnel clients to operate in co-action. This is thus not considered a plausible threat.

9.2.2.2. Traffic related risks

Tunnel encapsulation is recognized as being more resource demanding than mere packet forwarding. Given the same traffic load a Tunnel Server must thus be more generously provisioned than a corresponding router for it not to be more likely to get overthrown by large unexpected amounts of traffic than the router.

The authors have found no plausible treats to the tunnel service, due to large unexpected amounts of traffic needing encapsulation, which can be classified as a security threat rather than a case of under provision. This regardless of whether the traffic is due to a surge in the density of active tunnel clients or due to a surge in the traffic streams set-up by active clients.

9.2.2.3. Packet Delivery related threats

One potential risk related to packet delivery has been identified. This risk is the equivalent of the threat to routers in multi-access environments described in [11], [Section 4.3.2](#).

The risk is associated with the special case where the tunnel protocol requires special resource demanding and/or temporary state creation actions to be taken by the Tunnel Server for delivery of packets destined for not recently addressed Tunnel Clients. The situation where such actions must be performed for all packets at all times is considered to be unlikely. The actions required could be buffering of packets while the reachability of the destined node is being verified.

In case a malicious node (located either within or outside the zero-configuration site) is able to continuously send packets to continuously changing nodes, which by the Tunnel Server is perceived as being existing or potential client nodes, the malicious node may be able to exhaust the Tunnel Servers capability of delivering packets by saturating the packet buffering mechanism and the reachability state table as well as by keeping the Tunnel Server busy determining the reachability state of the ever changing client nodes.

9.3. Implications of Direct Tunneling

In case direct tunneling in between end-hosts is provided by the tunneling protocol, it will not (as described in [Section 9.2.1](#)) be possibly for end-hosts to filter out received Protocol-41 encapsulated packets based on whether the IPv4 source is an address belonging to a trusted Tunnel Server as such behavior evidently would break direct tunneling.

As other end-hosts generally are non-trusted, direct tunneling may thus open up for attacks against IPv6 ingress filtering.

Detailed analysis of the validity of this threat will have to depend on the particular zero-configuration solution.

[10. Acknowledgments](#)

Prior work by J. Mulahusic on the requirements for UE tunneling has been considered in the initial stage of the work.

This work has benefited from input and comments provided by Fred Templin in the initial phase of the work.

Thanks are due to Pekka Savola for many helpful comments and suggestions for improvements.

Corresponding work on assisted-tunneling, [5], has been an inspiration for the Zero-Configuration Tunneling work.

The authors would like to acknowledge the European Commission support in the co-funding of the Euro6IX project, where part of this work is being developed.

[11. Authors Addresses](#)

Mario Morelli
Telecom Italia Lab.
Via Guglielmo Reiss Romoli, 274
I-10148 Turin,
Italy

Phone: +39 011 228 7790
Fax: +39 011 228 5069
Email: mario.morelli@tilab.com

Karen Egede Nielsen
Ericsson
Skanderborgvej 232
8260 Viby J

Denmark

Nielsen, et. al.

[Page 15]

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

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

Jonne Soininen
Nokia
Linnoitustie 6
02600 ESPOO
Finland

Phone: +358 7180 08000
EMail: jonne.soininen@nokia.com

Juha Wiljakka
Nokia
Visiokatu 3
33720 TAMPERE
Finland

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

12. Changes from [draft-nielsen-v6ops-zeroconf-goals-00.txt](#)

- Aimed to clarify what is meant by full suite of native Ipv6 connectivity functions.
- Aimed to clarify what is meant when referring to IPv6-in-IPv4 encapsulation as defined in [1].
- Aimed to clarify the accepted limitations of the IPv6 connectivity service provided by Zero-Configuration Tunneling compared to the full suite of native IPv6

- Added that it is a prerequisite that Zero-Configuration Tunneling can work over IPv4 networks where Ipv4 multicast is not supported.
- Clarified goal, [Section 6.3](#), on usage of native IPv6 when available.
- Added goal on load sharing and scalability, [Section 6.10](#).
- Added [Section 8](#), discussing state or no state.
- Changed [Section 6.8](#) on reachability detection to refer to unicast Neighbor Discovery techniques rather than IPv6 NUD.
- Added material to the section on threats to tunnel servers, 9.2.2.
- Added section, [Section 9.3](#), discussing particular security considerations for usage of direct tunneling.
- Added text under Acknowledgements
- Specified references as informative.
- Performed various Editorial changes.

13. Informative References

- [1] Nordmark, E., Basic Transition Mechanisms for IPv6 Hosts and Routers, [draft-ietf-v6ops-mech-v2-04.txt](#) (work in progress), July 2004.
- [2] Lind, M., Scenarios and Analysis for Introducing IPv6 into ISP Networks, [draft-ietf-v6ops-isp-scenarios-analysis-03.txt](#) (work in progress), June 2004.
- [3] Huitema, C., Evaluation of Transition Mechanisms for Unmanaged Networks, [draft-ietf-v6ops-unmaneaval-03.txt](#) (work in progress), June 2004.
- [4] Wiljakka, J., Analysis on IPv6 Transition in 3GPP Networks, [draft-ietf-v6ops-3gpp-analysis-10.txt](#) (work in progress), May 2004.
- [5] Durand, A., Requirements for assisted tunneling, [draft-ietf-v6ops-assisted-tunneling-requirements-00.txt](#) (work in progress), June 2004.
- [6] Palet, J., Analysis of IPv6 Tunnel End-point Discovery Mechanisms, [draft-palet-v6ops-tun-auto-disc-01.txt](#) (work in progress), June 2004.
- [7] Wasserman, M., Recommendations for IPv6 in 3GPP standards, [RFC 3314](#).

- [8] Loughney, J., IPv6 Node Requirements, [draft-ietf-ipv6-node-requirements-10.txt](#) (work in progress), August 2004.

- [9] IAB, IESG, IAB/IESG Recommendations on IPv6 Address Allocations to Sites, [RFC 3177](#).
- [10] Savola, P., Security Considerations for 6to4, [draft-ietf-v6ops-6to4-security-04.txt](#) (work in progress), July 2004.
- [11] Nikander, P., IPv6 Neighbor Discovery (ND) Trust Models and Threats, [RFC 3756](#).
- [12] Narten, T., Privacy Extensions for Stateless Address Autoconfiguration in IPv6, [RFC 3041](#).
- [13] Narten, T., Neighbor Discovery for IP Version 6 (IPv6), [RFC 2461](#).
- [14] Jordi, P., [draft-palet-v6ops-solution-tun-auto-disc-00](#) (work in progress), September 2004.

Appendix A Out of Scope

[Editor's Note: This appendix can be removed in a future revision of this document]

The following issues have been considered as being out of scope of this work.

DNS:

DNS registration of the IPv6 addresses allocated to dual stack nodes while deploying Zero-Configuration Tunneling for IPv6 connectivity.

Mobile IPv6:

Support of Mobile IPv6 usage over the tunnel-link; here under potential mechanisms required to support MIPv6 movement detection as well as fast tunnel set-up for Mobile IPv6 session survivability.

Appendix B Open Issues

[Editor's Note: This appendix can be removed in a future revision of this document]

Allow NATs that support proto-41 forwarding:

Should the no NATs assumption be relaxed to allow only NATs which support proto-41 forwarding ?

Intellectual Property Statement

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 IETF's procedures with respect to rights in IETF Documents can be found in [RFC 3667 \(BCP 78\)](#) and [RFC 3668 \(BCP 79\)](#).

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.

Copyright Statement

Copyright (C) The Internet Society (2004). 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.

Disclaimer of Validity

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 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.

This Internet-Draft expires March 2, 2005.

