Internet Engineering Task Force P. Savola
Internet-Draft CSC/FUNET
Expires: September 30, 2004 J. Soininen
Nokia
Apr 2004

# Evaluation of v6ops Tunneling Scenarios and Mechanisms
**draft-savola-v6ops-tunneling-01.txt**

## Status of this Memo



## Copyright Notice



## Abstract

This memo analyses the v6ops scenarios/analysis work (Unmanaged, 3GPP, ISP and Enterprise) for their requirements for tunneling solutions, and analyses the proposed mechanisms on how they might fit in these requirements, and discusses possibilities for choosing solution(s).

Table of Contents

## 1. Introduction

This memo analyzes the v6ops scenarios/analysis work (Unmanaged [1], 3GPP [2], ISP [3], Enterprise [4]) at a bit more length, summarizes the exact requirements for tunneling in the scenarios, and analyzes how different mechanisms would fit into those specific tunneling requirements.

The mechanisms analyzed in this document are Teredo [5], ISATAP [6], STEP [7], and TSP [8] The latter two are examples of the tunnel server/broker concept [9]. Already-specified, and in many cases applicable tunneling mechanisms are 6to4 [10] and configured tunneling [11]. Some others include so-called "6over4" [12] and Layer 2 Tunneling Protocol (L2TP) [13].

## 2. The Selection Process

It is strongly desirable to be able to recommend as few mechanisms as reasonably possible. This is an important tradeoff to consider. However, it is likely that one has to give up some features to achieve that goal.

However, it is recognized that there are implementations out there: therefore, one can submit I-Ds specifying currently implemented (but only that; no additions) mechanisms to RFC-editor as invidial contributions for Informational or Experimental RFC. These documents will include a clear IESG Note and/or an applicability statement that these are not recommended mechanisms. The goal of this process is to improve interoperability of the implementations that already exist, and some have deemed fit to implement. XXX: anything about the timing when these can be submitted to the RFC-editor?

## 3. Scenarios

This section described the specific cases identified inside the scenarios where a form of tunneling is desirable.

### 3.1 3GPP Networks

There are two closely related cases:

1. Providing IPv6 connectivity to User Equipment (UE) when it roams to an another operator's network, and the other operator does not support IPv6 PDP contexts.

2. Providing IPv6 connectivity to UEs when the "home" 3GPP operator has not deployed even minimal IPv6 PDP context support yet in its network -- but would like to enable IPv6 connectivity through a

transition mechanism which can be transparently overlayed on the existing IPv4 3GPP network.

Note that the current document strongly recommends at least initial IPv6 PDP context support: this only requires support from HLR, SGSN and one GGSN.

The case where there is no IPv6 connectivity at all in the 3GPP network is considered out of scope (but a solution can be achieved using one of the Unmanaged connectivity mechanisms if appropriate).

In this scenario, it is assumed that the users are typically using private IPv4 addresses, but there is no NAT in the path between the different users in the same 3GPP network.

There are a couple of differences with the two cases described above: in the first case, the operator has deployed IPv6 PDP context support, as recommended, but some other ISP has not: waiting for all the operators to deploy IPv6 PDP context support prior to launching the service would be unacceptable -- so there has to be a solution to this.

Also, note that the tunnel from the foreign 3GPP network is terminated at the local 3GPP operator's network, inducing delay and a bottle-neck; "direct connectivity" optimization is not much different whether the tunnel would be terminated at a tunnel server, or communicating directly.  In the second case, one could argue that the deployment should only be temporary, with at least minimal IPv6 PDP context support being the goal.

In any case, the critical question in this specific scenario is, how desirable is it to have direct tunneling between the UEs in the same 3GPP network -- instead of a slightly longer "leg" through a tunnel server?  Clearly, this would be desirable -- but not a strict requirement.  Also, to re-iterate the (strongly) recommended deployment scenario in the 3GPP networks is the native (even if minimal) IPv6 PDP context support; accepting non-direct connectivity could be an acceptable tradeoff in the early adopter phase, also encouraging to move to proper IPv6 support.

## 3.2  Unmanaged Networks

The unmanaged scenarios seem to include two cases, with a subcase, totalling 3 different cases; these are derived from both unmanaged and ISP scenario documents.

1.  When the user's direct ISP does not offer any IPv6 service at all, and connectivity must be obtained automatically.

1. When the connectivity must be obtained with as little infrastructure as possible, without any signups or contracts, etc.

2. When the connectivity can be obtained from a third party, through a sign-up, contract, etc. -- for higher manageability, more control, increased security benefits, or for other reasons. (Whether such "3rd party connectivity" is feasible or good enough is a separate question.)

2. When the user's direct ISP would like to offer IPv6 service, but would require tunneling for some reason (e.g., access router, access link, or the gateway in the unmanaged network is incapable of IPv6). In this case the options are basically tunneling from the gateway, a separate IPv6 gateway or tunneling from the host(s).

It should be noted that a solution to problem 1.2) would solve problem 2) as well, but a solution to problem 2) would not necessarily be adequate for solving 1.2).

NAT traversal must be supported. Dynamic IPv4 addresses must be supported -- but the solution does not necessarily have to be better than dynamic IPv4 addresses are today; e.g., a dynamic IPv6 address would be acceptable as well.

When the gateway has been upgraded to support IPv6 (but access router or link has not, resulting in tunneling from the gateway, NAT traversal doesn't need to be used as often. That is, often the gateway has a public IPv4 address on its Internet-side interface, so a mechanism like 6to4 can be used on the gateway to provide "native-like" IPv6 support to the unmanaged network. However, there are a number of cases where this is not sufficient -- for example, when the gateway is connecting to a privately-addressed access network shared by multiple ISPs. In such case, the gateway may be unable to obtain IPv6 connectivity.

It seems obvious that direct tunneling between users is required at least when there is no ISP support -- to minimize the latency increase, and to decrease the bandwidth aggregation. However, it is not a strict requirement with hosts in the same ISP: in such a case, the slight increase in latency and concentration of traffic is probably manageable. It should be re-iterated that tunneling is not meant to be a permanent solution, and accepting that the connectivity may not be fully optimal should be acceptable.

One should note that direct connectivity that traverses NATs, as is requirement here, is a very difficult thing to do right; essentially,

this is what Teredo is doing. On the other hand, if the ISP is offering service which does not provide direct connectivity between hosts, the use of another mechanism (such as 6to4 or Teredo) "on the side" could perhaps provide at least partial direct connectivity.

### 3.3 Enterprise Scenarios

This scenario/analysis work is still incomplete, so it is not fully addressed in this memo.

The only scenario which is obvious at the moment is when the enterprise wishes to deploy IPv6 without changing the gear to be dual-stack, without injecting IPv6 into existing VLANs [14], or without adding additional IPv6 routers in the VLANs. In particular, this may be the case when the IPv6 deployment is "sparse" -- because if it was sufficiently "dense", it would make sense to expand the infrastructure in one of the several ways. It would be nice if the tunneling between the nodes is direct from node-to-node, but this is not a strict requirement.

There are at least three subcases of this scenario:

1. When the enterprise does not NAT between different parts of the internal network. This case is useful to separate from case 2 merely because it is generally the common case, where simplicity is the most highly valued.

2. When NATs exist within the enterprise network, e.g., to connect branch offices to the corporate network. Hence, some form of NAT traversal must be supported.

3. When the enterprise internally uses multicast applications/ protocols that they want to transition to IPv6. Here the enterprise has already deployed intra-domain IPv4 multicast to support this stage. As with case 1, there are no internal NATs in this case since IPv4 multicast itself does not cross NATs.

Case 3 is special because the number of enterprises with all-reaching IPv4 multicast deployment is low. While IPv4 multicast support is typically more commonplace than (even unicast) IPv6 support, the enterprises are typically in a better position than those which would not have IPv4 multicast at all. Therefore, native IPv6 multicast support is a possibility, while using tunneling mechanism(s) which support IPv6 multicast by tunneling is a simple short-term deployment fix.

The need for direct connectivity in all cases is strong due to two factors:

- o The high end-to-end bandwidth requirements for enterprise applications, e.g., many clients accessing various servers, such that bottlenecks occur if all traffic must be funneled through one or a few tunnel endpoints. However, note that heavy usage is quite a strong argument for native (or hierarchical/distributed) IPv6 deployment.

- o The use of collaborative applications, possibly including high-bandwidth streams such as audio/video.

Finally, most enterprise networks currently lack IPv6 expertise, and have a great need to keep costs low. Hence simplicity and low infrastructure costs are also required. (This is not really specific to enterprises, though.)

### 3.4 ISP Scenarios

ISPs have two possible requirements for tunneling: either inside their own infrastructure (e.g., through a non-upgraded core network) or to their peers or upstreams, or towards customers.

Internal tunneling requirements can be satisfied with configured tunneling or the use of [15] as described in [3] so there is no need to discuss that in this memo. Similarly, tunneling requirements towards peers or upstreams are satisfied by configured tunneling only.

As for tunneling towards customers, ISPs do not have specific scenarios which need to be addressed which haven't been already mentioned: some ISPs want to provide IPv6 connectivity, possibly over a tunnel, to their unmanaged or enterprise customers, or ISPs. However, these requirements have already been discussed; only two particular considerations should be explicitly mentioned:

1. Mechanisms used must be very simple if we want them to be adopted by many ISPs,

2. Very few ISPs want to be providing service, for free at least, to the users which are not their customers. Therefore being able to identify the user as your own customer is very important. (Whether this is done by explicit user authentication or basing on IP-address -based knowledge of whether the user is your customer is a detail.)

### 3.5 Additional Scenarios: IP Mobility

Recently, there some concerns have been raised about additional

scenarios work, which might be partially worked under v6ops WG. One such is work on the scenarios where IP nodes are not stationary, i.e., are expected to change IPv4 or IPv6 address relatively rapidly. In many cases this implies that either MIPv4 or MIPv6 is being run to ensure a relatively stable IP address being available, and to make the changes in IP addresses as transparent as possible.

This is not a formal scenario as such, but in these events it would be extremely desirable to have minimal amount of signalling when the attachment point (whether a care-of or home address) changes -- to ensure seamless IP mobility.

# 4. Scenarios and Mechanisms Evaluation

## 4.1 Scenarios Evaluation

| | NAT-T | Direct | ISP | Secure | Simple | Low Overhead | Mcast | Gateway |
|---|---|---|---|---|---|---|---|---|
| Unman 1.1 | * | * | N | - | - | - | - | - |
| Unman 1.2 | * | N | * | -/* | - | - | - | - |
| Unman 2 | * | - | * | - | * | - | - | - |
| 3GPP 1 | N | -/* | * | - | * | * | - | N |
| 3GPP 2 | N | -/* | * | - | * | * | - | N |
| Enterprise 1 | N | -/* | * | -/* | * | - | - | -/N |
| Enterprise 2 | * | -/* | * | - | * | - | - | -/N |
| Enterprise 3 | N | -/* | * | - | * | - | * | -/N |

Legend: *  = MUST; -  = Nice to have; N  = No

NAT traversal specifies whether NAT traversal is a requirement.

Direct tunneling states whether direct tunneling between the nodes using the same mechanism is a requirement.

"ISP" indicates whether the scenario assumes that the ISP provides IPv6 support.  That is, whether the mechanism must be able to operate (at least to a degree) without explicit support from an ISP which is identifying the user.

"Secure" is an overly simplistic term to evaluate whether the scenario requires some specific amount of security.  Here, security implies security issues which may not be trivially fixable, whether the operational mode or in the mechanism, such that it is difficult to secure, whether it creates new significant threats to either IPv4 or IPv6 infrastructures, etc.

Simple is a term which refers to the simplicity of the protocol or the operational model in which it operates; simple protocols and

specifications are easy to understand, evaluate, implement, and deploy, and thus preferable to more complex ones.

"Low overhead" refers to both minimal encapsulation -- as few added bytes in the messages or signalling as possible -- and signalling latency (e.g., the number of messages/round-trips to set-up, change, or tear down a tunnel).

Multicast states whether IPv6 multicast should be supported.

Gateway refers to whether the scenario also requires that an upgraded gateway, to provide IPv6 connectivity to the local link(s), must be supported.

### 4.2 Mechanisms Evaluation

One should note that we are not evaluating the specific version of the specification, but rather the mechanism in a more generic sense ("which features could this mechanism easily be made to work with?").

| | NAT-T | Direct | ISP | Secure | Simple | Low Overhead | Mcast | Gateway | Impl. | Depl |
|---|---|---|---|---|---|---|---|---|---|---|
| Teredo | Y | Y | N | Y | N | N | N | N | R | Y |
| ISATAP | N | Y@ | Y | N/R | R | Y | N | R | Y | R? |
| TSP | Y | N | Y? | Y | R | N | Y# | Y | R | R? |
| STEP | Y | N | Y | Y | Y/R | Y | Y# | R | N | N |
| L2TP | Y | N | Y | Y | N | N | Y# | R? | Y | Y |
| 6to4 | N | Y$ | N | N | Y | Y | N | Y | Y | Y |
| 6over4 | N | Y@ | Y | N/R | Y | Y | Y | R | R | N |

@: intra-site, when no NATs are in the path
#: in a non-optimal fashion
$: only between public IPv4 addresses

Legend: Y  = Yes; R  = Relatively good; N  = No

Notes: 6to4 does not work behind a NAT, so it is not applicable in 3GPP scenarios, and practically also not applicable in Enterprise scenarios.  6over4 requires IPv4 multicast infrastructure, so it is practically only applicable in Enterprise scenario 3.

NAT traversal states whether the mechanism is able to perform full NAT traversal.

Direct tunneling states whether the mechanism is able to provide direct tunneling between the nodes using the same mechanism.

"ISP" indicates whether ISP(s) must explicitly, identifying the user,

support the mechanism in order for the mechanism to operate (at least to a degree).

"Secure" is an overly simplistic term to evaluate whether the mechanism has obvious security issues which may not be trivially fixable, whether the operational mode or in the mechanism, such that it is difficult to secure, whether it creates new significant threats to either IPv4 or IPv6 infrastructures, etc.

Simple is a term which refers to the simplicity of the protocol or the operational model in which it operates; simple protocols and specifications are easy to understand, evaluate, implement, and deploy, and thus preferable to more complex ones.

"Low overhead" refers to both minimal encapsulation -- as few added bytes in the messages or signalling as possible -- and signalling latency (e.g., the number of messages/round-trips to set-up, change, or tear down a tunnel).

Multicast states whether the mechanism supports IPv6 multicast. Bidirectional tunnels support it but do not leverage the underlying link-layer for packet duplication, as e.g., so-called 6over4 does.

"Gateway" refers to whether the mechanism can be used on a router to provide "native-like" IPv6 connectivity to its link(s), or whether the mechanism can only be used at each individual host.

Implemented refers how widely the mechanism has been implemented (e.g., no implementations, one implementation, multiple interoperable implementations), and how large part of the mechanism has been implemented.

"Widely deployed" refers to how widely the mechanism has been deployed: is it in use as deployed by multiple vendors, or in many operational scenarios?

### 4.3  Evaluation Summary

By combining the two matrices, we obtain the following:

o  Unmanaged scenario 1.1) requires Teredo.

o  Unmanaged 1.2) requires STEP, TSP or L2TP.

o  Unmanaged 2) requires STEP or TSP.

o  3GPP 1) can be filled by STEP or ISATAP; if a higher level of security is required, ISATAP may not be applicable.

- o 3GPP 2) can be filled by STEP or ISATAP; only ISATAP if direct tunneling is a MUST requirement.

- o Enterprise 1) requires STEP, TSP, or ISATAP.

- o Enterprise 2) requires STEP or TSP.

- o Enterprise 3) requires STEP, TSP, or 6over4.

This evaluation indicates that Teredo is needed, no matter what. STEP can satisfy all the other requirements, and TSP almost all. So, the minimum number of recommended "new" mechanisms appears to be 2: Teredo and something else.

6to4 can (continue to) be used when the gateway has been upgraded in all the unmanaged scenarios; this does not work in every case, though.

On the other hand, if direct connectivity was a strict requirement, the minimum number of "new" mechanisms would have to be 3, as Teredo and ISATAP would not be sufficient on their own.

Actually, it would be preferable in some sense to combine TSP and STEP as they are best applicable in quite similar scenarios. This might allow one to combine the best features of both proposals. If an auto-discovery feature would be added to that mechanism, we would have a very powerful mechanism which would apply well in almost all scenarios.

## 4.4 Features of the Mechanisms

If we would assume that we'd prefer to recommend two solutions, in addition to 6to4 and configured tunneling, which are already there, we would have to judge, based on the scenario requirements, three critical features:

1. NAT traversal

2. Direct connectivity inside a site/ISP/enterprise

3. Operation with third party ISPs

In several scenarios, it's impossible to have both 1) and 2) without a relatively complex solution. Similarly, there is often a conflict with 2) and 3).

Therefore, we must be able to make a decision which features/ requirements we are willing to give up in which specific scenarios,

or whether we have to specify more mechanism to satisfy all the features.

For example, by picking Teredo and {TSP or STEP}, we would have to give up direct connectivity inside the 3GPP, Enterprise and the unmanaged scenarios where the ISP is offering service -- except where it would automatically come along where Teredo (or 6to4) service is active in any case.

On the other hand, ISATAP is unable to properly perform NAT traversal, and it is not designed to securely interact with third party ISPs. By picking Teredo and ISATAP, we would not have a solution to operate with 3rd party ISPs, ISPs which do not properly secure their borders, or in the cases where NAT traversal is required and Teredo is seen as too cumbersome a choice. However, Teredo does provide a means to interoperate with a specifically crafted, simplistic Tunnel Server implementation; if we assume that it would be acceptable to recommend implementation and deployment of Teredo to be able to use a tunnel server service, implementing a stub tunnel server could provide a means to achieve support in the 3rd party ISP case.

Direct connectivity inside (a vague definition of) a "site" is sometimes seen as attractive, e.g., to be done with ISATAP. There are two possible arguments: sparse and dense deployment. In the "sparse" deployment case, the requirements of the site can be met with an (auto-discovered) tunnel server solution. In the "dense" deployment case, when direct connectivity could be desirable because the tunnel servers would get overloaded or the bandwidth could be wasted, there is a strong reason to go for a dual-stack (whether partial, e.g., with hierarchical tunnel servers or full deployment) solution instead. Therefore, it seems that when other choices are feasible, "intra-site" direct connectivity is not a progressive way forward.

As Teredo is the only solution available for scenario 1.1), it seems that it cannot be left out.

## 5. Conclusions

There seems to be clear need for Teredo. There is also clear desire to keep using 6to4 in the cases where it's applicable.

There seems to be clear need for a tunnel server protocol which is able to traverse NATs and work with dynamic IPv4 addresses. This tunnel server should be able to automatically discover the server address if the service is provided by the ISP.

Direct connectivity is desirable but difficult to provide in a few specific scenarios when considering the other trade-offs.  However, global direct connectivity can be obtained with 6to4 and Teredo; local direct connectivity, inside 3GPP, ISP or enterprise is something that one could be able to live without.

(Further TBD.)

## 6.  Security Considerations

This memo analyses the tunneling scenario requirements and mechanisms trying to address these requirements.  As such, it does not have significant security considerations.  When considering which mechanism(s) to adopt, the security properties of the mechanisms vary considerably -- and this has to be taken in consideration in the evaluation and selection. However, mechanism-specific considerations are to be addressed in the respective documents.

## 7.  Acknowledgements

This memo was written from scratch at IETF59, and it has been improved based on feedback received both prior, during, and after the session by numerous people.

Alain Durand, Alain Baudot, Jeroen Massar, and Hesham Soliman sent substantive feedback, helping in improving this memo. Dave Thaler contributed the majority of enterprise scenarios, and provided other feedback as well.

## 8.  References

### 8.1  Normative References

[1]  Huitema, C., "Evaluation of Transition Mechanisms for Unmanaged Networks", draft-ietf-v6ops-unmaneval-01 (work in progress), February 2004.

[2]  Wiljakka, J., "Analysis on IPv6 Transition in 3GPP Networks", draft-ietf-v6ops-3gpp-analysis-09 (work in progress), March 2004.

[3]  Lind, M., "Scenarios and Analysis for Introducing IPv6 into ISP Networks", draft-ietf-v6ops-isp-scenarios-analysis-01 (work in progress), February 2004.

[4]  Bound, J., "IPv6 Enterprise Network Scenarios", draft-ietf-v6ops-ent-scenarios-01 (work in progress), February 2004.

[5]  Huitema, C., "Teredo: Tunneling IPv6 over UDP through NATs", draft-huitema-v6ops-teredo-01 (work in progress), February 2004.

[6]  Templin, F., Gleeson, T., Talwar, M. and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", draft-ietf-ngtrans-isatap-20 (work in progress), February 2004.

[7]  Savola, P., "Simple IPv6-in-IPv4 Tunnel Establishment Procedure (STEP)", draft-savola-v6ops-conftun-setup-02 (work in progress), January 2004.

[8]  Blanchet, M., "IPv6 Tunnel Broker with the Tunnel Setup Protocol(TSP)", draft-blanchet-v6ops-tunnelbroker-tsp-00 (work in progress), February 2004.

## 8.2 Informative References

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

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

[11]  Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", draft-ietf-v6ops-mech-v2-02 (work in progress), February 2004.

[12]  Carpenter, B. and C. Jung, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", RFC 2529, March 1999.

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

[14]  Chown, T., "Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks", draft-chown-v6ops-vlan-usage-00 (work in progress), October 2003.

[15]  Clercq, J., "Connecting IPv6 Islands across IPv4 Clouds with BGP", draft-ooms-v6ops-bgp-tunnel-02 (work in progress), March 2004.

## Authors' Addresses

Pekka Savola
CSC/FUNET

Espoo
Finland

EMail: psavola@funet.fi

Jonne Soininen
Nokia

EMail: jonne.soininen@nokia.com

## Intellectual Property Statement



## Disclaimer of Validity



## Copyright Statement



## Acknowledgment

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