

v6ops Working Group
Internet-Draft
Intended status: Informational
Expires: March 4, 2011

G. Van de Velde
O. Troan
Cisco Systems
T. Chown
University of Southampton
August 31, 2010

Non-Managed IPv6 Tunnels considered Harmful
<[draft-vandvelde-v6ops-harmful-tunnels-01.txt](#)>

Abstract

IPv6 is ongoing and natively being deployed by a growing community and it is important that the quality perception and traffic flows are as optimal as possible. Ideally it would be as good as the IPv4 perceptive experience.

This paper looks into a set of transitional technologies where the actual user has IPv6 connectivity through the means of IPv6-in-IPv4 tunnels. A subset of the available tunnels has the property of being non-managed (i.e. 6to4 [[RFC3056](#)] and Teredo [[RFC4380](#)]).

While native IPv6 deployments will keep growing it is uncertain or even expected that non-managed IPv6 tunnels will be providing the same user experience and operational quality as managed tunnels or native IPv6 connectivity.

This paper will detail some considerations around non-managed tunnels and will document the harmful element of these for the future growth of networks and the Internet.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on March 4, 2011.

Internet-Draft

Non-Managed Tunnels are Harmful

August 2010

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	4
2.	Managed Tunnelling Properties	4
3.	Tunnel User Experience Views	5
4.	Why do non-managed tunnels exist?	5
5.	Non-Managed Tunnelling Properties	6
5.1.	Performance	6
5.2.	Topological Considerations	7
5.3.	Operational Provisioning	7
5.4.	Operational Troubleshooting	7
5.5.	Security	8
5.6.	Content Services	8
6.	Conclusion	9
7.	IANA Considerations	9
8.	Security Considerations	9
9.	Acknowledgements	9
10.	References	9
10.1.	Normative References	9
10.2.	Informative References	9
	Authors' Addresses	10

1. Introduction

While the Internet and networks continue to grow, it is found that the deployment of IPv6 within these networks is an ongoing activity due to global IPv4 address pool depletion. An important aspect is that the quality, availability and security of the IPv6 connectivity is as good as possible, and when possible even more advanced as the IPv4 connectivity.

Historically IETF has been facilitating a variety of technologies and procedures to deploy IPv6 successfully in addition to existing IPv4 connectivity. In general and for the sake of this draft these procedures and technologies can be divided into three major groups: (1) native (dual-stack) IPv6, (2) Tunnelled IPv6 and (3) Translation. While native IPv6 deployments has been steadily growing, the value and the drawbacks of some tunnelling mechanisms can be investigated. Translational techniques provide a total different aspect of considerations and applicability and is beyond the scope of this paper. Transition techniques have been and still are in many cases important for the bootstrapping of IPv6, this paper will look into a range of property aspects of non-managed IPv6 tunnelling techniques. Areas of perverse traffic paths, security considerations, lack of business incentives to run tunnel relays/gateways, black holing and ownership of supportability will be analysed. Finally the paper will conclude that for the growth of IP connectivity, non-managed tunnelling techniques are considered harmful especially for those that want to access applications over the network through pervasive IPv6 connectivity and have no particular interest on how connectivity to the applications is established (IPv4, translation, IPv6, etc...)

2. Managed Tunnelling Properties

A managed tunnel is a tunnel has a few properties supporting the ownership and quality of the tunnel.

When using a managed service, there tends to be an administrative entity which provides quality assurance and can take action if users of the service are experiencing a degraded service. An example would be 6rd tunnels [[RFC5969](#)]

In addition there is a general trust awareness and agreement between the user of the managed tunnel service and the provider of the managed tunnel service.

3. Tunnel User Experience Views

The tunnel experience can be divided into three distinct segments: (1) the End-user view, (2) the Enterprise View and (3) the Service Provider View.

The End-user view exists mainly out of two different user profiles. The technical power user and the general user mainly trying to reach their favourite application on the network. The technical power user may have a particular interest to run IPv6 as a transport mechanism, and if his upstream service provider has no native IPv6 connectivity available, then non-managed tunneling mechanisms may provide a solution satisfying to the immediate needs of the technical power user. Alternatively, the general user trying to reach his favourite network application, may have no interest or awareness of his IPv6 usage, particularly when non-managed tunnels are utilized.

The Enterprise View is a more traffic flows and network oriented positioning. When the upstream service provider does not have an IPv6 offer, then the enterprise may start to rely upon a technology as 6to4 [[RFC3056](#)]. However this technology has the potential of creating quite perverse traffic paths when user want to reach

applications on the Internet. When user would like to reach other 6to4 [[RFC3056](#)] users, then more optimized traffic paths, generally following the IPv4 traffic paths are realized

The final view is how a Internet service provider looks into non-managed tunnel usage. A service provider may decide to deploy a 6to4 relay to increase the IPv6 quality of their customers. This a service which require resources (money, maintenance, etc...). Often the 6to4 relay service is not just (always) restricted to only the service providers customers, which as result provides often results in a demotivation to provide quality tunnel relay devices. From a content service provider perspective the usage of non-managed tunnel often results in measurable differences in RTT and reliability in some cases, and hence are reluctant to bring all services to mainstream IPv6 for all users 'just yet'.

[4.](#) Why do non-managed tunnels exist?

Non-managed tunnels exist due to a variety of reasons.

Early adopters: people and organisations with a desire to use new and potentially market disrupting technologies and applications may have a desire to use the latest IP even when the upstream provider doesn't have an available service offering.

Lock-step process to implement IPv6: It is not trivial to move a system or an organisation in lock-step towards IPv6 and the aid of tunnels help in this process.

The utilisation of tunnels aid in providing a de-coupling between infrastructure readiness and application readiness, and hence contribute to the development of both elements.

[5.](#) Non-Managed Tunnelling Properties

The properties of Non-managed tunnels span many different areas. In this section the properties are analysed and segmented within different areas of impact. In each case the comparison is made between native IPv6 connectivity and connectivity through a non-

managed tunnel. A common property of non-managed tunnels is that they often use well-known anycast addresses or other well known addresses and anticipate upon the goodwill of middleware (typically a relay or gateway) device to serve as a tunnel termination point. In some cases, for example a 6to4 relay can be provided by a connected responsible service provider, and hence good quality operation can be guaranteed.

Non-managed tunnels often have asymmetric behaviour. There is an outbound and an inbound connectivity behaviour from the tunnel initiator. It is possible to influence the good quality tunnel behaviour of the outbound connectivity (e.g. by explicit setting of the 6to4 relay), however, influencing good inbound connectivity is often an issue.

[5.1.](#) Performance

Deploying a tunnelling mechanism mostly results in encapsulation and de-capsulation efforts. Often this activity has a performance impact on the device, especially when the device does not use hardware acceleration for this functionality. If the performance impact is scoped into the device its lifetime through performance capacity management then the actual impact is predictive. Non-deterministic tunnels tend to have a non-predictive behaviour for capacity, and hence application and network performance is non-predictive. The key reason for this is the decoupling of the capacity management of the tunnel aggregation devices from the capacity desired by users of the aggregation devices.

During initial IPv6 deployment there have been mainly technical savvy people that have been using non-managed tunnel technologies and it has for many been working well. However, if non-managed tunnelling would be deployed in mass and especially when enabled by default by

CPE vendors or host vendors then those aggregation points could become overloaded and result in bad performance. There are a few measures that can be taken, i.e. upgrade the CPU power of the aggregation device or its bandwidth available, however this may not happen without the right motivation for the operator of the aggregation device (i.e. cash flows, reputation, competitive reasons, etc...).

[5.2.](#) Topological Considerations

Due to non-managed IPv6 tunnels the traffic flows may result in sub-optimal flows through the network topology between two communicating devices. The impact for example can cause increase of the RTT and packet loss, especially considering the availability (or better non-availability) of tunnel aggregation/de-aggregation points of certain topological areas or realms. The increase of non-managed tunnel usage would amplify the negative impact on good quality connectivity. For many operators of tunnel aggregation/de-aggregation devices there is little motivation to increase the quality and number of available devices within a topological area or logistical realm.

[5.3.](#) Operational Provisioning

Some elements regarding provisioning of both managed and non-managed tunnels can be controlled, while others are beyond control or influence of people and applications using tunnels. To make applications highly reliable and performing, all elements within the traffic path must provide an expected quality service and performance. For managed tunnels, the user or provider of the tunnel can exercise a degree of operational management and hence influence good quality behaviour upon the tunnel especially upon the aggregation and de-aggregation devices. In some cases even the traffic path between both aggregation and de-aggregation can be controlled. Non-managed tunnels however have less good quality behaviour of both tunnel aggregation and de-aggregation devices because often good quality behaviour is beyond the control or influence of the tunnel user. For non-managed tunnels the tunnel aggregator and/or tunnel de-aggregator are operated by a 3rd party which may have a conflicting interest with the user of the non-managed tunnel. An exception is where the use of the tunnel mechanism is all within one ISP, or ISPs who are 'well coupled', e.g. as happens between many NRENs.

[5.4.](#) Operational Troubleshooting

When one is using non-managed tunnels, then these tunnels may get aggregated or de-aggregated by a 3rd party or a device outside the control of a contracted service provider. Troubleshooting these

devices these devices will be pretty hard for the tunnel user or to

work around the issue.

Also some tools like traceroute don't work too well on asymmetric paths. Another aspect is that tunnels show as one hop in a traceroute, not indicating where problems may be.

[5.5.](#) Security

For an aggregating or de-aggregating tunnel device it is a non-trivial issue to separate the valid traffic from non-valid traffic because it is from the aggregation device perspective almost impossible to know -from- and -towards- about the tunnel traffic. This imposes potential attacks on the available resources of the aggregating/de-aggregating router. A detailed security analysis for 6to4 tunnels can be found in [[RFC3964](#)].

For the user of the non-managed IPv6 tunnel there is an underlying trust that the aggregating/de-aggregating device is a trustworthy device. However, some of the devices used are run by anonymous 3rd parties outside the trusted infrastructure from the user perspective, which is not an ideal situation. The usage of non-managed tunnels increases the risk of rogue aggregation/de-aggregation devices and may be open to malicious packet analyses or manipulation.

From the operator perspective, managing the aggregating/de-aggregating tunnel device, there is a trust assumption that no-one abuses the service. Abuse may impact preset or assumed service quality levels, and hence the quality provided can be impacted

There is also an impact caused by ipv4 firewalling upon non-managed tunnels. Common firewall policies recommend to block tunnels, especially non-managed tunnels, because there is no trust that the traffic within the tunnel is not of malicious intent. This restricts the applicability of some non-managed tunnel mechanisms (e.g. 6to4). Other tunnel mechanisms have found manners to avoid traditional firewall filtering (e.g. Teredo) and open the local network infrastructure for malicious influence (e.g. virus, worms, infrastructure attacks, etc..).

[5.6.](#) Content Services

When providing content services a very important related aspect is that these services are accessible with high reliability, are trustworthy and have a high performance. Using non-managed tunnels makes this a much harder equation and can result in all three elements to suffer negatively, without the ability to uniquely identify and resolve the root cause. The statistical impact of non-

managed tunnels has been measured by some Internet Content providers and is often an additional delay of 0(100msec) (need to add reference here)

This reduces the interest of content providers to provide content services over IPv6 when non-managed tunnels are used.

6. Conclusion

Non-managed tunnels have properties impacting the growth of networks and the Internet in a negative way. Consequences regarding black-holing, perverse traffic paths, lack of business incentive and operational management influence and security issues are a real pragmatic concern, while universal supportability for the tunnel relay services appear to be non-trivial. Due to these elements the usage of non-managed tunnelling can be considered harmful for the growth of networks and the Internet.

7. IANA Considerations

There are no extra IANA consideration for this document.

8. Security Considerations

There are no extra Security consideration for this document.

9. Acknowledgements

10. References

10.1. Normative References

10.2. Informative References

- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", [RFC 3056](#), February 2001.
- [RFC3964] Savola, P. and C. Patel, "Security Considerations for 6to4", [RFC 3964](#), December 2004.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through

- [RFC4798] De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", [RFC 4798](#), February 2007.
- [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", [RFC 5214](#), March 2008.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", [RFC 5969](#), August 2010.

Authors' Addresses

Gunter Van de Velde
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32 2704 5473
Email: gvandeve@cisco.com

Ole Troan
Cisco Systems
Folldalslia 17B
Bergen N-5239
Norway

Phone: +47 917 38519
Email: ot@cisco.com

Tim Chown
University of Southampton
Highfield
Southampton, S017 1BJ

United Kingdom

Phone: +44 23 8059 3257

Email: tjc@ecs.soton.ac.uk

Van de Velde, et al.

Expires March 4, 2011

[Page 10]