IPv6 Operations
Internet-Draft
Expires: January 17, 2005

T. Chown
University of Southampton

July 19, 2004

Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks
draft-chown-v6ops-vlan-usage-01

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.

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 to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

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

This Internet-Draft will expire on January 17, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

Ethernet VLANs are quite commonly used in enterprise networks for the purposes of traffic segregation. This document describes how such

VLANs can be readily used to deploy IPv6 networking in an enterprise, including the scenario of early deployment prior to availability of IPv6-capable switch-router equipment, where IPv6 may be routed in parallel with the existing IPv4 in the enterprise and delivered at Layer 2 via VLAN technology. The IPv6 connectivity to the enterprise may or may not enter the site via the same physical link.

Table of Contents

1. Introduction

Ethernet VLANs are quite commonly used in enterprise networks for the purposes of traffic segregation. This document describes how such VLANs can be readily used to deploy IPv6 networking in an enterprise, including the scenario of early deployment prior to availability of IPv6-capable switch-router equipment, where IPv6 may be routed in parallel with the existing IPv4 in the enterprise and delivered to the desired LANs via VLAN technology.

The IEEE 802.1Q VLAN standard allows separate LANs to be implemented over a single bridged LAN, by inserting "Virtual LAN" tagging or membership information into Ethernet frames. Hosts and switches that support VLANs effectively allow software-based reconfiguration of LANs through configuration of the tagging parameters. The software control means VLANs can be used to alter the LAN infrastructure without having to physically alter the wiring between the LAN segments and Layer 3 routers.

Many IPv4 enterprise networks are utilising VLAN technology. Where a site does not have IPv6-capable Layer 2/3 switch-router equipment, but VLANs are supported, a simple yet effective method exists to gradually introduce IPv6 to some or all of that site's network.

If such a site wishes to introduce IPv6, it may do so by deploying a parallel IPv6 routing infrastructure (which as described below may be a single PC-based IPv6 router), and then using VLAN technology to "overlay" IPv6 links onto existing IPv4 links. This can be achieved without needing any changes to the IPv4 configuration.

The IPv6 connectivity to the enterprise may or may not enter the site via the same physical link, and may be native or tunneled from the external provider to the IPv6 routing equipment.

This VLAN usage is a solution adopted by a number of sites already, and is referenced in our Campus Network IPv6 Transition [2] text.

2. Enabling IPv6 per link

The precise method by which IPv6 would be "injected" into the

existing IPv4 network is implementation specific. The general principle is that the IPv6 router device (e.g. performing IPv6 Router Advertisements [1] in the case of stateless autoconfiguration) is connected to the target link through the use of VLAN capable Layer 2 equipment.

2.1 IPv6 routing

In a typical scenario, one IPv6 router would be deployed, with both an external interface and one or more internal interfaces. The external interface connects to the wider IPv6 internet, and may be dual-stack if some tunnel mechanism is used for external connectivity, or IPv6-only if a native external connection is available.

By connecting the internal interface(s) directly to a VLAN-capable switch, and writing VLAN tags on the packets sent from the internal router interface to the switch, VLAN tagging on the switch can be used to carry tagged traffic across the internal VLAN-capable site infrastructure to IPv6 links that may be dispered widely across the site network.

It is not necessary to do VLAN tagging in all cases. On some Layer 3 switches, IPv6 traffic can directly be distributed to specific ports by adding them to the same protocol-based VLAN (in this case IPv6-based VLANs).

2.2 One VLAN per router interface

The VLAN marking may be done in different ways. Some sites may prefer to use one router interface per VLAN, e.g. if there are three internal IPv6 links, a PC-based IPv6 router with four Ethernet ports could be used, one for the external link and three for the internal links. In such a case one switch port would be needed per link, to receive the connectivity from each router port.

In such a deployment, the IPv6 routing could be cascaded through lower tier internal IPv6-only routers. Here, the internal facing ports on the IPv6 edge router may feed other IPv6 routers over IPv6-only links which in turn inject the IPv6 connectivity (the /64 size links and associated Router Advertisements) into the VLANs.

2.3 Collapsed VLANs on a single interface

Using multiple IPv6 routers and one port per IPv6 link (i.e. VLAN) may be unnecessary. Many devices now support VLAN tagging based on

virtual interfaces such that multiple IPv6 VLANs could be assigned from one physical router interface port. Thus it is possible to use just one router interface for "aggregated" VLAN trunking from a switch. This is a far more interesting case for a site planning the introduction of IPv6 to (part of) its site network.

This approach is viable while IPv6 traffic load is light. As traffic volume grows, the single collapsed interface could be extended to

utilise two or more physical ports, where the capacity of the IPv6 router device allows it.

2.4 Congruent IPv4 and IPv6 Subnets

Such a VLAN-based technique can be used to deploy IPv6-only VLANs in an enterprise network. However most enterprises will be interested in dual-stack IPv4-IPv6 networking.

In such a case the IPv6 connectivity may be injected into the existing IPv4 VLANs, such that the IPv4 and IPv6 subnets are congruent (i.e. they coincide exactly when superimposed). Such a method may have desirable administrative properties, e.g. the devices in each IPv4 subnet will be in the same IPv6 subnets also. This is the method being used in our Campus Network IPv6 Transition [2] text.

Further, IPv6-only devices may be gradually added into the subnet without any need to resize the IPv6 subnet (which may hold in effect an infinite number of hosts in a /64 in contrast to IPv4 where the subnet size is often relatively limited, or kept to a minimum possible due to address space usage concerns). The lack of requirement to periodically resize an IPv6 subnet is a useful administrative advantage for IPv6.

2.5 IPv6 Addressing

One site using this VLAN technique has chosen to number its IPv6 links with the format [Site IPv6 prefix]:[VLAN ID]::/64. This is not a recommended addressing plan, but some sites may wish to consider its usage.

2.6 Final IPv6 Deployment

The VLAN technique for IPv6 deployment offers a more structured alternative to opportunistic per-host intra-site tunnelling methods such as ISATAP [3]. It has the ability to offer a simple yet efficient method for early IPv6 deployment to an enterprise site.

When the site acquires IPv6-capable switch-router equipment, the

VLAN-based mathod can still be used for delivery of IPv6 links to physical switch interfaces, just as it is commonly today for IPv4 subnets, but with a common routing infrastructure.

3. Example VLAN topology

The following figure shows how a VLAN topology may be used to introduce IPv6 in an enterprise network, using a parallel IPv6 routing infrastructure and VLAN tagging.

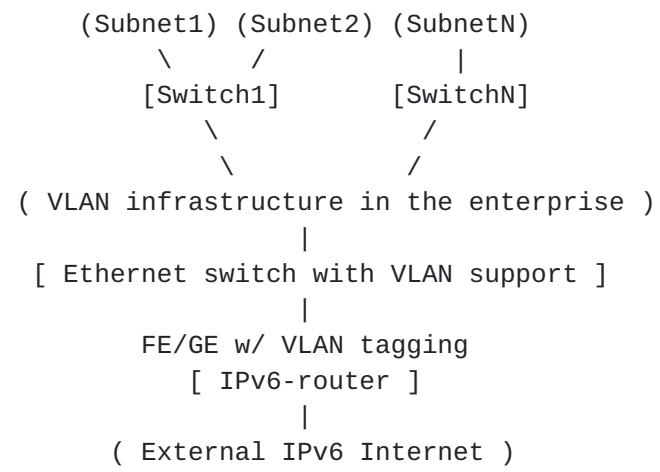

Figure 1: IPv6 deployment using VLANs

In this scenario, the router has one physical port facing towards the internal infrastructure, and is using the collapsed VLAN mechanism described above. It may have an additional interface towards the external infrastructure. The router can also function as a "one-handed" router.

A number of VLANs are handled by the internal-facing IPv6 router port; the VLANs are seen as logical subinterfaces of the physical interface. Therefore, the router acts as an IPv6 first-hop access router to the physical links, separately from the IPv4-first hop router. This technique allows a site to easily "inject" native IPv6 into all the links where a VLAN-capable infrastructure is available, enabling partial or full IPv6 deployment on the wire in a site.

4. Security Considerations

There are no additional security considerations particular to this

method of enabling IPv6 on a link.

Where the IPv6 connectivity is delivered into the enterprise network by a different path from the IPv4 connectivity, care should be given that equivalent application of security policy (e.g. firewalling) is made to the IPv6 path.

5. Acknowledgements

The author would like to thank colleagues on the 6NET project, where this technique for IPv4-IPv6 coexistence is widely deployed, including Janos Mohacsi (Hungarnet), Martin Dunmore and Chris Edwards (Lancaster University), Christian Strauf (JOIN Project, University of Muenster), Stig Venaas (UNINETT) and Pekka Savola (CSC/FUNET).

6 Informative References

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

[2] Chown, T., "IPv6 Campus Transition Scenario Description and Analysis", draft-chown-v6ops-campus-transition-00 (work in progress), July 2004.

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

Author's Address

Tim Chown
University of Southampton

Southampton, Hampshire SO17 1BJ
United Kingdom

EMail: tjc@ecs.soton.ac.uk

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 procedures with respect to rights in RFC documents can be found in BCP 78 and 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.

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.

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.

Acknowledgment

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