Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks
draft-chown-v6ops-vlan-usage-02

Status of this Memo

Copyright Notice

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, which focuses on the scenario of early deployment prior to availability of IPv6-capable switch-router equipment. In this method 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.

    It is expected that in the long run, sites migrating to dual-stack
    networking will either upgrade existing switch-router equipment to
    support IPv6 or procure new equipment that supports IPv6. If a site
    already has production routers deployed that support IPv6, the
    procedures described in this document are not required. In the
    interim however, a method is required for early IPv6 adopters that
    enables IPv6 to be deployed in a structured, managed way to some or
    all of an enterprise network which currently lacks IPv6 support in
    its core infrastructure.

    The IEEE 802.1Q VLAN standard allows separate LANs to be deployed
    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, in
    advance of the site's core infrastructure having dual-stack
    capability.

    If such a site wishes to introduce IPv6, it may do so by deploying a
    parallel IPv6 routing infrastructure (which is likely to be a
    different platform to the site's main infrastructure equipment, i.e.
    one that supports IPv6 where the existing equipment does not), 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 VLANs don't need to differentiate between IPv4

and IPv6; the deployment is just dual stack, as Ethernet is without
VLANs.

The IPv4 default route to the VLAN is provided by one (IPv4) router,
while the IPv6 default route to the VLAN is provided a different

(IPv6) router. The IPv6 router can provide native IPv6 connectivity
to the whole site with just a single physical interface, thanks to
VLAN tagging and trunking, as described below.

The IPv6 connectivity to the enterprise may or may not enter the site
via the same physical link as the IPv4 traffic, 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.

It should be noted that a parallel infrastructure will require
additional infrastructure and thus cost, and will often require a
separate link into the site (from an IPv6 provider), quite possibly
tunneled, that will require the site's security policy to be applied
(e.g. firewalling, and intrusion detection). For sites who believe
early adoption of IPv6 is important, that price is one they may be
quite willing to pay. However, this document focuses on the
technical issues of VLAN usage in such a scenario.

2. Enabling IPv6 per link

The precise method by which IPv6 would be "injected" into the
existing IPv4 network is deployment specific. For example, perhaps a
site has an IPv4-only router, connected to an Ethernet switch that
supports VLANs, and a number of hosts connected to that VLAN. Let's
further assume the site has a dozen of these setups which it wishes
to IPv6-enable immediately. This could be done by upgrading the
twelve routers to support IPv6, and turning IPv6 on on those routers.
However, this may not be practical for various reasons.

The simplest approach would be to connect an IPv6 router with one
interface to an ethernet switch, and connect that switch to other
switches, and then use VLAN tags between the switches and the IPv6
router to "reach" all the IPv4-only subnets from the IPv6 router.
Thus 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 over VLANs

In a typical scenario where connectivity is to be offered to a number
of existing IPv6 internal subnets, one IPv6 router could 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.

The internal interface(s) can be connected directly to a VLAN-capable
switch. It is then possible to write VLAN tags on the packets sent
from the internal router interface based on the target IPv6 link
prefix. The VLAN-tagged traffic is then transported across the
internal VLAN-capable site infrastructure to the target IPv6 links
(which may be dispersed widely across the site network).

Where the IPv6 router is unable to VLAN-tag the packets, a
protocol-based VLAN can be created on the VLAN-capable device
connected to the IPv6 router, causing IPv6 traffic to be tagged and
then redistributed on (congruent) IPv4 subnet links that lie in the
same VLAN.

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 BSD-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
(trunked) 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.

```
                    External IPv6 Internet
                              |
                              |
                      IPv6 Access Router
                              |
                              |
                  Switch-router with VLAN support
                              |
                              |
             +--------------+----------------+
             |Site enterprise infrastructure |
             |    with support for VLANs     |
             +----+--------------------+-----+
                  |                     |
                  |                     |
              VLAN switch A        VLAN switch B
                |       |              |
                |       |              |
            Subnet1    Subnet2       Subnet3
```
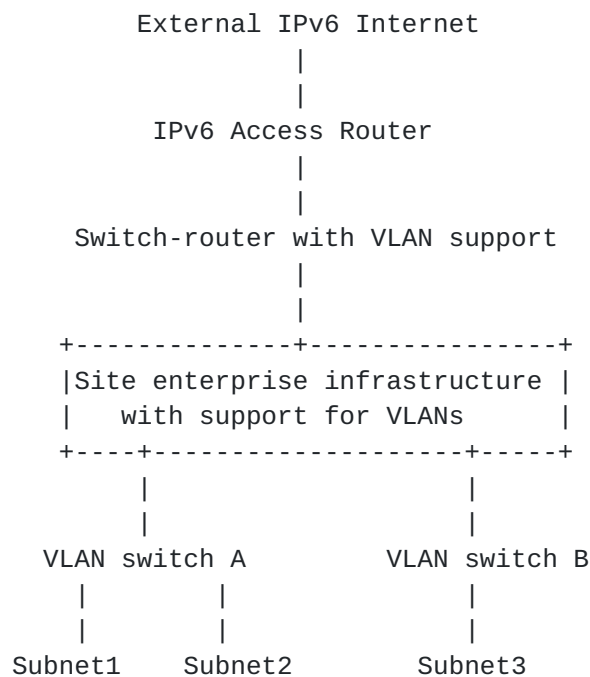
   Figure 1: IPv6 deployment using VLANs (physical diagram)

   In this scenario, the IPv6 access router has one physical port facing
   towards the internal infrastructure. In this example it need only be
   IPv6-enabled, as its purpose is solely to handle IPv6 traffic for the
   enterprise. The access router has an additional interface facing
   towards the external infrastructure, which in this example could be
   dual-stack if the external IPv6 connectivity is via a tunnel to an
   IPv6 ISP.

   A number of VLANs are handled by the internal-facing IPv6 router
   port; in this case IPv6 links Subnet1, Subnet2, Subnet3. The VLANs
   are seen as logical subinterfaces of the physical interface on the
   IPv6 access router, which is using the "collapsed VLAN" method

described above, tagging the inbound traffic with one of three VLAN
IDs depending on the target IPv6 Subnet prefix.

The following figure shows how the IPv6 view of the deployment looks; all IPv6 subnets are on-link to the IPv6 access router, whether they share the same physical links over the VLAN infrastructure or not.
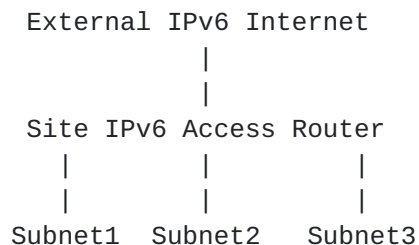

```
                    External IPv6 Internet
                              |
                              |
                    Site IPv6 Access Router
                      |        |          |
                      |        |          |
                   Subnet1  Subnet2    Subnet3
```


Figure 2: IPv6 view of the deployment (logical view)

In this example, 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, in particular Pekka Savola (CSC/FUNET), but also including Janos Mohacsi (Hungarnet), Martin Dunmore and Chris Edwards (Lancaster University), Christian Strauf (JOIN Project, University of Muenster) and Stig Venaas (UNINETT).

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

Appendix A.   Appendix: Configuration example

   In this section we describe a configuration example for BSD to deploy
   IPv6 networking across a number of IPv6 links on an enterprise (in
   this case, eight links), for a scenario similar to the one described
   above. Here the precise configuration may of course vary depending
   on the existing site VLAN deployment. This section highlights that
   the VLAN configuration must be manually configured; the support is
   not "automatic".

   In this example, the configuration is for an IPv6 BSD router
   connected directly to a site's external IPv6 access router. The BSD
   router has one interface (dc0) towards the site IPv6access router,
   and three interfaces (dc1, dc2, dc3) over which the internal routing
   is performed (the number of interfaces can be varied, three are used
   here to distribute the traffic load). The IPv6 documentation prefix
   (2001:db8::/32) is used in the example.

   --- Example IPv6 VLAN configuration, FreeBSD ---

```
#
# To IPv6 enable a vlan
#
# 1. Add a new vlan device to cloned_interfaces called vlanX
#
# 2. Add an ifconfig_vlanX line, the number in quotes is the vlan tag ID
#
# 3. Add vlanX to ipv6_network_interfaces
#
```

```
# 4. Add an ipv6_ifconfig_vlanX line, with a new unique prefix
#
# 5. Add vlanX to rtadvd_interface
#
# 6. Add vlanX to ipv6_router_flags


### Interfaces ###


# Bring physical interfaces up
ifconfig_dc0="up"
ifconfig_dc1="up"
ifconfig_dc2="up"
ifconfig_dc3="up"


# Create VLan interfaces
cloned_interfaces="vlan0 vlan1 vlan2 vlan3 vlan4 vlan5 vlan6 vlan7 vlan8"


# Upstream link to IPv6 Access Router
ifconfig_vlan0="vlan 37 vlandev dc0"


# Downstream interfaces, load balance over interfaces dc1,dc2,dc3
ifconfig_vlan1="vlan 11 vlandev dc1" # Subnet1
ifconfig_vlan2="vlan 17 vlandev dc2" # Subnet2
ifconfig_vlan3="vlan 24 vlandev dc3" # Subnet3
ifconfig_vlan4="vlan 25 vlandev dc1" # Subnet4
ifconfig_vlan5="vlan 34 vlandev dc2" # Subnet5
ifconfig_vlan6="vlan 14 vlandev dc3" # Subnet6
ifconfig_vlan7="vlan 19 vlandev dc1" # Subnet7
ifconfig_vlan8="vlan 13 vlandev dc2" # Subnet8


### IPv6 ###


# Enable ipv6
ipv6_enable="YES"


# Forwarding
ipv6_gateway_enable="YES"


# Define Interfaces
ipv6_network_interfaces="vlan0 vlan1 vlan2 vlan3 vlan4 vlan5 vlan6 vlan7
```

```
vlan8"
   # Define addresses
   ipv6_ifconfig_vlan0="2001:db8:d0:101::2 prefixlen 64" # Uplink
   ipv6_ifconfig_vlan1="2001:db8:d0:111::1 prefixlen 64" # Subnet1
   ipv6_ifconfig_vlan2="2001:db8:d0:112::1 prefixlen 64" # Subnet2
   ipv6_ifconfig_vlan3="2001:db8:d0:121::1 prefixlen 64" # Subnet3
   ipv6_ifconfig_vlan4="2001:db8:d0:113::1 prefixlen 64" # Subnet4
   ipv6_ifconfig_vlan5="2001:db8:d0:114::1 prefixlen 64" # Subnet5
   ipv6_ifconfig_vlan6="2001:db8:d0:115::1 prefixlen 64" # Subnet6
```

```
   ipv6_ifconfig_vlan7="2001:db8:d0:116::1 prefixlen 64" # Subnet7
   ipv6_ifconfig_vlan8="2001:db8:d0:117::1 prefixlen 64" # Subnet8


   # Router advs
   rtadvd_enable="YES"
   rtadvd_interfaces="-s vlan0 vlan1 vlan2 vlan3 vlan4 vlan5 vlan6 vlan7 vlan8"


   ### Routing ###


   # Multicast
   mroute6d_enable="YES"
   mroute6d_program="/sbin/pim6sd"


   # RIP-ng
   ipv6_router_enable="YES"
   ipv6_router_flags="-N
dc0,dc1,dc2,dc3,vlan1,vlan2,vlan3,vlan4,vlan5,vlan6,vlan7,vlan8"


   --- End of configuration ---
```

   Note that if there was only one internal facing interface, then again
   so long as the OS supported VLAN trunking, all the VLAN IDs could be
   associated to that interface (dc1, for example).

   The VLAN IDs need to be managed by the site administrator, but would
   probably already be assigned for existing IPv4 subnets (ones into
   which IPv6 is being introduced).

   For a large enterprise, a combination of internal tunnels and VLAN
   usage could be used; the whole site need not be enabled by VLAN
   tagging alone. This choice is one for the site administrator to
   make.

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment