

Network Working Group
Internet Draft
Expiration Date: December 2001
File name: [draft-shen-isis-ospf-p2p-over-lan-00.txt](#)

Naiming Shen
Acee Lindem
Jenny Yuan
Redback Networks
June 2001

Point-to-point operation over LAN
in link-state routing protocols

[draft-shen-isis-ospf-p2p-over-lan-00.txt](#)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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/lid-abstracts.html>

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

Abstract

Two different circuit types are commonly used by link state routing protocols: point-to-point and broadcast. It is important to identify the correct circuit type in forming adjacency with neighbors, in flooding link state database packets, in representation of the circuit subnet. This document describes a simple mechanism to treat the broadcast media as point-to-point connection from IP routing protocol point of view if there are only two devices on the LAN

media.

INTERNET DRAFT

P2P OVER LAN

June 2001

1. Introduction

Point-to-point and broadcast media are two most common circuit types used by link state routing protocols such as IS-IS [[ref1](#)] [[ref2](#)] and OSPF [[ref3](#)]. They are treated differently with respect to establishing neighbor adjacencies, link state database flooding, representation of the media subnet, SPF calculation and protocol packets. The most important difference is there is a designated router concept associated with the broadcast media and pseudo node is used to represent the information on the LAN media.

Compared with broadcast circuits, point-to-point circuits are afford more straightforward IGP operation. There is no designated router involved and there is no representation of the pseudo-node or network LSA in the link state database. For ISIS, there also no periodic database synchronization. Conversely, if there are more than two routers on the LAN media, the traditional view of the broadcast media will reduce the routing information in the network.

When there are only two routers on the broadcast media, it makes more sense to treat the connection between the two routers as a point-to-point one. This document describes the mechanism to allow link state routing protocols to operate using point-to-point connection over broadcast media under this condition. Some implication with forwarding IP packets on this type of circuit is also discussed. We will refer to this as p2p-over-lan circuit in this document.

2. Motivation

Even though the broadcast media is meant to handle more than two devices, there exist cases where only two routers are interconnected over the physical or logical broadcast media:

- o simply only two routers on the physical LAN.
- o two routers are connected directly back to back using broadcast media, mainly for long-haul operation.
- o only two routers exist on the virtual LAN.

In any of the above cases, the link state routing protocols will normally still treat the circuit as a broadcast type. Thus it will have the overhead involved with protocol LAN operation but without the benefit of reducing routing information designed for the LAN environment.

Even when there are multiple routers on the LAN an ISP may want to sub-group the routers into multiple VLANs since this allows them to assign different costs to IGP neighbors. When there are only two routers in some of the VLANs, this broadcast media can be viewed by the IGP as a mesh of point-to-point connections. As a side benefit, unnumbered interface can also be applied over LAN.

Shen, Lindem, Yuan

Expires December 2001

[Page 2]

INTERNET DRAFT

P2P OVER LAN

June 2001

[3.](#) Point-to-point connection over LAN media

The idea is very simple: provide a configuration mechanism to inform the IGP that the circuit is type point-to-point irrespective of the physical media type. For the IGP, this implies that it will send protocol packets with the appropriate point-to-point information and expects to receive protocol packets as they would be received on a point-to-point circuit. Over LAN media, the MAC header must contain the correct multicast MAC address to be received by the other side of the connection. For VLAN environments, the MAC header must also contain the proper VLAN ID.

[3.1](#) Operation of IS-IS

This p2p-over-lan circuit extension for IS-IS is only concerned in pure IP routing and forwarding operation.

Since physically the circuit is a broadcast one, the IS-IS packets need to have MAC addresses for this p2p-over-lan circuit. From link layer point of view, those packets are IS-IS LAN packets. The Multi-destination address, either AllL1ISs or AllL2ISs defined in [\[ref1\]](#), is used for the point-to-point IS-IS PDUs.

The circuit needs to have IP address(es) and the p2p IIH over this circuit MUST include the IP interface address(es) as defined in [\[ref2\]](#). The IP address(es) can be numbered or unnumbered. Note that the term "unnumbered" here means this interface sets the IP address to any one of the other IP addresses belong to this router.

If the circuit is configured as point-to-point type and receives LAN IIHs, it MUST discard the incoming packets; If the circuit is a LAN type and receive point-to-point IIHs, it MUST discard the incoming packets. If the system ID of incoming IIH does not match the system ID of already established adjacency over this p2p-over-lan circuit, it MUST discard the packet. The implementation should offer enough logging or debugging information to detect mis-configurations.

[3.2](#) Operation of OSPF

OSPF routers supporting the capabilities described herein must support an additional interface configuration parameter specifying the interface topology type. For LAN (i.e., broadcast capable) media, the interface may be viewed as a point-to-point interface. Both routers on the LAN will simply join the AllSPFRouters (224.0.0.5) multicast group and send all OSPF packets to 224.0.0.5. This is identical to operation over a physical point-to-point link as described in sections [8.1](#) and [8.2](#) of [\[ref3\]](#).

[3.3](#) IP forwarding and ARP

Unlike normal point-to-point IGP circuit, the IP nexthop for the routes using this p2p-over-lan circuit as outbound interface is not optional. The IP nexthop address has to be a valid interface or internal address on the adjacent router. This address is used by local router to obtain the MAC address for IP packet forwarding. Proxy ARP has to be enabled if the address is not the adjacent interface IP address.

In the case where unnumbered IP addresses are used for p2p-over-lan circuit, the source IA of ARP request and the target interface IA are usually on different subnets. The ARP should reply only if this circuit is a p2p-over-lan type and the source IA of the ARP request is the same as the neighbor's interface IP address at the other end. The neighbor's address is learned from IGP hello exchanges over this circuit.

[4.](#) Compatibility

Routers on both sides of the broadcast media connection have to support this p2p-over-lan extension in order to establish adjacency to each other. Otherwise, the traditional LAN model for the IGP has to be used on this media.

[5. Scalability Issues](#)

There is obvious advantage to use this extension if the broadcast media between two routers are connected back-to-back. To model the LAN as a number of vLANs with this extension does sacrifice the scalability property of the LAN representation for link-state routing protocols. It will in general increase the link-state database size, the amount of packets to be flooded and the route calculation time thus the network overall convergence time.

The network design engineers should carefully balance between the need of more precise routing control and the scalability of the network. The scalability impact is less of a concern if the LAN and routers involved are within a single link-state subdomain in hierarchical IGP routing.

[6. Security Issues](#)

This document does not introduce any new security issues to ISIS or OSPF. For ARP to support unnumbered IP interface addresses, it needs to verify the p2p-over-lan circuit type described in this document and to verify the ARP packet source IA to match the IGP adjacency interface IP address. This is due to normal ARP sanity check for common subnet can not be applied in this case.

Shen, Lindem, Yuan

Expires December 2001

[Page 4]

INTERNET DRAFT

P2P OVER LAN

June 2001

[7. Acknowledgments](#)

TBA.

[8. References](#)

- [ref1] ISO. Information Technology - Telecommunications and Information Exchange between Systems - Intermediate System to Intermediate System Routing Exchange Protocol for

Use in Conjunction with the Protocol for Providing the
Connectionless-Mode Network Service. ISO, 1990.

[ref2] R. Callon. Use of OSI ISIS for Routing in TCP/IP and Dual
Environments. INTERNET-RFC, Internet Engineering Task Force,
December 1990.

[ref3] J. Moy. OSPF Version 2. Technical Report [RFC2328](#) Internet
Engineering Task Force, 1998.

[9.](#) Authors' Addresses

Naiming Shen
Redback Networks
350 Holger Way
San Jose, CA, 95134 USA
naiming@redback.com

Acee Lindem
Redback Networks
102 Carric Bend Court
Apex, NC 27502 USA
acee@redback.com

Jenny Yuan
Redback Networks
350 Holger Way
San Jose, CA, 95134 USA
jenny@redback.com