

Point-to-Point Operation over LAN
in Link State Routing Protocols

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

The two predominant circuit types used by link state routing protocols are point-to-point and broadcast. It is important to identify the correct circuit type when forming adjacencies, flooding link state database packets, and representing the circuit topologically. This document describes a simple mechanism to treat the broadcast network as a point-to-point connection from the standpoint of IP routing.

1. Introduction

Point-to-point and broadcast are the two predominant circuit types used by link state routing protocols such as IS-IS [ISO10589] [RFC1195] and OSPF [RFC2328] [RFC5340]. They are treated differently with respect to establishing neighbor adjacencies, flooding link state information, representing the topology, and calculating the Shortest Path First (SPF) and protocol packets. The most important differences are that broadcast circuits utilize the concept of a designated router and are represented topologically as virtual nodes in the network topology graph.

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

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

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Motivation

Even though a broadcast circuit is meant to handle more than two devices, there are cases where only two routers are connected over either the physical or logical LAN segment:

1. The media itself is being used for point-to-point operation between two routers. This is mainly for long-haul operation.
2. There are only two routers on the physical LAN.
3. There are only two routers on a virtual LAN (vLAN).

In any of the above cases, the link state routing protocols will normally still treat the media as a broadcast circuit. Hence, they will have the overhead involved with protocol LAN operation without the benefits of reducing routing information and optimized flooding.

Being able to treat a LAN as a point-to-point circuit provides the benefit of reduction in the amount of information routing protocols must carry and manage. DR/DIS (Designated Router / Designated Intermediate System) election can be omitted. Flooding can be done as in p2p links without the need for using "LSA reflection" by the DR in OSPF or periodic Complete Sequence Number Packets (CSNPs) in IS-IS.

Also, if a broadcast segment wired as a point-to-point link can be treated as a point-to-point link, only the connection between the two routers would need to be advertised as a topological entity.

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 LAN can be viewed by the IGP as a mesh of point-to-point connections.

The IP unnumbered configuration is widely used in networks. It enables IP processing on a point-to-point interface without an explicit IP address. The IP unnumbered interface can "borrow" the IP address of another interface on the node. The advantages of unnumbered point-to-point links are obvious in the current IP addressing environment where addresses are a scarce resource. The unnumbered interface can also be applied over p2p-over-lan circuits. Separating the concept of network type from media type will allow LANs, e.g., ethernet, to be unnumbered and realize the IP address space savings. Another advantage is in simpler network management and configuration. In the case of an IPv6 network, a link local address used in IS-IS [RFC5308] and OSPFv3 [RFC5340] serves the same purpose.

3. IP Multi-Access Subnets

When an IP network includes multi-access segments, each segment is usually assigned a separate subnet, and each router connected to it is assigned a distinct IP address within that subnet. The role of the IP address assigned to a multi-access interface can be outlined as follows:

1. Source IP address - The interface address can be used by the router as the source IP address in locally originated IP packets that are destined for that subnet or have a best path next hop on that subnet.
2. Destination IP address - The interface address can be used by other devices in the network as a destination address for packets to router applications (examples include telnet, SMTP, TFTP, OSPF, BGP, etc).
3. Next-hop identifier - If other routers connected to the same segment need to forward traffic through the router, the corresponding routes in their routing tables will include the router's interface IP address. This address will be used to find the router's MAC (Media Access Control) address using the ARP/ND (Address Resolution Protocol / Neighbor Discovery) protocol. Effectively, the interface IP addresses help other routers find the data-link layer details that are required to specify the destination of the encapsulating data-link frame when it is sent on the segment.

The IP addressing scheme includes an option that allows the administrators to not assign any subnets to point-to-point links (links connecting only two devices and using protocols like PPP, SLIP, or HDLC for IP encapsulation). This is possible because the routers do not need next-hop identifiers on point-to-point links

(there is only one destination for any transmission), and an interface-independent IP address can be used as the source and destination. Using the unnumbered option for a point-to-point link essentially makes it a purely topological entity used only to reach other destinations.

4. 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 it 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.

In order to allow LAN links used to connect only two routers to be treated as unnumbered point-to-point interfaces, the MAC address resolution and nexthop IP address issues need to be addressed.

4.1. Operation of IS-IS

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

Since physically the circuit is a broadcast one, the IS-IS protocol packets need to have MAC addresses for this p2p-over-lan circuit. From a link-layer point of view, those packets are IS-IS LAN packets. The Multi-destination address including AllISs, AllL1ISs, and AllL2ISs, defined in [ISO10589], can be used for link-layer encapsulation; the use of AllISs is recommended.

The circuit needs to have IP address(es), and the p2p IS-IS Hello (IIH) over this circuit MUST include the IP interface address(es) as defined in [RFC1195]. The IPv4 address(es) included in the IIHs is either the IP address assigned to the interface in the case of a numbered interface or the interface-independent IP address in the case of an unnumbered interface. The IPv6 addresses are link-local IPv6 address(es) [RFC5308].

4.2. Operation of OSPF and OSPFv3

OSPF and OSPFv3 [RFC5340] routers supporting the capabilities described herein should support an additional interface configuration parameter specifying the interface topology type. For a LAN (i.e., broadcast-capable) interface, the interface may be viewed as a

point-to-point interface. Both routers on the LAN will simply join the AllSPFRRouters multicast group and send all OSPF packets with a destination address of AllSPFRRouters. AllSPFRRouters is 224.0.0.5 for OSPF and FF02::5 for OSPFv3. This is identical to operation over a physical point-to-point link as described in Sections 8.1 and 8.2 of [RFC2328].

4.3. ARP and ND

Unlike a normal point-to-point IGP circuit, the IP nexthop for the routes using this p2p-over-lan circuit as an 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 a local router to obtain the MAC address for IP packet forwarding. The ARP process has to be able to resolve the internal IPv4 address used for the unnumbered p2p-over-lan circuits. For the ARP implementation (which checks that the subnet of the source address of the ARP request matches the local interface address), this check needs to be relaxed for the unnumbered p2p-over-lan circuits. The misconfiguration detection is handled by the IGPs and is described in Section 4.5. In the IPv6 case, the ND resolves the MAC for the link-local address on the p2p-over-lan circuit, which is part of the IPv6 neighbor discovery process [RFC4861].

4.4. Other MAC Address Resolution Mechanisms

In more general cases, while p2p-over-lan circuit is used as an unnumbered link, other MAC address resolution mechanisms are needed for IP packet forwarding; for example, if link state IGP is not configured over this p2p-over-lan link, or if the mechanism described in Section 4.3 is not possible. The following techniques can be used to acquire the MAC address and/or the next-hop IP address of the remote device on an unnumbered point-to-point LAN link.

1. Static configuration. A router can be statically configured with the MAC address that should be used as the destination MAC address when sending data out of the interface.
2. MAC address gleaning. If a dynamic routing protocol is running between the routers connected to the link, the MAC address of the remote device can be taken from a data-link frame carrying a packet of the corresponding routing protocol.

4.5. Detection of Misconfiguration

With this p2p-over-lan extension, the difference between a LAN and a point-to-point circuit can be made purely by configuration. It is important to implement the mechanisms for early detection of misconfiguration.

If the circuit is configured as the point-to-point type and receives LAN hello packets, the router MUST discard the incoming packets; if the circuit is a LAN type and receives point-to-point hello packets, it MUST discard the incoming packets. If the system ID or the router ID of an incoming hello packet does not match the system ID or the router ID for an established adjacency over a p2p-over-lan circuit, the packet MUST be discarded. Furthermore, if OSPF hello suppression (as described in [RFC1793]) is active for the adjacency, the hello suppression MUST be terminated for a period of RouterIntervalSeconds. After this interval, either the neighbor adjacency will time out and an adjacency may be formed with a neighbor with a different router ID, or hello suppression may be renegotiated. The implementation should offer logging and debugging information of the above events.

5. Compatibility Considerations

Both routers on a LAN must support the p2p-over-lan extension and both must have the LAN segment configured as a p2p-over-lan circuit for successful operation. Both routers SHOULD support at least one of the above listed methods for mapping IP addresses on the link to MAC address. If a proprietary method of IP address to MAC address resolution is used by one router, both routers must be capable of using the same method. Otherwise, the link should be configured as a standard LAN link, with traditional IGP LAN models used.

6. Scalability and Deployment Considerations

While there is advantage to using this extension on the LANs that are connected back to back or only contain two routers, there are trade offs when modeling a LAN as multiple VLANs and using this extension since one does sacrifice the inherent scalability benefits of multi-access networks. In general, it will increase the link state database size, the amount of packets flooded, and the route calculation overhead.

Deployment of the described technique brings noticeable benefits from the perspective of IP address usage: the network management and the router configuration. Note, however, that use of the IP unnumbered

option for point-to-point LAN links inherits the same problems as those present for serial links, i.e., not being able to ping or monitor a specific interface between routers.

7. Security Considerations

This document does not introduce any new security issues to IS-IS, OSPF, ARP, or ND. Implementations may have 'source address subnet checks' that need to be relaxed as described in Section 4.3. These are used to manage misconfigurations, not so much to secure ARP -- if an attacker would be attached to the LAN, (s)he could pick a subnet-wise correct address as well.

If one router on a link thinks that a LAN should be either broadcast or p2p-over-lan, and the other router has a different opinion, the adjacencies will never form, as specified in Section 4.5. There are no fallbacks at either end to resolve the situation, except by a manual configuration change.

8. Acknowledgments

The authors would like to acknowledge the following individuals (in alphabetical order by last name): Pedro Marques, Christian Martin, Danny McPherson, Ajay Patel, Jeff Parker, Tony Przygienda, Alvaro Retana, and Pekka Savola.

9. Normative References

- [ISO10589] ISO, "Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", International Standard 10589:2002, Second Edition, 2002.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.
- [RFC1793] Moy, J., "Extending OSPF to Support Demand Circuits", RFC 1793, April 1995.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

[RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, October 2008.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

Contributors

The following individuals are the authors that contributed to the contents of this document.

Acee Lindem
Cisco Systems
7025 Kit Creek Road
Research Triangle Park, NC 27709
USA
EMail: acee@cisco.com

Jenny Yuan
Cisco Systems
225 West Tasman Drive
San Jose, CA 95134
USA
EMail: jenny@cisco.com

Russ White
Cisco Systems, Inc.
7025 Kit Creek Rd.
Research Triangle Park, NC 27709
EMail: riw@cisco.com

Stefano Previdi
Cisco Systems, Inc.
De Kleetlaan 6A
1831 Diegem - Belgium
EMail: sprevidi@cisco.com

Editors' Addresses

Naiming Shen
Cisco Systems
225 West Tasman Drive
San Jose, CA 95134
USA
EMail: naiming@cisco.com

Alex Zinin
Alcatel-Lucent
750D Chai Chee Rd, #06-06
Technopark@ChaiChee
Singapore 469004

EMail: alex.zinin@alcatel-lucent.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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.

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, THE IETF TRUST 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.

Intellectual Property

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.