

Guidelines for Running OSPF Over Frame Relay Networks

Status of this Memo

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

Abstract

This memo specifies guidelines for implementors and users of the Open Shortest Path First (OSPF) routing protocol to bring about improvements in how the protocol runs over frame relay networks. We show how to configure frame relay interfaces in a way that obviates the "full-mesh" connectivity required by current OSPF implementations. This allows for simpler, more economic network designs. These guidelines do not require any protocol changes; they only provide recommendations for how OSPF should be implemented and configured to use frame relay networks efficiently.

Acknowledgements

This memo is the result of work done in the OSPF Working Group of the IETF. Comments and contributions from several sources, especially Fred Baker of ACC, John Moy of Proteon, and Bala Rajagopalan of AT&T Bell Laboratories are included in this work.

1. Introduction

A frame relay (FR) network provides virtual circuits (VCs) to interconnect attached devices. Each VC is uniquely identified at each FR interface by a Data Link Connection Identifier (DLCI). RFC 1294 specifies the encapsulation of multiprotocol traffic over FR [1]. The devices on a FR network may either be fully interconnected with a "mesh" of VCs, or partially interconnected. OSPF characterizes FR networks as non-broadcast multiple access (NBMA) because they can support more than two attached routers, but do not have a broadcast capability [2]. Under the NBMA model, the physical FR interface on a router corresponds to a single OSPF interface through which the router is connected to one or more neighbors on the FR network; all the neighboring routers must also be directly connected to each other

over the FR network. Hence OSPF implementations that use the NBMA model for FR do not work when the routers are partially interconnected. Further, the topological representation of a multiple access network has each attached router bi-directionally connected to the network vertex with a single link metric assigned to the edge directed into the vertex.

We see that the NBMA model becomes more restrictive as the number of routers connected to the network increases. First, the number of VCs required for full-mesh connectivity increases quadratically with the number of routers. Public FR services typically offer performance guarantees for each VC provisioned by the service. This means that real physical resources in the FR network are devoted to each VC, and for this the customer eventually pays. The expense for full-mesh connectivity thus grows quadratically with the number of interconnected routers. We need to build OSPF implementations that allow for partial connectivity over FR. Second, using a single link metric (per TOS) for the FR interface does not allow OSPF to weigh some VCs more heavily than others according to the performance characteristics of each connection. To make efficient use of the FR network resources, it should be possible to assign different link metrics to different VCs.

These limitations of the current OSPF model for FR become more severe as the network size increases, and render FR technology less cost effective than it could be for large networks. We propose solutions to these problems that do not increase complexity by much and do not require any changes to the OSPF protocol.

2. Summary of Recommendations

We propose expanding the general view of an OSPF interface to account for its functional type (point-to-point, broadcast, NBMA) rather than its physical type. In most instances, the physical network can only serve one function and can only be defined as one type of OSPF interface. For multiplexed interfaces such as FR however, logical connections between routers can serve different functions. Hence one VC on a FR interface can be viewed distinctly from other VCs on the same physical interface. The solution requires that OSPF be able to support logical interfaces (networks) as well as physical interfaces. Each logical network can be either point-to-point, that is, a single VC, or NBMA, that is, a collection of VCs. It is not necessary to define new interface types for logical networks, since the operation of the protocol over logical point-to-point networks and logical NBMA networks remains the same as for the corresponding physical networks. For instance, logical point-to-point links could be numbered or unnumbered. It is only necessary for implementations to provide the hooks that give users the ability to configure an individual VC as a

logical point-to-point network or a collection of VCs as a logical NBMA network.

The NBMA model does provide some economy in OSPF protocol processing and overhead and is the recommended mode of operation for small homogeneous networks. Other than the Designated Router (DR) and the backup Designated Router (BDR), each router maintains only two adjacencies, one each with the DR and BDR, regardless of the size of the NBMA network. When FR VCs are configured as point-to-point links, a router would have many more adjacencies to maintain, resulting in increased protocol overhead. If all VCs were to have comparable performance characteristics as well, there may not be compelling reasons to assign a different link metric to each VC.

3. Implementing OSPF over FR

We recommend that OSPF router implementations be built so that administrators can configure network layer interfaces that consist of one or more FR VCs within a single physical interface. Each logical network interface could then be configured as the appropriate type of OSPF interface, that is, point-to-point for a single VC, or NBMA for a collection of VCs. This capability would allow a router to belong to one or more distinct IP subnets on a single physical FR interface. Thus, it is necessary that the router be able to support multiple IP addresses on a single physical FR interface. As with physical NBMA networks, logical NBMA networks must be full-mesh connected. While logical point-to-point links can be either numbered or unnumbered, we show that it is easier to implement routers to handle numbered logical point-to-point links.

3.1 Numbered Logical Interfaces

The router administrator should be able to configure numbered logical interfaces over FR as follows:

STEP 1: Configure the physical interface specifying relevant parameters such as the slot, connector, and port numbers, physical frame format, encoding, and clock mode. In its internal interface MIB [3], the router should create a new ifEntry in the ifTable, assign the physical interface an ifIndex, and increment the ifNumber by one.

STEP 2: Configure the data-link layer over the interface, specifying frame relay as the encapsulation method. Parameters such as the DLCI encoding type and length, maximum frame size, management interface (Annex D, LMI), and address resolution procedure (manual, inverse ARP). If a management interface is not supported, FR VCs must be

configured manually.

STEP 3: Configure the IP network layer for the interface by specifying the number of logical interfaces and the IP address and subnet mask for each numbered logical interface. Specify the VCs (by DLCI) associated with each logical network interface if there is more than one. If an address resolution protocol such as Inverse ARP [4] is being used, it should suffice to specify a list of IP addresses for the FR interface and have Inverse ARP create the DLCI-IP address binding.

STEP 4: Configure OSPF to run over each logical interface as appropriate, specifying the necessary interface parameters such as area ID, link metric, protocol timers and intervals, DR priority, and list of neighbors (for the DR). OSPF interfaces consisting of one VC can be treated as point-to-point links while multi-VC OSPF interfaces are treated as NBMA subnets. In its internal OSPF MIB [5], the router should create additional entries in the ospfIfTable each with the appropriate ospfIfType (nbma or pointToPoint).

3.2 Unnumbered Point-to-Point Logical Interfaces

OSPF uses the IP address to instance each numbered interface. However, since an unnumbered point-to-point link does not have an IP address, the ifIndex from the interface MIB is used instead [5]. This is straightforward for a physical point-to-point network, since the ifIndex is assigned when the interface is configured. Logical interfaces over FR however, do not have distinct and unique values for ifIndex. To allow OSPF to instance unnumbered logical point-to-point links, it is necessary to assign each such link a unique ifIndex in STEP 3 above. This could lead to some confusion in the interfaces table since a new ifTable entry would have to be created for each logical point-to-point link. This type of departure from the standard practice of creating interface table entries only for physical interfaces could be viewed as an unnecessary complication.

Alternatively, it is possible to build a private MIB that contains data structures to instance unnumbered logical links. However, making recommendations for the structure and use of such a private MIB is beyond the scope of this work. Even if unnumbered point-to-point logical links were implemented in this manner, it would still be necessary to allow a FR interface to be configured with multiple IP addresses when a router is connected to multiple NBMA subnets through a single physical interface. Hence, while it is possible to define unnumbered logical point-to-point links in OSPF, we find this

alternative less attractive than using numbered logical point-to-point links.

4. Using OSPF over FR

The ability to configure distinct logical interfaces over FR gives users a great deal of flexibility in designing FR networks for use with OSPF. Because routers can be partially interconnected over FR, it is possible to design networks more cost-effectively than before. The issues to consider are the price/cost structure for VCs (fixed, distance-sensitive, banded) and ports, performance guarantees provided, traffic distribution (local, long-haul), and protocol efficiency. We have mentioned that the NBMA model provides some economy in OSPF protocol processing and overhead and is recommended for small homogeneous networks. In general, users should configure their networks to contain several small "NBMA clusters," which are in turn interconnected by long-haul VCs. The best choices for the number of routers in each cluster and the size of the long-haul logical point-to-point links depends on the factors mentioned above. If it is necessary to architect a more "flat" network, the ability to assign different link metrics to different (groups of) VCs allows for greater efficiency in using FR resources since VCs with better performance characteristics (throughput, delay) could be assigned lower link metrics.

5. Conclusion

We have specified guidelines for OSPF implementors and users to bring about improvements in how the protocol runs over frame relay networks. These recommendations do not require any protocol changes and allow for simpler, more efficient and cost-effective network designs. We recommend that OSPF implementations be able to support logical interfaces, each consisting of one or more virtual circuits and used as numbered logical point-to-point links (one VC) or logical NBMA networks (more than one VC). The current NBMA model for frame relay should continue to be used for small homogeneous networks consisting of a few routers.

6. References

- [1] Bradley, T., Brown, C., and A. Malis, "Multiprotocol Interconnect over Frame Relay", RFC 1294, Wellfleet Communications, Inc., BBN Communications, January 1992.
- [2] Moy, J., "OSPF Version 2", RFC 1583, Proteon, Inc., March 1994.
- [3] McCloghrie, K., and M. Rose, Editors, "Management Information Base for Network Management of TCP/IP-based Internets: MIB-II", STD 17, RFC 1213, Hughes LAN Systems, Inc., Performance Systems International, March 1991.
- [4] Bradley, T., and C. Brown, "Inverse Address Resolution Protocol", RFC 1293, Wellfleet Communications, Inc., January 1992.
- [5] Baker, F., and R. Coltun, "OSPF Version 2 Management Information Base", RFC 1253, ACC, Computer Science Center, August 1991.

Security Considerations

Security issues are not discussed in this memo.

7. Authors' Addresses

Osmund S. deSouza
AT&T Bell Laboratories
Room 1K-606
101 Crawfords Corner Road
Holmdel, NJ 07733

Phone: (908) 949-1393
EMail: osmund.desouza@att.com

Manoel A. Rodrigues
Room 1K-608
AT&T Bell Laboratories
101 Crawfords Corner Road
Holmdel, NJ 07733

Phone: (908) 949-4655
EMail: manoel.rodrigues@att.com

