

Network Working Group
Internet Draft
Expiration Date: July 2004
File name: [draft-ietf-ospf-dc-07.txt](#)

Sira Panduranga Rao
UTA
Alex Zinin
Alcatel
Abhay Roy
Cisco Systems

January 2004

Detecting Inactive Neighbors over OSPF Demand Circuits
draft-ietf-ospf-dc-07.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. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "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>.

Abstract

OSPF is a link-state intra-domain routing protocol used in IP networks. OSPF behavior over demand circuits is optimized in [RFC1793](#) to minimize the amount of overhead traffic. A part of OSPF demand circuit extensions is the Hello suppression mechanism. This technique allows a demand circuit to go down when no interesting traffic is going through the link. However, it also introduces a problem, where it becomes impossible to detect a OSPF-inactive neighbor over such a link. This memo introduces a new mechanism called "neighbor probing" to address the above problem.

1. Motivation

In some situations, when operating over demand circuits, the remote neighbor may be unable to run OSPF, and, as a possible result, unable to route application traffic. Possible scenarios include:

- o The OSPF process might have died on the remote neighbor.
- o Oversubscription ([Section 7 of \[RFC1793\]](#)) may cause a continuous drop of application data at the link level.

Problem here is that the local router cannot identify the problems such as this, since Hello exchange is suppressed on demand circuits. If the topology of the network is such that other routers cannot communicate their knowledge about the remote neighbor via flooding, the local router and all routers behind it will never know about the problem, so application traffic may continue being forwarded to the OSPF-incapable router.

This memo describes a backward-compatible neighbor probing mechanism based on the details of the standard flooding procedure followed by OSPF routers.

2. Proposed Solution

The solution this document proposes uses the link-state update packets to detect whether the OSPF process is operational on the remote neighbor. We call this process "Neighbor probing". The idea behind this technique is to allow either of the two neighbors connected over a demand circuit to test the remote neighbor at any time (see [Section 2.1](#)).

The routers across the demand circuit can be connected by either a point-to-point link, a virtual link, or a point-to-multipoint interface. The case of routers connected by broadcast networks or NBMA is not considered, since Hello suppression is not used in these cases ([Section 3.2 \[RFC1793\]](#)).

The neighbor probing mechanism is used as follows. After a router has synchronized the LSDB with its neighbor over the demand circuit, the demand circuit may be torn down if there is no more application traffic. When application traffic starts going over the link, the link is brought up. If `ospfIfDemandNbrProbe` is enabled, the routers SHOULD probe each other. While the link is up, the routers may also periodically probe each other every `ospfIfDemandNbrProbeInterval`. Neighbor probing should not be considered as interesting traffic and do not cause the demand circuit to remain up (relevant details of implementation are outside of the scope of this document).

The case when one or more of the router's links are oversubscribed (see [section 7 of \[RFC1793\]](#)) should be considered by the implementations. In such a situation even if the link status is up and application data is being sent on the link, only a limited number of neighbors is really reachable. To make sure temporarily unreachable neighbors are not mistakenly declared down, Neighbor probing should be restricted to those neighbors that are actually reachable (i.e., there is a circuit established with the neighbor at the moment the probing procedure needs to be initiated). This check itself is also considered an implementation detail.

2.1 Neighbor Probing

The neighbor probing method described in this section is completely compatible with standard OSPF implementations, because it is based on standard behavior that must be followed by OSPF implementations in order to keep their LSDBs synchronized.

When a router needs to verify OSPF capability of a neighbor reachable through a demand circuit, it should flood to the neighbor any LSA in its LSDB that would normally be sent to the neighbor during the initial LSDB synchronization process (in most cases such an LSA must have already been flooded to the neighbor by the time the probing procedure starts). For example, the router may flood its own router-LSA (without originating a new version), or the neighbor's own router-LSA. If the neighbor is still alive and OSPF-capable, it replies with a link state acknowledgement or a link state update (an implied acknowledgement) and the LSA is removed from the neighbor's retransmission list. The implementations should limit the number of times an LSA can be retransmitted to `ospfIfDemandNbrProbeRetxLimit`, when used for neighbor probing. If no acknowledgement (explicit or implicit) is received for a predefined period of time, the probing router should treat this as evidence of the neighbor's unreachability (proving wrong the assumption of reachability used in [\[RFC1793\]](#)) and should bring the adjacency down.

Note that when the neighbor being probed receives such a link state update packet, the received LSA has the same contents as the LSA in the neighbor's LSDB, and hence should normally not cause any additional flooding. However, since LSA refreshes are not flooded over demand circuits, the received LSA may have a higher Sequence Number. This will result in the first probe LSA being flooded further by the neighbor. Note that if the current version of the probe LSA has already been flooded to the neighbor, it will not be propagated any further by the neighbor. Also note that in any case subsequent (non-first) probe LSAs will not cause further flooding until the

LSA's sequence number is incremented.

Again, the implementation should insure (through internal mechanisms) that OSPF link state update packets sent over the demand circuit for the purpose of neighbor probing do not prevent that circuit from being torn down.

3. Support of Virtual Links and Point-to-multipoint Interfaces

Virtual links can be treated analogous to point-to-point links and so the techniques described in this memo are applicable to virtual links as well. The case of point-to-multipoint interface running as demand circuit ([section 3.5 \[RFC1793\]](#)) can be treated as individual point-to-point links, for which the solution has been described in [section 2](#).

4. Compatibility issues

All mechanisms described in this document are backward-compatible with standard OSPF implementations.

5. Considerations

In addition to the lost functionality mentioned in [Section 6 of \[RFC1793\]](#), there is an added overhead in terms of the amount of data (link state updates and acknowledgements) being transmitted due to neighbor probing whenever the link is up and thereby increasing the overall cost.

6. Acknowledgements

The original idea of limiting the number of LSA retransmissions on demand circuits (used as part of the solution described in this document) and its implementation belong to Padma Pillay-Esnault and Derek Yeung.

The authors would like to thank John Moy, Vijayapal Reddy Patil, SVR Anand, and Peter Psenak for their comments on this work.

A significant portion of Sira's work was carried out as part of the HFCL-IIISc Research Project (HIRP), Bangalore, India. He would like to thank the team for their insightful discussions.

7. Security Considerations

The mechanism described in this document does not modify security aspects of the OSPF routing protocol.

8. Normative References

[RFC2328]

Moy, J., "OSPF Version 2", [RFC 2328](#), April 1998.

[RFC1793]

Moy, J., "Extending OSPF to Support Demand Circuits", [RFC 1793](#), April 1995.

Appendix A. Configurable Parameters

This memo defines the following additional configuration parameters for OSPF interfaces.

`ospfIfDemandNbrProbe`

Indicates whether or not neighbor probing is enabled to determine whether or not the neighbor is inactive. Neighbor probing is disabled by default.

`ospfIfDemandNbrProbeRetxLimit`

The number of consecutive LSA retransmissions before the neighbor is deemed inactive and the neighbor adjacency is brought down. Sample value is 10 consecutive LSA retransmissions.

`ospfIfDemandNbrProbeInterval`

Defines how often the neighbor will be probed. Sample value is 2 minutes.

Authors' addresses

Sira Panduranga Rao
The University of Texas at Arlington
Arlington, TX 76013
Email: siraprao@hotmail.com

Alex Zinin
Alcatel
Sunnyvale, CA
E-mail: zinin@psg.com

Abhay Roy
Cisco Systems
170 W. Tasman Dr.
San Jose, CA 95134
E-mail: akr@cisco.com

