

Network Working Group
Request for Comments: 3883
Updates: 1793
Category: Standards Track

S. Rao
UTA
A. Zinin
Alcatel
A. Roy
Cisco Systems
October 2004

Detecting Inactive Neighbors over OSPF Demand Circuits (DC)

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2004).

Abstract

OSPF is a link-state intra-domain routing protocol used in IP networks. OSPF behavior over demand circuits (DC) is optimized in RFC 1793 to minimize the amount of overhead traffic. A part of the 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 an 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 [RFC2328], 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.

The problem here is that the local router cannot identify problems such as this, since the 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 the 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 Non-Broadcast Multi-Access (NBMA) links 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 Link State Database (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 should 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 are 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 the 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 analogously to point-to-point links, so the techniques described in this memo are applicable to virtual links as well. The case of point-to-multipoint interface running as a 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. Deployment Considerations

In addition to the lost functionality mentioned in Section 6 of [RFC1793], there is additional overhead in terms of the amount of data (link state updates and acknowledgements) being transmitted due to neighbor probing whenever the link is up, 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-IISc 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", STD 54, 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 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. The sample value is 2 minutes.

Authors' Addresses

Sira Panduranga Rao
The University of Texas at Arlington
416 Yates Street, 300 Nedderman Hall
Arlington, TX 76019

EEmail: siraprao@hotmail.com

Alex Zinin
Alcatel
701 E Middlefield Rd
Mountain View, CA 94043

EEmail: zinin@psg.com

Abhay Roy
Cisco Systems
170 W. Tasman Dr.
San Jose, CA 95134

EEmail: akr@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2004).

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/S HE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY 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 IETF's procedures with respect to rights in IETF 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

