Network Working Group
Internet Draft
Expiration Date: January 2001
File name: draft-zinin-ospf-stub-adv-00.txt

Alex Zinin
Alvaro Retana
Liem Nguyen
Russ White
Cisco Systems
July 2000

OSPF Stub Router Advertisement
draft-zinin-ospf-stub-adv-00.txt

Status of this Memo

Abstract

In some cases, it is desirable not to route transit traffic via a specific OSPF router. However, OSPF [RFC2328] does not specify a standard way to accomplish this. This memo describes a backward-compatible technique that may be used by OSPF implementations to advertise unavailability to forward transit traffic or to lower the preference level for the paths through such a router.

1 Motivation

In some situations, it may be advantageous to inform routers in a network not to use a specific router as a transit point, but still route to it. Possible situations include the following.

- The router is in a critical condition (for example, has very high CPU load or does not have enough memory to store all LSAs or build the routing table).
- Graceful introduction and removal of the router to/from the network.
- Other (administrative or traffic engineering) reasons.

Note that the proposed solution does not remove the router from the topology view of the network (as could be done by just flushing that router's router-LSA), but prevents other routers from using it for transit routing, while still routing packets to router's own IP addresses, i.e., the router is announced as stub.

It must be emphasized that the proposed solution provides real benefits in networks designed with at least some level of redundancy so that traffic can be routed around the stub router. Otherwise, traffic destined for the networks reachable through such a stub router will either be still routed through it or black-holed (see Section 4 for more details).

2 Proposed Solution

The solution described in this document solves three challenges associated with the outlined problem. In the description below, router X is the router announcing itself as a stub.

1) Making other routers prefer routes around router X while performing the Dijkstra calculation.

2) Allowing other routers to reach IP prefixes directly connected to router X

3) Allowing router X itself to use its links to reach remote routers and destinations.

Note that it would be easy to address issue 1) alone by just flushing router X's router-LSA from the domain. However, it does not solve problems 2) and 3), since other routers will not be able to use links to router X in Dijkstra (no back link), and because router X will not have links to its neighbors.

To address all three problems, router X announces its router-LSA to

the neighbors as follows.

- costs of all non-stub links (links of the types other than 3) are set to LSInfinity (16-bit value 0xFFFF, rather than 24-bit value 0xFFFFFF used in summary and AS-external LSAs).
- costs of stub links (type 3) are set to the interface output cost.

This addresses issues 1) and 2). To address issue 3), router X must use real link costs in its Dijkstra calculation. This will allow it to use its links to reach remote routers.

3 Implementation details

To simplify the implementation of the described technique, it may be useful for OSPF routers to keep two versions of their router-LSAs. One version (public) is installed in the LSDB and sent to the neighbors, while another version (internal) is used for local routing table calculation. When the router announces itself as stub, it constructs the public version as indicated in Section 2, but the internal version is constructed as in standard OSPF [RFC2328]. When the router does not announce itself as stub, both versions are constructed as in standard OSPF and would both yield the same result, hence it is not necessary (though acceptable) to keep two versions of the LSAs at this point.

4 Compatibility issues

Some inconsistency may be seen when the network is constructed of the routers that perform intra-area Dijkstra calculation as specified in [RFC1247] (discarding link records in router-LSAs that have LSInfinity cost value) and routers that perform it as specified in [RFC1583] and higher (do not treat links with LSInfinity cost as unreachable). Note that this inconsistency will not lead to routing loops, because if there are some alternate paths in the network, both types of routers will agree on using them rather than the path through the stub router. If the path through the stub router is the only one, the routers of the first type will not use the stub router for transit (which is the desired behavior), while the routers of the second type will still use this path.

5 Acknowledgements

The authors of this document do not make any claims on the originality of the ideas described. Among other people, we would like to acknowledge Henk Smit for being part of one of the initial

discussions around this topic.

6 Security Considerations

The technique described in this document does not introduce any new security issues into OSPF protocol.

7 References

[RFC2328]
J. Moy. OSPF version 2. Technical Report RFC 2328, Internet Engineering Task Force, 1998. ftp://ftp.isi.edu/in-notes/rfc2328.txt.

[RFC1247]
J. Moy. OSPF version 2. RFC 1247, Internet Engineering Task Force, 1991. ftp://ftp.isi.edu/in-notes/rfc1247.txt.

[RFC1583]
J. Moy. OSPF version 2. RFC 1583, Internet Engineering Task Force, 1994. ftp://ftp.isi.edu/in-notes/rfc1583.txt.

8 Authors' addresses

Alex D. Zinin
Cisco Systems
150 West Tasman Dr.
San Jose,CA 95134
E-mail: azinin@cisco.com

Liem Nguyen
7025 Kit Creek Rd.
Research Triangle Park, NC 27709
USA
e-mail: lhnguyen@cisco.com

Alvaro Retana
7025 Kit Creek Rd.
Research Triangle Park, NC 27709
USA
e-mail: aretana@cisco.com

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