Network Working Group
Internet Draft
Expiration Date: August 1998
File name: draft-ietf-ospf-stdreport-03.txt

J. Moy
Ascend Communications, Inc.

February 1998

OSPF Standardization Report

Status of this Memo

Abstract

This memo documents how the requirements for advancing a routing protocol to Full Standard, set out in [Ref2], have been met for OSPFv2.

Please send comments to ospf@gated.cornell.edu.

Table of Contents

1. Introduction

OSPFv2, herein abbreviated simply as OSPF, is an IPv4 routing protocol documented in [Ref8]. OSPF is a link-state routing protocol. It is designed to be run internal to a single Autonomous System. Each OSPF router maintains an identical database describing the Autonomous System's topology. From this database, a routing table is calculated by constructing a shortest-path tree. OSPF features include the following:

- OSPF responds quickly to topology changes, expending a minimum of network bandwidth in the process.
- Support for CIDR addressing.
- OSPF routing exchanges can be authenticated, providing routing security.
- Equal-cost multipath.
- An area routing capability is provided, enabling an Autonomous system to be split into a two level hierarchy to further reduce the amount of routing protocol traffic.
- OSPF allows import of external routing information into the Autonomous System, including a tagging feature that can be exploited to exchange extra information at the AS boundary (see [Ref7]).

An analysis of OSPF together with a more detailed description of OSPF features was originally provided in [Ref6], as a part of promoting OSPF to Draft Standard status. The analysis of OSPF remains unchanged. Two additional major features have been developed for OSPF since the protocol achieved Draft Standard status: the Point-to-MultiPoint interface and Cryptographic Authentication. These features are described in Sections 2.1 and 2.2 respectively of this memo.

The OSPF MIB is documented in [Ref4]. It is currently at Draft Standard status.

2. Modifications since Draft Standard status

OSPF became a Draft Standard with the release of RFC 1583 [Ref3]. Implementations of the new specification in [Ref8] are backward-compatible with RFC 1583. The differences between the two documents are described in the Appendix Gs of [Ref1] and [Ref8]. These differences are listed briefly below. Two major features were also

added, the Point-to-MultiPoint interface and Cryptographic Authentication, which are described in separate sections.

- Configuration requirements for OSPF area address ranges have been relaxed to allow greater flexibility in area assignment. See Section G.3 of [Ref1] for details.

- The OSPF flooding algorithm was modified to a) improve database convergence in networks with low speed links b) resolve a problem where unnecessary LSA retransmissions could occur as a result of differing clock granularities, c) remove race conditions between the flooding of MaxAge LSAs and the Database Exchange process, d) clarify the use of the MinLSArrival constant, and e) rate-limit the response to less recent LSAs received via flooding. See Sections G.4 and G.5 of [Ref1] and Section G.1 of [Ref8] for details.

- To resolve the long-standing confusion regarding representation of point-to-point links in OSPF, the specification now optionally allows advertisement of a stub link to a point-to-point link's subnet, ala RIP. See Section G.6 of [Ref1].

- Several problems involving advertising the same external route from multiple areas were found and fixed, as described in Section G.7 of [Ref1] and Section G.2 of [Ref8]. Without the fixes, persistent routing loops could form in certain such configurations. Note that one of the fixes was not backward-compatible, in that mixing routers implementing the fixes with those implementing just RFC 1583 could cause loops not present in an RFC 1583-only configuration. This caused an RFC1583Compatibility global configuration parameter to be added, as described in Section C.1 of [Ref1].

- In order to deal with high delay links, retransmissions of initial Database Description packets no longer reset an OSPF adjacency.

- In order to detect link MTU mismatches, which can cause problems both in IP forwarding and in the OSPF routing protocol itself, MTU was added to OSPF's Database Description packets. Neighboring routers refuse to bring up an OSPF adjacency unless they agree on their common link's MTU.

- The TOS routing option was deleted from OSPF. However, for backward compatibility the formats of OSPF's various LSAs remain unchanged, maintaining the ability to specify TOS metrics in router-LSAs, summary-LSAs, ASBR-summary-LSAs, and AS-external-LSAs.

- o OSPF's routing table lookup algorithm was changed to reflect current practice. The "best match" routing table entry is now always selected to be the one providing the most specific (longest) match. See Section G.4 of [Ref8] for details.

2.1. Point-to-MultiPoint interface

The Point-to-MultiPoint interface was added as an alternative to OSPF's NBMA interface when running OSPF over non-broadcast subnets. Unlike the NBMA interface, Point-to-MultiPoint does not require full mesh connectivity over the non-broadcast subnet. Point-to-MultiPoint is less efficient than NBMA, but is easier to configure (in fact, it can be self-configuring) and is more robust than NBMA, tolerating all failures within the non-broadcast subnet. For more information on the Point-to-MultiPoint interface, see Section G.2 of [Ref1].

There are at least six independent implementations of the Point-to-MultiPoint interface. Interoperability has been demonstrated between at least two pairs of implementations: between 3com and Bay Networks, and between cisco and Cascade.

2.2. Cryptographic Authentication

Non-trivial authentication was added to OSPF with the development of the Cryptographic Authentication type. This authentication type uses any keyed message digest algorithm, with explicit instructions included for the use of MD5. For more information on OSPF authentication, see Section 4.

There are at least three independent implementations of the OSPF Cryptographic authentication type. Interoperability has been demonstrated between the implementations from cisco and Cascade.

3. Updated implementation and deployment experience

When OSPF was promoted to Draft Standard Status, a report was issued documenting current implementation and deployment experience (see [Ref6]). That report is now quite dated. In an attempt to get more current data, a questionnaire was sent to OSPF mailing list in January 1996. Twelve responses were received, from 11 router vendors and 1 manufacturer of test equipment. These responses represented 6 independent implementations. A tabulation of the results are presented below.

Table 1 indicates the implementation, interoperability and deployment of the major OSPF functions. The number in each column represents the number of responses in the affirmative.

Feature	Imple-mented	Inter-operated	Deployed
OSPF areas	10	10	10
Stub areas	10	10	9
Virtual links	10	9	8
Equal-cost multipath	10	7	8
NBMA support	9	8	7
CIDR addressing	8	5	6
OSPF MIB	8	5	5
Cryptographic auth.	3	1	1
Point-to-Multipoint ifc.	6	3	4

Table 1: Implementation of OSPF features

Table 2 indicates the size of the OSPF routing domains that vendors have tested. For each size parameter, the number of responders and the range of responses (minimum, mode, mean and maximum) are listed.

Parameter	Responses	Min	Mode	Mean	Max
Max routers in domain	7	30	240	460	1600
Max routers in single area	7	20	240	380	1600
Max areas in domain	7	1	10	16	60
Max AS-external-LSAs	9	50	10K	10K	30K

Table 2: OSPF domain sizes tested

Table 3 indicates the size of the OSPF routing domains that vendors have deployed in real networks. For each size parameter, the number of responders and the range of responses (minimum, mode, mean and maximum) are listed.

Parameter	Responses	Min	Mode	Mean	Max
Max routers in domain	8	20	350	510	1000
Max routers in single area	8	20	100	160	350
Max areas in domain	7	1	15	23	60
Max AS-external-LSAs	6	50	1K	2K	5K

Table 3: OSPF domain sizes deployed

In an attempt to ascertain the extent to which OSPF is currently deployed, vendors were also asked in January 1998 to provide deployment estimates. Four vendors of OSPF routers responded, with a total estimate of 182,000 OSPF routers in service, organized into 4300 separate OSPF routing domains.

4. Protocol Security

All OSPF protocol exchanges are authenticated. OSPF supports multiple types of authentication; the type of authentication in use can be configured on a per network segment basis. One of OSPF's authentication types, namely the Cryptographic authentication option, is believed to be secure against passive attacks and provide significant protection against active attacks. When using the Cryptographic authentication option, each router appends a "message digest" to its transmitted OSPF packets. Receivers then use the shared secret key and received digest to verify that each received OSPF packet is authentic.

The quality of the security provided by the Cryptographic authentication option depends completely on the strength of the message digest algorithm (MD5 is currently the only message digest algorithm specified), the strength of the key being used, and the correct implementation of the security mechanism in all communicating OSPF implementations. It also requires that all parties maintain the secrecy of the shared secret key.

None of the OSPF authentication types provide confidentiality. Nor do they protect against traffic analysis. Key management is also not addressed by the OSPF specification.

For more information, see Sections 8.1, 8.2, and Appendix D of [Ref1].

References

[Ref1] Moy, J., "OSPF Version 2", RFC 2178, July 1997.

[Ref2] Hinden, B., "Internet Routing Protocol Standardization Criteria", RFC 1264, October 1991.

[Ref3] Moy, J., "OSPF Version 2", RFC 1583, March 1994.

[Ref4] Baker, F,, and R. Coltun, "OSPF Version 2 Management Information Base", RFC 1850, November 1995.

[Ref5] Moy, J., "OSPF Protocol Analysis", RFC 1245, August 1991.

[Ref6] Moy, J., "Experience with the OSPF Protocol", RFC 1246, August 1991.

[Ref7] Varadhan, K., S. Hares and Y. Rekhter, "BGP4/IDRP for IP--- OSPF Interaction", RFC 1745, December 1994.

[Ref8] Moy, J., "OSPF Version 2", Internet Draft, <draft-ietf-ospf-vers2-02.txt>, January 1998.

Security Considerations

Security considerations are addressed in Section 4 of this memo.

Author's Address

John Moy
Ascend Communications, Inc.
1 Robbins Road
Westford, MA 01886

Phone: 978-952-1367
Fax: 978-392-2075
Email: jmoy@casc.com

This document expires in August 1998.