

Network Working Group
Internet Draft
Expiration Date: May 2003

Acee Lindem (Redback Networks)
Naiming Shen (Redback Networks)
Rahul Aggarwal (Redback Networks)
Scott Shaffer (Genuity, Inc.)
JP Vasseur (Cisco Systems, Inc)

Extensions to IS-IS and OSPF for Advertising
Optional Router Capabilities

[draft-raggarwa-igp-cap-01.txt](#)

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#), except that the right to produce derivative works is not granted.

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 and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``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>.

2. Abstract

It is useful for routers in a domain to know of the capabilities of their IGP neighbors, and/or of other routers in the domain. This draft proposes extensions to IS-IS and OSPF for advertising optional router capabilities. We define an optional Router Capability TLV for IS-IS, while for OSPF we define an optional Router Information Opaque LSA.

3. Motivation

It is useful for routers in a domain to know of the capabilities of their IGP neighbors, and/or of other routers in the domain. This can be useful for various purposes:

- o In MPLS Traffic Engineering (TE) as a TE discovery mechanism [15, 16] to announce a LSR's TE capabilities like Path Computation Server capability (Capability of a LSR to be a Path Computation Server for TE LSP path computation) or the intention of a LSR to be part of a particular MPLS TE mesh group.
- o For network management and troubleshooting. It gives operators a network wide view of IGP capabilities on different routers in the network. The presence of a capability on a given router implies that the software version supports the capability and the router is configured to support it. On the other hand the absence of an expected capability on a particular router can imply either mis-configuration or an incorrect software version. Hence this capability information can be used to track problems resulting from mis-configuration or an incorrect software version.

There is no existing mechanism in IS-IS to advertise optional router capabilities. On the other hand OSPF uses the options field in the hello packet to advertise optional router capabilities [2]. However this attribute is not extensible for advertising optional capabilities such as hitless graceful restart or MPLS TE capabilities. We propose extensions to IS-IS and OSPF for advertising these optional capabilities. For current IS-IS and OSPF capabilities this advertisement will be used primarily for MPLS TE and informational purposes. Conceivably, future IS-IS and OSPF capability advertisements could be used for other purposes.

4. IS-IS Router Capability TLV

IS-IS [7] routers may optionally advertise their router capabilities in the TLV with code type 242. This TLV specifies the router ID of the router that originates the TLV, defines the flooding scope of the TLV, specifies the router capability bits in the first sub-TLV and certain capability related information in other sub-TLVs. This draft does not specify how an application may use the Router Capability TLV and such specification is outside the scope of this draft.

The router ID is a 32 bit unsigned integer to represent the router that originated this capability TLV. This is needed since this TLV can be flooded over the entire domain, hence the router ID of the originating router must be kept.

The capability bits are defined in a mandatory sub-TLV with code 1. It starts as a 32 bits flag, where each bit can represent a router capability. This flag can be expanded as needed to include more capabilities.

Some of the router capabilities may require more information than a single bit. The extra capability information can be encoded as sub-TLVs under this router capability TLV. The definition of these sub-TLVs is outside the scope of this draft.

If a router does not advertise this TLV, it does not imply that the router does not support one or more of the defined capabilities. If this TLV is included in the LSP, the router SHOULD set all the defined bits corresponding to the capabilities which the software supports, unless they are explicitly configured off.

4.1 Flooding Scope of the Router Capability TLV

There are three bits currently defined for this TLV in the information flag to control the flooding scope of the TLV. The Flooding bit, the Transit bit and the Down bit.

There are two flooding types defined for this router capability TLV's flooding scope. One is the domain wide flooding scope and the other is the intra-area flooding scope. The F bit if set indicates this TLV has the domain wide flooding scope.

The Transit bit can be used to signal the routers on the edge of the IGP routing domain to redistribute this TLV information into another routing process. How this is done is an application specific issue and is outside the scope of this document.

The L1/L2 routers MUST observe the Down bit to avoid TLV leak looping. This Down bit is not set when the router first originates this TLV and it MUST be set when leaking into a lower level or into another area of the same level. When the Down bit is set, this TLV can no longer be leaked to a higher level or into another area of the same level. This capability TLV MUST be preserved at the level boundary during TLV leaking. The L1/L2 router SHOULD NOT leak the TLV back into the same area which originated this TLV. It MAY be able to alter certain capability contents during TLV leaking when specified by applications.

4.2 Encoding of the Router Capability TLV

The following figure depicts the structure of this IS-IS Router Capability TLV.

- x CODE - 242
- x LENGTH - total length of the value field in this TLV
- x VALUE - 4-octet information flag, 4-octet router ID, 1-octet sub-tlv length, the mandatory sub-TLV code 1 for capability flags, and optional sub-TLVs for extra capability information, structured as follows:

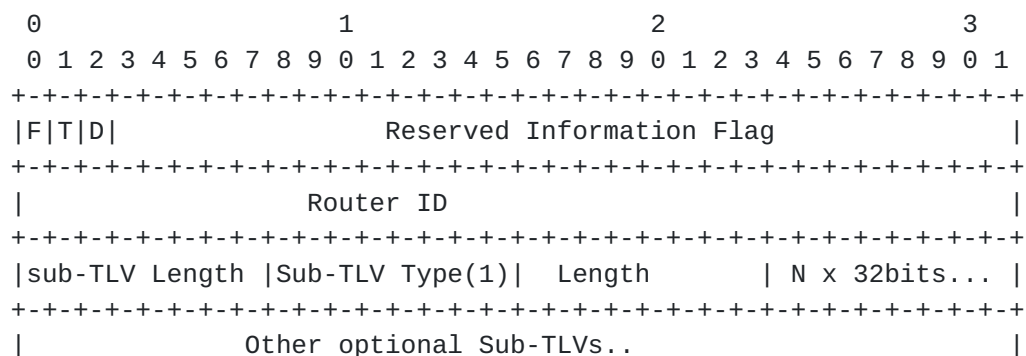


Figure 1. IS-IS Router Capability TLV

The first field is the 4-octet information flag, which consists of the F, T and D bits, the reserved information bits.

Bit F represents the Flooding scope of the TLV. If set, this TLV SHOULD be flooded to entire IGP domain. Otherwise, it SHOULD NOT be leaked into the other level or another area in the same level.

Bit T determines the Transit behavior into other routing domains. For example, if this bit is set, a router can leak this capability information into another routing protocol.

Bit D represents Down/Up behavior during the TLV leaking. When the capability is leaked from level 2 into level 1 or it is leaked into another area of the same level, this D bit MUST be set. Otherwise this bit MUST be cleared.

Router ID is an unsigned 32 bit number representing the router that originates this router capability TLV.

The next octet of the TLV is the total sub-TLV length of this router capability TLV. This sub-TLV length includes the first mandatory sub-TLV. The minimum value of this field is 6.

The first sub-TLV with code 1 is a mandatory sub-TLV, the router capability flag sub-TLV. The length is the length of this sub-TLV. Its set to $N \times 4$ octets. N starts from 1 and can be increased when there is a need. Each 4 octets are referred to as a capability flag. For each capability flag the bits are indexed from the most significant to the least significant, where each bit represents one router capability.

There can be other sub-TLVs after the first sub-TLV to include extra information describing certain router capabilities. The description of those sub-TLVs is outside the scope of this draft.

The above data structure can be replicated within this TLV, but can not exceed the maximum length of 255 octets. If no other sub-TLVs are used and the capability flag is the minimum 4 octets, this encoding can contain up to 17 router capability TLVs where each have a minimum of 15 octets of data (4 byte information flag, 4 byte router-id, 1 byte total sub-tlv length, 6 byte capability flag).

4.3 Reserved IS-IS Router Capability Bits

We have assigned some pre-determined bits to the first capability flag.

Bit	Capabilities
0-3	Reserved
4	IS-IS hitless graceful restart capable [9]
5	IS-IS and BGP blackhole avoidance capable [11]
6	IS-IS wide metric processing capable [8]
7	IS-IS short metric processing capable [6]
8	IS-IS hmac-md5 authentication capable [10]
9	IS-IS Traffic Engineering support [8]
10	IS-IS point-to-point over LAN [12]
11	IS-IS Path Computation Server discovery [16]
12	M-ISIS capable [13]
13	IS-IS IPv6 capable [14]
14-31	For future assignments

5. OSPF Router Information LSA

OSPF routers will optionally advertise their optional capabilities in an area-scoped, local scope, or AS-scoped Opaque-LSA [1]. If a router does not advertise this LSA, it does not imply that the router does not support one or more of the defined capabilities. For current OSPF capabilities, the advertisement will be used for MPLS TE [15] and information purposes. Conceivably, future OSPF capabilities could require other capability LSA advertisement. The Router Information LSA will be originated at startup and re-originated when router capabilities change or when periodically refreshed.

The Router Information LSA will have an Opaque type of 4 and Opaque ID of 0.

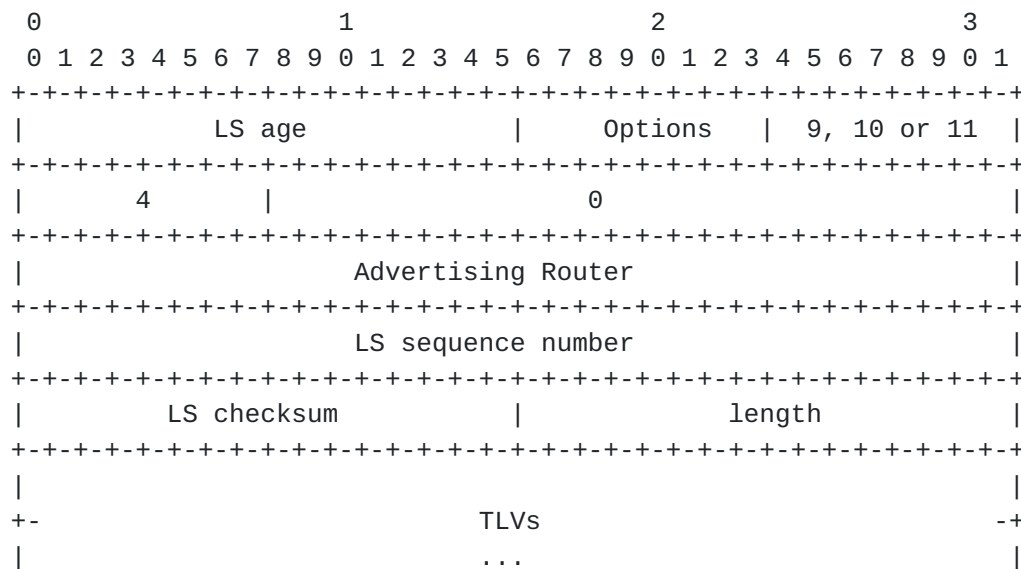


Figure 2. OSPF Router Information LSA

The format of the TLVs within the body of a Router Information LSA is the same as the TLV format used by the Traffic Engineering Extensions to OSPF [3]. The TLV header consists of a 16-bit Type field and a 16-bit length field, and is followed by zero or more bytes of value. The length field indicates the length of the value portion in bytes. The value portion is padded to four-octet alignment, but the padding is not included in the length field. For example, a one byte value would have the length field set to 1, and three bytes of padding would be added to the end of the value portion of the TLV.




Figure 3. OSPF TLV Format

5.1 OSPF Router Capability TLV

The first TLV in the body of a Router Information LSA is the Router Capability TLV. It MUST be included. A router advertising an optional Router Information LSA SHOULD set the supported optional capabilities, unless they are explicitly configured off, in the Router Capability TLV.

The format of the Router Capability TLV is as follows :

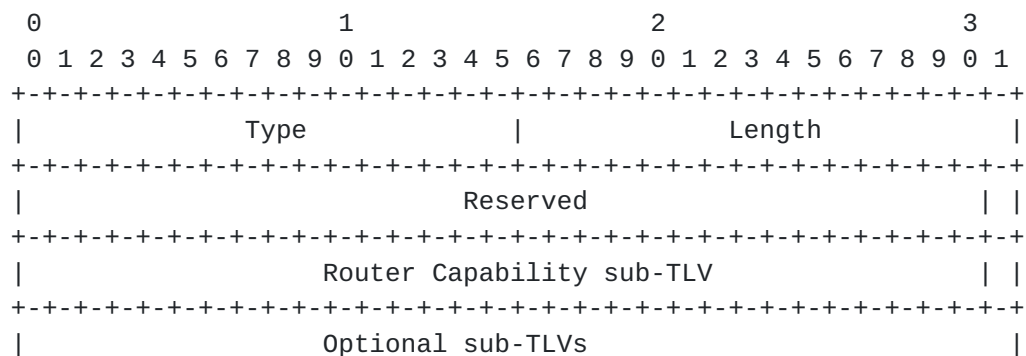


Figure 4. OSPF Router Capability TLV

Type	A 16 bit field set to 1.
Length	A 16 bit field that indicates the length of the TLV, other than the Type and the Length fields in bytes.

The first four bytes of the TLV are reserved. This is followed by a Router Capability sub-TLV that **MUST** be included. The format of the Router Capability sub-TLV is as follows :




Figure 5. OSPF Router Capability Sub-TLV

Type	A 16 bit field set to 1.
Length	A 16 bit field that indicates the length of the value portion in bytes. Its set to N x 4 octets. N starts from 1 and can be increased when there is a need. Each 4 octets are referred to as a capability flag.
Value	This comprises one or more capability flags. For each 4 octets, the bits are indexed from the most significant to the least significant, where each bit represents one router capability. When the first 32 capabilities are defined, a new capability flag will be used to accommodate the next capability.

The Router Capability sub-TLV MAY be followed by optional sub-TLVs. In some cases it may be desirable to advertise additional information for a particular capability. This can be done by including other sub-TLVs.

5.2 Reserved OSPF Router Capability Bits

We have assigned some pre-determined bits to the first capability flag.

Bit	Capabilities
0-3	Reserved
4	Hitless graceful restart capable [4]
5	OSPF hitless graceful restart helper [4]
6	Stub Router support [5]
7	Traffic Engineering support [3]
8	OSPF point-to-point over LAN [12]
9	OSPF Path Computation Server discovery [15, 16]
10-31	Future assignments

5.3 Flooding Scope of the Router Information LSA

The flooding scope of the Router Information LSA is determined by the LSA type. A local scope i.e. Type 9 LSA is flooded on a link, an area-scoped i.e. Type 10 LSA is flooded throughout the area, while an AS-scoped i.e. Type 11 LSA is flooded throughout the AS. Choice of the flooding scope is made by the advertising router and is a matter of local policy. A router information LSA must be announced using one flooding mode. A router may announce more than one router information LSA for local scope, intra-area scope or domain scope capabilities.

6. Security Consideration

This document does not introduce new security issues. The security considerations pertaining to the original IS-IS and OSPF protocols remain relevant.

7. Acknowledgments

The idea for this work grew out of a conversation with Andrew Partan and we would like to thank him for his contribution.

8. References

- [1] Coltun, R., "The OSPF Opaque LSA Option", [RFC 2370](#), July 1998.
- [2] Moy, J., "OSPF Version 2", [RFC 2328](#), April 1998.
- [3] Katz, D., D. Yeung and K. Kompella, "Traffic Engineering Extensions to OSPF", Internet Draft, work in progress.
- [4] Moy, J., "OSPF Hitless OSPF Restart", Internet Draft, work in progress.
- [5] Retana, A., et al, "OSPF Stub Router Advertisement", [RFC 3137](#), June 2001.
- [6] Callon, R., "OSI IS-IS for IP and Dual Environment", [RFC 1195](#), December 1990.

- [7] ISO, "Intermediate system to Intermediate system routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)," ISO/IEC 10589:1992.
- [8] Li, T. et al, "IS-IS Extensions for Traffic Engineering", Internet Draft, work in Progress.
- [9] Shand, M., "Restart Signaling for IS-IS", Internet Draft, work in Progress.
- [10] Li, T., "IS-IS Cryptographic Authentication", Internet Draft, work in progress.
- [11] McPherson, D., "IS-IS Transient Blackhole Avoidance", Internet Draft, work in progress.
- [12] N. Shen, et al, "Point-to-point operation over LAN in link-state-routing protocols", Internet Draft, work in progress.
- [13] T. Przygienda, N. Shen, N. Sheth, "M-ISIS: Multi Topology (MT) Routing in IS-IS", Internet Draft, work in progress.
- [14] C. Hopps, "Routing IPv6 with IS-IS", Internet Draft, work in progress.
- [15] Vasseur, Psenak, "Traffic Engineering Capability TLV for OSPF", Internet Draft, work in progress.
- [16] Vasseur et al, "RSVP Path computation request and reply " messages", [draft-vasseur-mpls-computation-rsvp-te-03.txt](#), work in progress

9. Author Information

Acee Lindem
Redback Networks
350 Holger Way
San Jose, CA 95134
e-mail: acee@redback.com

Naiming Shen
Redback Networks
350 Holger Way
San Jose, CA 95134
e-mail: naiming@redback.com

Rahul Aggarwal
Redback Networks
350 Holger Way
San Jose, CA 95134
e-mail: rahul@redback.com

Scott Shaffer
Genuity, Inc.
3 Van de Graaff Drive
PO Box 3073
Burlington, MA 01803
e-mail: sshaffer@genuity.com

JP Vasseur
Cisco Systems, Inc.
300 Apollo Drive
Chelmsford, MA 01824
e-mail: jpv@cisco.com