OSPF Transport Instance Extensions
draft-ietf-ospf-transport-instance-11.txt

Abstract

   OSPFv2 and OSPFv3 include a reliable flooding mechanism to
   disseminate routing topology and Traffic Engineering (TE) information
   within a routing domain. Given the effectiveness of these
   mechanisms, it is convenient to envision using the same mechanism for
   dissemination of other types of information within the domain.
   However, burdening OSPF with this additional information will impact
   intra-domain routing convergence and possibly jeopardize the
   stability of the OSPF routing domain. This document presents
   mechanism to relegate this ancillary information to a separate OSPF
   instance and minimize the impact.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

OSPFv2 [OSPFV2] and OSPFv3 [OSPFV3] include a reliable flooding
mechanism to disseminate routing topology and Traffic Engineering
(TE) information within a routing domain. Given the effectiveness of
these mechanisms, it is convenient to envision using the same
mechanism for dissemination of other types of information within the
domain. However, burdening OSPF with this additional information
will impact intra-domain routing convergence and possibly jeopardize
the stability of the OSPF routing domain. This document presents
mechanism to relegate this ancillary information to a separate OSPF
instance and minimize the impact.

1.1. Requirements notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
document are to be interpreted as described in [RFC-KEYWORDS].

[2](#). **OSPF Transport Instance**

   In order to isolate the effects of flooding and processing of non-
   routing information, it will be relegated to a separate protocol
   instance. This instance should be given lower priority when
   contending for router resources including processing, backplane
   bandwidth, and line card bandwidth. How that is realized is an
   implementation issue and is outside the scope of this document.

   Throughout the document, non-routing refers to routing information
   that is not used for IP or IPv6 routing calculations. The OSPF
   transport instance is ideally suited for dissemination of routing
   information for other protocols and layers.

[2.1](#). **OSPFv2 Transport Instance Packets Differentiation**

   OSPFv2 currently doesn't offer a mechanism to differentiate Transport
   instance packets from normal instance packets sent and received on
   the same interface. However, the [MULTI-INST] provides the necessary
   packet encoding to support multiple OSPF protocol instances.

[2.2](#). **OSPFv3 Transport Instance Packets Differentiation**

   Fortunately, OSPFv3 already supports separate instances within the
   packet encodings. The existing OSPFv3 packet header instance ID
   field will be used to differentiate packets received on the same link
   (refer to section 2.4 in [OSPFV3]).

[2.3](#). **Instance Relationship to Normal OSPF Instances**

   There are basically two alternatives for the relationship between a
   normal OSPF instance and an OSPF transport instance. In both cases,
   we must guarantee that any information we've received is treated as
   valid if and only if the router sending it is reachable. We'll refer
   to this as the "condition of reachability" in this document.

   1. Ships in the Night - The OSPF transport instance has no
       relationship or dependency on any other OSPF instance.

   2. Child Instance - The OSPF transport instance has a child-parent
       relationship with a normal OSPF instance and is dependent on this
       for topology information and assuring the "condition of
       reachability".

2.3.1. Ships in the Night Relationship to Normal OSPF Instances

   In this mode, the OSPF transport instance is not dependent on any
   other OSPF instance. It does, however, have much of the same as
   topology information must be advertised to satisfy the "condition of
   reachability".

   Prefix information does not need to be advertised. This implies that
   for OSPFv2, only router-LSAs, network-LSAs, and type 4 summary-LSAs
   need to be advertised. In the router-LSAs, the stub (type 3) links
   may be suppressed. For OSPFv3, this implies that router-LSAs,
   network-LSAs, and inter-area-router-LSAs must be advertised.

2.3.2. Tighter Coupling with Normal OSPF Instances

   Further optimizations and coupling between an OSPF transport instance
   and a normal OSPF instance are beyond the scope of this document.
   This is an area for future study.

2.4. Network Prioritization

   While OSPFv2 (section 4.3 in [OSPFV2]) are normally sent with IP
   precedence Internetwork Control, any packets sent by an OSPF
   transport instance will be sent with IP precedence Flash (B'011').
   This is only appropriate given that this is a pretty flashy
   mechanism.

   Similarly, OSPFv3 transport instance packets will be sent with the
   traffic class mapped to flash (B'011') as specified in ([OSPFV3].

   By setting the IP/IPv6 precedence differently for OSPF transport
   instance packets, normal OSPF routing instances can be given priority
   during both packet transmission and reception. In fact, some router
   implementations map the IP precedence directly to their internal
   packet priority. However, internal router implementation decisions
   are beyond the scope of this document.

2.5. OSPF Transport Instance Omission of Routing Calculation

   Since the whole point of the transport instance is to separate the
   routing and non-routing processing and fate sharing, a transport
   instance SHOULD NOT install any IP or IPv6 routes. OSPF routers
   SHOULD NOT advertise any transport instance LSAs containing IP or
   IPv6 prefixes and OSPF routers receiving LSAs advertising IP or IPv6
   prefixes SHOULD ignore them. This implies that an OSPFv2 transport
   instance Link State Database should not include any summary-LSAs
   (type 3) , AS-external-LSAs (type 5), or NSSA-LSAs (type 7) and the
   router-LSAs should not include any stub (type 3) links. An OSPFv3

transport instance Link State database should not include any inter-
area-prefix-LSAs (type 0x2003), AS-external-LSAs (0x4005), NSSA-LSAs
(type 0x2007), or intra-area-prefix-LSAs (type 0x2009). If they are
erroneously advertised, they will be flooded as per standard OSPF but
MUST be ignored by OSPF routers supporting this specification.

2.6. Non-routing Instance Separation

It has been suggested that an implementation could obtain the same
level of separation between IP routing information and non-routing
information in a single instance with slight modifications to the
OSPF protocol. The authors refute this contention for the following
reasons:

o Adding internal and external mechanisms to prioritize routing
   information over non-routing information are much more complex
   than simply relegating the non-routing information to a separate
   instance as proposed in this specification.

o The instance boundary offers much better separation for allocation
   of finite resources such as buffers, memory, processor cores,
   sockets, and bandwidth.

o The instance boundary decreases the level of fate sharing for
   failures. Each instance may be implemented as a separate process
   or task.

o With non-routing information, many times not every router in the
   OSPF routing domain requires knowledge of every piece of non-
   routing information. In these cases, groups of routers which need
   to share information can be segregated into sparse topologies
   greatly reducing the amount of non-routing information any single
   router needs to maintain.

2.7. Non-Routing Sparse Topologies

With non-routing information, many times not every router in the OSPF
routing domain requires knowledge of every piece of non-routing
information. In these cases, groups of routers which need to share
information can be segregated into sparse topologies. This will
greatly reduce the amount of information any single router needs to
maintain with the core routers possibly not requiring any non-routing
information at all.

With normal OSPF, every router in an OSPF area must have every piece
of topological information and every intra-area IP or IPv6 prefix.
With non-routing information, only the routers needing to share a set
of information need be part of the corresponding sparse topology.

For directly attached routers, one only needs to configure the
desired topologies on the interfaces with routers requiring the non-
routing information. When the routers making up the sparse topology
are not part of a uniconnected graph, two alternatives exist. The
first alternative is configure tunnels to form a fully connected
graph including only those routers in the sparse topology. The
second alternative is use remote neighbors as described in
Section 2.7.1.

2.7.1. Remote OSPF Neighbor

With sparse topologies, OSPF routers sharing non-routing information
may not be directly connected. OSPF adjacencies with remote
neighbors are formed exactly as they are with regular OSPF neighbors.
The main difference is that a remote OSPF neighbor's address is
configured and IP routing is used to deliver OSPF protocol packets to
the remote neighbor. Other salient feature of the remote neighbor
include:

o All OSPF packets have the remote neighbor's configured IP address
   as the IP destination address.

o The adjacency is represented in the router Router-LSA as a router
   (type-1) link with the link data set to the remote neighbor's
   configured IP address.

o Similar to NBMA networks, a poll-interval is configured to
   determine if the remote neighbor is reachable. This value is
   normally much higher than the hello interval with 40 seconds
   RECOMMENDED as the default.

[3](#). **OSPF Transport Instance Information Encoding**

   The format of the TLVs within the body of an LSA containing non-
   routing information is the same as the format used by the Traffic
   Engineering Extensions to OSPF [[TE](#)]. The LSA payload consists of one
   or more nested Type/Length/Value (TLV) triplets. The format of each
   TLV is:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Type              |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Value...                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

                           TLV Format

   However, each unique application using the mechanisms defined in this
   document will have it's own unique ID. Whether to encode this ID as
   the top-level TLV or make it part of the OSPF LSA ID is open for
   debate.

   The specific TLVs and sub-TLVs relating to a given application and
   the corresponding IANA considerations MUST for standard applications
   MUST be specified in the document corresponding to that application.

[3.1](#). **OSPFv2 Transport Instance Information Encoding**

   Application specific information will be flooded in opaque LSAs as
   specified in [[OPAQUE](#)].

[3.2](#). **OSPFv3 Transport Instance Information Encoding**

   Application specific information will be flooded in separate LSAs
   with separate function codes. Refer to section A.4.2.1 of [[OSPFV3](#)]
   for information on the LS Type encoding in OSPFv3.

4. Security Considerations

The security considerations for the Transport Instance will not be different for those for OSPFv2 [OSPFV2] and OSPFv3 [OSPFV3].

5. IANA Considerations

No IANA actions are required.

6. References

6.1. Normative References

   [MULTI-INST]
             Lindem, A., Mirtorabi, S., and A. Roy, "OSPF Multi-
             Instance Extensions", RFC 6549, March 2012.

   [OPAQUE]   Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The
             OSPF Opaque LSA Option", RFC 5250, July 2008.

   [OSPFV2]   Moy, J., "OSPF Version 2", RFC 2328, April 1998.

   [OSPFV3]   Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF
             for IPv6", RFC 5340, July 2008.

   [RFC-KEYWORDS]
             Bradner, S., "Key words for use in RFCs to Indicate
             Requirement Levels", RFC 2119, March 1997.

   [TE]   Katz, D., Yeung, D., and K. Kompella, "Traffic Engineering
             Extensions to OSPF", RFC 3630, September 2003.

6.2. Informative References

   [ISIS-GEN-APP]
             Ginsberg, L., Previdi, S., and M. Shand, "Advertising
             Generic Information in IS-IS", RFC 6823, December 2012.

Appendix A. Acknowledgments

The RFC text was produced using Marshall Rose's xml2rfc tool.

Although very different mechanisms are utilized, the concept of using a separate instance to advertise non-routing information in an IGP was first introduced in "Advertising Generic Information in IS-IS" [ISIS-GEN-APP].

Thanks to Jonathan Sadler for comments on the document.

Authors' Addresses

   Acee Lindem
   Ericsson
   301 Midenhall Way
   Cary, NC 27513
   USA

   Email: acee.lindem@ericsson.com

   Abhay Roy
   Cisco Systems
   225 West Tasman Drive
   San Jose, CA 95134
   USA

   Email: akr@cisco.com

   Sina Mirtorabi
   Cisco Systems
   3 West Plumeria Drive
   San Jose, CA 95134
   USA

   Email: sina@cisco.com