

Network Working Group
Request for Comments: 4973
Category: Experimental

P. Srisuresh
Kazeon Systems
P. Joseph
Consultant
July 2007

OSPF-xTE: Experimental Extension to OSPF for Traffic Engineering

Status of This Memo

This memo defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document defines OSPF-xTE, an experimental traffic engineering (TE) extension to the link-state routing protocol OSPF. OSPF-xTE defines new TE Link State Advertisements (LSAs) to disseminate TE metrics within an autonomous System (AS), which may consist of multiple areas. When an AS consists of TE and non-TE nodes, OSPF-xTE ensures that non-TE nodes in the AS are unaffected by the TE LSAs. OSPF-xTE generates a stand-alone TE Link State Database (TE-LSDB), distinct from the native OSPF LSDB, for computation of TE circuit paths. OSPF-xTE is versatile and extendible to non-packet networks such as Synchronous Optical Network (SONET) / Time Division Multiplexing (TDM) and optical networks.

IESG Note

The content of this RFC was at one time considered by the IETF, and therefore it may resemble a current IETF work in progress or a published IETF work. This RFC is not a candidate for any level of Internet Standard. The IETF disclaims any knowledge of the fitness of this RFC for any purpose and in particular notes that the decision to publish is not based on IETF review for such things as security, congestion control, or inappropriate interaction with deployed protocols. The RFC Editor has chosen to publish this document at its discretion. Readers of this RFC should exercise caution in evaluating its value for implementation and deployment. See RFC 3932 for more information.

See RFC 3630 for the IETF consensus protocol for OSPF Traffic Engineering. The OSPF WG position at the time of publication is that although this proposal has some useful properties, the protocol in RFC 3630 is sufficient for the traffic engineering needs that have been identified so far, and the cost of migrating to this proposal exceeds its benefits.

Table of Contents

1. Introduction	3
2. Principles of Traffic Engineering	3
3. Terminology	5
3.1. Native OSPF Terms	5
3.2. OSPF-xTE Terms	6
4. Motivations behind the Design of OSPF-xTE	9
4.1. Scalable Design	9
4.2. Operable in Mixed and Peer Networks	9
4.3. Efficient in Flooding Reach	9
4.4. Ability to Reserve TE-Exclusive Links	10
4.5. Extensible Design	11
4.6. Unified for Packet and Non-Packet Networks	11
4.7. Networks Benefiting from the OSPF-xTE Design	11
5. OSPF-xTE Solution Overview	12
5.1. OSPF-xTE Solution	12
5.2. Assumptions	13
6. Strategy for Transition of Opaque LSAs to OSPF-xTE	14
7. OSPF-xTE Router Adjacency -- TE Topology Discovery	14
7.1. The OSPF-xTE Router Adjacency	14
7.2. The Hello Protocol	15
7.3. The Designated Router	15
7.4. The Backup Designated Router	15
7.5. Flooding and the Synchronization of Databases	16
7.6. The Graph of Adjacencies	16
8. TE LSAs for Packet Network	18
8.1. TE-Router LSA (0x81)	18
8.1.1. Router-TE Flags: TE Capabilities of the Router	19
8.1.2. Router-TE TLVs	20
8.1.3. Link-TE Flags: TE Capabilities of a Link	22
8.1.4. Link-TE TLVs	23
8.2. TE-Incremental-Link-Update LSA (0x8d)	26
8.3. TE-Circuit-Path LSA (0x8C)	28
8.4. TE-Summary LSAs	31
8.4.1. TE-Summary Network LSA (0x83)	32
8.4.2. TE-Summary Router LSA (0x84)	33
8.5. TE-AS-external LSAs (0x85)	34
9. TE LSAs for Non-Packet Network	36
9.1. TE-Router LSA (0x81)	36
9.1.1. Router-TE flags - TE Capabilities of a Router	37

9.1.2. Link-TE Options: TE Capabilities of a TE Link	38
9.2. TE-positional-ring-network LSA (0x82)	38
9.3. TE-Router-Proxy LSA (0x8e)	40
10. Abstract Topology Representation with TE Support	42
11. Changes to Data Structures in OSPF-xTE Nodes	44
11.1. Changes to Router Data Structure	44
11.2. Two Sets of Neighbors	44
11.3. Changes to Interface Data Structure	44
12. IANA Considerations	45
12.1. TE LSA Type Values	45
12.2. TE TLV Tag Values	46
13. Acknowledgements	46
14. Security Considerations	47
15. Normative References	48
16. Informative References	48

1. Introduction

This document defines OSPF-xTE, an experimental traffic engineering (TE) extension to the link-state routing protocol OSPF. The objective of OSPF-xTE is to discover TE network topology and disseminate TE metrics within an autonomous system (AS). A stand-alone TE Link State Database (TE-LSDB), different from the native OSPF LSDB, is created to facilitate computation of TE circuit paths. Devising algorithms to compute TE circuit paths is not an objective of this document.

OSPF-xTE is different from the Opaque-LSA-based approach outlined in [OPQLSA-TE]. Section 4 describes the motivations behind the design of OSPF-xTE. Section 6 outlines a transition path for those currently using [OPQLSA-TE] for intra-area and wish to extend this using OSPF-xTE across the AS.

Readers interested in TE extensions for packet networks alone may skip section 9.0.

2. Principles of Traffic Engineering

The objective of traffic engineering (TE) is to set up circuit path(s) between a pair of nodes or links and to forward traffic of a certain forwarding equivalency class (FEC) through the circuit path. Only unicast circuit paths are considered in this section; multicast variations are outside the scope.

A traffic engineered circuit path is unidirectional and may be identified by the tuple: (FEC, TE circuit parameters, origin node/link, destination node/link).

A forwarding equivalency class (FEC) is a grouping of traffic that is forwarded in the same manner by a node. An FEC may be classified based on a number of criteria, as follows:

- a) traffic arriving on a specific interface,
- b) traffic arriving at a certain time of day,
- c) traffic meeting a certain packet based classification criteria (ex: based on a match of the fields in the IP and transport headers within a packet),
- d) traffic in a certain priority class,
- e) traffic arriving on a specific set of TDM (Synchronous Transport Signal (STS)) circuits on an interface, or
- f) traffic arriving on a certain wavelength of an interface.

Discerning traffic based on the FEC criteria is mandatory for Label Edge Routers (LERs). The intermediate Label-Switched Routers (LSRs) are transparent to the traffic content. LSRs are only responsible for maintaining the circuit for its lifetime. This document will not address definition of FEC criteria, the mapping of an FEC to circuit, or the associated signaling to set up circuits. [MPLS-TE] and [GMPLS-TE] address the FEC criteria. [RSVP-TE] and [CR-LDP] address signaling protocols to set up circuits.

This document is concerned with the collection of TE metrics for all the TE enforceable nodes and links within an autonomous system. TE metrics for a node may include the following.

- a) Ability to perform traffic prioritization,
- b) Ability to provision bandwidth on interfaces,
- c) Support for Constrained Shortest Path First (CSPF) algorithms,
- d) Support for certain TE-Circuit switch type, and
- e) Support for a certain type of automatic protection switching.

TE metrics for a link may include the following.

- a) available bandwidth,
- b) reliability of the link,
- c) color assigned to the link,
- d) cost of bandwidth usage on the link, and
- e) membership in a Shared Risk Link Group (SRLG).

A number of CSPF (Constraint-based Shortest Path First) algorithms may be used to dynamically set up TE circuit paths in a TE network.

OSPF-xTE mandates that the originating and the terminating entities of a TE circuit path be identifiable by IP addresses.

3. Terminology

Definitions of the majority of the terms used in the context of the OSPF protocol may be found in [OSPF-V2]. MPLS and traffic engineering terms may be found in [MPLS-ARCH]. RSVP-TE and CR-LDP signaling-specific terms may be found in [RSVP-TE] and [CR-LDP], respectively.

The following subsections describe the native OSPF terms and the OSPF-xTE terms used within this document.

3.1. Native OSPF Terms

- o Native node (Non-TE node)

A native or non-TE node is an OSPF router that is capable of IP packet forwarding but does not take part in a TE network. A native OSPF node forwards IP traffic using the shortest-path forwarding algorithm and does not run the OSPF-xTE extensions.

- o Native link (Non-TE link)

A native (or non-TE) link is a network attachment to a TE or non-TE node used for IP packet traversal.

- o Native OSPF network (Non-TE network)

A native OSPF network refers to an OSPF network that does not support TE. "Non-TE network", "native-OSPF network", and "non-TE topology" are used synonymously throughout the document.

- o LSP

LSP stands for "Label-Switched Path". An LSP is a TE circuit path in a packet network. The terms "LSP" and "TE circuit path" are used synonymously in the context of packet networks.

- o LSA

LSA stands for OSPF "Link State Advertisement".

- o LSDB

LSDB stands for "Link State Database". An LSDB contains a representation of the topology of a network. A native LSDB, constituted of native OSPF LSAs, represents the topology of a native IP network. The TE-LSDB, on the other hand, is constituted of TE LSAs and is a representation of the TE network topology.

3.2. OSPF-xTE Terms

- o TE node

A TE node is a node in the traffic engineering (TE) network. A TE node has a minimum of one TE link attached to it. Associated with each TE node is a set of supported TE metrics. A TE node may also participate in a native IP network.

In a SONET/TDM or photonic cross-connect network, a TE node is not required to be an OSPF-xTE node. An external OSPF-xTE node may act as proxy for the TE nodes that cannot be routers themselves.

- o TE link

A TE link is a network attachment point to a TE node and is intended for traffic engineering use. Associated with each TE link is a set of supported TE metrics. A TE link may also optionally carry native IP traffic.

Of the various links attached to a TE node, only the links that take part in a traffic-engineered network are called TE links.

- o TE circuit path

A TE circuit path is a unidirectional data path that is defined by a list of TE nodes connected to each other through TE links. A TE circuit path is also often referred simply as a circuit path or a circuit.

For the purposes of OSPF-xTE, the originating and terminating entities of a TE circuit path must be identifiable by their IP addresses. As a general rule, all nodes and links party to a traffic-engineered network should be uniquely identifiable by an IP address.

- o OSPF-xTE node (OSPF-xTE router)

An OSPF-xTE node is a TE node that runs the OSPF routing protocol and the OSPF-xTE extensions described in this document. An autonomous system (AS) may consist of a combination of native and OSPF-xTE nodes.

- o TE Control network

The IP network used by the OSPF-xTE nodes for OSPF-xTE communication is referred as the TE control network or simply the control network. The control network can be independent of the TE data network.

- o TE network (TE topology)

A TE network is a network of connected TE nodes and TE links, for the purpose of setting up one or more TE circuit paths. The terms "TE network", "TE data network", and "TE topology" are used synonymously throughout the document.

- o Packet-TE network (Packet network)

A packet-TE network is a TE network in which the nodes switch MPLS packets. An MPLS packet is defined in [MPLS-TE] as a packet with an MPLS header, followed by data octets. The intermediary node(s) of a circuit path in a packet-TE network perform MPLS label swapping to emulate the circuit.

Unless specified otherwise, the term "packet network" is used throughout the document to refer to a packet-TE network.

- o Non-packet-TE network (Non-packet network)

A non-packet-TE network is a TE network in which the nodes switch non-packet entities such as STS time slots, Lambda wavelengths, or simply interfaces.

SONET/TDM and fiber cross-connect networks are examples of non-packet-TE networks. Circuit emulation in these networks is accomplished by the switch fabric in the intermediary nodes (based on TDM time slot, fiber interface, or Lambda).

Unless specified otherwise, the term non-packet network is used throughout the document to refer a non-packet-TE network.

- o Mixed network

A mixed network is a network that is constituted of both packet-TE and non-TE networks. Traffic in the network is strictly datagram oriented, i.e., IP datagrams or MPLS packets. Routers in a mixed network may be TE or native nodes.

OSPF-xTE is usable within a packet network or a mixed network.

- o Peer network

A peer network is a network that is constituted of packet-TE and non-packet-TE networks combined. In a peer network, a TE node could potentially support TE links for the packet as well as non-packet data.


OSPF-xTE is usable within a packet network or a non-packet network or a peer network, which is a combination of the two.

- CSPF

CSPF stands for "Constrained Shortest Path First". Given a TE LSDB and a set of constraints that must be satisfied to form a circuit path, there may be several CSPF algorithms to obtain a TE circuit path that meets the criteria.

- TLV

A TLV stands for a data object in the form: Tag-Length-Value. All TLVs are assumed to be of the following format, unless specified otherwise. The Tag and Length are 16 bits wide each. The Length includes the 4 octets required for Tag and Length specification. All TLVs described in this document are padded to 32-bit alignment. Any padding required for alignment will not be a part of the length field, however. TLVs are used to describe traffic engineering characteristics of the TE nodes, TE links, and TE circuit paths.



- o Router-TE TLVs (Router TLVs)

TLVs used to describe the TE capabilities of a TE node.

- o Link-TE TLVs (Link TLVs)

TLVs used to describe the TE capabilities of a TE link.

4. Motivations behind the Design of OSPF-xTE

There are several motivations that led to the design of OSPF-xTE. OSPF-xTE is scalable, efficient, and usable across a variety of network topologies. These motivations are explained in detail in the following subsections. The last subsection lists real-world network scenarios that benefit from the OSPF-xTE.

4.1. Scalable Design

In OSPF-xTE, an area-level abstraction provides the scaling required for the TE topology in a large autonomous system (AS). An OSPF-xTE area border router will advertise summary LSAs for TE and non-TE topologies independent of each other. Readers may refer to section 10 for a topological view of the AS from the perspective of a OSPF-xTE node in an area.

[OPQLSA-TE], on the other hand, is designed for intra-area and is not scalable to AS-wide scope.

4.2. Operable in Mixed and Peer Networks

OSPF-xTE assumes that an AS may be constituted of coexisting TE and non-TE networks. OSPF-xTE dynamically discovers TE topology and the associated TE metrics of the nodes and links that form the TE network. As such, OSPF-xTE generates a stand-alone TE-LSDB that is fully representative of the TE network. Stand-alone TE-LSDB allows for speedy TE computations.

[OPQLSA-TE] is designed for packet networks and is not suitable for mixes and peer networks. TE-LSDB in [OPQLSA-TE] is derived from the combination of Opaque LSAs and native LSDB. Further, the TE-LSDB thus derived has no knowledge of the TE capabilities of the routers in the network.

4.3. Efficient in Flooding Reach

OSPF-xTE is able to identify the TE topology in a mixed network and to limit the flooding of TE LSAs to only the TE nodes. Non-TE nodes are not bombarded with TE LSAs.

In a TE network, a subset of the TE metrics may be prone to rapid change, while others remain largely unchanged. Changes in TE metrics must be communicated at the earliest throughout the network to ensure that the TE-LSDB is up-to-date within the network. As a general rule, a TE network is likely to generate significantly more control traffic than a native network. The excess traffic is almost directly proportional to the rate at which TE circuits are set up and torn down within the TE network. The TE database synchronization should occur much quicker compared to the aggregate circuit set up and tear-down rates. OSPF-xTE defines TE-Incremental-Link-update LSA (section 8.2) to advertise only a subset of the metrics that are prone to rapid changes.

The more frequent and wider the flooding, the larger the number of retransmissions and acknowledgements. The same information (needed or not) may reach a router through multiple links. Even if the router did not forward the information past the node, it would still have to send acknowledgements across all the various links on which the LSAs tried to converge. It is undesirable to flood non-TE nodes with TE information.

4.4. Ability to Reserve TE-Exclusive Links

OSPF-xTE draws a clear distinction between TE and non-TE links. A TE link may be configured to permit TE traffic alone, and not permit best-effort IP traffic on the link. This permits TE enforceability on the TE links.

When links of a TE topology do not overlap the links of a native IP network, OSPF-xTE allows for virtual isolation of the two networks. Best-effort IP network and TE network often have different service requirements. Keeping the two networks physically isolated can be expensive. Combining the two networks into a single physically connected network will bring economies of scale, while service enforceability can be maintained individually for each of the TE and non-TE sections of the network.

[OPQLSA-TE] does not support the ability to isolate best-effort IP traffic from TE traffic on a link. All links are subject to best-effort IP traffic. An OSPF router could potentially select a TE link to be its least cost link and inundate the link with best-effort IP traffic, thereby rendering the link unusable for TE purposes.

4.5. Extensible Design

The OSPF-xTE design is based on the tried-and-tested OSPF paradigm, and it inherits all the benefits of OSPF, present and future. TE LSAs are extensible, just as the native OSPF on which OSPF-xTE is founded are extensible.

4.6. Unified for Packet and Non-Packet Networks

OSPF-xTE is usable within a packet network or a non-packet network or a combination peer network.

Signaling protocols such as RSVP and LDP work the same across packet and non-packet networks. Signaling protocols merely need the TE characteristics of nodes and links so they can signal the nodes to formulate TE circuit paths. In a peer network, the underlying control protocol must be capable of providing a unified LSDB for all TE nodes (nodes with packet-TE links as well as non-packet-TE links) in the network. OSPF-xTE meets this requirement.

4.7. Networks Benefiting from the OSPF-xTE Design

Below are examples of some real-world network scenarios that benefit from OSPF-xTE.

- o IP providers transitioning to provide TE services

Providers needing to support MPLS-based TE in their IP network may choose to transition gradually. They may add new TE links or convert existing links into TE links within an area first and progressively advance to offering MPLS in the entire AS.

Not all routers will support TE extensions at the same time during the migration process. Use of TE-specific LSAs and their flooding to OSPF-xTE only nodes will allow the vendor to introduce MPLS TE without destabilizing the existing network. The native OSPF-LSDB will remain undisturbed while newer TE links are added to the network.

- o Providers offering best-effort-IP & TE services

Providers choosing to offer both best-effort-IP and TE based packet services simultaneously on the same physically connected network will benefit from the OSPF-xTE design. By maintaining independent LSDBs for each type of service, TE links are not cannibalized in a mixed network.

- o Large TE networks

The OSPF-xTE design is advantageous in large TE networks that require the AS to be sub-divided into multiple areas. OSPF-xTE permits inter-area exchange of TE information, which ensures that all nodes in the AS have up-to-date, AS-wide, TE reachability knowledge. This in turn will make TE circuit setup predictable and computationally bounded.

- o Non-Packet Networks and Peer Networks

Vendors may also use OSPF-xTE for their non-packet TE networks. OSPF-xTE defines the following functions in support of non-packet TE networks.

- (a) "Positional-Ring" type network LSAs.
- (b) Router proxying -- allowing a router to advertise on behalf of other nodes (that are not packet/OSPF-capable).

5. OSPF-xTE Solution Overview

5.1. OSPF-xTE Solution

Locally-scoped Opaque LSA (type 9) is used to discover the TE topology within a network. Section 7.1 describes in detail the use of type 9 Opaque LSA for TE topology discovery. TE LSAs are designed for use by the OSPF-xTE nodes. Section 8.0 describes the TE LSAs in detail. Changes required of the OSPF data structures to support OSPF-xTE are described in section 11.0. A new TE-neighbors data structure will be used to advertise TE LSAs along TE topology.

An OSPF-xTE node will have a native LSDB and a TE-LSDB, while a native OSPF node will have just a native LSDB. Consider the OSPF area, constituted of OSPF-xTE and native OSPF routers, shown in Figure 1. Nodes RT1, RT2, RT3, and RT6 are OSPF-xTE routers with TE and non-TE link attachments. Nodes RT4 and RT5 are native OSPF routers with no TE links. When the LSA database is synchronized, all nodes will share the same native LSDB. OSPF-xTE nodes alone will have the additional TE-LSDB.




Figure 1. A (TE + native) OSPF Network Topology

5.2. Assumptions

OSPF-xTE is an extension to the native OSPF protocol and does not mandate changes to the existing OSPF. OSPF-xTE design makes the following assumptions.

- (1) An OSPF-xTE node will need to establish router adjacency with at least one other OSPF-xTE node in the area in order for the router's TE database to be synchronized within the area. Failing this, the OSPF router will not be in the TE calculations of other TE routers in the area.

It is the responsibility of the network administrator(s) to ensure connectedness of the TE network. Otherwise, there can be disjoint TE topologies within a network.

- (2) OSPF-xTE nodes must advertise the link state of its TE links. TE links are not obligated to support native IP traffic. Hence, an OSPF-xTE node cannot be required to synchronize its link-state database with neighbors on all its links. The only requirement is to have the TE LSDB synchronized across all OSPF-xTE nodes in the area.
- (3) A link in a packet network may be designated as a TE link or a native-IP link or both. For example, a link may be used for both TE and non-TE traffic, as long as the link is under subscribed in bandwidth for TE traffic (for example, 50% of the link capacity is set aside for TE traffic).
- (4) Non-packet TE sub-topologies must have a minimum of one node running OSPF-xTE protocol. For example, a SONET/SDH TDM ring must have a minimum of one Gateway Network Element (GNE) running OSPF-xTE. The OSPF-xTE node will advertise on behalf of all the TE nodes in the ring.

6. Strategy for Transition of Opaque LSAs to OSPF-xTE

Below is a strategy to transition implementations currently using Opaque LSAs ([OPQLSA-TE]) within an area to adapt OSPF-xTE in a gradual fashion across the AS.

- (1) Use [OPQLSA-TE] within an area. Derive TE topology within the area from the combination of Opaque LSAs and native LSDB.
- (2) Use TE-Summary LSAs and TE-AS-external LSAs for inter-area communication. Use the TE topology within an area to summarize the TE networks in the area and advertise the same to all TE nodes in the backbone. The TE-ABRs (TE area border routers) on the backbone area will in turn advertise these summaries within their connected areas.

7. OSPF-xTE Router Adjacency -- TE Topology Discovery

OSPF creates adjacencies between neighboring routers for the purpose of exchanging routing information. The following subsections describe the use of locally-scoped Opaque LSAs to discover OSPF-xTE neighboring routers. The capability is used as the basis to build a TE topology.

7.1. The OSPF-xTE Router Adjacency

OSPF uses the options field in the Hello packet to advertise optional router capabilities [OSPF-V2]. However, all the bits in this field have been allocated and there is no way to advertise OSPF-xTE

capability using the options field at this time. This document proposes using local-scope Opaque LSA (OPAQUE-9 LSA) to advertise support for OSPF-xTE and establish OSPF-xTE adjacency. In order to exchange Opaque LSAs, the neighboring routers must have the O-bit (Opaque option bit) set in the options field.

[OSPF-CAP] proposes a format for exchanging router capabilities via OPAQUE-9 LSA. Routers supporting OSPF-xTE will be required to set the "OSPF Experimental TE" bit within the "router capabilities" field. Two routers will not become TE neighbors unless they share a common network link on which both routers advertise support for OSPF-xTE. Routers that do not support OSPF-xTE may simply ignore the advertisement.

7.2. The Hello Protocol

The Hello protocol is primarily responsible for dynamically establishing and maintaining neighbor adjacencies. In a TE network, it is not required for all links and neighbors to establish adjacency using this protocol. OSPF-xTE router adjacency between two routers is established using the method described in the previous section.

For non-broadcast multi-access (NBMA) and broadcast networks, the HELLO protocol is responsible for electing the Designated Router and the Backup Designated Router. Routers supporting the TE option shall be given a higher precedence for becoming a designated router over those that do not support TE.

7.3. The Designated Router

When a router's non-TE link first becomes functional, it checks to see whether there is currently a Designated Router for the network. If there is one, it accepts that Designated Router, regardless of its router priority, so long as the current designated router is TE compliant. Otherwise, the router itself becomes Designated Router if it has the highest Router Priority on the network and is TE compliant.

OSPF-xTE must be implemented on the most robust routers, as they become likely candidates to take on the role as Designated Router.

7.4. The Backup Designated Router

The Backup Designated Router is also elected by the Hello Protocol. Each Hello Packet has a field that specifies the Backup Designated Router for the network. Once again, TE-compliance must be weighed in conjunction with router priority in electing the Backup Designated Router.

7.5. Flooding and the Synchronization of Databases

In OSPF, adjacent routers within an area are required to synchronize their databases. However, a more concise requirement is that all routers in an area must converge on the same LSDB. As stated in item 2 of section 5.2, a basic assertion of OSPF-xTE is that the links used by the OSPF-xTE control network for flooding must not be required to match the links used by the data network for real-time data forwarding. For instance, it should not be required to send OSPF-xTE messages over a TE link that is configured to reject non-TE traffic. However, the control network must be set up such that a minimum of one path exists between any two OSPF or OSPF-xTE routers within the network, for flooding purposes. This revised control network connectivity requirement does not jeopardize convergence of LSDB within an area.

In a mixed network, where some of the neighbors are TE compliant and others are not, the designated OSPF-xTE router will exchange different sets of LSAs with its neighbors. TE LSAs are exchanged only with the TE neighbors. Native LSAs are exchanged with all neighbors (TE and non-TE alike). Restricting the scope of TE LSA flooding to just the OSPF-xTE nodes will not affect the native nodes that coexist with the OSPF-xTE nodes.

The control traffic for a TE network (i.e., TE LSA advertisement) is likely to be higher than that of a native OSPF network. This is because the TE metrics may vary with each TE circuit setup and the corresponding state change must be advertised at the earliest, not exceeding the MinLSInterval of 5 seconds. To minimize advertising repetitive content, OSPF-xTE defines a new TE-incremental-Link-update LSA (section 8.2) that would advertise just the TLVs that changed for a link.

The OSPFIGP-TE well-known multicast address 224.0.0.24 has been assigned by IANA for the exchange of TE-compliant database descriptors during database synchronization.


7.6. The Graph of Adjacencies

If two routers have multiple networks in common, they may have multiple adjacencies between them. The adjacency may be one of two types - native OSPF adjacency and TE adjacency. OSPF-xTE routers will form both types of adjacency.

Two types of adjacency graphs are possible, depending on whether a Designated Router is elected for the network. On physical point-to-point networks, point-to-multipoint networks, and virtual links, neighboring routers become adjacent whenever they can communicate

directly. The adjacency can be either (a) TE-compliant or (b) native. In contrast, on broadcast and NBMA networks the designated router and the backup designated router may maintain two sets of adjacency. The remaining routers will form either TE-compliant or native adjacency.

In the broadcast network in Figure 2, routers RT7 and RT3 are chosen as the Designated and Backup Designated Routers, respectively. Routers RT3, RT4 and RT7 are TE-compliant, but RT5 and RT6 are not. So RT4 will have TE-compliant adjacency with the designated and backup routers, while RT5 and RT6 will only have native adjacency with the Designated and Backup Designated Routers.



Adjacency Legend:

```

----- Native adjacency (primary)
***** Native adjacency (backup)
::::: TE-compliant adjacency (primary)
;;;;; TE-compliant adjacency (backup)

```

Figure 2. Two Adjacency Graphs with TE-Compliant Routers

8. TE LSAs for Packet Network

The OSPFv2 protocol currently has a total of 11 LSA types. LSA types 1 through 5 are defined in [OSPF-V2]. LSA types 6, 7, and 8 are defined in [MOSPF], [NSSA], and [BGP-OSPF], respectively. LSA types 9 through 11 are defined in [OPAQUE].

Each LSA type has a unique flooding scope. Opaque LSA types 9 through 11 are general purpose LSAs, with flooding scope set to link-local, area-local, and AS-wide (except stub areas) respectively.

In the following subsections, we define new LSAs for traffic engineering (TE) use. The values for the new TE LSA types are assigned with the high bit of the LSA-type octet set to 1. The new TE LSAs are largely modeled after the existing LSAs for content format and have a unique flooding scope.


TE-router LSA is defined to advertise TE characteristics of an OSPF-xTE router and all the TE links attached to the router. TE-incremental-Link-Update LSA is defined to advertise incremental updates to the metrics of a TE link. Flooding scope for both these LSAs is restricted to an area.

TE-Summary network and router LSAs are defined to advertise the reachability of area-specific TE networks and area border routers (along with router TE characteristics) to external areas. Flooding scope of the TE-Summary LSAs is the TE topology in the entire AS less the non-backbone area for which the advertising router is an ABR. Just as with native OSPF summary LSAs, the TE-Summary LSAs do not reveal the topological details of an area to external areas.

TE-AS-external LSA and TE-Circuit-Path LSA are defined to advertise AS external network reachability and pre-engineered TE circuits, respectively. While flooding scope for both these LSAs can be the entire AS, flooding scope for the pre-engineered TE circuit LSA may optionally be restricted to just the TE topology within an area.


8.1. TE-Router LSA (0x81)

The TE-router LSA (0x81) is modeled after the router LSA and has the same flooding scope as the router LSA. However, the scope is restricted to only the OSPF-xTE nodes within the area. The TE router LSA describes the TE metrics of the router as well as the TE links attached to the router. Below is the format of the TE-router LSA. Unless specified explicitly otherwise, the fields carry the same meaning as they do in a router LSA. Only the differences are explained below. Router-TE flags, Router-TE TLVs, Link-TE options, and Link-TE TLVs are each described in the following sub-sections.



8.1.1. Router-TE Flags: TE Capabilities of the Router

The following flags are used to describe the TE capabilities of an OSPF-xTE router. The remaining bits of the 32-bit word are reserved for future use.



Bit LSR - When set, the router is considered to have LSR (Label-Switched Router) capability.

Bit LER - When set, the router is considered to have LER capability. All MPLS border routers will be required to have LER capability. Setting both the LER and E bits indicates an AS Boundary router with LER capability. Setting both the LER and B bits indicates an area border router with LER capability.

Bit PSC - Indicates the node is packet-switch capable.

Bit LSP - An MPLS Label switch TLV TE-NODE-TLV-MPLS-SWITCHING follows. This is applicable only when the PSC flag is set.

Bit SIG - An MPLS Signaling-protocol-support TLV TE-NODE-TLV-MPLS-SIG-PROTOCOLS follows.

BIT CSPF - A CSPF algorithm support TLV TE-NODE-TLV-CSPF-ALG follows.

8.1.2. Router-TE TLVs

The following Router-TE TLVs are defined.

8.1.2.1. TE-NODE-TLV-MPLS-SWITCHING

MPLS switching TLV is applicable only for packet switched nodes. The TLV specifies the MPLS packet switching capabilities of the TE node.

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								
Tag = 0x8001										Length = 6																													
Label Depth										QOS																													

Label Depth is the depth of label stack the node is capable of processing on its ingress interfaces. An octet is used to represent label depth. A default value of 1 is assumed when the TLV is not listed. Label depth is relevant when an LER has to pop multiple labels off the MPLS stack.

QOS is a single-octet field that may be assigned '1' or '0'. Nodes supporting QOS are able to interpret the EXP bits in the MPLS header to prioritize multiple classes of traffic through the same LSP.

8.1.2.2. TE-NODE-TLV-MPLS-SIG-PROTOCOLS

MPLS signaling protocols TLV lists all the signaling protocol supported by the node. An octet is used to list each signaling protocol supported.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Tag = 0x8002                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Protocol-1   |   ...   |   ....   |                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

RSVP-TE protocol is represented as 1, CR-LDP as 2, and LDP as 3. These are the only permitted signaling protocols at this time.

8.1.2.3. TE-NODE-TLV-CSPF-ALGORITHMS

The CSPF algorithms TLV lists all the CSPF algorithm codes supported. Support for CSPF algorithms makes the node eligible to compute complete or partial circuit paths. Support for CSPF algorithms can also be beneficial in knowing whether or not a node is capable of expanding loose routes (in an MPLS signaling request) into a detailed circuit path.

Two octets are used to list each CSPF algorithm code. The algorithm codes may be vendor defined and unique within an Autonomous System. If the node supports 'n' CSPF algorithms, the Length would be $(4 + 4 * ((n+1)/2))$ octets.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Tag = 0x8003                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               CSPF-1                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               CSPF-n                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

8.1.2.4. TE-NODE-TLV-NULL

When a TE-Router or a TE link has multiple TLVs to describe the metrics, the NULL TLV is used to terminate the TLV list.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Tag = 0x8888                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

8.1.3. Link-TE Flags: TE Capabilities of a Link

The following flags are used to describe the TE capabilities of a link. The remaining bits of the 32-bit word are reserved for future use.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| T | N | P |   |   |   | D |                               | S | L | B | C |
| E | T | K |   |   |   | B |                               | R | U | W | O |
|   | E | T |   |   |   | S |                               | L | G |   | L |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| <---- Boolean TE flags -----> | <- TE flags pointing to TLVs -> |

```

- Bit TE - Indicates whether TE is permitted on the link. A link can be denied for TE use by setting the flag to 0.
- Bit NTE - Indicates whether non-TE traffic is permitted on the TE link. This flag is relevant only when the TE flag is set.
- Bit PKT - Indicates whether or not the link is capable of IP packet processing.
- Bit DBS - Indicates whether or not database synchronization is permitted on this link.
- Bit SRLG - Shared Risk Link Group TLV TE-LINK-TLV-SRLG follows.
- Bit LUG - Link Usage Cost Metric TLV TE-LINK-TLV-LUG follows.
- Bit BW - One or more Link Bandwidth TLVs follow.
- Bit COL - Link Color TLV TE-LINK-TLV-COLOR follows.

8.1.4. Link-TE TLVs

8.1.4.1. TE-LINK-TLV-SRLG

The SRLG describes the list of Shared Risk Link Groups (SRLG) the link belongs to. Two octets are used to list each SRLG. If the link belongs to 'n' SRLGs, the Length would be $(4 + 4 * ((n+1)/2))$ octets.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Tag = 0x0001                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               SRLG-1                                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               SRLG-n                                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

8.1.4.2 TE-LINK-TLV-BANDWIDTH-MAX

The Bandwidth TLV specifies the maximum bandwidth of the link, as follows.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Tag = 0x0002                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Maximum Bandwidth                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Bandwidth is expressed in units of 32 bytes/sec (256 bits/sec). A 32-bit field for bandwidth would permit specification not exceeding 1 terabit/sec.

Maximum Bandwidth is the maximum link capacity expressed in bandwidth units. Portions or all of this bandwidth may be used for TE use.

8.1.4.3. TE-LINK-TLV-BANDWIDTH-MAX-FOR-TE

The Bandwidth TLV specifies the maximum bandwidth available for TE use, as follows.

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	2	3	4	5	6	7	8	9	0	1																								
+																+																+																+																	
																Tag = 0x0003																																Length = 8																	
+																+																+																+																	
																Maximum Bandwidth available for TE use																																																	
+																+																+																+																	

Bandwidth is expressed in units of 32 bytes/sec (256 bits/sec). A 32-bit field for bandwidth would permit specification not exceeding 1 terabit/sec.

"Maximum Bandwidth available for TE use" is the total reservable bandwidth on the link for use by all the TE circuit paths traversing the link. The link is oversubscribed when this field is more than the Maximum Bandwidth. When the field is less than the Maximum Bandwidth, the remaining bandwidth on the link may be used for non-TE traffic in a mixed network.

8.1.4.4. TE-LINK-TLV-BANDWIDTH-TE

The Bandwidth TLV specifies the bandwidth reserved for TE as follows.


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										
+										+										+										+											
										Tag = 0x0004																				Length = 8											
+										+										+										+											
										TE Bandwidth subscribed																															
+										+										+										+											

Bandwidth is expressed in units of 32 bytes/sec (256 bits/sec). A 32-bit field for bandwidth would permit specification not exceeding 1 terabit/sec.

"TE Bandwidth subscribed" is the bandwidth that is currently subscribed from of the link. "TE Bandwidth subscribed" must be less than the "Maximum bandwidth available for TE use". New TE circuit paths are able to claim no more than the difference between the two bandwidths for reservation.

8.1.4.5. TE-LINK-TLV-LUG

The link usage cost TLV specifies bandwidth unit usage cost, TE circuit set-up cost, and any time constraints for setup and teardown of TE circuits on the link.



Circuit Setup time constraint

This 64-bit number specifies the time at or after which a TE-circuit path may be set up on the link. The set-up time constraint is specified as the number of seconds from the start of January 1, 1970 UTC. A reserved value of 0 implies no circuit setup time constraint.

Circuit Teardown time constraint

This 64-bit number specifies the time at or before which all TE-circuit paths using the link must be torn down. The teardown time constraint is specified as the number of seconds from the start of January 1 1970 UTC. A reserved value of 0 implies no circuit teardown time constraint.

8.1.4.6. TE-LINK-TLV-COLOR

The color TLV is similar to the SRLG TLV, in that an Autonomous System may choose to issue colors to a TE link meeting certain criteria. The color TLV can be used to specify one or more colors assigned to the link as follows. Two octets are used to list each color. If the link belongs to 'n' number of colors, the Length would be $(4 + 4 * ((n+1)/2))$ octets.

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								
Tag = 0x0006										Length = $4(1 + (n+1)/2)$																													
Color-1																																						
Color-n																																							

8.1.4.7. TE-LINK-TLV-NULL

When a TE link has multiple TLVs to describe its metrics, the NULL TLV is used to terminate the TLV list. The TE-LINK-TLV-NULL is same as the TE-NODE-TLV-NULL described in section 8.1.2.4

8.2. TE-Incremental-Link-Update LSA (0x8d)

A significant difference between a native OSPF network and a TE network is that the latter may be subject to frequent real-time circuit pinning and is likely to undergo TE-state updates. Some links might undergo changes more frequently than others. Flooding the network with TE-router LSAs at the aggregated speed of all link metric changes is simply not desirable. A smaller in size TE-incremental-link-update LSA is designed to advertise only the incremental link updates.

A TE-incremental-link-update LSA will be advertised as frequently as the link state is changed (not exceeding once every MinLSInterval seconds). The TE link sequence is largely the advertisement of a sub-portion of router LSA. The sequence number on this will be incremented with the TE-router LSA's sequence as the basis. When an updated TE-router LSA is advertised within 30 minutes of the previous advertisement, the updated TE-router LSA will assume a sequence number that is larger than the most frequently updated of its links.

Below is the format of the TE-incremental-link-update LSA.

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																																						
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										LS age																				Options																				0x8d																			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										Link State ID (same as Link ID)																																																											
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										Advertising Router																																																											
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										LS sequence number																																																											
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										LS checksum																				length																																							
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										Link Data																																																											
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										Type																				0																				Link-TE options																			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										Link-TE options																				Zero or more Link-TE TLVs																																							
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										# TOS																				metric																																							
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										...																																																											
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										TOS																				0																				TOS metric																			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							

Link State ID

This would be exactly the same as would have been specified for Link ID, for a link within the router LSA.

Link Data

This specifies the router ID the link belongs to. In majority of cases, this would be same as the advertising router. This choice for Link Data is primarily to facilitate proxy advertisement for incremental link updates.

Suppose that a proxy router LSA was used to advertise the TE-router LSA of a SONET/TDM node, and that the proxy router is now required to advertise incremental-link-update for the same SONET/TDM node. Specifying the actual router-ID to which the link in the incremental-link-update LSA belongs helps receiving nodes in finding the exact match for the LSA in their database.

The tuple of (LS Type, LSA ID, Advertising router) uniquely identifies the LSA and replaces LSAs of the same tuple with an older sequence number. However, there is an exception to this rule in the context of TE-link-update LSA. TE-Link-update LSA will initially assume the sequence number of the TE-router LSA it belongs to. Further, when a new TE-router LSA update with a larger sequence number is advertised, the newer sequence number is assumed by all the link LSAs.

8.3. TE-Circuit-Path LSA (0x8C)

TE-Circuit-path LSA (next page) may be used to advertise the availability of pre-engineered TE circuit path(s) originating from any router in the network. The flooding scope may be area-wide or AS-wide. Fields are as follows.


Link State ID

The ID of the far-end router or the far-end link-ID to which the TE circuit path(s) is being advertised.

TE-circuit-path(s) flags

- Bit G - When set, the flooding scope is set to be AS-wide.
Otherwise, the flooding scope is set to be area-wide.
- Bit E - When set, the advertised Link-State ID is an AS boundary router (E is for external). The advertising router and the Link State ID belong to the same area.
- Bit B - When set, the advertised Link State ID is an area border router (B is for Border)
- Bit D - When set, this indicates that the duration of circuit path validity follows.
- Bit S - When set, this indicates that setup time of the circuit path follows.
- Bit T - When set, this indicates that teardown time of the circuit path follows.
- CktType - This 4-bit field specifies the circuit type of the Forward Equivalency Class (FEC).

0x01 - Origin is Router, Destination is Router.
 0x02 - Origin is Link, Destination is Link.
 0x04 - Origin is Router, Destination is Link.
 0x08 - Origin is Link, Destination is Router.



Circuit Duration (Optional)

This 64-bit number specifies the seconds from the time of the LSA advertisement for which the pre-engineered circuit path will be valid. This field is specified only when the D-bit is set in the TE-circuit-path flags.

Circuit Setup time (Optional)

This 64-bit number specifies the time at which the TE circuit path may be set up. This field is specified only when the S-bit is set in the TE-circuit-path flags. The set-up time is specified as the number of seconds from the start of January 1, 1970 UTC.

Circuit Teardown time (Optional)

This 64-bit number specifies the time at which the TE circuit path may be torn down. This field is specified only when the T-bit is set in the TE-circuit-path flags. The teardown time is specified as the number of seconds from the start of January 1 1970 UTC.

No. of TE Circuit Paths

This specifies the number of pre-engineered TE circuit paths between the advertising router and the router specified in the Link State ID.

Circuit-TE ID

This is the ID of the far-end router for a given TE circuit path segment.

Circuit-TE Data

This is the virtual link identifier on the near-end router for a given TE circuit path segment. This can be a private interface or handle the near-end router uses to identify the virtual link.

The sequence of (Circuit-TE ID, Circuit-TE Data) pairs lists the end-point nodes and links in the LSA as a series.

Circuit-TE flags

This lists the zero or more TE-link TLVs that all member elements of the LSP meet.

8.4. TE-Summary LSAs

TE-Summary LSAs are Type 0x83 and 0x84 LSAs. These LSAs are originated by area border routers. A TE-Summary-network LSA (0x83) describes the reachability of TE networks in a non-backbone area, advertised by the area border router. A Type 0x84 summary LSA describes the reachability of area border routers and AS border routers and their TE capabilities.

One of the benefits of having multiple areas within an AS is that frequent TE advertisements within the area do not impact outside the area. Only the TE abstractions befitting the external areas are advertised.

8.4.1. TE-Summary Network LSA (0x83)

A TE-Summary network LSA may be used to advertise reachability of TE-networks accessible to areas external to the originating area. The content and the flooding scope of a TE-Summary LSA is different from that of a native Summary LSA.

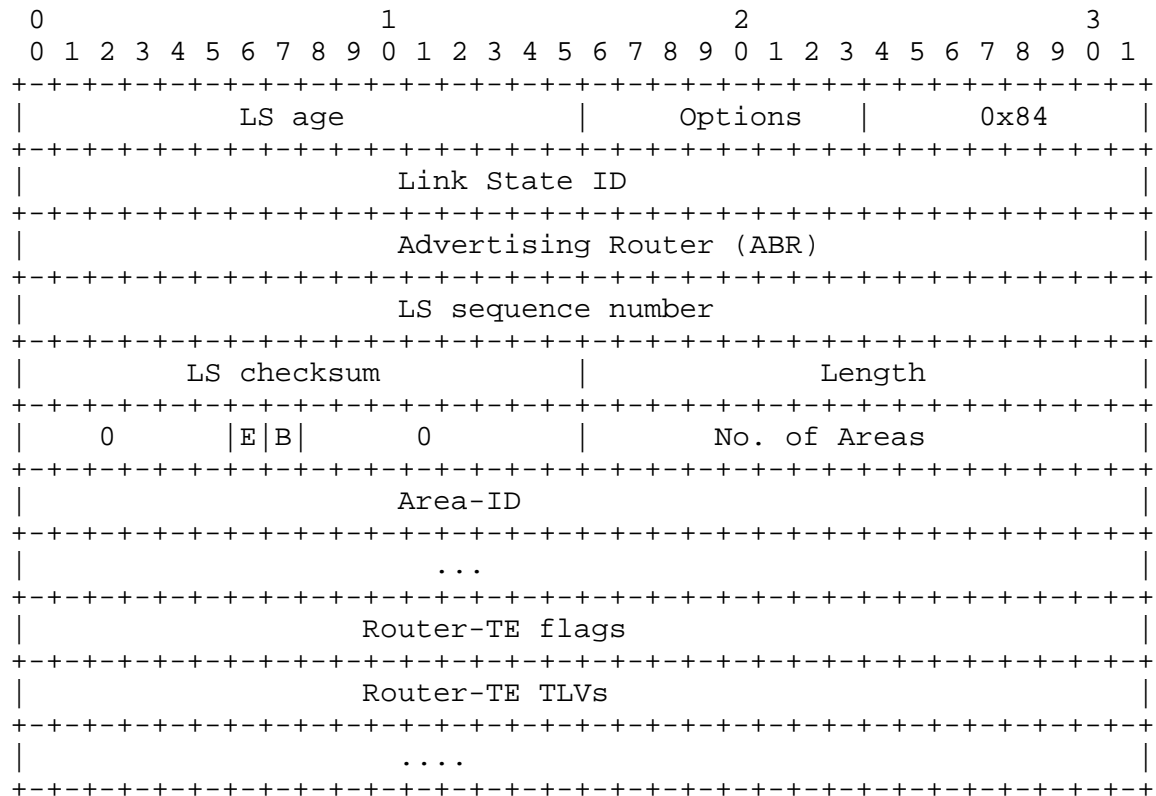
The scope of flooding for a TE-Summary network LSA is AS-wide, with the exception of the originating area and the stub areas. The area border router for each non-backbone area is responsible for advertising the reachability of backbone networks into the area.

Unlike a native-summary network LSA, a TE-Summary network LSA does not advertise summary costs to reach networks within an area. This is because TE parameters are not necessarily additive or comparable. The parameters can be varied in their expression. For example, a TE-Summary network LSA will not summarize a network whose links do not fall under an SRLG (Shared-Risk Link Group). This way, the TE-Summary LSA merely advertises the reachability of TE networks within an area. The specific circuit paths can be computed by the ABR. Pre-engineered circuit paths are advertised using TE-Circuit-path LSAs (refer to Section 8.3).

[illegible]

8.4.2. TE-Summary Router LSA (0x84)

A TE-Summary router LSA may be used to advertise the availability of area border routers (ABRs) and AS border routers (ASBRs) that are TE-capable. The TE-Summary router LSAs are originated by the Area Border Routers. The scope of flooding for the TE-Summary router LSA is the non-backbone area the advertising ABR belongs to.



Link State ID

The ID of the area border router or the AS border router whose TE capability is being advertised.

Advertising Router

The ABR that advertises its TE capabilities (and the OSPF areas it belongs to) or the TE capabilities of an ASBR within one of the areas for which the ABR is a border router.

No. of Areas

Specifies the number of OSPF areas the link state ID belongs to.

Area-ID

Specifies the OSPF area(s) the link state ID belongs to. When the link state ID is same as the advertising router ID, the Area-ID lists all the areas the ABR belongs to. In the case the link state ID is an ASBR, the Area-ID simply lists the area the ASBR belongs to. The advertising router is assumed to be the ABR from the same area the ASBR is located in.

Summary-router-TE flags

Bit E - When set, the advertised Link-State ID is an AS boundary router (E is for external). The advertising router and the Link State ID belong to the same area.

Bit B - When set, the advertised Link state ID is an Area border router (B is for Border)

Router-TE flags, Router-TE TLVs

TE capabilities of the link-state-ID router.

TE Flags and TE TLVs are as applicable to the ABR/ASBR specified in the link state ID. The semantics is same as specified in the Router-TE LSA.

8.5. TE-AS-external LSAs (0x85)

TE-AS-external LSAs are the Type 0x85 LSAs. This is modeled after AS-external LSA format and flooding scope. TE-AS-external LSAs are originated by AS boundary routers with TE extensions, and describe the TE networks and pre-engineered circuit paths external to the AS. As with AS-external LSA, the flooding scope of the TE-AS-external LSA is AS-wide, with the exception of stub areas.

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																																						
-----										-----										-----										-----																																							
										LS age																				Options																				0x85																			
										Link State ID																																																											
										Advertising Router																																																											
										LS sequence number																																																											
										LS checksum																				length																																							
										Network Mask																																																											
										Forwarding address																																																											
										External Route Tag																																																											
										# of Virtual TE links																				0																																							
										Link-TE flags																																																											
										Link-TE TLVs																																																											
										...																																																											
										TE-Forwarding address																																																											
										External Route TE Tag																																																											
										...																																																											

Network Mask

The IP address mask for the advertised TE destination. For example, this can be used to specify access to a specific TE node or TE link with an mask of 0xffffffff. This can also be used to specify access to an aggregated set of destinations using a different mask. ex: 0xff000000.

Link-TE flags, Link-TE TLVs

The TE attributes of this route. These fields are optional and are provided only when one or more pre-engineered circuits can be specified with the advertisement. Without these fields, the LSA will simply state TE reachability info.

Forwarding address

Data traffic for the advertised destination will be forwarded to this address. If the Forwarding address is set to 0.0.0.0, data traffic will be forwarded instead to the LSA's originator (i.e., the responsible AS boundary router).

External Route Tag

A 32-bit field attached to each external route. This is not used by the OSPF protocol itself. It may be used to communicate information between AS boundary routers; the precise nature of such information is outside the scope of this specification.

9. TE LSAs for Non-Packet Network

A non-packet network would use the TE LSAs described in the previous section for a packet network with some variations. These variations are described in the following subsections.

Two new LSAs, TE-Positional-ring-network LSA and TE-Router-Proxy LSA are defined for use in non-packet TE networks.

Readers may refer to [SONET-SDH] for a detailed description of the terms used in the context of SONET/SDH TDM networks,

9.1. TE-Router LSA (0x81)

The following fields are used to describe each router link (i.e., interface). Each router link is typed (see the below Type field). The Type field indicates the kind of link being described.

Type

A new link type "Positional-Ring Type" (value 5) is defined. This is essentially a connection to a TDM-Ring. TDM ring network is different from LAN/NBMA transit network in that nodes on the TDM ring do not necessarily have a terminating path between themselves. Second, the order of links is important in determining the circuit path. Third, the protection switching and the number of fibers from a node going into a ring are determined by the ring characteristics, for example, 2-fiber vs. 4-fiber ring and Unidirectional Path Switched Ring (UPSR) vs. Bidirectional Line Switched Ring (BLSR).

Type	Description
1	Point-to-point connection to another router
2	Connection to a transit network
3	Connection to a stub network
4	Virtual link
5	Positional-Ring type.

Link ID

Identifies the object that this router link connects to. Value depends on the link's Type. For a positional-ring type, the Link ID shall be IP Network/Subnet number just as the case with a broadcast transit network. The following table summarizes the updated Link ID values.

Type	Link ID
1	Neighboring router's Router ID
2	IP address of Designated Router
3	IP network/subnet number
4	Neighboring router's Router ID
5	IP network/subnet number

Link Data

This depends on the link's Type field. For type-5 links, this specifies the router interface's IP address.

9.1.1 Router-TE flags - TE Capabilities of a Router

Flags specific to non-packet TE nodes are described below.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|L|L|P|T|L|F|                                     |S|S|S|C|
|S|E|S|D|S|S|                                     |T|E|I|S|
|R|R|C|M|C|C|                                     |A|L|G|P|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|<---- Boolean TE flags ----->|<- TE flags pointing to TLVs ->|

```

Bit TDM - Indicates the node is TDM circuit switch capable.

Bit LSC - Indicates the node is capable of Lambda switching.

Bit FSC - Indicates the node is capable of fiber-switching (can also be a non-fiber link type).

9.1.2 Link-TE Options: TE Capabilities of a TE Link

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|T|N|P|T|L|F|D|                                     |S|L|B|C|
|E|T|K|D|S|S|B|                                     |R|U|W|O|
| |E|T|M|C|C|S|                                     |L|G|A|L|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|<---- Boolean TE flags ----->|<- TE flags pointing to TLVs ->|

```


TDM, LSC, FSC bits - Same as defined for router TE options.

9.2. TE-positional-ring-network LSA (0x82)

Network LSA is adequate for packet TE networks. A new TE-positional-ring-network LSA is defined to represent type-5 link networks, found in non-packet networks such as SONET/SDH TDM rings. A type-5 ring is a collection of network elements (NEs) forming a closed loop. Each NE is connected to two adjacent NEs via a duplex connection to provide redundancy in the ring. The sequence in which the NEs are placed on the Ring is pertinent. The NE that provides the OSPF-xTE functionality is termed the Gateway Network Element (GNE). The GNE selection criteria is outside the scope of this document. The GNE is also termed the Designated Router for the ring.

The TE-positional-ring-network LSA (0x82) is modeled after the network LSA and has the same flooding scope as the network LSA amongst the OSPF-xTE nodes within the area. Below is the format of the TE-Positional-Ring-network LSA. Unless specified explicitly otherwise, the fields carry the same meaning as they do in a network LSA. Only the differences are explained below.

A TE-positional-ring-network LSA is originated for each Positional-Ring type network in the area. The tuple of (Link State ID, Network Mask) below uniquely represents a ring. The TE option must be set in the Options flag while propagating the LSA.



Link State ID

This is the IP interface address of the network's Gateway Network Element, which is also the designated router.

Advertising Router

Router ID of the network's Designated Router.

Ring type

There are 8 types of SONET/SDH rings defined as follows.

- 1 - A Unidirectional Line Switched 2-fiber ring (2-fiber ULSR)
- 2 - A Bidirectional Line switched 2-fiber ring (2-fiber BLSR)
- 3 - A Unidirectional Path Switched 2-fiber ring (2-fiber UPSR)
- 4 - A Bidirectional Path switched 2-fiber ring (2-fiber BPSR)
- 5 - A Unidirectional Line Switched 4-fiber ring (4-fiber ULSR)
- 6 - A Bidirectional Line switched 4-fiber ring (4-fiber BLSR)
- 7 - A Unidirectional Path Switched 4-fiber ring (4-fiber UPSR)
- 8 - A Bidirectional Path switched 4-fiber ring (4-fiber BPSR)

Capacity Unit

Two units are currently defined, as follows.

- 1 - Synchronous Transport Signal (STS), which is the basic signal rate for SONET signals. The rate of an STS signal is 51.84 Mbps
- 2 - Synchronous Transport Multiplexer (STM), which is the basic signal rate for SDH signals. The rate of an STM signal is 155.52 Mbps

Ring capacity

Ring capacity expressed in number of Capacity Units.

Network Element Node Id

The Router ID of each of the routers in the positional-ring network. The list must start with the designated router as the first element. The Network Elements (NEs) must be listed in strict clockwise order as they appear on the ring, starting with the Gateway Network Element (GNE). The number of NEs in the ring can be deduced from the LSA header's length field.

9.3. TE-Router-Proxy LSA (0x8e)

This is a variation to the TE-router LSA in that the TE-router LSA is not advertised by the network element, but rather by a trusted TE-router Proxy. This is typically the scenario in a non-packet TE network, where some of the nodes do not have OSPF functionality and count on a helper node to do the advertisement for them. One such example would be the SONET/SDH Add-Drop Multiplexer (ADM) nodes in a TDM ring. The nodes may principally depend upon the GNE (Gateway Network Element) to do the advertisement for them. TE-router-Proxy LSA shall not be used to advertise area border routers and/or AS border routers.

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																																						
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										LS age																				Options																				0x8e																			
										Link State ID										(Router ID of the TE Network Element)																																																	
										Advertising Router																																																											
										LS sequence number																																																											
										LS checksum																				length																																							
										0																				Router-TE flags																																							
										Router-TE flags (contd.)																				Router-TE TLVs																																							
																																																																				
																													# of TE links																																							
										Link ID																																																											
										Link Data																																																											
										Type																				0																				Link-TE options																			
										Link-TE flags																				Zero or more Link-TE TLVs																																							
										Link ID																																																											
										Link Data																																																											
										...																																																											

10. Abstract Topology Representation with TE Support

Below, we consider a TE network composed of three OSPF areas, Area-1, Area-2 and Area-3, attached together through the backbone area. Area-1 has a single area border router, ABR-A1 and no ASBRs. Area-2 has an area border router ABR-A2 and an AS border router ASBR-S1. Area-3 has two area border routers ABR-A2 and ABR-A3 and an AS border router ASBR-S2. The following network also assumes a pre-engineered TE circuit path between ABR-A1 and ABR-A2; between ABR-A1 and ABR-A3; between ABR-A2 and ASBR-S1; and between ABR-A3 and ASBR-S2.

The following figure is an inter-area topology abstraction from the perspective of routers in Area-1. The abstraction illustrates reachability of TE networks and nodes within area to the external areas in the same AS and to the external ASes. The abstraction also illustrates pre-engineered TE circuit paths advertised by ABRs and ASBRs.

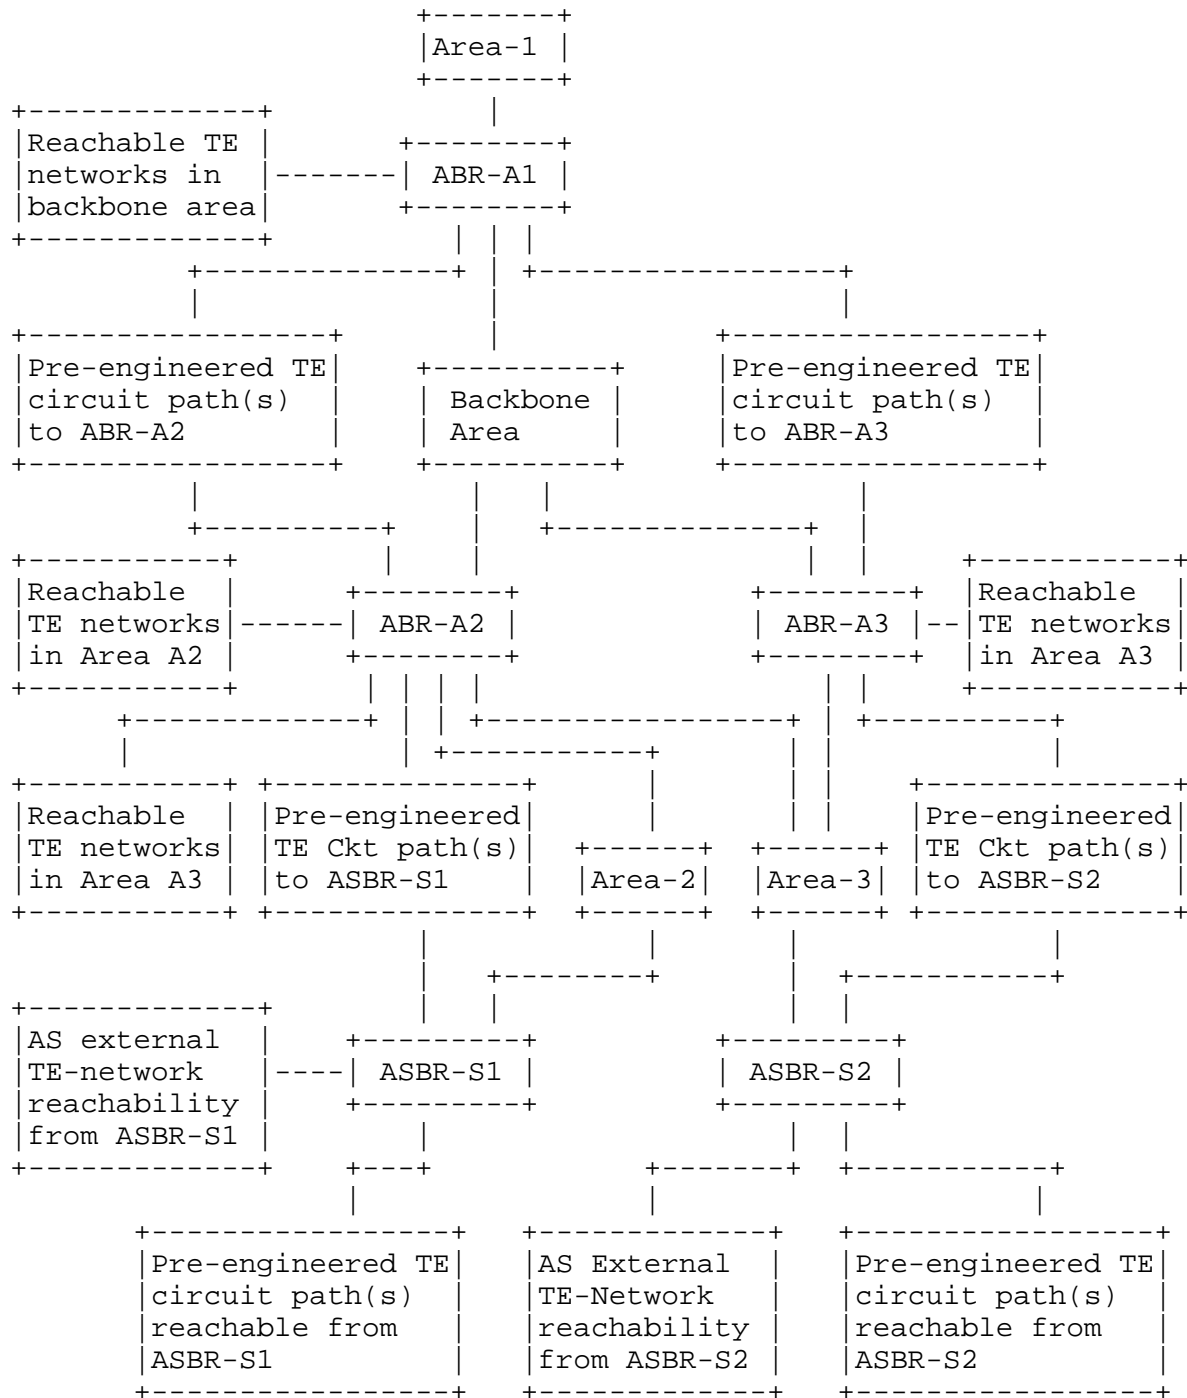


Figure 3: Inter-Area Abstraction as viewed by Area-1 TE-routers

11. Changes to Data Structures in OSPF-xTE Nodes

11.1. Changes to Router Data Structure

An OSPF-xTE router must be able to include the router-TE capabilities (as specified in section 8.1) in the router data structure. OSPF-xTE routers providing proxy service to other TE routers must also track the router and associated interface data structures for all the TE client nodes for which the proxy service is being provided. Presumably, the interaction between the Proxy server and the proxy clients is out-of-band.

11.2. Two Sets of Neighbors

Two sets of neighbor data structures are required. TE-neighbors set is used to advertise TE LSAs. Only the TE nodes will be members of the TE-neighbor set. Native neighbors set will be used to advertise native LSAs. All neighboring nodes supporting non-TE links are part of the Native neighbors set.

11.3. Changes to Interface Data Structure

The following new fields are introduced to the interface data structure.

TePermitted

If the value of the flag is TRUE, the interface may be advertised as a TE-enabled interface.

NonTePermitted

If the value of the flag is TRUE, the interface permits non-TE traffic on the interface. Specifically, this is applicable to packet networks, where data links may permit both TE and IP packets. For FSC and LSC TE networks, this flag is set to FALSE.

FloodingPermitted

If the value of the flag is TRUE, the interface may be used for OSPF and OSPF-xTE packet exchange to synchronize the LSDB across all adjacent neighbors. This is TRUE by default to all NonTePermitted interfaces that are enabled for OSPF. However, it is possible to set this to FALSE for some of the interfaces.

TE-TLVs

Each interface may define any number of TLVS that describe the link characteristics.

The following existing fields in Interface data structure will take on additional values to support TE extensions.

Type

The OSPF interface type can also be of type "Positional-Ring". The Positional-Ring type is different from other types (such as broadcast and NBMA) in that the exact location of the nodes on the ring is relevant, even though they are all on the same ring. SONET ADM ring is a good example of this. Complete ring positional-ring description may be provided by the GNE on a ring as a TE-network LSA for the ring.

List of Neighbors

The list may be statically defined for an interface without requiring the use of Hello protocol.

12. IANA Considerations

The IANA has assigned multicast address 224.0.0.24 to OSPFIGP-TE for the exchange of TE database descriptors.

TE LSA types and TE TLVs will be maintained by the IANA, using the following criteria.

12.1. TE LSA Type Values

LSA type is an 8-bit field required by each LSA. TE LSA types will have the high bit set to 1. TE LSAs can range from 0x80 through 0xFF. The following values are defined in sections 8.0 and 9.0. The remaining values are available for assignment by the IANA with IETF Consensus [RFC2434].

TE LSA Type	Value
TE-Router LSA	0x81
TE-Positional-ring-network LSA	0x82
TE-Summary Network LSA	0x83
TE-Summary router LSA	0x84
TE-AS-external LSAs	0x85
TE-Circuit-paths LSA	0x8C
TE-incremental-link-Update LSA	0x8d
TE-Router-Proxy LSA	0x8e

12.2. TE TLV Tag Values

TLV type is a 16-bit field required by each TE TLV. TLV type shall be unique across the router and link TLVs. A TLV type can range from 0x0001 through 0xFFFF. TLV type 0 is reserved and unassigned. The following TLV types are defined in sections 8.0 and 9.0. The remaining values are available for assignment by the IANA with IETF Consensus [RFC2434].

TE TLV Tag	Reference Section	Value
TE-LINK-TLV-SRLG	Section 8.1.4.1	0x0001
TE-LINK-TLV-BANDWIDTH-MAX	Section 8.1.4.2	0x0002
TE-LINK-TLV-BANDWIDTH-MAX-FOR-TE	Section 8.1.4.3	0x0003
TE-LINK-TLV-BANDWIDTH-TE	Section 8.1.4.4	0x0004
TE-LINK-TLV-LUG	Section 8.1.4.5	0x0005
TE-LINK-TLV-COLOR	Section 8.1.4.6	0x0006
TE-LINK-TLV-NULL	Section 8.1.4.7	0x8888
TE-NODE-TLV-MPLS-SWITCHING	Section 8.1.2.1	0x8001
TE-NODE-TLV-MPLS-SIG-PROTOCOLS	Section 8.1.2.2	0x8002
TE-NODE-TLV-CSPF-ALG	Section 8.1.2.3	0x8003
TE-NODE-TLV-NULL	Section 8.1.2.4	0x8888

13. Acknowledgements

The authors wish to specially thank Chitti Babu and his team for implementing the protocol specified in a packet network and verifying several portions of the specification in a mixed packet network. The authors also wish to thank Vishwas Manral, Riyad Hartani, and Tricci So for their valuable comments and feedback on the document. Lastly, the authors wish to thank Alex Zinin and Mike Shand for their document (now defunct) titled "Flooding optimizations in link state routing protocols". The document provided inspiration to the authors to be sensitive to the high flooding rate, likely in TE networks.

14. Security Considerations

Security considerations for the base OSPF protocol are covered in [OSPF-V2] and [SEC-OSPF]. This memo does not create any new security issues for the OSPF protocol. Security measures applied to the native OSPF (refer [SEC-OSPF]) are directly applicable to the TE LSAs described in the document. Discussed below are the security considerations in processing TE LSAs.

Secure communication between OSPF-xTE nodes has a number of components. Authorization, authentication, integrity and confidentiality. Authorization refers to whether a particular OSPF-xTE node is authorized to receive or propagate the TE LSAs to its neighbors. Failing the authorization process might indicate a resource theft attempt or unauthorized resource advertisement. In either case, the OSPF-xTE nodes should take proper measures to audit/log such attempts so as to alert the administrator to take necessary action. OSPF-xTE nodes may refuse to communicate with the neighboring nodes that fail to prompt the required credentials.

Authentication refers to confirming the identity of an originator for the datagrams received from the originator. Lack of strong credentials for authentication of OSPF-xTE LSAs can seriously jeopardize the TE service rendered by the network. A consequence of not authenticating a neighbor would be that an attacker could spoof the identity of a "legitimate" OSPF-xTE node and manipulate the state, and the TE database including the topology and metrics collected. This could potentially cause denial-of-service on the TE network. Another consequence of not authenticating is that an attacker could pose as OSPF-xTE neighbor and respond in a manner that would divert TE data to the attacker.

Integrity is required to ensure that an OSPF-xTE message has not been accidentally or maliciously altered or destroyed. The result of a lack of data integrity enforcement in an untrusted environment could be that an imposter will alter the messages sent by a legitimate adjacent neighbor and bring the OSPF-xTE on a node and the whole network to a halt or cause a denial of service for the TE circuit paths effected by the alteration.

Confidentiality of OSPF-xTE messages ensures that the TE LSAs are accessible only to the authorized entities. When OSPF-xTE is deployed in an untrusted environment, lack of confidentiality will allow an intruder to perform traffic flow analysis and snoop the TE control network to monitor the traffic metrics and the rate at which circuit paths are being setup and torn-down. The intruder could cannibalize a lesser secure OSPF-xTE node and destroy or compromise the state and TE-LSDB on the node. Needless to say, the least secure

OSPF-xTE will become the Achilles heel and make the TE network vulnerable to security attacks.

15. Normative References

- [MPLS-ARCH] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [MPLS-TE] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, September 1999.
- [OSPF-V2] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [SEC-OSPF] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", RFC 2154, June 1997.
- [OSPF-CAP] Lindem, A., Ed., Shen, N., Vasseur, J., Aggarwal, R., and S. Schaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 4970, July 2007.
- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

16. Informative References

- [BGP-OSPF] Ferguson, D., "The OSPF External Attribute LSA", unpublished.
- [CR-LDP] Jamoussi, B., Andersson, L., Callon, R., Dantu, R., Wu, L., Doolan, P., Worster, T., Feldman, N., Fredette, A., Girish, M., Gray, E., Heenanen, J., Kilty, T., and A. Malis, "Constraint-Based LSP Setup using LDP", RFC 3212, January 2002.
- [GMPLS-TE] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [MOSPF] Moy, J., "Multicast Extensions to OSPF", RFC 1584, March 1994.
- [NSSA] Murphy, P., "The OSPF Not-So-Stubby Area (NSSA) Option", RFC 3101, January 2003.
- [OPAQUE] Coltun, R., "The OSPF Opaque LSA Option", RFC 2370, July 1998.

- [OPQLSA-TE] Katz, D., Yeung, D., and K. Kompella, "Traffic Engineering Extensions to OSPF", RFC 3630, September 2003.
- [RSVP-TE] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [SONET-SDH] Chow, M.-C., "Understanding SONET/SDH Standards and Applications", Holmdel, N.J.: Andan Publisher, 1995.

Authors' Addresses

Pyda Srisuresh
Kazeon Systems, Inc.
1161 San Antonio Rd.
Mountain View, CA 94043
U.S.A.

Phone: (408) 836-4773
EMail: srisuresh@yahoo.com

Paul Joseph
Consultant
10100 Torre Avenue, # 121
Cupertino, CA 95014
U.S.A.

Phone: (408) 777-8493
EMail: paul_95014@yahoo.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78 and at www.rfc-editor.org/copyright.html, 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/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST 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 procedures with respect to rights in RFC 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.

