

INTERNET-DRAFT  
Intended Status: Proposed Standard  
Expires: October 13, 2011

A. Malis, ed.  
Verizon Communications  
A. Lindem, ed.  
Ericsson  
April 11, 2011

**Updates to ASON Routing for OSPFv2 Protocols (RFC 5787bis)  
draft-malis-ccamp-rfc5787bis-03.txt**

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at  
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Abstract

The ITU-T has defined an architecture and requirements for operating an Automatically Switched Optical Network (ASON).

The Generalized Multiprotocol Label Switching (GMPLS) protocol suite is designed to provide a control plane for a range of network technologies including optical networks such as time division multiplexing (TDM) networks including SONET/SDH and Optical Transport Networks (OTNs), and lambda switching optical networks.

The requirements for GMPLS routing to satisfy the requirements of ASON routing, and an evaluation of existing GMPLS routing protocols are provided in other documents. This document defines extensions to the OSPFv2 Link State Routing Protocol to meet the requirements for routing in an ASON.

Note that this work is scoped to the requirements and evaluation expressed in [RFC 4258](#) and [RFC 4652](#) and the ITU-T Recommendations current when those documents were written. Future extensions of revisions of this work may be necessary if the ITU-T Recommendations are revised or if new requirements are introduced into a revision of [RFC 4258](#).

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">4</a>
<a href="#">1.1.</a>	<a href="#">Conventions Used in This Document</a>	<a href="#">5</a>
<a href="#">2.</a>	<a href="#">Routing Areas, OSPF Areas, and Protocol Instances</a>	<a href="#">5</a>
<a href="#">3.</a>	<a href="#">Terminology and Identification</a>	<a href="#">6</a>
<a href="#">4.</a>	<a href="#">Reachability</a>	<a href="#">6</a>
<a href="#">5.</a>	<a href="#">Link Attribute</a>	<a href="#">7</a>
<a href="#">5.1.</a>	<a href="#">Local Adaptation</a>	<a href="#">7</a>
<a href="#">5.2.</a>	<a href="#">Bandwidth Accounting</a>	<a href="#">8</a>
<a href="#">6.</a>	<a href="#">Routing Information Scope</a>	<a href="#">8</a>
<a href="#">6.1.</a>	<a href="#">Link Advertisement (Local and Remote TE Router ID Sub-TLV)</a>	<a href="#">9</a>
<a href="#">6.2.</a>	<a href="#">Reachability Advertisement (Local TE Router ID sub-TLV)</a>	<a href="#">10</a>
<a href="#">7.</a>	<a href="#">Routing Information Dissemination</a>	<a href="#">11</a>
<a href="#">7.1</a>	<a href="#">Import/Export Rules</a>	<a href="#">11</a>
<a href="#">7.2</a>	<a href="#">Loop Prevention</a>	<a href="#">11</a>
<a href="#">7.2.1</a>	<a href="#">Inter-RA Export Upward/Downward Sub-TLVs</a>	<a href="#">12</a>
<a href="#">7.2.2</a>	<a href="#">Inter-RA Export Upward/Downward Sub-TLV Processing</a>	<a href="#">13</a>
<a href="#">8.</a>	<a href="#">OSPFv2 Scalability</a>	<a href="#">13</a>
<a href="#">9.</a>	<a href="#">Security Considerations</a>	<a href="#">14</a>
<a href="#">10.</a>	<a href="#">IANA Considerations</a>	<a href="#">14</a>
<a href="#">10.1.</a>	<a href="#">Sub-TLVs of the Link TLV</a>	<a href="#">14</a>



<a href="#">10.2.</a>	Sub-TLVs of the Node Attribute TLV . . . . .	<a href="#">15</a>
<a href="#">10.3.</a>	Sub-TLVs of the Router Address TLV . . . . .	<a href="#">15</a>
<a href="#">11.</a>	Management Considerations . . . . .	<a href="#">16</a>
<a href="#">11.1.</a>	Routing Area (RA) Isolation . . . . .	<a href="#">16</a>
<a href="#">11.2</a>	Routing Area (RA) Topology/Configuration Changes . . . .	<a href="#">16</a>
<a href="#">12.</a>	Comparison to Requirements in <a href="#">RFC 4258</a> . . . . .	<a href="#">16</a>
<a href="#">13.</a>	References . . . . .	<a href="#">22</a>
<a href="#">13.1.</a>	Normative References . . . . .	<a href="#">22</a>
<a href="#">13.2.</a>	Informative References . . . . .	<a href="#">23</a>
<a href="#">14.</a>	Acknowledgements . . . . .	<a href="#">24</a>
<a href="#">Appendix A.</a>	ASON Terminology . . . . .	<a href="#">25</a>
<a href="#">Appendix B.</a>	ASON Routing Terminology . . . . .	<a href="#">26</a>
Authors' Addresses	. . . . .	<a href="#">27</a>

## 1. Introduction

The Generalized Multiprotocol Label Switching (GMPLS) [[RFC3945](#)] protocol suite is designed to provide a control plane for a range of network technologies including optical networks such as time division multiplexing (TDM) networks including SONET/SDH and Optical Transport Networks (OTNs), and lambda switching optical networks.

The ITU-T defines the architecture of the Automatically Switched Optical Network (ASON) in [[G.8080](#)].

[RFC4258] describes the routing requirements for the GMPLS suite of routing protocols to support the capabilities and functionality of ASON control planes identified in [[G.7715](#)] and in [[G.7715.1](#)].

[RFC4652] evaluates the IETF Link State routing protocols against the requirements identified in [[RFC4258](#)]. [Section 7.1 of \[RFC4652\]](#) summarizes the capabilities to be provided by OSPFv2 [[RFC2328](#)] in support of ASON routing. This document describes the OSPFv2 specifics for ASON routing.

Multi-layer transport networks are constructed from multiple networks of different technologies operating in a client-server relationship. The ASON routing model includes the definition of routing levels that provide scaling and confidentiality benefits. In multi-level routing, domains called routing areas (RAs) are arranged in a hierarchical relationship. Note that as described in [[RFC4652](#)], there is no implied relationship between multi-layer transport networks and multi-level routing. The multi-level routing mechanisms described in this document work for both single-layer and multi-layer networks.

Implementations may support a hierarchical routing topology (multi-level) for multiple transport network layers and/or a hierarchical routing topology for a single transport network layer.

This document describes the processing of the generic (technology-independent) link attributes that are defined in [[RFC3630](#)], [[RFC4202](#)], and [[RFC4203](#)] and that are extended in this document. As described in [Section 5.2](#), technology-specific traffic engineering attributes and their processing may be defined in other documents that complement this document.

Note that this work is scoped to the requirements and evaluation expressed in [[RFC4258](#)] and [[RFC4652](#)] and the ITU-T Recommendations current when those documents were written. Future extensions of revisions of this work may be necessary if the ITU-T Recommendations are revised or if new requirements are introduced into a revision of



[[RFC4258](#)].

### **1.1. Conventions Used in This Document**

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 2119](#) [[RFC2119](#)].

The reader is assumed to be familiar with the terminology and requirements developed in [[RFC4258](#)] and the evaluation outcomes described in [[RFC4652](#)].

General ASON terminology is provided in [Appendix A](#). ASON routing terminology is described in [Appendix B](#).

## **2. Routing Areas, OSPF Areas, and Protocol Instances**

An ASON routing area (RA) represents a partition of the data plane, and its identifier is used within the control plane as the representation of this partition.

RAs are hierarchically contained: a higher-level (parent) RA contains lower-level (child) RAs that in turn MAY also contain RAs, etc. Thus, RAs contain RAs that recursively define successive hierarchical RA levels. Routing information may be exchanged between levels of the RA hierarchy, i.e., Level N+1 and N, where Level N represents the RAs contained by Level N+1. The links connecting RAs may be viewed as external links (inter-RA links), and the links representing connectivity within an RA may be viewed as internal links (intra-RA links). The external links to an RA at one level of the hierarchy may be internal links in the parent RA. Intra-RA links of a child RA MAY be hidden from the parent RA's view. [[RFC4258](#)]

An ASON RA can be mapped to an OSPF area, but the hierarchy of ASON RA levels does not map to the hierarchy of OSPF areas. Instead, successive hierarchical levels of RAs MUST be represented by separate instances of the protocol. Thus, inter-level routing information exchange (as described in [Section 7](#)) involves the export and import of routing information between protocol instances.

An ASON RA may therefore be identified by the combination of its OSPF instance identifier and its OSPF area identifier. With proper and careful network-wide configuration, this can be achieved using just the OSPF area identifier, and this process is RECOMMENDED in this document. These concepts are discussed in [Section 7](#).

A key ASON requirement is the support of multiple transport planes or layers. Each transport node has associated topology (links and





reachability) which is used for ASON routing.

### **3. Terminology and Identification**

This section describes the mapping of key ASON entities to OSPF entities. [Appendix A](#) contains a complete glossary of ASON routing terminology.

There are three categories of identifiers used for ASON routing (G7715.1): transport plane names, control plane identifiers for components, and SCN addresses. This section discusses the mapping between ASON routing identifiers and corresponding identifiers defined for GMPLS routing, and how these support the physical (or logical) separation of transport plane entities and control plane components. GMPLS supports this separation of identifiers and planes.

In the context of OSPF Traffic Engineering (TE), an ASON transport node corresponds to a unique OSPF TE node. An OSPF TE node is uniquely identified by the TE Router Address TLV [[RFC3630](#)]. In this document, this TE Router Address is referred to as the TE Router ID, which is in the ASON transport plane name space. The TE Router ID should not be confused with the OSPF Router ID which uniquely identifies an OSPF router within an OSPF routing domain [[RFC2328](#)] and is in a name space for control plane components.

Note: The Router Address top-level TLV definition, processing, and usage are unchanged from [[RFC3630](#)]. This TLV specifies a stable OSPF TE node IP address, i.e., the IP address is always reachable when there is IP connectivity to the associated OSPF TE node.

ASON defines a Routing Controller (RC) as an entity that handles (abstract) information needed for routing and the routing information exchange with peering RCs by operating on the Routing Database (RDB). ASON defines a Protocol Controller (PC) as an entity that handles protocol-specific message exchanges according to the reference point over which the information is exchanged (e.g., E-NNI, I-NNI), and internal exchanges with the Routing Controller (RC) [[RFC4258](#)]. In this document, an OSPF router advertising ASON TE topology information will perform both the functions of the RC and PC. Each OSPF router is uniquely identified by its OSPF Router ID [[RFC2328](#)].

### **4. Reachability**

Reachability in ASON refers to the set of endpoints reachable in the transport plane by a node or the reachable endpoints of a level N. Reachable entities are identified in the transport plane name space



(ASON SNPP name space). In order to advertise blocks of reachable address prefixes, a summarization mechanism is introduced that is based on the techniques described in [\[RFC5786\]](#). For ASON reachability advertisement, blocks of reachable address prefixes are advertised together with the associated data plane node. The data plane node is identified in the control plane by its TE Router ID, as discussed in [section 6](#).

In order to support ASON reachability advertisement, the Node Attribute TLV defined in [\[RFC5786\]](#) is used to advertise the combination of a TE Router ID and its set of associated reachable address prefixes. The Node Attribute TLV can contain the following sub-TLVs:

- TE Router ID sub-TLV: Length: 4; Defined in [Section 6.2](#)
- Node IPv4 Local Address sub-TLV: Length: variable; [\[RFC5786\]](#)
- Node IPv6 Local Address sub-TLV: Length: variable; [\[RFC5786\]](#)

A router may support multiple transport nodes as discussed in [section 6](#), and, as a result, may be required to advertise reachability (ASON TRIs) separately for each transport node. As a consequence, it MUST be possible for the router to originate more than one TE LSA containing the Node Attribute TLV when used for ASON reachability advertisement.

Hence, the Node Attribute TLV [\[RFC5786\]](#) advertisement rules must be relaxed for ASON. A Node Attribute TLV MAY appear in more than one TE LSA originated by the RC when the RC is advertising reachability information for a different transport node identified by the Local TE Router Sub-TLV (refer to [section 6.1](#)).

## 5. Link Attribute

With the exception of local adaptation (described below), the mapping of link attributes and characteristics to OSPF TE Link TLV Sub-TLVs is unchanged [\[RFC4652\]](#). OSPF TE Link TLV Sub-TLVs are described in [\[RFC3630\]](#) and [\[RFC4203\]](#). Advertisement of this information SHOULD be supported on a per-layer basis, i.e., one TE LSA per unique switching capability and bandwidth granularity combination.

### 5.1. Local Adaptation

Local adaptation is defined as a TE link attribute (i.e., sub-TLV) that describes the cross/inter-layer relationships.

The Interface Switching Capability Descriptor (ISCD) TE Attribute [\[RFC4202\]](#) identifies the ability of the TE link to support cross-connection to another link within the same layer. When advertising



link adaptation, it also identifies the ability to use a locally terminated connection that belongs to one layer as a data link for another layer (adaptation capability). However, the information associated with the ability to terminate connections within that layer (referred to as the termination capability) is advertised with the adaptation capability.

For instance, a link between two optical cross-connects will contain at least one ISCD attribute describing the Lambda Switching Capable (LSC) switching capability. Conversely, a link between an optical cross-connect and an IP/MPLS Label Switching Router (LSR) will contain at least two ISCD attributes, one for the description of the LSC termination capability and one for the Packet Switching Capable (PSC) adaptation capability.

In OSPFv2, the Interface Switching Capability Descriptor (ISCD) is a sub-TLV (type 15) of the top-level Link TLV (type 2) [RFC4203]. The adaptation and termination capabilities are advertised using two separate ISCD sub-TLVs within the same top-level Link TLV.

An interface MAY have more than one ISCD sub-TLV, [RFC4202] and [RFC4203]. Hence, the corresponding advertisements should not result in any compatibility issues.

## **5.2. Bandwidth Accounting**

GMPLS routing defines an Interface Switching Capability Descriptor (ISCD) that provides, among other things, the available (maximum/minimum) bandwidth per priority available for Label Switched Path (LSPs). One or more ISCD sub-TLVs can be associated with an interface, [RFC4202] and [RFC4203]. This information, combined with the Unreserved Bandwidth Link TLV sub-TLV [RFC3630], provides the basis for bandwidth accounting.

In the ASON context, additional information may be included when the representation and information in the other advertised fields are not sufficient for a specific technology, e.g., SDH. The definition of technology-specific information elements is beyond the scope of this document. Some technologies will not require additional information beyond what is already defined in [RFC3630], [RFC4202], and [RFC4203].

## **6. Routing Information Scope**

For ASON routing, the control plane component routing adjacency topology (i.e., the associated Protocol Controller (PC) connectivity) and the transport topology are NOT assumed to be congruent [RFC4258]. Hence, a single OSPF router (i.e., the PC) MUST be able to advertise



on behalf of multiple transport layer nodes. The OSPF routers are identified by OSPF Router ID and the transport nodes are identified by TE Router ID.

The Router Address TLV [RFC3630] is used to advertise the TE Router ID associated with the advertising Routing Controller. TE Router IDs for additional transport nodes are advertised through specification of the Local TE Router Identifier in the Local and Remote TE Router TE sub-TLV and the Local TE Router Identifier sub-TLV described in the sections below. These Local TE Router Identifiers are typically used as the local endpoints for TE Label Switched Paths (LSPs) terminating on the associated transport node.

It MAY be feasible for multiple OSPF Routers to advertise TE information for the same transport node. However, this is not considered a required use case and is not discussed further.

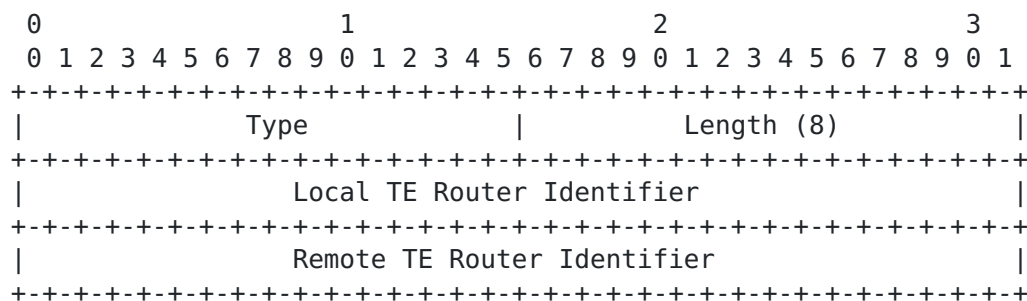
### **6.1. Link Advertisement (Local and Remote TE Router ID Sub-TLV)**

An OSPF router advertising on behalf of multiple transport nodes will require additional information to distinguish the link endpoints amongst the subsumed transport nodes. In order to unambiguously specify the transport topology, the local and remote transport nodes MUST be identified by TE router ID.

For this purpose, a new sub-TLV of the OSPFv2 TE LSA top-level Link TLV is introduced that defines the Local and Remote TE Router ID.

The Type field of the Local and Remote TE Router ID sub-TLV is assigned a value TBD. The Length field takes the value 8. The Value field of this sub-TLV contains 4 octets of the Local TE Router Identifier followed by 4 octets of the Remote TE Router Identifier. The value of the Local and Remote TE Router Identifier SHOULD NOT be set to 0.

The format of the Local and Remote TE Router ID sub-TLV is:



This sub-TLV MUST be included as a sub-TLV of the top-level Link TLV





if the OSPF router is advertising on behalf of one or more transport nodes having TE Router IDs different from the TE Router ID advertised in the Router Address TLV. Therefore, it **MUST** be included if the OSPF router is advertising on behalf of multiple transport nodes.

Note: The Link ID sub-TLV identifies the other end of the link (i.e., Router ID of the neighbor for point-to-point links) [[RFC3630](#)]. When the Local and Remote TE Router ID Sub-TLV is present, it **MUST** be used to identify local and remote transport node endpoints for the link and the Link-ID sub-TLV **MUST** be ignored. The Local and Remote ID sub-TLV, if specified, **MUST** only be specified once.

## 6.2. Reachability Advertisement (Local TE Router ID sub-TLV)

When an OSPF router is advertising on behalf of multiple transport nodes, the routing protocol **MUST** be able to associate the advertised reachability information with the correct transport node.

For this purpose, a new sub-TLV of the OSPFv2 TE LSA top-level Node Attribute TLV is introduced. This TLV associates the local prefixes (see above) to a given transport node identified by TE Router ID.

The Type field of the Local TE Router ID sub-TLV is assigned a value TBD. The Length field takes the value 4. The Value field of this sub-TLV contains the Local TE Router Identifier [[RFC3630](#)] encoded over 4 octets.

The format of the Local TE Router ID sub-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 (4)
|                                     |                                     |
|                                     Local TE Router Identifier             |
|                                     |                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

This sub-TLV **MUST** be included as a sub-TLV of the top-level Node Attribute TLV if the OSPF router is advertising on behalf of one or more transport nodes having TE Router IDs different from the TE Router ID advertised in the Router Address TLV. Therefore, it **MUST** be included if the OSPF router is advertising on behalf of multiple transport nodes.



## **7. Routing Information Dissemination**

An ASON routing area (RA) represents a partition of the data plane, and its identifier is used within the control plane as the representation of this partition. An RA may contain smaller RAs inter-connected by links. ASON RA levels do not map directly to OSPF areas. Rather, hierarchical levels of RAs are represented by separate OSPF protocol instances.

Routing controllers (RCs) supporting multiple RAs disseminate information downward and upward in this ASON hierarchy. The vertical routing information dissemination mechanisms described in this section do not introduce or imply hierarchical OSPF areas. RCs supporting RAs at multiple levels are structured as separate OSPF instances with routing information exchange between levels described by import and export rules between these instances. The functionality described herein does not pertain to OSPF areas or OSPF Area Border Router (ABR) functionality.

### **7.1 Import/Export Rules**

RCs supporting RAs disseminate information upward and downward in the hierarchy by importing/exporting routing information as TE LSAs. TE LSAs are area-scoped opaque LSAs with opaque type 1 [[RFC3630](#)]. The information that MAY be exchanged between adjacent levels includes the Router Address, Link, and Node Attribute top-level TLVs.

The imported/exported routing information content MAY be transformed, e.g., filtered or aggregated, as long as the resulting routing information is consistent. In particular, when more than one RC is bound to adjacent levels and both are allowed to import/export routing information, it is expected that these transformations are performed in a consistent manner. Definition of these policy-based mechanisms is outside the scope of this document.

In practice, and in order to avoid scalability and processing overhead, routing information imported/exported downward/upward in the hierarchy is expected to include reachability information (see [Section 4](#)) and, upon strict policy control, link topology information.

### **7.2 Loop Prevention**

When more than one RC is bound to an adjacent level of the ASON hierarchy, and is configured to export routing information upward or downward, a specific mechanism is required to avoid looping of routing information. Looping is the re-advertisement of routing information into an RA that had previously advertised that routing



information upward or downward into an upper or lower level RA in the ASON hierarchy. For example, without loop prevention mechanisms, this could happen when the RC advertising routing information downward in the hierarchy is not the same one that advertises routing information upward in the hierarchy.

### **7.2.1 Inter-RA Export Upward/Downward Sub-TLVs**


The Inter-RA Export Sub-TLVs can be used to prevent the re-advertisement of OSPF TE routing information into an RA which previously advertised that information. The type value TBD will indicate that the associated routing information has been exported downward. The type value TBD will indicate that the associated routing information has been exported upward. While it is not required for routing information exported downward, both Sub-TLVs will include the Routing Area (RA) ID from the which the routing information was exported. This RA is not necessarily the RA originating the routing information but RA from which the information was immediately exported.

These additional Sub-TLVs MAY be included in TE LSAs that include any of the following top-level TLVs:

- Router Address top-level TLV
- Link top-level TLV
- Node Attribute top-level TLV

The Type field of the Inter-RA Export Upward and Inter-RA Export Downward sub-TLVs are respectively assigned the values TBD1 and TBD2. The Length of the Associated RA ID TLV is 4 octets. The Value field in these sub-TLVs contains the associated RA ID. The RA ID value must be a unique identifier for the RA within the ASON routing domain.

The format of the Inter-RA Export Upward and Inter-RA Export Downward Sub-TLVs is graphically depicted below:





### **7.2.2 Inter-RA Export Upward/Downward Sub-TLV Processing**

TE LSAs MAY be imported or exported downward or upward in the ASON routing hierarchy. The direction and advertising RA ID are advertised in an Inter-RA Export Upward/Downward Sub-TLV. They MUST be retained and advertised in the receiving RA with the associated routing information.

When exporting routing information upward in the ASON routing hierarchy, any information received from a level above, i.e., tagged with an Inter-RA Export Downward Sub-TLV, MUST NOT be exported upward. Since an RA at level N is contained by a single RA at level N+1, this is the only checking that is necessary and the associated RA ID is used solely for informational purposes.

When exporting routing information downward in the ASON routing hierarchy, any information received from a level below, i.e., tagged with an Inter-RA Export Upward Sub-TLV MUST NOT be exported downward if the target RA ID matches the RA ID associated with the routing information. This additional checking is required for routing information exported downward since a single RA at level N+1 may contain multiple RAs at level N in the ASON routing hierarchy. In order words, routing information MUST NOT be exported downward into the RA from which it was received.

## **8. OSPFv2 Scalability**

The extensions described herein are only applicable to ASON routing domains and it is not expected that the attendant reachability (see [Section 4](#)) and link information will ever be mixed with global or local IP routing information. If there were ever a requirement for a given RC to participate in both domains, separate OSPFv2 instances would be utilized. However, in a multi-level ASON hierarchy, the potential volume of information could be quite large and the recommendations in this section SHOULD be followed by RCs implementing this specification.

- Routing information exchange upward/downward in the hierarchy between adjacent RAs SHOULD, by default, be limited to reachability information. In addition, several transformations such as prefix aggregation are RECOMMENDED to reduce the amount of information imported/exported by a given RC when such transformations will not impact consistency.
- Routing information exchange upward/downward in the ASON hierarchy involving TE attributes MUST be under strict policy control. Pacing and min/max thresholds for triggered updates are strongly RECOMMENDED.





- The number of routing levels MUST be maintained under strict policy control.

## **9. Security Considerations**

This document specifies the contents and processing of OSPFv2 TE LSAs [[RFC3630](#)] and [[RFC4202](#)]. The TE LSA extensions defined in this document are not used for SPF computation, and have no direct effect on IP routing. Additionally, ASON routing domains are delimited by the usual administrative domain boundaries.

Any mechanisms used for securing the exchange of normal OSPF LSAs can be applied equally to all TE LSAs used in the ASON context. Authentication of OSPFv2 LSA exchanges (such as OSPF cryptographic authentication [[RFC2328](#)] and [[RFC5709](#)]) can be used to secure against passive attacks and provide significant protection against active attacks. [[RFC5709](#)] defines a mechanism for authenticating OSPFv2 packets by making use of the HMAC algorithm in conjunction with the SHA family of cryptographic hash functions.

If a stronger authentication were believed to be required, then the use of a full digital signature [[RFC2154](#)] would be an approach that should be seriously considered. Use of full digital signatures would enable precise authentication of the OSPF router originating each OSPF link-state advertisement, and thereby provide much stronger integrity protection for the OSPF routing domain.

## **10. IANA Considerations**

This document is classified as Standards Track. It defines new sub-TLVs for inclusion in OSPF TE LSAs. According to the assignment policies for the registries of code points for these sub-TLVs, values must be assigned by IANA [[RFC3630](#)].

The following subsections summarize the required sub-TLVs.

### **10.1. Sub-TLVs of the Link TLV**

This document defines the following sub-TLVs of the Link TLV advertised in the OSPF TE LSA:

- Local and Remote TE Router ID sub-TLV
- Associated RA ID sub-TLV
- Inter-RA Export Upward sub-TLV
- Inter-RA Export Downward sub-TLV

Codepoints for these Sub-TLVs should be allocated from the "Types for sub-TLVs of TE Link TLV (Value 2)" registry standards action range (0



- 32767) [[RFC3630](#)].

Note that the same values for the Associated RA ID sub-TLV, Inter-RA Export Upward sub-TLV, and Inter-RA Export Downward Sub-TLV MUST be used when they appear in the Link TLV, Node Attribute TLV, and Router Address TLV.

### **10.2. Sub-TLVs of the Node Attribute TLV**

This document defines the following sub-TLVs of the Node Attribute TLV advertised in the OSPF TE LSA:

- Local TE Router ID sub-TLV
- Associated RA ID sub-TLV
- Inter-RA Export Upward sub-TLV
- Inter-RA Export Downward sub-TLV

Codepoints for these Sub-TLVs should be assigned from the "Types for sub-TLVs of TE Node Attribute TLV (Value 5)" registry standards action range (0 - 32767) [[RFC5786](#)].

Note that the same values for the Associated RA ID sub-TLV, Inter-RA Export Upward sub-TLV, and Inter-RA Export Downward Sub-TLV MUST be used when they appear in the Link TLV, Node Attribute TLV, and Router Address TLV.

### **10.3. Sub-TLVs of the Router Address TLV**

The Router Address TLV is advertised in the OSPF TE LSA [[RFC3630](#)]. Since this TLV currently has no Sub-TLVs defined, a "Types for sub-TLVs of Router Address TLV (Value 1)" registry must be defined.

The registry guidelines for the assignment of types for sub-TLVs of the Router Address TLV are as follows:

- o Types in the range 0-32767 are to be assigned via Standards Action.
- o Types in the range 32768-32777 are for experimental use; these will not be registered with IANA, and MUST NOT be mentioned by RFCs.
- o Types in the range 32778-65535 are not to be assigned at this time. Before any assignments can be made in this range, there MUST be a Standards Track RFC that specifies IANA Considerations that covers the range being assigned.

This document defines the following sub-TLVs for inclusion in the



Router Address TLV:

- Associated RA ID sub-TLV
- Inter-RA Export Upward sub-TLV
- Inter-RA Export Downward sub-TLV

Codepoints for these Sub-TLVs should be allocated from the "Types for sub-TLVs of Router Address TLV (Value 1)" registry standards action range (0 - 32767).

Note that the same values for the Associated RA ID sub-TLV, Inter-RA Export Upward sub-TLV, and Inter-RA Export Downward Sub-TLV MUST be used when they appear in the Link TLV, Node Attribute TLV, and Router Address TLV.

## **11. Management Considerations**

### **11.1. Routing Area (RA) Isolation**

If the RA Identifier is mapped to the OSPF Area ID as recommended in [section 2.0](#), OSPF [[RFC2328](#)] implicitly provides isolation. On any intra-RA link, packets will only be accepted if the area-id in the OSPF packet header matches the area ID for the OSPF interface on which the packet was received. Hence, RCs will only establish adjacencies and exchange reachability information (see [Section 4.0](#)) with RCs in the same RC. Other mechanisms for RA isolation are beyond the scope of this document.

### **11.2 Routing Area (RA) Topology/Configuration Changes**

The GMPLS Routing for ASON requirements [[RFC4258](#)] dictate that the routing protocol MUST support reconfiguration and SHOULD support architectural evolution. OSPF [[RFC2328](#)] includes support for the dynamic introduction or removal of ASON reachability information through the flooding and purging of OSPF opaque LSAs [[RFC5250](#)]. Also, when an RA is partitioned or an RC fails, stale LSAs SHOULD NOT be used unless the advertising RC is reachable. The configuration of OSPF RAs and the policies governing the redistribution of ASON reachability information between RAs are implementation issues outside of the OSPF routing protocol and beyond the scope of this document.

## **12. Comparison to Requirements in [RFC 4258](#)**

The following table shows how this draft complies with the requirements in [[RFC4258](#)]. The first column contains a requirements number (1-30) and the relevant section in [RFC 4258](#). The second column describes the requirement, the third column discusses the compliance



to that requirement, and the fourth column lists the relevant section in draft, and/or another RFC that already satisfies the requirement.

<a href="#">RFC 4258</a> Section (Req. Number)	<a href="#">RFC 4258</a> Requirement	Compliance	Reference
3.0 (1)	The failure of an RC, or the failure of communications between RCs, and the subsequent recovery from the failure condition MUST NOT disrupt call in progress.	Implied by separation of transport and control plane.	Not an attribute of routing protocol.
3.1 (2)	Multiple Hierarchical Level of ASON Routing Areas (RAs).	Yes	Sections <a href="#">2</a> and <a href="#">3</a>
3.1 (3)	Prior to establishing communications, RCs MUST verify that they are bound to the same parent RA.	Yes when RA maps to OSPF Area ID.	<a href="#">Section 11.1</a>
3.1 (4)	The RC ID MUST be unique within its containing RA.	Yes	<a href="#">RFC 2328</a> and <a href="#">Section 3</a> .
3.1 (5)	Each RA within a carrier's network SHALL be uniquely identifiable. RA IDs MAY be associated with a transport plane name space, whereas RC IDs are associated with a control plane name space.	Yes - although uniqueness is the operator's responsibility.	Sections <a href="#">2</a> , <a href="#">3</a> , and <a href="#">11.1</a>
3.2 (6)	Hierarchical Routing Information Dissemination	Yes	<a href="#">Section 7</a>
3.2 (7)	Routing Information exchanged between levels N and N+1 via separate instances and import/export.	Yes	<a href="#">Section 7.1</a>





3.2 (8)	Routing Information exchanged between levels N and N+1 via external link (inter-RA links).	No - Not described.	
3.2 (9)	Routing information exchange MUST include reachability information and MAY include, upon policy decision, node and link topology.	Yes	Sections <a href="#">4</a> , <a href="#">6</a> , <a href="#">6.1</a> , <a href="#">6.2</a> , and <a href="#">8</a>
3.2 (10)	There SHOULD NOT be any dependencies on the different routing protocols used within an RA or in different RAs.	Yes - separate instances.	Sections <a href="#">2</a> and <a href="#">3</a>
3.2 (11)	The routing protocol SHALL differentiate the routing information originated at a given-level RA from derived routing information (received from external RAs), even when this information is forwarded by another RC at the same level.	Yes	<a href="#">Section 7.2</a>
3.2 (12)	The routing protocol MUST provide a mechanism to prevent information propagated from a Level N+1 RA's RC into the Level N RA's RC from being re-introduced into the Level N+1 RA's RC.	Yes	<a href="#">Section 7.2</a>
3.2 (13)	The routing protocol MUST provide a mechanism to prevent information propagated from a Level N-1 RA's RC into the Level N RA's RC from being re-introduced into the Level N-1 RA's RC.	Yes	<a href="#">Section 7.2</a>



3.2 (14)	Instance of a Level N routing function and an instance of a Level N+1 routing function in the same system.	Yes	Sections <a href="#">2</a> , <a href="#">3</a> , and <a href="#">7</a>
3.2 (15)	The Level N routing function is on a separate system the Level N+1 routing function.	Not described but possible.	N/A
3.3 (16)	The RC MUST support static (i.e., operator assisted) and MAY support automated configuration of the information describing its relationship to its parent and its child within the hierarchical structure (including RA ID and RC ID).	Yes - automation requirement is ambiguous.	Sections <a href="#">2</a> and <a href="#">3</a> . Config is product specific.
3.3 (17)	The RC MUST support static (i.e., operator assisted) and MAY support automated configuration of the information describing its associated adjacencies to other RCs within an RA.	Yes - when OSPF area maps to RA discovery is automatic.	<a href="#">RFC 2328</a> and <a href="#">Section 11.1</a>
3.3 (18)	The routing protocol SHOULD support all the types of RC adjacencies described in Section 9 of [ <a href="#">G.7715</a> ]. The latter includes congruent topology (with distributed RC) and hubbed topology (e.g., note that the latter does not automatically imply a designated RC).	Yes	<a href="#">RFC 2328</a>



3.4 (19)	The routing protocol SHOULD be capable of supporting architectural evolution in terms of the number of hierarchical levels of RAs, as well as the aggregation and segmentation of RAs.	Yes	<a href="#">RFC 2328</a> , RFC 5250, and <a href="#">Section 11.2</a> .
3.5.2 (20)	Advertisements MAY contain the following common set of information regardless of whether they are link or node related: - RA ID of the RA to which the advertisement is bounded - RC ID of the entity generating the advertisement - Information to uniquely identify advertisements - Information to determine whether an advertisement has been updated - Information to indicate when an advertisement has been derived from a different level RA	Yes  Yes  Yes  No - Must compare to old  Yes	<a href="#">Section 7.2.1</a>  <a href="#">RFC 2328</a>  <a href="#">RFC 2328</a> , RFC 5250  <a href="#">Section 7.2.1</a>
3.5.3 (21)	The Node Attributes Node ID and Reachability must be advertised. It MAY be advertised as a set of associated external (e.g., User Network Interface (UNI)) address/address prefixes or a set of associated SNPP link IDs/SNPP ID prefixes, the selection of which MUST be consistent within the applicable scope.	Yes - Prefixes only for reachability	<a href="#">RFC 5786</a> , <a href="#">Section 4</a> and 6



3.5.4 (22)	The Link Attributes Local SNPP link ID, Remote SNPP link ID, and layer specific characteristics must be advertised.	Yes	<a href="#">Section 6.1</a>
3.5.4 (23)	Link Signaling Attributes other than Local Adaptation (Signal Type, Link Weight, Resource Class, Local Connection Types, Link Capacity, Link Availability, Diversity Support)	Yes	<a href="#">Section 5</a> , <a href="#">RFC 4652</a> - <a href="#">Section 5.3.1</a>
3.5.4 (24)	Link Signaling Local Adaptation	Yes	<a href="#">Section 5.1</a>
5 (25)	The routing adjacency topology (i.e., the associated PC connectivity topology) and the transport network topology SHALL NOT be assumed to be congruent.	Yes	<a href="#">Section 2</a> , 3, and 6
5 (26)	The routing topology SHALL support multiple links between nodes and RAs.	Yes	<a href="#">RFC 2328</a> , RFC 3630
5 (27)	The routing protocol SHALL converge such that the distributed RDBs become synchronized after a period of time.	Yes	<a href="#">RFC 2328</a> , RFC 5250
5 (28)	Self-consistent information at the receiving level resulting from any transformation (filter, summarize, etc.) and forwarding of information from one RC to RC(s) at different levels when multiple RCs are bound to a single RA.	Yes - However, this is not a routing protocol function.	<a href="#">Section 7.1</a>





5 (29)	In order to support operator-assisted changes in the containment relationships of RAs, the routing protocol SHALL support evolution in terms of the number of hierarchical levels of RAs. For example: support of non-disruptive operations such as adding and removing RAs at the top/bottom of the hierarchy, adding or removing a hierarchical level of RAs in or from the middle of the hierarchy, as well as aggregation and segmentation of RAs.	Partial - OSPF supports the purging of stale advertisements and origination of new. The non-disruptive behavior is implementation specific.	<a href="#">RFC 2328</a> and <a href="#">RFC 5250</a>
5 (30)	A collection of links and nodes such as a subnetwork or RA MUST be able to represent itself to the wider network as a single logical entity with only its external links visible to the topology database.	Yes - Within an RA it must be consistent.	Sections <a href="#">4</a> and <a href="#">6</a>

## [13. References](#)

### [13.1. Normative References](#)

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), September 2003.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", [RFC 3945](#), October 2004.
- [RFC4202] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol



Label Switching (GMPLS)", [RFC 4202](#), October 2005.

- [RFC4203] Kompella, K., Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4203](#), October 2005.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", [RFC 5250](#), July 2008.
- [RFC5786] Aggarwal, R. and K. Kompella, "Advertising a Router's Local Addresses in OSPF TE Extensions", [RFC 5786](#), March 2010.

### **[13.2.](#) Informative References**

- [RFC2154] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", [RFC 2154](#), June 1997.
- [RFC4258] Brungard, D., Ed., "Requirements for Generalized Multi-Protocol Label Switching (GMPLS) Routing for the Automatically Switched Optical Network (ASON)", [RFC 4258](#), November 2005.
- [RFC4652] Papadimitriou, D., Ed., Ong, L., Sadler, J., Shew, S., and D. Ward, "Evaluation of Existing Routing Protocols against Automatic Switched Optical Network (ASON) Routing Requirements", [RFC 4652](#), October 2006.
- [RFC5709] Bhatia, M., Manral, V., Fanto, M., White, R., Barnes, M., Li, T., and R. Atkinson, "OSPFv2 HMAC-SHA Cryptographic Authentication", [RFC 5709](#), October 2009.

For information on the availability of ITU Documents, please see <http://www.itu.int>.

- [G.7715] ITU-T Rec. G.7715/Y.1306, "Architecture and Requirements for the Automatically Switched Optical Network (ASON)", June 2002.
- [G.7715.1] ITU-T Draft Rec. G.7715.1/Y.1706.1, "ASON Routing Architecture and Requirements for Link State Protocols", November 2003.
- [G.805] ITU-T Rec. G.805, "Generic functional architecture of transport networks)", March 2000.



[G.8080] ITU-T Rec. G.8080/Y.1304, "Architecture for the Automatically Switched Optical Network (ASON)," 2006 (and Amendments 1 and 2).

#### **14. Acknowledgements**

The editors would like to thank Dimitri Papadimitriou for editing [RFC 5787](#), from which this document is derived, and Lyndon Ong, Remi Theillaud, Stephen Shew, Jonathan Sadler, Deborah Brungard, and Lou Berger for their useful comments and suggestions.

## **Appendix A. ASON Terminology**

This document makes use of the following terms:

**Administrative domain:** (See Recommendation [G.805].) For the purposes of [G7715.1], an administrative domain represents the extent of resources that belong to a single player such as a network operator, a service provider, or an end-user. Administrative domains of different players do not overlap amongst themselves.

**Control plane:** performs the call control and connection control functions. Through signaling, the control plane sets up and releases connections, and may restore a connection in case of a failure.

**(Control) Domain:** represents a collection of (control) entities that are grouped for a particular purpose. The control plane is subdivided into domains matching administrative domains. Within an administrative domain, further subdivisions of the control plane are recursively applied. A routing control domain is an abstract entity that hides the details of the RC distribution.

**External NNI (E-NNI):** interfaces located between protocol controllers between control domains.

**Internal NNI (I-NNI):** interfaces located between protocol controllers within control domains.

**Link:** (See Recommendation G.805.) A "topological component" that describes a fixed relationship between a "subnetwork" or "access group" and another "subnetwork" or "access group". Links are not limited to being provided by a single server trail.

**Management plane:** performs management functions for the transport plane, the control plane, and the system as a whole. It also provides coordination between all the planes. The following management functional areas are performed in the management plane: performance, fault, configuration, accounting, and security management.

**Management domain:** (See Recommendation G.805.) A management domain defines a collection of managed objects that are grouped to meet organizational requirements according to geography, technology, policy, or other structure, and for a number of functional areas such as configuration, security, (FCAPS), for the purpose of providing control in a consistent manner. Management domains can be disjoint, contained, or overlapping. As such, the resources



within an administrative domain can be distributed into several possible overlapping management domains. The same resource can therefore belong to several management domains simultaneously, but a management domain shall not cross the border of an administrative domain.

**Subnetwork Point (SNP):** The SNP is a control plane abstraction that represents an actual or potential transport plane resource. SNPs (in different subnetwork partitions) may represent the same transport resource. A one-to-one correspondence should not be assumed.

**Subnetwork Point Pool (SNPP):** A set of SNPs that are grouped together for the purposes of routing.

**Termination Connection Point (TCP):** A TCP represents the output of a Trail Termination function or the input to a Trail Termination Sink function.

**Transport plane:** provides bidirectional or unidirectional transfer of user information, from one location to another. It can also provide transfer of some control and network management information. The transport plane is layered; it is equivalent to the Transport Network defined in Recommendation G.805.

**User Network Interface (UNI):** interfaces are located between protocol controllers between a user and a control domain. Note: There is no routing function associated with a UNI reference point.

## **Appendix B. ASON Routing Terminology**

This document makes use of the following terms:

**Routing Area (RA):** an RA represents a partition of the data plane, and its identifier is used within the control plane as the representation of this partition. Per [[G.8080](#)], an RA is defined by a set of sub-networks, the links that interconnect them, and the interfaces representing the ends of the links exiting that RA. An RA may contain smaller RAs inter-connected by links. The limit of subdivision results in an RA that contains two sub-networks interconnected by a single link.

**Routing Database (RDB):** a repository for the local topology, network topology, reachability, and other routing information that is updated as part of the routing information exchange and may additionally contain information that is configured. The RDB may contain routing information for more than one routing area (RA).





Routing Components: ASON routing architecture functions. These functions can be classified as protocol independent (Link Resource Manager or LRM, Routing Controller or RC) or protocol specific (Protocol Controller or PC).

Routing Controller (RC): handles (abstract) information needed for routing and the routing information exchange with peering RCs by operating on the RDB. The RC has access to a view of the RDB. The RC is protocol independent.

Note: Since the RDB may contain routing information pertaining to multiple RAs (and possibly to multiple layer networks), the RCs accessing the RDB may share the routing information.

Link Resource Manager (LRM): supplies all the relevant component and TE link information to the RC. It informs the RC about any state changes of the link resources it controls.

Protocol Controller (PC): handles protocol-specific message exchanges according to the reference point over which the information is exchanged (e.g., E-NNI, I-NNI), and internal exchanges with the RC. The PC function is protocol dependent.

#### Authors' Addresses

Andrew G. Malis  
Verizon Communications  
117 West St.  
Waltham MA 02451 USA

EMail: [andrew.g.malis@verizon.com](mailto:andrew.g.malis@verizon.com)

Acee Lindem  
Ericsson  
102 Carric Bend Court  
Cary, NC 27519

EMail: [acee.lindem@ericsson.com](mailto:acee.lindem@ericsson.com)