Extending OSPF Routing Protocol for Shared Mesh Restoration

H. Liu, E. Bouillet, D. Pendarakis, N. Komaee, J. Labourdette, and S. Chaudhuri
liu1999@ieee.org

(Extended Abstract)

1. Introduction

Recently, distributed control architectures have been proposed for optical mesh networks as a means to automate operations, enhance interoperability and scalability as well as facilitate the deployment of new applications [1]. With the application of distributed control plane, a connection request can be sent to the ingress node either from the client directly connected to the optical network or by the management plane. The ingress node computes the explicit path and then initiates the path establishment with signaling messages. Path computation must take into account various requirements and constraints, including bandwidth, recovery and survivability, optimal utilization of network resources. This requires dissemination of information about the network topology and various link attributes to every node in the network using routing protocols.

End-to-end path protection and restoration techniques are commonly used in optical networks in order to support high service availability and quick recovery from network failure. In shared mesh restoration [2, 3], the restoration path is pre-computed and its resource is reserved along the path through signaling protocols [4]. However no cross-connections are performed along the restoration path. The complete establishment of the restoration path occurs only after the working path fails, and requires some additional signaling. Multiple restoration paths can share the same reserved network resource if their corresponding working paths are mutually diverse. This scheme achieves efficient utilization of network resource by sharing the restoration resource. It can perform reasonably fast switching time and guarantees successful recovery from a single failure. The resource reserved for restoration can even be used for low priority pre-emptible traffic in normal network operating mode (i.e. while there are no failure along the working paths). Of course, the restoration path needs to be activated via signaling when the working path fails. It may result in a recovery slower than the dedicated mesh protection [2].

Efficient methods for aggregation and dissemination of link resource availability and sharing information are required for routing shared mesh-restored paths. This paper describes extensions to the OSPF routing protocol in support of path computation for shared mesh restoration [5]. It is based on the OSPF routing extensions required to support Traffic Engineering (TE) and Generalized Multi-Protocol Label Switching (GMPLS) [6, 7]. New optional sub-TLVs (Type/Length/Value) are added to the link TLV of the TE Link State Advertisements (LSA) so that the sharing information of the restoration resource on the TE link reserved for shared mesh restoration is disseminated.

2. Required Information to Support Path Computation for Shared Mesh Restoration

The term Shared Risk Link Group (SRLG) is used to indicate a group of optical lines that share a common risk of failure. In shared mesh restoration, multiple restoration paths can share the same reserved restoration resource only if the sets of SRLGs traversed by their respective working paths are disjoint in order to guarantee recovery from a single SRLG failure. This imposes additional constraints on the path computation. To compute the restoration path for the shared mesh restoration, the path computation algorithm needs to have the restoration resource sharing information of the links in the network [8]. In general, the more detail information is available, the better results the path computation algorithms can achieve. On the other hand, in order to reduce the amount of information handled by routing protocol and improve the scalability, it is desirable to aggregate the routing information. For optical networks there may be multiple data links between a pair of nodes. Data links between the same pair of nodes, with similar characteristics, can then be bundled together and advertised as a single link bundle or a TE link into the routing protocol [9]. To support path computation for shared mesh restoration, all or some information below should be disseminated by routing protocol.

(1) Summarized information about the restoration resource sharing on a TE link for shared mesh restoration, including the total number of restoration paths sharing the restoration resource reserved on the TE link, the total number of SRLGs protected by the reserved restoration resource, and the total sharable restoration bandwidth.

(2) The list of SRLGs protected by the reserved restoration resource on the TE link and their respective sharable restoration bandwidth.

The list of SRLGs protected by the TE link is defined as the union of SRLGs traversed by all the working paths whose respective restoration paths share the reserved restoration resource on this TE link [10]. The sharable restoration bandwidth for a SRLG indicates the available restoration bandwidth on the TE link that can be reserved for recovering this SRLG failure. If a working path only traverses one SRLG, the available restoration bandwidth that its restoration path can share on this TE link is the sharable restoration bandwidth for this SRLG. When a working path traverses multiple SRLGs, the sharable restoration bandwidth available for its restoration path may become smaller on this TE link. The total sharable restoration bandwidth is the bandwidth reserved on the TE link for restoration, which is the union of the sharable restoration bandwidth for all SRLGs.

3. OSPF Extensions

OSPF traffic engineering extensions and GMPLS extensions make use of the Opaque LSA [11]. An Opaque LSA, called Traffic Engineering LSA has been defined to carry the additional attributes related to traffic engineering and GMPLS links, and standard link-state database flooding mechanisms are used for distributing TE LSAs. The LSA payload consists of one or more nested Type/Length/Value (TLV) triplets for extensibility. A link TLV is constructed of a set of sub-TLVs that specify the link attributes, including link type, traffic engineering metric, available resource, administrative group, local and remote link identifiers, link protection type, and interface switching capability descriptor. The information in the TE LSAs is used to build an extended TE link state database for the explicit path computation just as the router LSAs are used to build a regular link state database for packet forwarding.

The extensions in support of carrying link state information for the path computation of shared mesh restoration can be based upon the OSPF-TE and GMPLS extensions. The sub-TLVs carrying the sharing information of the restoration resource on a TE link can be added to the link TLV of the TE LSA so that the information can be used by the path computation algorithm to compute the restoration path. Specifically, we defined the two new sub-TLVs, Restoration Information Summary sub-TLV and SRLG Sharable Restoration Bandwidth sub-TLV.

Restoration Information Summary sub-TLV specifies the sharing information of the restoration resource reserved for the shared mesh restoration on the TE link. As shown in Fig. 1, it contains a 16-bit Type field to distinguish it from other sub-TLVs and a 16-bit Length field to specify the length of this sub-TLV, which follows the same format as other TLVs. The field of #Shared Restoration Path (16 bits) specifies the number of restoration paths sharing the bandwidth reserved for the mesh restoration on this TE link. The field of #SRLGs Restored (16 bits) indicates the number of SRLGs protected by the TE link. The SRLGs here includes all distinct SRLGs traversed by all the working paths whose respective restoration paths share the reserved restoration bandwidth on this TE link. Similarly, #Nodes Restored (16 bits) indicates the number of nodes protected by the TE link. Total Sharable Restoration Bandwidth specifies the bandwidth that has been allocated for shared mesh restoration at each of the eight priority levels. The bandwidth might have been reserved by one or more shared restoration path.

SRLG Sharable Restoration Bandwidth sub-TLV identifies the sharable restoration bandwidth range for a protected SRLG on this TE link. By specifying a range instead of a fixed value, it may reduce the amount of information handled by the routing protocol. Its format is shown in Fig. 2. Besides the standard Type and Length fields, it contains a Priority field (8 bits) that indicates the priority of sharable restoration bandwidth. The SRLG sharable restoration bandwidth can be encoded per priority. If the SRLG sharable restoration bandwidth is not encoded per priority, the value of the priority field can be set to be 0xFF, which means the sharable restoration bandwidth is the same for all priorities. Sharable Restoration Bandwidth Lower Bound and Upper Bound fields (32 bits) specify the range for sharable restoration bandwidth. The sharable restoration bandwidth for each of SRLGs listed in the SRLG fields falls into this range, i.e. Sharable Restoration Bandwidth Lower Bound <= sharable restoration bandwidth for SRLG N < Sharable Restoration Bandwidth Upper Bound. If the lower bound is equal to the upper bound, a fixed value of sharable restoration bandwidth is specified and the sharable restoration bandwidth for each of SRLGs in the following list is equal to this value.

The Restoration Information Summary sub-TLV and the SRLG Sharable Restoration Bandwidth sub-TLV are optional. If a TE LSA doesn't carry them, it means that the information is unknown. There may be more than one

SRLG Sharable Restoration Bandwidth sub-TLVs in the TE LSA. However if there are one or more SRLG Sharable Restoration Bandwidth sub-TLVs in the TE LSA, the sharable restoration bandwidth for the SRLGs not listed in these sub-TLVs is assumed to be equal to the total sharable restoration bandwidth at that priority on this TE link.

<table>
<tr><td>Type</td><td>Length</td></tr>
<tr><td>#Shared Restoration Paths</td><td>#SRLGs Restored</td></tr>
<tr><td>#Nodes Restored</td><td>Reserved</td></tr>
<tr><td colspan="2">Total Sharable Restoration Bandwidth at Priority 0</td></tr>
<tr><td colspan="2">Total Sharable Restoration Bandwidth at Priority 1</td></tr>
<tr><td colspan="2">Total Sharable Restoration Bandwidth at Priority 2</td></tr>
<tr><td colspan="2">Total Sharable Restoration Bandwidth at Priority 3</td></tr>
<tr><td colspan="2">Total Sharable Restoration Bandwidth at Priority 4</td></tr>
<tr><td colspan="2">Total Sharable Restoration Bandwidth at Priority 5</td></tr>
<tr><td colspan="2">Total Sharable Restoration Bandwidth at Priority 6</td></tr>
<tr><td colspan="2">Total Sharable Restoration Bandwidth at Priority 7</td></tr>
</table>

Figure 1. The format of Restoration Information Summary sub-TLV

<table>
<tr><td colspan="2">Type</td><td>Length</td></tr>
<tr><td>Priority</td><td colspan="2">Reserved</td></tr>
<tr><td colspan="3">Sharable Restoration Bandwidth Lower Bound</td></tr>
<tr><td colspan="3">Sharable Restoration Bandwidth Upper Bound</td></tr>
<tr><td colspan="3">SRLG 1</td></tr>
<tr><td colspan="3">…</td></tr>
<tr><td colspan="3">SRLG N</td></tr>
</table>

Figure 2. The format of SRLG Sharable Restoration Bandwidth sub-TLV

4. Conclusions

Efficient methods are needed to aggregate and disseminate the routing information including optical link resource availability and sharing so that the amount advertised by the routing protocol is minimized but information necessary for path computation is not lost. We proposed to extend the GMPLS OSPF-TE routing protocol to carry necessary sharing information of the reserved resource on a TE link in support of computing shared mesh restored paths.

5. References

[1] ITU-T Recommendation G.8080, “Architecture for the ASON,” 2001.
[2] B. Doshi, et al., “Optical Network Design and Restoration,” Bell-labs Technical Journal, January-March 1999.
[3] S. Datta, et al., “Efficient Channel Reservation for Backup Paths in Optical Mesh Networks,” IEEE GLOBECOM 2001, San Antonio, TX, November 2001.
[4] L. Berger, et al., "Generalized MPLS Signaling - RSVP-TE Extensions," RFC 3473, January 2003.
[5] J. Moy, "OSPF Version 2," RFC 2328, April 1998.
[6] D. Katz, et al., “Traffic Engineering Extensions to OSPF Version 2,” Internet Draft, Work in Progress, October 2002.
[7] K. Kompella, et al., “OSPF Extensions in Support of Generalized MPLS,” Internet Draft, Work in Progress, December 2002.
[8] E. Bouillet, et al., "Stochastic Approaches to Route Shared Mesh Restored Lightpaths in Optical Mesh Networks," Proc. of IEEE Infocom 2002, New York, NY, June 2002.
[9] K. Kompella, et al., “Link Bundling in MPLS Traffic Engineering,” Internet Draft, Work in Progress, January 2003.
[10] H. Liu, et al., “OSPF-TE Extensions in Support of Shared Mesh Restoration,” Internet Draft, Work in Progress, February 2003.
[11] R. Coltun, “The OSPF Opaque LSA Option,” RFC 2370, 1999.

Extending OSPF Routing Protocol for Shared Mesh Restoration

H. Liu, E. Bouillet, D. Pendarakis, N. Komaee,
J. Labourdette, and S. Chaudhuri

Traditional Management Plane for Optical Transport Networks

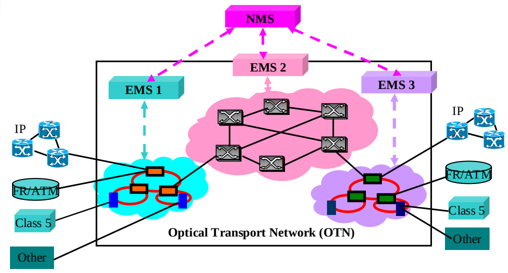

- A lot of manual operations
- Integration of different EMS and NMS is complex
  - multiple types of equipment from different vendors with different technologies
- Automatic end-to-end provisioning is not easy
  - planning, path computation, connection establishment

Distributed Control Plane

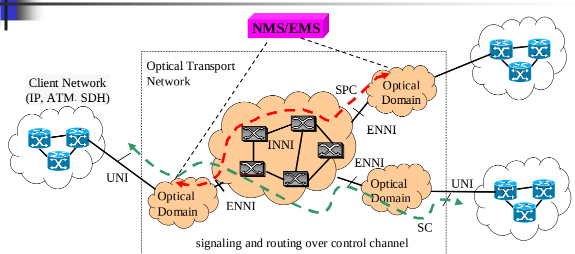

- Distributed control plane offers
  - automatic neighbor and topology discovery
  - automatic end-to-end provisioning and connection modification
  - scalability and interoperability
  - unified traffic engineering and protection/restoration
    - for example, in an environment where IP router networks are interconnected via a mesh optical network

Connection Provisioning through Distributed Control Plane

- Connection provisioning
  - Connection request received from a client or a management agent at ingress node
  - Ingress node computes the explicit path from ingress to egress by taking into account a set of constraints (bandwidth, resource availability, protection and restoration, and traffic engineering)
    - ***Require routing protocol to disseminate network topology and link state information***
  - Signaling the connection establishment along the path
- Standardization Activities
  - IETF: Generalized Multi-Protocol Label Switching (GMPLS) Protocols
    - Extends MPLS/IP protocols based on generalized interface requirements
      - signaling (RSVP-TE and CR-LDP with GMPLS extensions)
      - routing (OSPF-TE and IS-IS with GMPLS extensions)
      - link management and neighbor discovery (LMP)
  - ITU: Automatically Switched Optical Network (ASON)
    - Architecture and components including UNI and NNI interfaces
  - OIF: Focuses on application of IETF protocols in an overlay model and generates implementation agreements
    - Optical UNI 1.0 (approved in October 2001) and NNI (in progress)

Shared Mesh Restoration

- Optical transport networks need fast end-to-end protection and restoration
- Shared mesh restoration is very efficient in term of network resource utilization
  - restoration path is pre-computed and pre-provisioned with signaling protocol
  - resource is reserved on the links but no cross-connections are created along the restoration path
  - the complete establishment of the restoration path occurs only after the primary path fails
  - the common restoration resource reserved on a link may be shared by multiple restoration paths to restore multiple primary paths
  - the resource reserved for restoration can be used for low priority pre-emptible traffic in normal operating mode

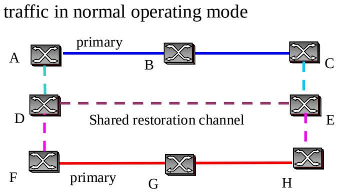

Routing Protocol

- Disseminates network topology and link resource availability over control channel
- Manages the link state database and routing tables
  - makes routing decision
- Provides path computation algorithm with the routing information to obtain explicit route
- Traffic Engineering (TE) and GMPLS routing extensions
  - support multiple types of GMPLS Traffic Engineering links
  - extends OSPF and IS-IS to carry new link attributes
    - link type, traffic engineering metric, available resource, administrative group, local and remote link identifiers or interface IP addresses, link protection type, interface switching capability descriptor, shared risk link group (SRLG), etc.
- draft-katz-yeung-ospf-traffic-09.txt, draft-ietf-ccamp-gmpls-routing-05.txt, draft-ietf-ccamp-ospf-gmpls-extensions-09.txt

Link Bundling

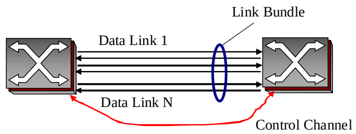

- Neighboring nodes (e.g. OXCs) connected by multiple parallel links
  - For standard OSPF, each physical link between a pair of nodes forms a routing adjacency
    - not scale well
- To improve routing scalability and reduce the amount of information handled by routing protocol, in GMPLS routing protocol
  - aggregates and abstracts the attributes of the links with similar characteristics between a pair of nodes
  - advertises as a single link bundle or Traffic Engineering (TE) link
- Control channel and data link may be separated

OSPF Extensions to Support Shared Mesh Restoration

- Link resource sharing information is required to compute shared mesh restored paths.
  - In order to avoid contention for the reserved restoration resource during a single SRLG failure
    - two restoration paths may share the common reserved restoration resource ***only*** if their respective working paths are mutually SRLG disjoint.
    - One failure then does not disrupt both working paths simultaneously.
- Standard GMPLS OSPF-TE routing protocol does not carry enough link sharing information
- Extensions to the GMPLS OSPF-TE protocol are necessary in support of path computation for shared mesh restored paths.
  - Define two optional sub-TLVs (Type/Length/Value) in the Traffic Engineering Link State Advertisements (LSA)
    - Restoration Information Summary sub-TLV
    - SRLG Sharable Restoration Bandwidth sub-TLV

OSPF Extensions to Support Shared Mesh Restoration

- Restoration Information Summary Sub-TLV

Type	Length
#Shared Restoration Paths	#SRLGs Restored
#Nodes Restored	Reserved
Total Sharable Restoration Bandwidth at Priority 0	
Total Sharable Restoration Bandwidth at Priority 1	
Total Sharable Restoration Bandwidth at Priority 2	
Total Sharable Restoration Bandwidth at Priority 3	
Total Sharable Restoration Bandwidth at Priority 4	
Total Sharable Restoration Bandwidth at Priority 5	
Total Sharable Restoration Bandwidth at Priority 6	
Total Sharable Restoration Bandwidth at Priority 7	

- Note:
  - #Shared Restoration Paths: the total number of restoration paths sharing the restoration resource reserved on the TE link
  - #SRLGs Restored: the total number of SRLGs protected by the reserved restoration resource

OSPF Extensions to Support Shared Mesh Restoration

- SRLG Sharable Restoration Bandwidth Sub-TLV

Type		Length
Priority	Reserved	
Sharable Restoration Bandwidth Lower Bound		
Sharable Restoration Bandwidth Upper Bound		
SRLG 1		
…		
SRLG N		

- Note:
  - The SRLG sharable restoration bandwidth is encoded per priority.
    - Set to be 0xFF if the sharable restoration bandwidth is the same for all priorities.
  - Sharable Restoration Bandwidth Lower Bound and Upper Bound fields specify the range for sharable restoration bandwidth.
    - Sharable Restoration Bandwidth Lower Bound <= sharable restoration bandwidth for SRLG N < Sharable Restoration Bandwidth Upper Bound.
    - If the lower bound is equal to the upper bound, a fixed value of sharable restoration bandwidth is specified.
      - The sharable restoration bandwidth for each of SRLGs in the list is equal to this value.

Performance Results

- Bandwidth usage for router LSAs and TE LSAs with different abstraction levels of restoration resource information in various sizes of networks

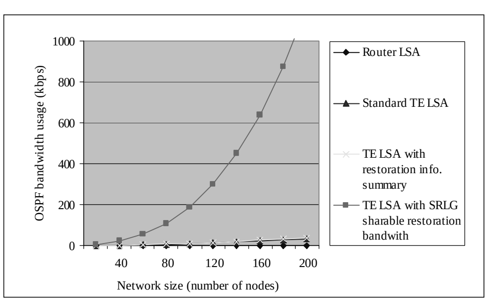

Conclusions

- Efficient methods are required to aggregate and disseminate the routing information
- Optical link resource sharing information is needed to compute shared mesh restored path
- We proposed to extend the GMPLS OSPF-TE routing protocol to carry necessary sharing information of the reserved resource on a TE link in support of computing shared mesh restored paths.
  - Define two optional sub-TLVs in the Traffic Engineering Link State Advertisements
    - Restoration Information Summary sub-TLV
    - SRLG Sharable Restoration Bandwidth sub-TLV