

Link Bundling in MPLS Traffic Engineering

[draft-ietf-mpls-bundle-01.txt](#)

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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

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

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

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

2. Abstract

For the purpose of GMPLS signaling in certain cases a combination of <link identifier, label> is not sufficient to unambiguously identify the appropriate resource used by an LSP. Such cases are handled by using the link bundling construct which is described in this document.

3. Link Bundling

As defined in [GMPLS-ROUTING], a TE link is a logical construct that represents a way to group/map the information about certain physical resources (and their properties) that interconnect LSRs into the information that is used by Constrained SPF for the purpose of path computation, and by GMPLS signaling.

As further stated in [GMPLS-ROUTING], depending on the nature of resources that form a particular TE link, for the purpose of GMPLS signaling in some cases a combination of <link identifier, label> is sufficient to unambiguously identify the appropriate resource used by an LSP. In other cases, a combination of <link identifier, label> is not sufficient. Such cases are handled by using the link bundling construct which is described in this document.

Consider a TE link such that for the purpose of GMPLS signaling a combination of <link identifier, label> is not sufficient to unambiguously identify the appropriate resources used by an LSP. In this situation the link bundling construct assumes that the set of resources that form the TE link could be partitioned into disjoint subsets, such that (a) the partition is minimal, and (b) within each subset a label is sufficient to unambiguously identify the appropriate resources used by an LSP. We refer to such subsets as "component links", and to the whole TE link as a "bundled link". On a bundled link a combination of <(bundled) link identifier, component link identifier, label> is sufficient to unambiguously identify the appropriate resources used by an LSP.

Since within each component link a label is sufficient to unambiguously identify the resources used by an LSP, one could also say that a component link is a TE link, and a bundled link is a collection of TE links.

The partition of resources that form a bundled link into component links has to be done consistently at both ends of the bundled link.

The purpose of link bundling is to improve routing scalability by reducing the amount of information that has to be handled by OSPF and/or IS-IS. This reduction is accomplished by performing information aggregation/abstraction. As with any other information aggregation/abstraction, this results in losing some of the information. To limit the amount of losses one needs to restrict the type of the information that can be aggregated/abstracted.

[3.1. Restrictions on Bundling](#)

All component links in a bundle must begin and end on the same pair of LSRs, have the same Link Type (i.e., point-to-point or multi-access), the same Traffic Engineering metric, and the same set of resource classes at each end of the links.

A Forwarding Adjacency may be a component link; in fact, a bundle can consist of a mix of point-to-point links and FAs.

If the component links are all multi-access links, the set of IS-IS or OSPF routers connected to each component link must be the same, and the Designated Router for each component link must be the same. If these conditions cannot be enforced, multi-access links must not be bundled.

[3.2. Routing Considerations](#)

A component link may be either numbered or unnumbered. A bundled link may itself be numbered or unnumbered independent of whether the component links of that bundled link are numbered or not.

Handling identifiers for unnumbered component links, including the case where a link is formed by a Forwarding Adjacency, follows the same rules as for an unnumbered TE link (see [Section 4](#) of [RSVP-UNNUM]/[CRLDP-UNNUM]). Furthermore, link local identifiers for all unnumbered links of a given LSR (whether component links, Forwarding Adjacencies or bundled links) MUST be unique in the context of that LSR.

The "liveness" of the bundled link is determined by the liveness of each of the component links within the bundled link - a bundled link is alive when at least one its component links is determined to be alive. The liveness of a component link can be determined by any of several means: IS-IS or OSPF hellos over the component link, or RSVP Hello, or LMP hellos (see [[LMP](#)]), or from layer 1 or layer 2 indications.

Once a bundled link is determined to be alive, it can be advertised as a TE link and the TE information can be flooded. If IS-IS/OSPF hellos are run over the component links, IS-IS/OSPF flooding can be restricted to just one of the component links [[ZININ](#)] [[MOY](#)].

In the future, as new Traffic Engineering parameters are added to IS-IS and OSPF, they should be accompanied by descriptions as to how they can be bundled, and possible restrictions on bundling.

3.3. Signaling Considerations

Typically, an LSP's ERO will choose the bundled link to be used for the LSP, but not the component link, since information about the bundled link is flooded, but information about the component links is not. If the ERO chooses the component link by means outside the scope of this document, this section does not apply. Otherwise, the choice of the component link for the LSP is a local matter between the two LSRs at each end of the bundled link.

Signaling must identify both the component link to use and the label to use. The choice of the component link to use is always made by the sender of the Path/REQUEST message (if an LSP is bidirectional [[GMPLS-SIG](#)], the sender chooses a component link in each direction). For unidirectional LSPs, and the forward direction of bidirectional LSPs, the sender of a Resv/MAPPING message chooses the label. For the reverse direction of a bidirectional LSP, the sender of the Path/REQUEST message selects the upstream label.

With RSVP the choice of the component link is indicated by the sender of the Path message by including the IF_ID RSVP_HOP object in the Path message, as described in section 8 of [[GMPLS-RSVP](#)]. With CR-LDP the choice of the component link is indicated by the sender of the REQUEST message by including the IF_ID TLV in the REQUEST message, as described in section 8 of [[GMPLS-CRLDP](#)].

If the component link is numbered, the IF_ID RSVP_HOP object, or IF_ID TLV carries either Type 1 (IPv4 address) or Type 2 (IPv6 address) TLVs (see [[GMPLS-SIG](#)]). The address carried in the TLV identifies the link for which label allocation must be done.

If the component link is unnumbered, the IF_ID RSVP_HOP object, or IF_ID TLV carries Type 3 (IF_INDEX) TLV (see [[GMPLS-SIG](#)]). The value carried in Type 3 TLV contains the identifier of the selected component link assigned to the link by the sender of the Path/REQUEST message. Processing this object is the same as specified in [Section 6.1](#) of [RSVP-UNNUM]/[CRLDP-UNNUM].

For the purpose of processing the IF_ID RSVP_HOP object or IF_ID TLV, an unnumbered component link formed by a Forwarding Adjacency is treated the same way as an unnumbered TE link formed by a Forwarding Adjacency (see [Section 5](#) of [RSVP-UNNUM]/[CDLDP-UNNUM]).

[4. Traffic Engineering Parameters for Bundled Links](#)

In this section, we define the Traffic Engineering parameters to be advertised for a bundled link, based on the configuration of the component links and of the bundled link. The definition of these parameters for component links was undertaken in [[ISIS-TE](#)] and [[OSPF-TE](#)]; we use the terminology from [[OSPF-TE](#)].

[4.1. OSPF Link Type](#)

The Link Type of a bundled link is the (unique) Link Type of the component links. (Note: this parameter is not present in IS-IS.)

[4.2. OSPF Link ID](#)

For point-to-point links, the Link ID of a bundled link is the (unique) Router ID of the neighbor. For multi-access links, this is the interface address of the (unique) Designated Router. (Note: this parameter is not present in IS-IS.)

[4.3. Local and Remote Interface IP Address](#)

(Note: in IS-IS, these are known as IPv4 Interface Address and IPv4 Neighbor Address, respectively.)

If the bundled link is numbered, the Local Interface IP Address is the local address of the bundled link; similarly, the Remote Interface IP Address is the remote address of the bundled link.

[4.4. Local and Remote Identifiers](#)

If the bundled link is unnumbered, the link local identifier is set to the identifier chosen for the bundle by the advertising LSR. The link remote identifier is set to the identifier chosen by the neighboring LSR for the reverse link corresponding to this bundle, if known; otherwise, this is set to 0.

[4.5. Traffic Engineering Metric](#)

The Traffic Engineering Metric for a bundled link is that of the component links.

[4.6.](#) Maximum Link Bandwidth

This TLV is not used. The maximum LSP Bandwidth (as described below) replaces the maximum link bandwidth for bundled links.

[4.7.](#) Total Reservable Bandwidth

We assume that for a given bundled link either each of its component links is configured with the Total Reservable Bandwidth, or the bundled link is configured with the Total Reservable Bandwidth. In the former case, the Total Reservable Bandwidth of the bundled link is set to the sum of the Total Reservable Bandwidths of all component links associated with the bundled link.

[4.8.](#) Unreserved Bandwidth

The unreserved bandwidth of a bundled link at priority *p* is the sum of the unreserved bandwidths at priority *p* of all the component links associated with the bundled link.

[4.9.](#) Resource Classes (Administrative Groups)

The Resource Classes for a bundled link are the same as those of the component links.

[4.10.](#) Maximum LSP Bandwidth

The Maximum LSP Bandwidth takes the place of the Maximum Link Bandwidth. For an unbundled link the Maximum Link Bandwidth is defined in [[GMPLS-ROUTING](#)]. The Maximum LSP Bandwidth of a bundled link at priority *p* is defined to be the maximum of the Maximum LSP Bandwidth at priority *p* of all of its component links.

The details of how Maximum LSP Bandwidth is carried in IS-IS is given in [[GMPLS-ISIS](#)]. The details of how Maximum LSP Bandwidth is carried in OSPF is given in [[GMPLS-OSPF](#)].

5. Bandwidth Accounting

The RSVP (or CR-LDP) Traffic Control module, or its equivalent, on an LSR with bundled links must apply admission control on a per-component link basis. An LSP with a bandwidth requirement b and setup priority p fits in a bundled link if at least one component link has maximum LSP bandwidth $\geq b$ at priority p . If there are several such links, the choice of which link is used for the LSP is up to the implementation.

In order to know the maximum LSP bandwidth (per priority) of each component link, the Traffic Control module must track the unreserved bandwidth (per priority) for each component link.

A change in the unreserved bandwidth of a component link results in a change in the unreserved bandwidth of the bundled link. It also potentially results in a change in the maximum LSP bandwidth of the bundle; thus, the maximum LSP bandwidth should be recomputed.

If one of the component links goes down, the associated bundled link remains up and continues to be advertised, provided that at least one component link associated with the bundled link is up. The unreserved bandwidth of the component link that is down is set to zero, and the unreserved bandwidth and maximum LSP bandwidth of the bundle must be recomputed. If all the component links associated with a given bundled link are down, the bundled link MUST not be advertised into OSPF/IS-IS.

6. Security Considerations

This document raises no new security issues for RSVP or CR-LDP.

7. References

[GMPLS-ISIS] Kompella, K., Rekhter, Y., Banerjee, A. et al, "IS-IS Extensions in Support of Generalized MPLS", [draft-ietf-isis-gmpls-extensions-02.txt](#) (work in progress)

[GMPLS-OSPF] Kompella, K., Rekhter, Y., Banerjee, A. et al, "OSPF Extensions in Support of Generalized MPLS", [draft-ietf-ccamp-ospf-gmpls-extensions-00.txt](#) (work in progress)

[GMPLS-ROUTING] Kompella, K., Rekhter, Y., Banerjee, A. et al, "Routing Extensions in Support of Generalized MPLS", [draft-ietf-ccamp-gmpls-routing-00.txt](#)

[GMPLS-SIG] Ashwood, P., et al., "Generalized MPLS - Signalling Functional Description", [draft-ietf-generalized-mpls-signalling-05.txt](#)

[GMPLS-RSVP] Ashwood, P., et al., "Generalized MPLS Signalling RSVP-TE Extensions", [draft-ietf-mpls-generalized-rsvp-te-04.txt](#)

[GMPLS-CRLDP] Ashwood, P., et al., "Generalized MPLS Signaling - CR-LDP Extensions", [draft-ietf-mpls-generalized-cr-ldp-04.txt](#)

[ISIS-TE] Smit, H., Li, T., "IS-IS extensions for Traffic Engineering", [draft-ietf-isis-traffic-02.txt](#) (work in progress)

[LMP] Lang, J., Mitra, K., et al., "Link Management Protocol (LMP)", [draft-ietf-ccamp-lmp-00.txt](#) (work in progress)

[MOY] Moy, J., [draft-ietf-ospf-ppp-flood-00.txt](#) (work in progress)

[OSPF-TE] Katz, D., Yeung, D., "Traffic Engineering Extensions to OSPF", [draft-katz-yeung-ospf-traffic-04.txt](#) (work in progress)

[UNNUM-CRLDP] Kompella, K., Rekhter, Y., Kullberg, A., "Signalling Unnumbered Links in CR-LDP", [draft-ietf-mpls-crldp-unnum-01.txt](#) (work in progress)

[UNNUM-RSVP] Kompella, K., Rekhter, Y., "Signalling Unnumbered Links in RSVP-TE", [draft-ietf-mpls-rsvp-unnum-01.txt](#) (work in progress)

[ZININ] Zinin, A., Shand, M., "Flooding optimizations in link-state routing protocols", [draft-ietf-ospf-isis-flood-opt-00.txt](#) (work in progress)

8. Author Information

Kireeti Kompella
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
Email: kireeti@juniper.net

Yakov Rekhter
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
Email: yakov@juniper.net

Lou Berger
Movaz Networks, Inc.
Voice: +1 301 468 9228
Email: lberger@movaz.com