

Working Group
Internet-Draft
Date Created: June, 2009
Expiration Date: January, 2010
Intended Status: Informational

Peter Ashwood-Smith
Don Fedyk
David Allan
Jerome Chiabaut
Nigel Bragg

Preliminary

Shortest Path Bridging and Backbone Bridging with IS-IS
[draft-fedyk-isis-spb-00.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/ietf/lid-abstracts.txt>.

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

This Internet-Draft will expire in January 6th 2010.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Internet Draft [draft-fedyk-isis-spb-00.txt](#)

July 2009

Abstract

Several techniques are being developed which use IS-IS to deliver link state based layer 2 forwarding. The superset of the extensions proposed to IS-IS to allow these capabilities is found in [IS-IS-L2]. One technique for layer 2 forwarding is being specified in the IEEE 802.1aq task group, under the over-arching title of "Shortest Path Bridging" (SPB). SPB however only requires a subset of the proposed IS-IS extensions in [[IS-IS-L2](#)]. For clarity this informational draft documents only the subset required by SPB. In addition a high level introduction, describing how these TLVs are used is provided for those who do not follow the IEEE work in detail. A reference is also given to the normative IEEE 802.1aq document. The ordering of material in this document follows that of Clause 28 of IEEE 802.1aq, to aid cross-referencing.

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 [[RFC2119](#)].

Table of Contents

1.	Terminology.....	3
2.	Introduction.....	4
3.	New TLVs for SPB and SPBB.....	6
3.1	Base Vlan-Identifiers sub-TLV.....	7
3.2	SPB/SPBB Instance sub-TLV.....	8
3.3	SPB Link Metric sub-TLV.....	9
3.4	The Group MAC Address sub-TLV.....	9
3.5	The Service Identifier and Unicast Address sub-TLV.....	10
4.	Security Considerations.....	10
5.	IANA Considerations / ISO Considerations.....	10
6.	References.....	11
6.1	Normative References.....	11
6.2	Informative References.....	11

7.	Acknowledgments.....	12
8.	Author's Addresses.....	12

Internet Draft [draft-fedyk-isis-spb-00.txt](#)

July 2009

1. Terminology

In addition to well understood IS-IS terms, this memo uses terminology from IEEE 802.1 and introduces a few new terms:

802.1ah	Provider Backbone Bridges a.k.a. Mac-in-Mac encapsulation
802.1aq	Shortest Path Bridging (SPB)
B-DA	Backbone Destination Address in 802.1ah Mac-in-Mac header
B-MAC	Backbone MAC Address
B-SA	Backbone Source address in 802.1ah Mac-in-Mac header
B-VID	Backbone VLAN ID in 802.1ah Mac-in-Mac header
B-VLAN	Backbone Virtual LAN
C-MAC	Customer MAC. Inner MAC in 802.1ah Mac-in-Mac header
C-VID	Customer VLAN ID
C-VLAN	Customer Virtual LAN
DA	Destination Address
FIB	Forwarding information base (B-DA/B-VID to next hop(s))
ISID	802.1ah: service membership in datapath(not always=SID)
MAC	Media Access Control
PBB	Provider Backbone Bridges as specified in 802.1ah
M-IS-IS	Multi Topology IS-IS as used in [MT]
MT	Multi Topology. As used in [MT]
MT-ID	Multi Topology Identifier (12 bits). As used in [MT]
SPSourceID	20 bit nodal identifier used with SID forms mcast DA.
PATHID	The unique identifier for a path used for symmetric tie breaking
PBB	Provider Backbone Bridge
PBT	Provider Backbone Transport
SA	Source Address
SID	Service Identifier at control plane, groups many ISIDs
UNI	User Network Interface: Customer/Backbone attach point
VID	VLAN ID

2. Introduction

Shortest Path Bridging (SPB) and Shortest Path Backbone Bridging (SPBB) both apply the [\[IS-IS\]](#) protocol to the utilization of mesh topologies for native Ethernet bridging. While 802.1aq has the umbrella title of shortest path bridging, this memo adopts the convention of ascribing unique attributes to the terms SPB and SPBB individually; SPB uses an IEEE 802.1Q forwarding paradigm, whilst SPBB extends the FIB population techniques specified in 802.1Qay combined with IEEE 802.1ah [\[PBB\]](#) adaptation. Both SPB and SPBB forward packets on shortest path trees with minimum path cost as a first order tie-breaker, where for any pair of nodes A and B, the unicast path for A to B is the exact reverse of the path from B to A (reverse path congruency), and all multicast traffic between the two nodes follows the unicast path (multicast and unicast congruency). These are direct extensions to fundamental Ethernet forwarding properties in IEEE bridged networks.

In SPB, conventional bridge learning is used to associate (customer) MAC addresses to ports and hence routes through the SPB region. The source-rooted tree associated with each node is assigned a unique VLAN ID (the SPVID) to identify it.

The 802.1ah [\[PBB\]](#) MAC-in-MAC encapsulation used by SPBB permits the isolation of customer Ethernet addressing from backbone Ethernet

addressing in the core of a network. This has an important consequence; the association between a customer MAC (C-MAC) and a B-MAC to resolve forwarding across the core is required only at the edge of the network. Flooding is only done by the edge adaptation functions to learn which B-MAC reaches a given C-MAC using the normal PBB C-MAC learning behavior, flooding at the C-MAC layer resolving to a service specific multicast tree at the B-MAC layer.

A minimum of one B-VID MUST be assigned to each instance of SPBB (IS-IS MT-ID). The B-VID MUST be unique backbone network wide. Two B-VIDs MAY be used by a single SPBB instance (MT-ID) when it is desired to use more than one equal cost shortest path permutation.

With Spanning Tree Protocol (STP), Rapid STP or Multiple STP[802.1Q] Ethernet networks use a shared spanning tree (or a small number of shared trees) to route traffic based on the VLAN ID and then on learned MAC addresses. Per service multicast is instantiated as a (*,G) multicast address which is a proper subset of the VID.

In order to move from a shared spanning tree to mesh connectivity, SPB and SPBB use one or more broadcast trees per source as a template for per service multicast trees i.e.(S,G).

SPB uses the SPVID when flooding, with conventional pruning as C-MAC addresses are learned.

SPBB adapts client flooding onto service specific multicast trees instantiated by encoding S,G in the destination MAC address. SPBB therefore uses both B-VID and B-MAC when forwarding.

SPBB uses algorithmically constructed addresses for the multicast DA. Multicast addresses are local to a PBB domain. The multicast B-MAC DA address constructed for a particular node/service consists of a unique nodal identifier called a SPSourceID (20 bits), combined with the service identifier (SID) (24 bits). Every SPBB node(per SPBB instance/MT-ID) therefore advertises this unique SPSourceID which can be correlated to the node's ISIS SYSID by all nodes in the SPBB network.

In SPB, a VID (called the Base VID) is also defined for interoperation with xSTP regions. This VID identifies the Common and Internal Spanning Tree (CIST) in the Shortest Path Region. It is created as a conventional Spanning Tree within the SPB region, and Group MAC addresses, installed by conventional Ethernet mechanisms,

may be supported on this VID.

A node performing SPB or SPBB calculations (for a given instance/MT-ID) use the IS-IS topology, and link metrics to compute which leaves of each shortest path tree require transit of the local node, and this node can then do pair-wise comparison of services of interest between the root and the leaves to populate the FIB accordingly. The link metrics are forced to be equal in both directions by defaulting to the largest of the two unidirectional metrics specified for a link as seen during the establishment of a standard IS-IS adjacency.

A node performing SPB or SPBB for a given MT instance MUST specify the NLPID for SPB and SPBB (IANA allocation 0xC1 pending) in its IS-IS hello message. Links that do not support this NLPID for this MT instance must be excluded from the shortest path computations by the given MT instance.

Both SPB and SPBB employ a transitive symmetric tie breaking algorithm which chooses deterministically and progressively between equal cost alternatives by ranking the paths according to the sorted list of node identifiers that make up the path. This sorted list of node identifiers (SYSIDs) is called the Path Identifier (PATHID). For SPB, and when only one B-VID is used in SPBB, the tie breaking algorithm will always pick from a set of equal cost shortest path alternatives by choosing the path with the lowest PATHID. Two B-VID instances MAY be used for an SPBB instance, when it will assign the lowest PATHID paths to one B-VID and the highest PATHID paths to the other B-VID. In this manner an SPBB instance will use more of the available paths in a network but will do so by assignment of packets at the head ends to one of different SPBB B-VIDs.

3. New TLVs for SPB and SPBB

SPB and SPBB require a subset of the [MT] TLVs.

SPB and SPBB inherit the Multi Topology mechanisms from [MT] to allow multiple logical bridging instances to exist within a single IS-IS control instance. Top level TLV's used by SPB and SPBB therefore begin with the Multi Topology Identifier (MT-ID) fields as defined in [MT] and the new sub-TLVs identified here may be used as

sub-TLVs of the corresponding new top level Multi Topology TLV's defined in [\[MT\]](#).

In the IEEE the Multiple Spanning tree protocol allowed multiple VLANs to represent a single VLAN. The single "logical" VLAN was identified by a Base VID so that bridges external to the region could have a consistent VID to identify the VLAN. This concept of a Base VID extends easily to SPB and SPBB. In this way a Base VID can be used to identify a topology instance and to correlate VLANs that are allocated to a particular shortest path VLAN instance.

A Shortest Path Region also has the property that viewed from outside the Region it appears as a single 802.1Q or 802.1ah bridge, irrespective of how the VLAN is implemented.

In SPB, each node uses a unique VID as its source identifier (an SPVID), and each SPVID is correlated to a Base VID. The set of SPVIDs that map to a given Base VID form the SPB region. In SPBB, this concept is reused however there is a slight modification. A single B-VID maps to a Base VID for all bridges. However multiple B-VIDs (currently two) can map to the same Base VID allowing multiple trees within the SPBB VLAN. A typical use for multiple trees is to instantiate equal cost paths and provide the opportunity to load spread services.

A Base VID identifies a traditional spanning tree in both SPB and SPBB that can be used to represent the VLAN for proper bridge behavior when viewed by bridges outside the shortest path region. In essence the VLAN identified by this Base VID can appear as a single bridge allowing proper spanning tree behavior.

The following sections introduce the new [\[MT\]](#) TLVs which are used by SPB and SPBB, and give an overview of their use from an SPB and SPBB perspective.

The Base Vlan-Identifier sub-TLV (section 2.3.5 of [\[IS-IS-L2\]](#)) is carried within the Multi Topology aware Port Capability TLV ([section 2.3](#) of [\[IS-IS-L2\]](#)), and this is carried in an IIH PDU.

The purpose of this sub-TLV is to check for compatible configuration of SPB or SPBB mode of operation and then of SPB or SPBB parameters as bridges form adjacencies, and to prevent adjacency formation when incompatible configurations are detected.

In informal terms this requires :

- agreement on the NLPID for SPB and SPBB
- specification of the bridging mode to be used (SPB or SPBB)
- binding of individual VIDs to the Base VID, and specification of the shortest path bridging algorithm to be used for each VID
- use (or not) of auto allocation capability for SPVIDs in SPB and SPSourceIDs in SPBB.

It was mentioned earlier that each SPB node is assigned a unique VID (a Shortest Path VID, or SPVID) as its source identifier for all traffic it transmits. The set of SPVIDs, one for each SPB node in the region, are bound to the Base VID, configured to execute a specific tiebreaking algorithm, and collectively provide the shortest path trees to support the VLAN. A control flag in the sub-TLV determines whether the SPVIDs are provisioned, or auto-allocated by the procedure in [801.1aq]

In SPB, the Base VID is also used to identify a VLAN providing peer inter-working with other non-SPB bridges outside the SPB Region. This VLAN forms a spanning tree across the region to achieve this.

In SPBB, the VID(s) on which forwarding is performed are Region-wide assignments. At present, the use of one or two VIDs is defined, with the latter capability available for edge-based load spreading using Equal Cost Multiple Trees generated via the symmetric tie-breaking variations.

Three algorithms are currently available to SPB and SPBB :

- spanning tree algorithm, which constructs a tree which is the same as would be constructed by Spanning Tree Protocol [\[802.1D\]](#)
- shortest path trees, selecting the low PATHID as a tie-breaker
- shortest path trees, selecting the high PATHID as a tie-breaker

Internet Draft [draft-fedyk-isis-spb-00.txt](#)

July 2009

3.2 SPB/SPBB Instance sub-TLV

The SPB/SPBB Instance sub-TLV (section 2.5.1 of [[IS-IS-L2](#)]) is carried within the Multi Topology Aware Capability TLV (section 2.5 of [[IS-IS-L2](#)]).

The purpose of this sub-TLV is to announce configuration and other parameters to the entire SPB or SPBB Region. The instance sub-TLV carries some information elements common to the Base VLAN-Identifiers sub-TLV described in the previous section, which are :

- binding of individual VIDs to the Base VID, and specification of the shortest path bridging algorithm to be used for each VID
- use (or not) of auto allocation capability for SPVIDs in SPB and SPSourceIDs in SPBB.

The SPB/SPBB Instance sub-TLV carries further information :

- the SPSourceID - the 20 bit network wide unique identifier used in the higher order bits of the SPBB multicast DA for packets originating at this node
- various Spanning Tree parameters for inter-working with non-SPB regions

In SPB, the multicast tree built off each SPB node is uniquely associated with an SPVID which thereby identifies the source. The required (S,G) trees, and loop avoidance checking, may be directly implemented using this SPVID by standard Ethernet forwarding. The SPVID can be configured from a pool or it can be auto allocated.

In SPBB multicast, the same capability is achieved by encoding the service-specific (S,G) tree in the multicast Destination Address. This is achieved by concatenating the PBB Service Identifier with the nodal SPSourceID. The distribution of SPSourceID therefore allows all SPBB nodes to compute the forwarding state they need to install, based only on topology and service endpoint locations. The computed SPBB multicast DA looks like this:

```

+--+--+--+-----+-----+-----+-----+-----+-----+
|M/L| A | SPSourceID (20 bit) | I-SID (24 bit) |
+--+--+--+-----+-----+-----+-----+-----+-----+

```

Where M/L = multicast/local bits
 (2 bits - both set to 1)
 A = SPSourceID allocation style

(2 bits - both 0 initially)

Figure 1: SPBB multicast MAC address construction

The SPSourceID can be provisioned, or auto allocated.

The Spanning tree inter-working parameters for SPB comprise :

- a Spanning tree compatible Bridge identifier, configured exactly as specified in [802.1D]. This allows SPB to build a compatible Spanning tree using link state.
- The Base VID identifies a VLAN capable of covering multiple Regions, SPB and non-SPB. In SPB, this is known as the Common and Internal Spanning Tree (CIST). At SPB Region boundaries, the CIST Root Identifier and the CIST External Root Path Cost may be imported from xSTP and flooded by IS-IS as a part of SPB's Spanning Tree emulation.

3.3 SPB Link Metric sub-TLV

SPB Link Metric sub-TLV (section 2.6 of [IS-IS-L2]) is carried within the Extended Reachability TLV, or the Multi Topology Intermediate System TLV.

The purpose of this sub-TLV is to announce SPB link metrics, in a form which enables SPB to build xSTP compatible Spanning Trees, typically to create the SPB component of the CIST (above) :

- indicates the administrative cost or weight of using a link
- a standard IEEE port identifier used to build a spanning tree associated with this link

3.4 The Group MAC Address sub-TLV

The Group MAC Address sub-TLV (section 2.2.1 of [IS-IS-L2]) is carried within the Group Address TLV (section 2.2 of [IS-IS-L2]), which is in turn carried within the Multicast Group Level 1 link state PDU.

This sub-TLV is used only by SPB. SPBB builds and installs per-service per-source multicast addresses algorithmically, as described earlier, using the SPSourceID and PBB Service Identifier

information announced in other sub-TLVs.

By default, SPB nodes broadcast traffic to all other nodes in their region. When inter-working with non-SPB regions over the CIST, multicast group membership may be signaled over the CIST using mechanisms such as MMRP [[802.1ak](#)]. The Group MAC Address sub-TLV allows such registrations to be imported and announced by IS-IS.

The sub-TLV carries the following information for SPB :

- the VID with which all subsequent MAC addresses are associated
- sets of group records, each consisting of a multicast group address and a list of unicast MAC addresses known to be sources of that group.

3.5 The Service Identifier and Unicast Address sub-TLV

The Service Identifier and Unicast Address sub-TLV (section 2.5.2 of [[IS-IS-L2](#)]) is carried within the Multi Topology Aware Capability TLV (section 2.5 of [[IS-IS-L2](#)]).

The purpose of this sub-TLV is to announce service group membership (using the PBB Service Identifier or I-SID) on the originating node, also to advertise an additional B-MAC unicast address present on or reachable by the node. It is applicable to SPBB only.

This sub-TLV carries:

- the unicast B-MAC address which must be used to send to the set of PBB I-SIDs announced in the sub-TLV, and which this node will use as its source B-MAC when transmitting these I-SIDs
- the unicast VID which must be used to send to the set of PBB I-SIDs announced in the sub-TLV, and which this node will use when transmitting these I-SIDs
- a list of PBB I-SIDs and their transmit and receive properties.

Announcement of I-SIDs in this way allows all SPBB nodes to see all service endpoints, and allows nodes not terminating a particular service to algorithmically determine the per-service per-source forwarding state which they must install if they lie on the shortest path between two or more service end-points.

The advertisement of the B-MAC unicast address to be used to reach the set of services allows different granularities of addressing to be used within the SPBB node, without compromising inter-working between nodes of different types. It also has application in some resiliency schemes.

4. Security Considerations

This document adds no additional security risks to IS-IS.

SPBB assumes that the link state bridged subnetwork consists of trusted devices and that the UNI ports to the domain are untrusted. Care is required to ensure untrusted access to the trusted domain does not occur.

5. IANA Considerations / ISO Considerations

See the subset of [[IS-IS-L2](#)] cited by this document and also the NLPID assignments requested by [[NLPID](#)].

6. References

6.1 Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [MT] Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs), [RFC 5120](#), February 2008
- [802.1D] "IEEE draft Standard for Local and Metropolitan Networks, Media Access Control (MAC) Bridges", IEEE 802.D June 2004
- [802.1aq] "IEEE draft Standard for Local and Metropolitan Networks, Virtual Bridged Local Area Networks, Amendment 9: Shortest Path Bridging",

IEEE 802.1aq D2.0, June 2009

- [802.1ak] "IEEE Standard for Local and Metropolitan Networks, Virtual Bridged Local Area Networks, Amendment 7: Multiple Registration Protocol"
IEEE Std 802.1ak - 2007 amendment to IEEE 802.Q - 2005
- [PBB] "IEEE Standard for Local and Metropolitan Networks, Virtual Bridged Local Area Networks, Amendment 7: Provider Backbone Bridges"
IEEE Std 802.1ah - 2008 amendment to IEEE 802.Q - 2005
- [IS-IS-L2] Extensions to IS-IS for Layer-2 Systems, IETF, Internet Draft, [draft-ietf-isis-layer2-00.txt](#), Work in Progress, June 2009
- [NLPID] IANA Considerations for NLPIDs, IETF, Internet Draft, Draft-eastlake-nlpid-iana-considerations-00.txt, Work in Progress, June 23, 2009

6.2 Informative References

- [IS-IS] ISO/IEC 10589:2002, "Intermediate system to Intermediate system routing information exchange protocol" ISO/IEC 10589:2002.

Fedyk et al.

[Page 11]

Internet Draft [draft-fedyk-isis-spb-00.txt](#) July 2009

7. Acknowledgments

The authors would like to thank Antonela Paraschiv, Daniel Joyal, Paul Unbehagen and Gautam Khera for their detailed review of this work.

8. Author's Addresses

Don Fedyk

Alcatel-Lucent
220 Hayden Road
Groton, MA, USA
Email donald.fedyk@alcatel-lucent.com

Peter Ashwood-Smith
Huawei Technologies Canada
411 Leggget Drive, Suite 503
Kanata, Ontario, K2k3C9, Canada
Email: peter.ashwoodsmith@huawei.com

Nigel Bragg
Nortel Networks
London Road, Harlow,
Essex CM17 9NA, UK
Email: nbragg@nortel.com

David Allan
Nortel Networks
3500 Carling Ave.
Ottawa, ON, Canada
K1Y4H7
Email: dallan@nortel.com

Jerome Chiabaut
Nortel Networks
3500 Carling Ave.
Ottawa, ON, Canada
K1Y 4H7
Email: chiabaut@nortel.com