

LSR WG
Internet-Draft
Intended status: Standards Track
Expires: 7 February 2026

C. Barth
T. Li
V. P. Beeram
R. Bonica
HPE
6 August 2025

Using IS-IS To Advertise Power Group Membership
draft-many-lsr-power-group-01

Abstract

This document introduces Power Groups. A Power Group is a hierarchical abstraction of power consumed by hardware components. In IS-IS, interfaces can reference the Power Group to which they belong. Therefore, Power Groups provide a method of organizing interfaces into groups by power characteristics.

The TE path placement algorithm can use Power Group membership information to implement TE policy. Power Group information is particularly useful when implementing TE policies that support power-savings and sustainability.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 7 February 2026.

Copyright Notice

Copyright (c) 2025 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Conventions and Definitions	3
3. Example Architecture	3
4. Power Groups	4
5. Interfaces and Power Groups	6
6. Power-Save Capability and Power Group Hierarchies	7
7. Link State Database Elements	7
7.1. The Power Group TLV	7
7.2. The Sleeping Adjacency TLV	8
7.3. Interface Extensions	9
7.3.1. The Power Group Member Sub-TLV	9
7.3.2. The Interface Power Sub-TLV	10
7.3.3. Unidirectional Sleeping Bandwidth Sub-TLV	10
7.3.4. The Power-Sleep Capable Bit	11
8. Operational Considerations	11
9. Security Considerations	11
10. IANA Considerations	12
11. Acknowledgements	13
12. References	13
12.1. Normative References	13
12.2. Informative References	13
Authors' Addresses	13

1. Introduction

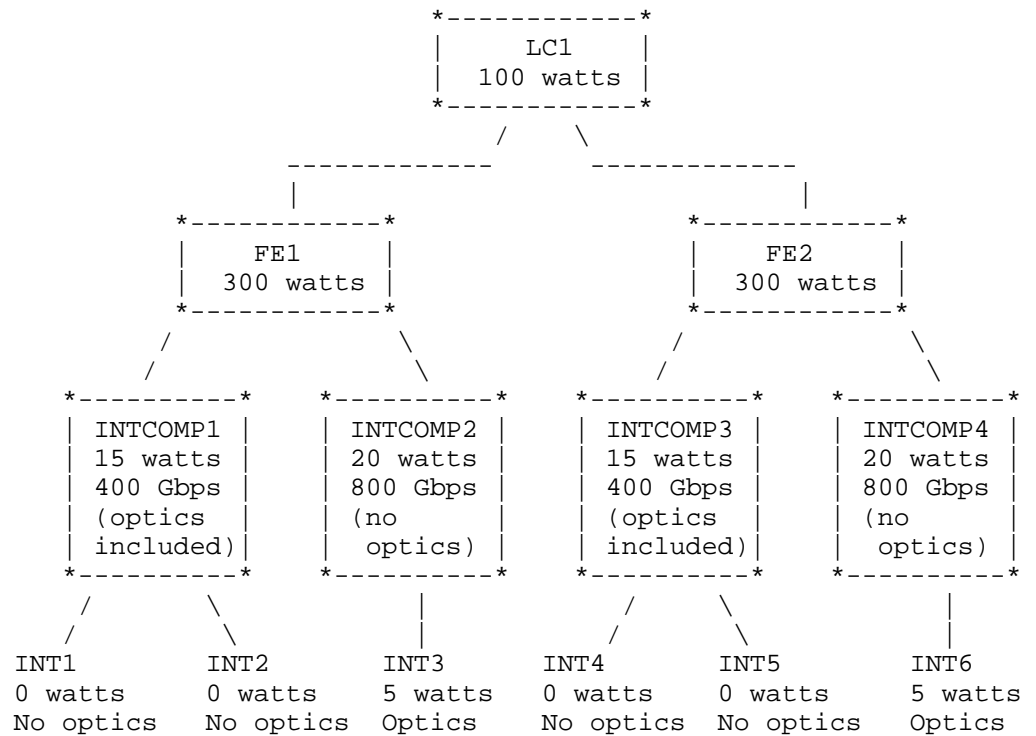
This document introduces Power Groups. A Power Group is a hierarchical abstraction of power consumed by hardware components. In IS-IS, interfaces can reference the Power Group to which they belong. Therefore, Power Groups provide a method of organizing interfaces into groups by power characteristics.

The TE path placement algorithm can use Power Group membership information to implement TE policy. Power Group information is particularly useful when implementing TE policies that support power-savings and sustainability.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Example Architecture



Line Card 1 (LC1) consumes 100 watts

Figure 1: Line Card 1

Figure 1 depicts a line card (LC1). LC1 contains two forwarding engines (FE1 and FE2) and four interface complexes (INTCOMP1 through INTCOMP4). INTCOMP1 supports in two interfaces (INT1 and INT2). Likewise, INTCOMP3 supports in two interfaces (INT4 and INT5). INTCOMP2 and INTCOMP4 support one interface each (INT3 and INT6).

An interface complex includes PHY, MAC, encryption, gearbox, and other related circuitry. INTCOMP1 and INTCOMP3 also contain optics. INTCOMP2 and INTCOMP4 do not contain optics. Therefore, the interfaces that they support have their own optics.

INTCOMP1 and INTCOMP3 provide 400 Gbps of forwarding capacity each, while INCOMP2 and INTCOMP4 provide 800 Gbps of forwarding capacity each.

Each hardware component consumes power. LC1 consumes 100 watts while FE1 and FE2 consume 300 watts each. INTCOMP1 and INTCOMP3 consume 15 watts each, while INTCOMP2 and INTCOMP4 consume 20 watts each. INT3 and INT6 contain optics that consume 5 watts each. INT1, INT2, INT4 and INT5 do not do separate optics. Therefore, they do not consume power beyond what is consumed by the complex.

INT1 and INT2 depend upon INTCOMP1. If INTCOMP1 fails, so do INT1 and INT2. Likewise, INT3 depends upon INTCOMP2. If INTCOMP2 fails, so does INT3.

INTCOMP1 and INTCOMP2 depend on FE1. If FE1 fails, so do INTCOMP1, INTCOMP2, INT1, INT2, and INT3. Likewise, INTCOMP3 and INTCOMP4 depend on FE2. If FE2 fails, so do INTCOMP3, INTCOMP4, INT4, INT5, and INT6.

FE1 and FE2 depend on LC1. If LC1 fails, so do all of the forwarding engines, interface complexes, and interfaces in the diagram.

4. Power Groups

A Power Group is a hierarchical abstraction of power consumed by hardware components. Each Power Group, except for the one at the top of the hierarchy, has exactly one parent. The Power Group at the top of the hierarchy does not have a parent. Many Power Groups can have the same parent.

Each Power Group has one or more components and each component consumes power. The power consumed by a Power Group is equal to the sum of the power consumed by each of its components. The power consumed by a Power Group does not include the power consumed by its ancestors or by its children.

The parent-child relationship reflects dependency. One Power Group is the child of another if any one of the child components depends upon any one of the parent components.

A network device's power consumption characteristics can be described by any number of equivalent Power Group hierarchies. The paragraphs below demonstrate how two equivalent Power Group hierarchies can describe the power consumption characteristics of the line card in Figure 1.

Identifier	Parent	Power Consumption	Hardware Components
1	None	100 watts	LC1
2	1	300 watts	FE1
3	1	300 watts	FE2
4	2	15 watts	INTCOMP1
5	2	20 watts	INTCOMP2
6	3	15 watts	INTCOMP3
7	3	20 watts	INTCOMP4
8	5	5 watts	INT3
9	7	5 watts	INT6

Table 1: A Granular Power Group Hierarchy

Table 1 describes the power consumption characteristics of the line card in Figure 1 using a granular Power Group hierarchy. We call it granular because each Power Group contains only one component. The power consumed by each Power Group is equal to the power consumed by its component.

In Table 1, Power Group 7 is the child of Power Group 3 because INTCOMP4 depends upon FE2. Likewise, Power Group 3 is the child of Power Group 1 because FE2 depends on LC1. Furthermore, Power Group 8 is the child of Power Group 5 because INT3 depends upon INCOMP2. Likewise, Power Group 9 is the child of Power Group 7 because INT6 depends on INTCOMP4.

Identifier	Parent	Power Consumption	Hardware Components
1	None	700 watts	LC1, FE1, FE2
2	1	15 watts	INTCOMP1
3	1	20 watts	INTCOMP2
4	1	15 watts	INTCOMP3
5	1	20 watts	INTCOMP4
6	1	5 watts	INT3
7	1	5 watts	INT6

Table 2: A Less Granular Power Group Hierarchy

Table 2 describes the power consumption characteristics of the line card in Figure 1 using a less granular Power Group hierarchy. We call it less granular because Power Group 1 contains three components (LC1, FE1 and FE2). Its power consumption is equal to the sum of the power consumed by LC1, FE1 and FE2 (i.e., 700 watts).

Power Group 2 and Power Group 3 are children of Power Group 1 because INTCOMP1 and INTCOMP2 depend on FE1. Likewise, Power Group 4 and Power Group 5 are children of Power Group 1 because INTCOMP3 and INTCOMP4 depend on FE2. Finally, Power Group 6 and Power Group 7 are children of Power Group 1 because INT3 and INT6 depend on INCOMP2 and INCOMP4..

Section 6 describes how a network device's power-save capability determines which of the equivalent Power Group hierarchies it should advertise.

Section 7.3.2 describes how IS-IS advertises Power Group information.

5. Interfaces and Power Groups

An interface is not part of a Power Group, even if it contains optics and consumes power. However, an interface can reference a Power Group. When it references a Power Group, it MUST reference the Power Group that contains the interface complex that supports it. See Section 7.3.1.

Therefore, Power Groups can be used to associate interfaces that depend on a common set of hardware components and have common power consumption characteristics.

A Link Aggregation Group (LAG) interface requires support from multiple interface complexes. Therefore a LAG interface references every Power Group that contains an interface complex that supports it.

Section 7.3.2 describes how an interface can advertise the power that it consumes.

6. Power-Save Capability and Power Group Hierarchies

A network device SHOULD advertise the least granular Power Group hierarchy that can exercise its complete power-savings capability.

Assume that a network contains line cards that are power-save capable. Those line cards contain forwarding engines and interface complexes that are also power-save capable. This means that the line cards, forwarding engines and interface complexes can be powered on and off independently of the chassis.

In order to exercise its complete power savings capability, information regarding line card, forwarding engine and interface complex dependencies is required. Therefore, the line card must advertise the granular Power Group hierarchy in Table 1.

Now assume that another network contains line cards that are power-save capable. Those line cards contain interface complexes that are also power-save capable. However, the forwarding engines are not power-save capable.

In order to exercise its complete power savings capability, information regarding line card, and interface complex dependencies is required. However, information regarding forwarding engine dependencies is not required. Therefore, the line card could advertise either the granular Power Group hierarchy in Table 1 or the less granular Power Group hierarchy in Table 2.

7. Link State Database Elements

7.1. The Power Group TLV

The Power Group is a top level TLV that describes a Power Group.

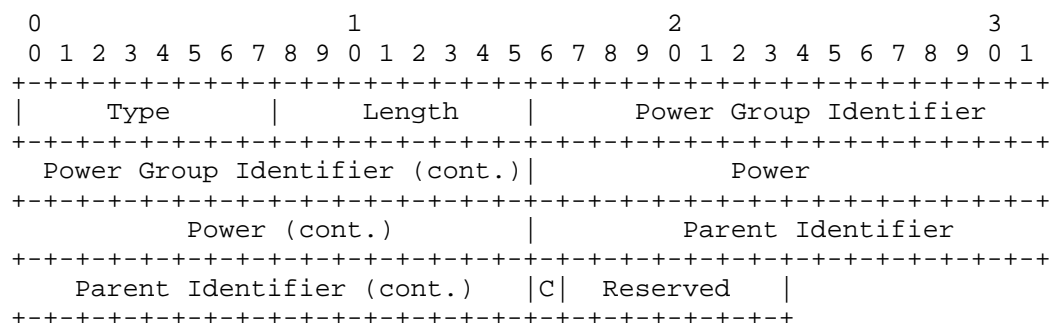


Figure 2: Power Group TLV

Where:

- * Type: 1 octet, value TBD1
- * Length: 1 octet, unsigned integer. Value 13.
- * Power Group Identifier: 4 octets, selector. MUST NOT be equal to 0.
- * Power: 4 octets, unsigned integer. The power consumed by the Power Group, in milliwatts.
- * Parent Identifier: 4 octets, selector.
- * C: 1 bit. Indicates that every component in the Power Group is power-save capable.
- * Reserved: 7 bits. MUST be set to 0 by the sender and MUST be ignored by the receiver.

The Power Group Identifier has node-local significance. If the Parent Identifier is equal to 0, the Power Group has no parent (i.e., it is the root of a Power Group hierarchy). Otherwise, the Parent Identifier MUST NOT be set to 0.

7.2. The Sleeping Adjacency TLV

The Sleeping Adjacency TLV is a top level TLV. If an adjacency is in the power-sleep mode, the TLVs that represent it appear only in the Sleeping Adjacency TLV. They do not also appear as top-level TLVs.

The Sleeping Adjacency TLV can include TLVs 22, 23, 141, 222 and 223.




Figure 3: Sleeping Adjacencies TLV

Where:

- * Type: 1 octet, value TBD2
- * Length: 1 octet, unsigned integer. The length of the TLV, measured in octets, not including the type and length fields.
- * Sleeping Adjacencies: A list of adjacency TLVs of type 22, 23, 141, 222 and 223. These TLVs represent adjacencies that are in the power-sleep mode.

7.3. Interface Extensions

7.3.1. The Power Group Member Sub-TLV

This sub-TLV is found in TLVs for advertising neighbor information.

This sub-TLV advertises a Power Group to which the interface belongs. Because a LAG interface can belong to many Power Groups, many instances of this sub-TLV may be advertised.




Figure 4: Power Group Member Sub-TLV

Where:

- * Type: 1 octet, value TBD3
- * Length: 1 octet, unsigned integer. Value 4,

- * Power Group Identifier: 4 octets, selector.

7.3.2. The Interface Power Sub-TLV

This sub-TLV is found in TLVs for advertising neighbor information.

This sub-TLV advertises the power consumed by an interface. A dynamic value might cause unnecessary churn in the Link State Database (LSDB), so a static value should be used.




Figure 5: Interface Power Sub-TLV

Where:

- * Type: 1 octet, value TBD4
- * Length: 1 octet, unsigned integer. Value 4
- * Power: 4 octets, unsigned integer. The power consumed by the interface, in milliwatts.

7.3.3. Unidirectional Sleeping Bandwidth Sub-TLV

This sub-TLV is found in TLVs for advertising neighbor information.

This sub-TLV advertises the sleeping bandwidth between two directly connected IS-IS neighbors. The sleeping bandwidth advertised by this sub-TLV MUST be the sleeping bandwidth from the system originating the Link State Advertisement (LSA) to its neighbor.




Figure 6: Unidirectional Sleeping Bandwidth Sub-TLV

Where:

- * Type: 1 octet, value TBD5
- * Length: 1 octet, unsigned integer. Value 4.
- * Sleeping Bandwidth: 4 octets, IEEE floating-point format measured in bytes per second.

The Sleeping bandwidth field carries the sleeping bandwidth on a link, forwarding adjacency [RFC4206], or bundled link. For a link or forwarding adjacency, sleeping bandwidth is defined as the maximum bandwidth [RFC5305] minus the bandwidth currently allocated to RSVP-TE label switched paths that was transitioned to power-sleep. For a bundled link, sleeping bandwidth is defined to be the sum of the component link sleeping bandwidths.

7.3.4. The Power-Sleep Capable Bit

This is a bit in the Link Attribute Sub-TLV (19). Presence of this bit indicates that the link may be put into power-sleep mode. The position of this bit is TBD5.

8. Operational Considerations

Network operators must exercise care when configuring interfaces to be power-sleep capable.

The TE path placement algorithm can use Power Groups to implement TE policies that support power-savings. In this case, the path placement algorithm identifies a Power Group in which all interfaces are power-sleep capable. It then diverts traffic from those interfaces. When traffic has been diverted, power can be removed from every hardware component in the Power Group.

Removing power from those components may cause the network to be insufficiently redundant. The subsequent failure of a single link may bisect the network.

Therefore, network operators must configure selected interfaces so that they are not power-sleep capable and will never be powered down.

9. Security Considerations

TBD

10. IANA Considerations

IANA is requested to add the following entries to the IS-IS Top-Level TLV Codepoints registry (<https://www.iana.org/assignments/isis-tlv-codepoints/isis-tlv-codepoints.xhtml#tlv-codepoints>):

Value	Name	IIH	LSP	SNP	Purge	MP	Status Reference
TBD1	Power Group	N	Y	N	N	N	This document
TBD2	Sleeping Adjacencies	N	Y	N	N	N	This document

Table 3: IS-IS Top-Level TLV Codepoints

IANA is also requested to add the following entries to the IS-IS Sub-TLVs for TLVs Advertising Neighbor Information registry (<https://www.iana.org/assignments/isis-tlv-codepoints/isis-tlv-codepoints.xhtml#isis-tlv-codepoints-advertising-neighbor-information>):

Type	Description	22	23	25	141	222	223	MP	Reference
TBD3	Power Group Member	Y	Y	Y(s)	Y	Y	Y	N	This document
TBD4	Interface Power	Y	Y	Y(s)	Y	Y	Y	N	This document
TBD5	Unidirectional Sleeping Bandwith	Y	Y	Y(s)	Y	Y	Y	N	This document

Table 4: IS-IS Sub-TLVs for TLVs Advertising Neighbor Information

IANA is also requested to add the following entry to the IS-IS Neighbor Link-Attribute Bit Values registry (<https://www.iana.org/assignments/isis-tlv-codepoints/isis-tlv-codepoints.xhtml#isis-tlv-codepoints-19of22>):

Value	Name	L2BM	Reference
TBD5	Power-Sleep Capable	N	This document

Table 5: IS-IS Neighbor Link-Attribute Bit Values

11. Acknowledgements

TBD

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/rfc/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.

12.2. Informative References

- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, DOI 10.17487/RFC4206, October 2005, <<https://www.rfc-editor.org/rfc/rfc4206>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/rfc/rfc5305>>.

Authors' Addresses

Colby Barth
HPE
United States of America
Email: cbarth@juniper.net

Tony Li
HPE
United States of America

Email: tli@juniper.net

Vishnu Pavan Beeram
HPE
United States of America
Email: vbeeram@juniper.net

Ron Bonica
HPE
United States of America
Email: rbonica@juniper.net