Network Working Group
Internet Draft
Intended status: Standards Track
Expires: September 03, 2025

Y. Liu
China Mobile
C. Lin
New H3C Technologies
March 03, 2025

BGP-LS extensions for energy efficiency
draft-liu-idr-bgp-ls-energy-efficiency-00

Abstract

[draft-liu-spring-sr-policy-energy-efficiency-00] describes the types of energy consumption information, how to collect energy consumption information, and the framework for path selection based on energy consumption information.

This document elaborates on extending BGP-LS to carry energy consumption information and transmit it to the controller, enabling the controller to perform routing calculations based on energy consumption metrics.

Table of Contents

1. Introduction

The importance of energy efficiency in modern networks is increasingly evident. In addition to reducing the power consumption of devices, network technologies can be leveraged to redirect traffic towards energy-efficient devices and paths, effectively lowering the energy consumption of network communications.

[draft-cx-green-green-metrics] described various metrics can be used to evaluate energy consumption.

[draft-liu-spring-sr-policy-energy-efficiency-00] describes how energy-consumption information is utilized in SR networks. It elaborates on the use of energy consumption information for path computation in SR networks.

The metrics for energy consumption include:

Node Maximum Energy Consumption: The power consumption of a node at full load, measured in watts.

Node Real-Time Energy Consumption: The real-time power consumption of a node, measured in watts.

Node Maximum Unit Energy Consumption: The power consumption of a node at full load divided by traffic, measured in watts per gigabyte (W/GB).

Node Real-Time Unit Energy Consumption: The real-time power consumption of a node divided by real-time traffic, measured in watts per gigabyte (W/GB).

Node Average Unit Energy Consumption: The change in power consumption of a node over a measurement period divided by the change in traffic, measured in watts per gigabyte (W/GB).

Interface Maximum Unit Energy Consumption: The power consumption of an interface at full load divided by traffic, measured in watts per gigabyte (W/GB).

Interface Real-Time Unit Energy Consumption: The real-time power consumption of an interface divided by real-time traffic, measured in watts per gigabyte (W/GB).

Interface Average Unit Energy Consumption: The change in power consumption of an interface over a measurement period divided by the change in traffic, measured in watts per gigabyte (W/GB).

This document describes how to extend BGP-LS to carry energy consumption information and transmit it to the controller.

1.1. Requirements Language

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. BGP-LS Extensions for energy consumption

2.1. Link Attribute TLV

New Link Attribute TLVs called "Link energy consumption" is requested from "Link Attribute TLVs" in this document for Link energy consumption associated with BGP-LS Attribute of a Link NLRI.

2.1.1. Interface Maximum Unit Energy Consumption

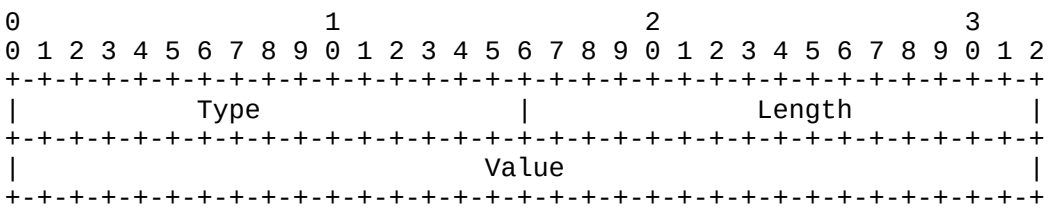

Type (2 octets): TBD1, Interface Maximum Unit Energy Consumption
Length (2 octets): 4.
Value (4 octet): Interface Maximum Unit Energy Consumption value range (0 - 4,294,967,295)

Figure 1: Link Interface Maximum Unit Energy Consumption

2.1.2. Interface Real-Time Unit Energy Consumption

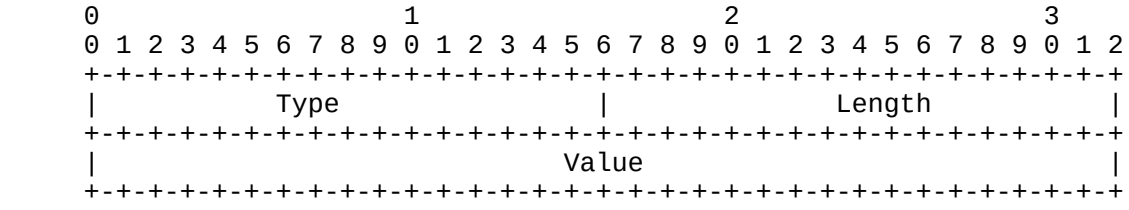

Type (2 octets): TBD2, Interface Real-Time Unit Energy Consumption
Length (2 octets): 4.
Value (4 octet): Interface Real-Time Unit Energy Consumption value range (0 - 4,294,967,295)

Figure 1: Link Interface Real-Time Unit Energy Consumption

2.1.3. Interface Average Unit Energy Consumption

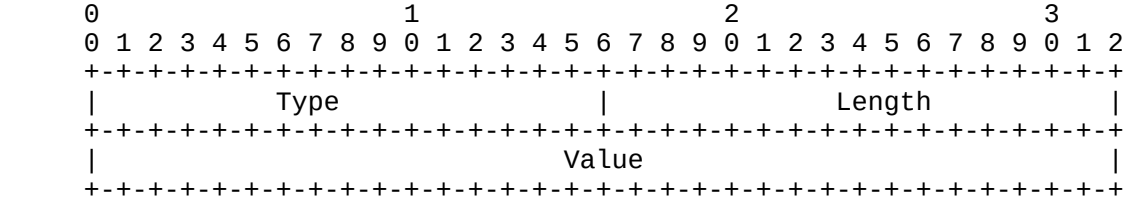

Type (2 octets): TBD3, Interface Average Unit Energy Consumption
Length (2 octets): 4.
Value (4 octet): Interface Average Unit Energy Consumption value range (0 - 4,294,967,295)

Figure 3: Link Interface Average Unit Energy Consumption

2.2. Node Attributes TLVs

New Node Attribute TLV called "Node energy consumption" is requested from "Node Attribute TLVs" in this document for energy consumption associated with BGP-LS Attribute of a Node NLRI.

2.2.1. Node Maximum Energy Consumption

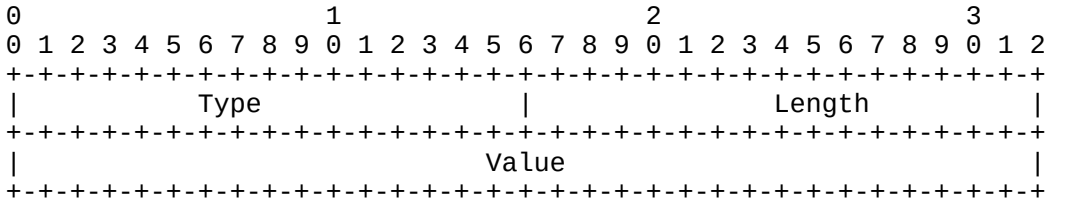

Type (2 octets): TBD4 Node Maximum Energy Consumption
Length (2 octets): 4
Value (4 octet): Node Maximum Energy Consumption value range (0 - 4,294,967,295)

Figure 6 Node Maximum Energy Consumption

2.2.2. Node Real-Time Energy Consumption

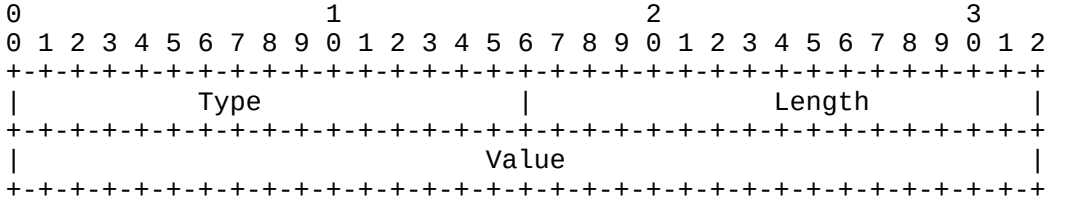

Type (2 octets): TBD5 Node Real-Time Energy Consumption
Length (2 octets): 4
Value (4 octet): Node Real-Time Energy Consumption value range (0 - 4,294,967,295)

Figure 7 Node Real-Time Energy Consumption

2.2.3. Node Maximum Unit Energy Consumption

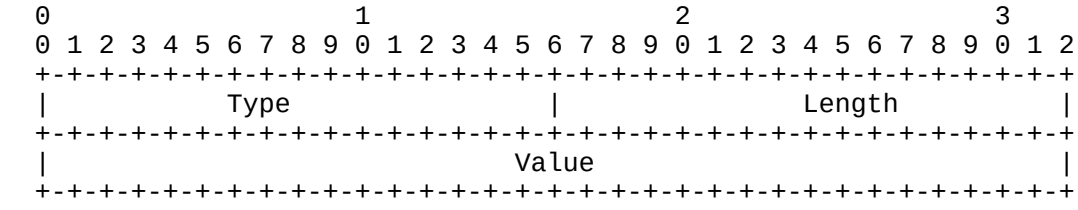

Type (2 octets): TBD6 Node Real-Time energy consumption
Length (2 octets): 4
Value (4 octet): Real-Time energy consumption value range (0 - 4,294,967,295)

Figure 8 Node Real-Time energy consumption

2.2.4. Node Real-Time Unit Energy Consumption

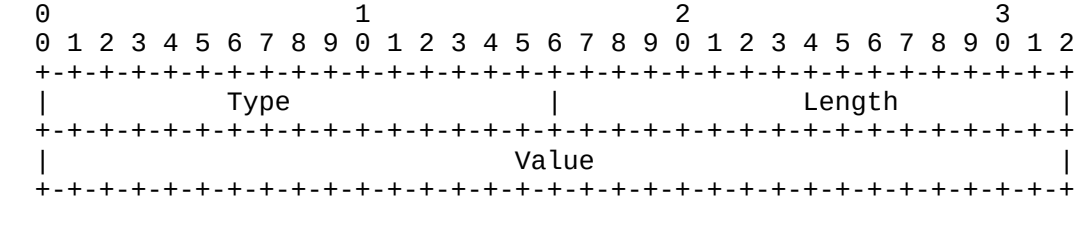

Type (2 octets): TBD7 Node Real-Time Unit Energy Consumption
Length (2 octets): 4
Value (4 octet): Node Real-Time Unit Energy Consumption value range (0 - 4,294,967,295)

Figure 9 Node Real-Time Unit Energy Consumption

2.2.5. Node Average Unit Energy Consumption

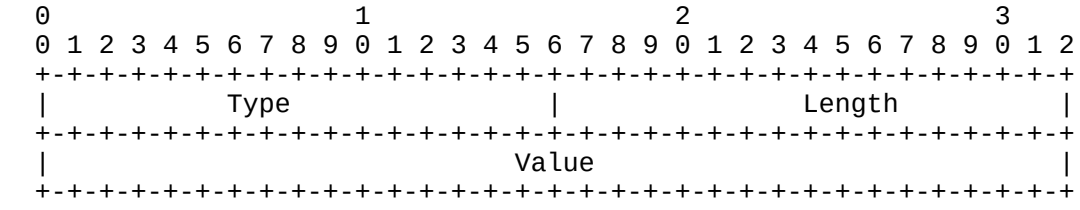

Type (2 octets): TBD8 Node Average Unit Energy Consumption
Length (2 octets): 4
Value (4 octet): Node Average Unit Energy Consumption value range (0 - 4,294,967,295)

Figure 10 Node Average Unit Energy Consumption

2.3. BGP-LS Board NLRI

For energy consumption information, devices can directly report it to the controller via BGP-LS. Since the energy consumption data of a device is typically collected on a per-board basis, this document introduces a new BGP-LS NLRI type, BOARD NLRI, to report board-related information.

[RFC7752] defines the NLRI format for BGP-LS.

2.3.1. Board NLRI Type

Type	NLRI Type
1	Node NLRI
2	Link NLRI
3	IPv4 Topology Prefix NLRI
4	IPv6 Topology Prefix NLRI
TBD9	Board NLRI

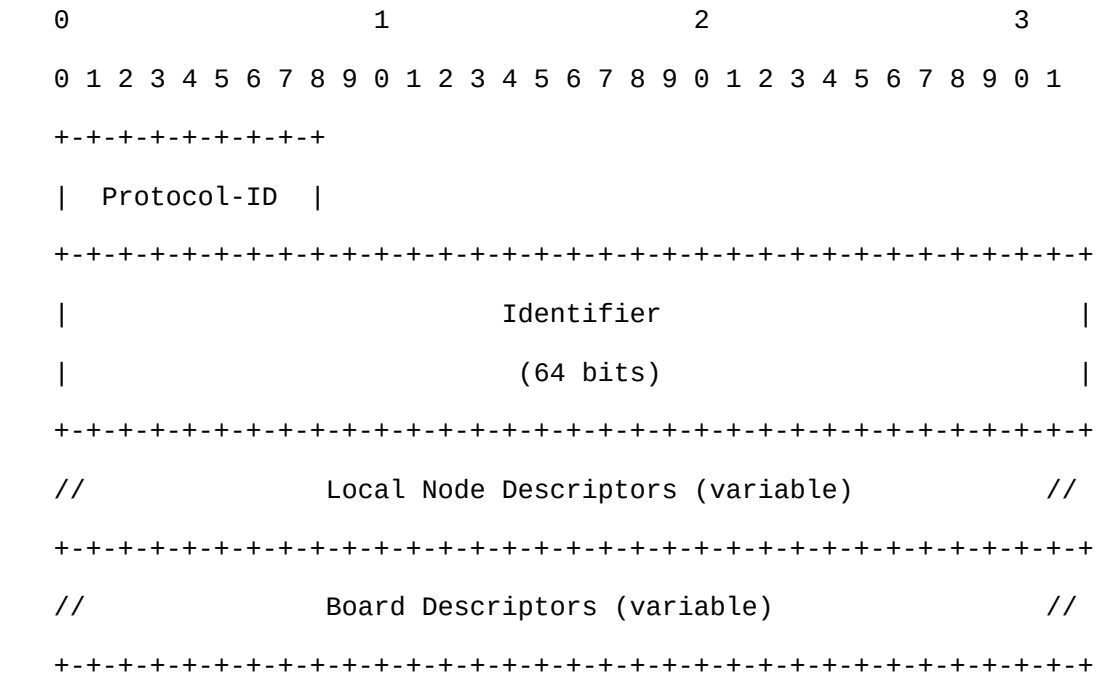

Figure 11: The Board NLRI Format

2.3.2. Board Descriptors

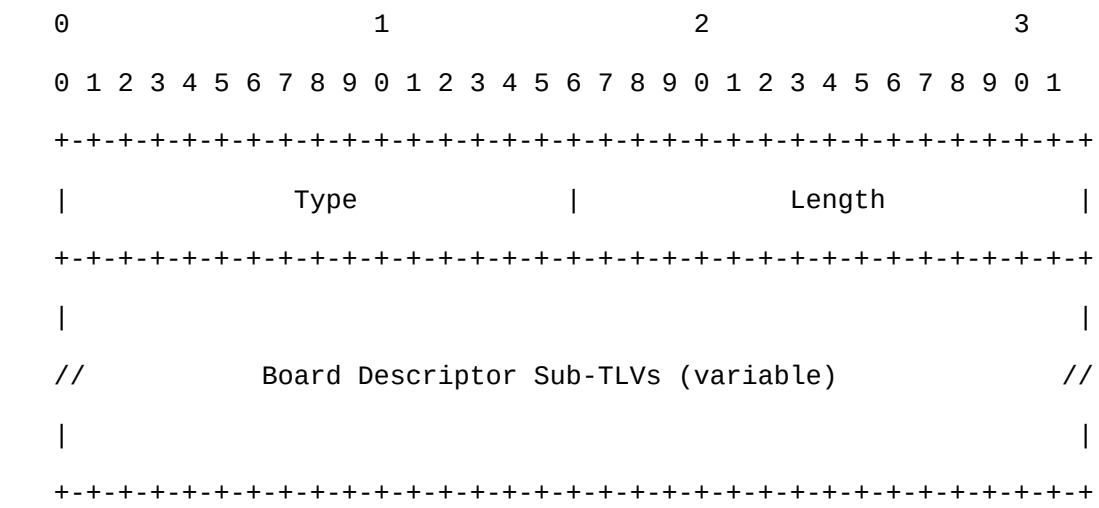

Figure 12: Board Descriptors TLV Format

Sub-TLV Code Point	Description	Length
TBD10	Board ID	4

A new Board ID Sub-TLV is defined, which contains a 4-byte Board ID.

2.4. BGP-LS Board Attributes TLVs

New Board Attribute TLVs called "Board energy consumption" is requested from "Board Attribute TLVs" in this document for Board energy consumption associated with BGP-LS Attribute of a Board NLRI.

2.4.1. Board Maximum Energy Consumption

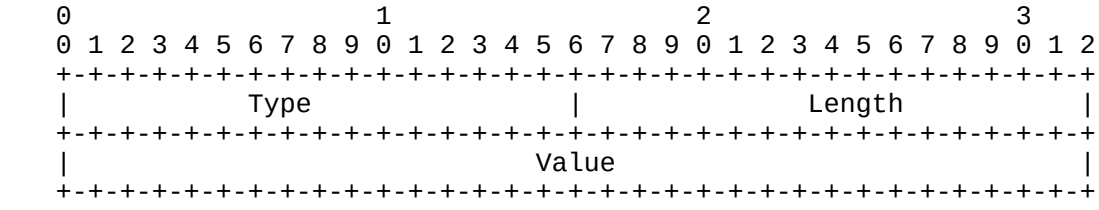

Type (2 octets): TBD11, Maximum Energy Consumption
Length (2 octets): 4.
Value (4 octet): Maximum Energy Consumption value range (0 - 4,294,967,295)

Figure 13: Board Maximum Energy Consumption

2.4.2. Board Real-Time Energy Consumption

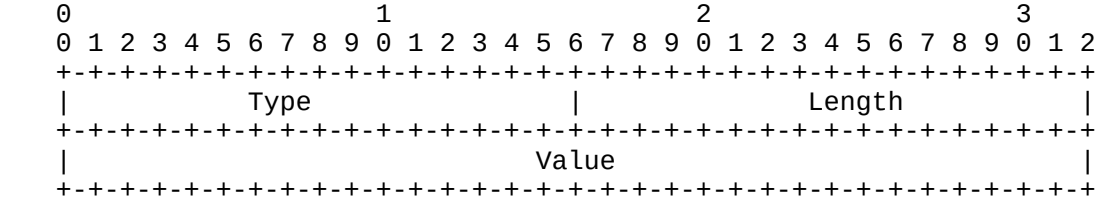

Type (2 octets): TBD12, Real-Time Energy Consumption
Length (2 octets): 4.
Value (4 octet): Real-Time Energy Consumption value range (0 - 4,294,967,295)

Figure 14: Board Real-Time Energy Consumption

2.4.3. Board Maximum Unit Energy Consumption

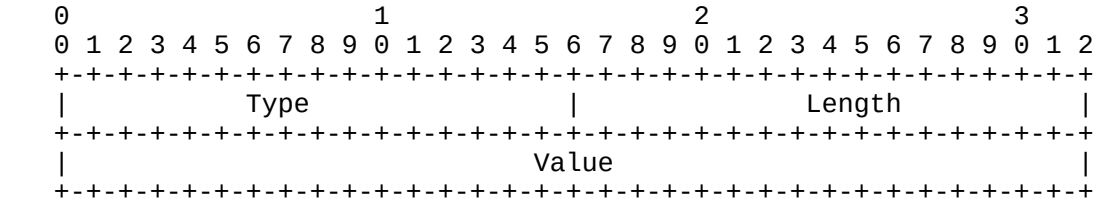

Type (2 octets): TBD13, Maximum Unit Energy Consumption
Length (2 octets): 4.
Value (4 octet): Maximum Unit Energy Consumption value range (0 - 4,294,967,295)

Figure 15: Board Maximum Unit Energy Consumption consumption

2.4.4. Real-Time Unit Energy Consumption

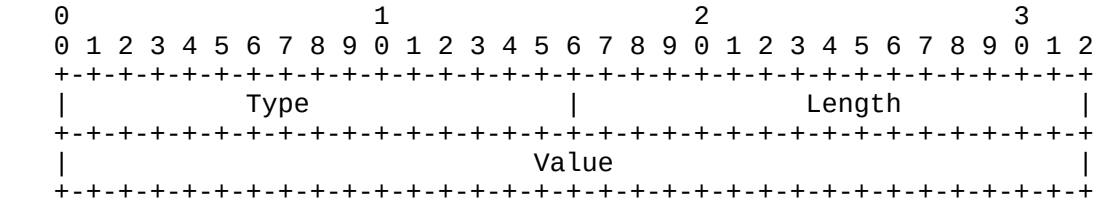

Type (2 octets): TBD14, Real-Time Unit Energy Consumption
Length (2 octets): 4.
Value (4 octet): Real-Time Unit Energy Consumption value range (0 - 4,294,967,295)

Figure 16: Board Real-Time Unit Energy Consumption

2.4.5. Board Average Unit Energy Consumption

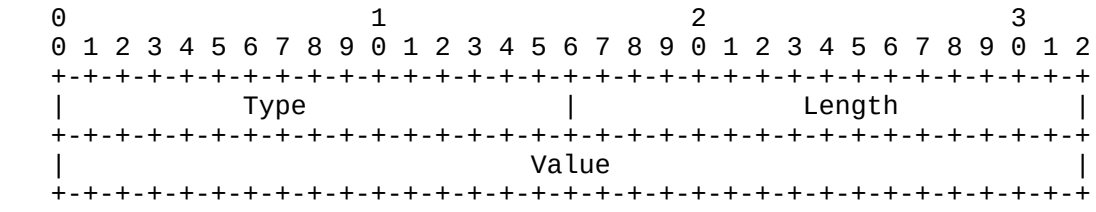

Type (2 octets): TBD15, Average Unit Energy Consumption
Length (2 octets): 4.
Value (4 octet): Average Unit Energy Consumption value range (0 - 4,294,967,295)

Figure 17: Board Average Unit Energy Consumption

3. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the "Security Considerations" section of [RFC4271] for a discussion of BGP security.

Security considerations for acquiring and distributing BGP-LS information are discussed in [RFC7752].

4. IANA Considerations

4.1. Link energy consumption TLV

A new Link Attribute TLV called "Link energy consumption" is requested from "Link Attribute TLVs" in this document for Link energy consumption associated with BGP-LS Attribute of a Link NLRI.

Value	Type
Description	Link energy consumption
Reference	[This Document]

TBD1, Interface Maximum Unit Energy Consumption
TBD2, Interface Real-Time Unit Energy Consumption
TBD3, Interface Average Unit Energy Consumption

4.2. Node energy consumption TLV

New Node Attribute TLVs called "Node energy consumption" is requested from "Node Attribute TLVs" in this document for energy consumption associated with BGP-LS Attribute of a Node NLRI.

Value	TBD6, TBD7, TBD10
Description	Node energy consumption
Reference	[This Document]

TBD4: Node Maximum Energy Consumption
TBD5, Node Real-Time Energy Consumption
TBD6, Node Maximum Unit Energy Consumption
TBD7, Node Real-Time Unit Energy Consumption
TBD8, Node Average Unit Energy Consumption

4.3. Board NLRI Type

Type	NLRI Type
1	Node NLRI
2	Link NLRI
3	IPv4 Topology Prefix NLRI
4	IPv6 Topology Prefix NLRI
TBD9	Board NLRK

4.4. Board ID Sub-TLV Type

A New Board ID Sub-TLV Type is request from

Sub-TLV Code	Description
TBD10	Board ID

4.5. Board energy consumption TLV

Value	Type
Description	Board energy consumption
Reference	[This Document]

TBD11: Maximum energy consumption
TBD12, nominal energy consumption
TBD13, Real-time energy consumption
TBD14, energy consumption change rate
TBD15, Average energy consumption change rate

4.6. IGP energy consumption Metric-Type for BGP-LS

Although the Metric-Type values advertised in IGP protocols do not need to be completely consistent with those advertised in BGP-LS, we recommend maintaining consistency with the values in IGP protocols when defining new values in BGP-LS. This ensures easier mapping when adding new types in the future.

Type	Description	Reference	Allowed in generic-metric
0	IGP Metric	[RFC9350] Section 5.1	No
1	Min Unidirectional Link Delay as defined	[RFC9350] Section 5.1	No

	in [RFC8570, Section 4.2],and [RFC7471, Section 4.2]	
2	Traffic Engineering Default Metric as defined in [RFC5305,Section 3.7], and [RFC3630, Section 2.5.5]	No Section 5.1

TBD16	Node Maximum Energy Consumption
TBD17	Node Real-Time Energy Consumption
TBD18	Node Maximum Unit Energy Consumption
TBD19	Node Real-Time Unit Energy Consumption
TBD20	Node Average Unit Energy Consumption
TBD21	Interface Maximum Unit Energy Consumption
TBD22	Interface Real-Time Unit Energy Consumption
TBD23	Interface Average Unit Energy Consumption

5. References

5.1. Normative References

[RFC5305] T. Li, Redback Networks, Inc., H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <https://www.rfc-editor.org/info/rfc5305>.

[RFC5311] D. McPherson, Ed., Arbor Networks, L. Ginsberg, S. Previdi, M. Shand, Cisco Systems, "Simplified Extension of Link State PDU (LSP) Space for IS-IS", RFC 5311, DOI 10.17487/RFC5311, February 2009, <https://www.rfc-editor.org/info/rfc5311>.

[draft-liu-spring-sr-policy-energy-consumption-00] Yisong Liu,C. Lin, "Computing Energy Consumption Path in Segment Routing Networks", draft-liu-spring-sr-energy-consumption-00, DOI 10.17487/draft-lin-spring-sr-for-green-metric-00, Febauary 2025, <https://www.rfc-editor.org/info/DOI 10.17487/draft-liu-spring-sr-energy-consumption-00>.

[RFC7752] H. Gredler, Ed., Individual Contributor, J. Medved, S. Previdi, Cisco Systems, Inc., A. Farrel, Juniper Networks, Inc, S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <https://www.rfc-editor.org/info/rfc7752>.

5.2. Informational References

TBD

Authors' Addresses

Yisong Liu
China Mobile
China
Email: liuyisong@chinamobile.com

Changwang Lin
New H3C Technologies
China
Email: linchangwang.04414@h3c.com