spring
Internet-Draft
Intended status: Informational
Expires: 8 January 2026

R. Arco
Nokia
L. Contreras
Telefonica

7 July 2025

Energy-aware Routing Using Flex-Algo in Segment Routing
draft-ralc-lsr-energy-aware-routing-00

Abstract

This document proposes enhancements to the Segment Routing (SR) Flexible Algorithm (Flex-Algo) framework by integrating power consumption metrics into routing decisions. It introduces metrics encompassing both node-level and link-level energy consumption attributes and proposes dynamic energy-aware path computation. By incorporating these metrics alongside traditional routing parameters, the enhanced Flex-Algo framework can enable networks to optimize routing paths for energy efficiency, and then leverage on advances router capabilities to further reduce operational costs as well as supporting sustainability objectives.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction 2
1.1. Related work 3
1.2. Terminology 4
2. Proposed Modifications for the reporting of metrics to be used by Flex-Algo 5
2.1. Additional information to be reported 6
2.1.1. Adjustment factors 6
2.1.2. Configurable default value 7
2.1.3. Power modes 7
3. Proposed TLV formats 7
3.1. Value field for reporting measured metrics 8
3.2. Value field for reporting average metrics 9
4. Integration into Path Computation 10
5. Open points for further discussion 10
6. Security Considerations 11
7. IANA Considerations 11
8. Acknowledgements 11
9. References 11
9.1. Normative References 11
9.2. Informative References 12
Authors’ Addresses 13

1. Introduction

The increasing significance of energy efficiency in networks, driven by both environmental sustainability goals and the need to reduce operational costs, highlights the necessity to develop routing mechanisms that enable path selection based on energy consumption. This includes the ability to define routing policies that leverage emerging capabilities in routers, such as power state management,to progressively reduce energy usage over time. But at the same time is also important to support legacy devices on the overall management of the network in which refers to energy-aware routing.

This document proposes enhancements to Flexible Algorithm (Flex-Algo) and its applicability to Segment Routing (SR) [RFC9350] by integrating power consumption metrics into routing decisions [I-D.ietf-idr-sr-policy-metric][RFC9256], to be populated by IGP protocols [RFC8491][RFC8665][RFC8667].

The motivation of this document is to address the growing importance of energy efficiency in networks, driven by environmental concerns and operational cost savings. The inherent flexibility and scalability of Segment Routing make it suitable for implementing energy-aware routing by leveraging energy consumption metrics to guide path selection. This document details extensions to previous works for energy data collection, path computation policies, and path issuance, aiming to reduce network carbon footprint and support sustainable network growth amid increasing traffic demands. The approach supports both MPLS SR and SRv6 data planes and includes mechanisms to avoid traffic oscillation by setting thresholds for path switching.

1.1. Related work

The relevance of energy-aware routing in general, but also in relation with Segment Routing, is reflected on some existing documents.

[I-D.li-lsr-flex-algo-energy-efficiency] proposes extensions to IGP protocols, specifically IS-IS and OSPF, to advertise energy consumption information within a network. It defines new TLVs and sub-TLVs for conveying various energy consumption metrics such as node maximum energy, real-time energy, unit energy consumption, and interface energy consumption values. These metrics enable energy-aware routing by allowing network controllers to compute paths that optimize for energy efficiency, balancing performance with operational costs. The document also introduces Flex-Algorithm constraints to exclude links or nodes from routing computations if their energy consumption exceeds specified thresholds, thereby facilitating the selection of low-power paths. These extensions aim to improve network sustainability by integrating energy consumption into routing decisions.

[I-D.liu-spring-sr-policy-energy-efficiency] proposes a method for computing energy consumption paths in Segment Routing (SR) networks to optimize traffic routing for improved energy efficiency. It outlines a framework where a network controller centrally collects energy consumption data from nodes and links within an SR domain using IGP and BGP-LS extensions or direct reporting mechanisms like NETCONF. Based on such information, the controller computes energy-efficient paths considering metrics such as maximum and real-time energy consumption at node and interface levels, and distributes optimized SR policies to head-end devices for forwarding. The approach includes mechanisms to avoid traffic oscillation by setting thresholds for path switching.

Furthermore, some other docuemnts address aspects such as energy-related metrics or data models.

Regarding metrics, [I-D.cx-green-green-metrics] define metrics focusing on attributes like power consumption, energy efficiency, and carbon footprint. It categorizes green metrics at different levels, such as equipment, flow, path, and network-wide. The metrics can be used for network optimization and benchmarking.

In terms of models, [I-D.cwbgp-green-common-energy-management] defines a common YANG data model module intended for reuse across various energy efficiency-related network management modules. It introduces identities, types, and groupings that represent energy-saving modes, power states, and energy consumption parameters. The module supports standardized monitoring and control of power and energy consumption at device and component levels. On the other hand, [I-D.cwbgp-green-common-energy-management] specifies a YANG network topology data model for energy efficiency management, augmenting existing network topology models to include energy consumption and energy-saving information at both the device and component levels. Finally, [I-D.cwbgp-green-inventory-energy-management] defines a YANG network inventory data model for energy efficiency management that captures static capability information related to energy saving modes and methods in network elements and components. It extends the network inventory base model to include energy management capabilities and power parameters.

Importantly, [I-D.belmq-green-framework] proposes a structured approach for energy efficiency management in network devices and their components. The framework describes physical power topologies, relationships among devices and components, and the semantics of power states. It can be expected that such framework could allow the energy-aware routing described by this document.

(Note. Additional documents will be considered in this section, as previous work is flourishing).

1.2. Terminology

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 RFC 2119 [RFC2119].

In addition, this document uses the terms defined in [I-D.bclp-green-terminology].

2. Proposed Modifications for the reporting of metrics to be used by Flex-Algo

To enable energy-aware routing using Segment Routing Flex-Algo it is necessary to consider metrics reporting real-time data about energy consumption at the node and link level, allowing routing decisions to optimize paths based on energy efficiency and/ or overall energy consumption.

[I-D.li-lsr-flex-algo-energy-efficiency] and [I-D.cx-green-green-metrics] propose different set of metrics.

[I-D.li-lsr-flex-algo-energy-efficiency] defines the following energy consumption metrics:

* 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).

From the set of metrics defined for node and interface (or port), the ones reporting maximum values can be considered static and then no necessary to be reported periodically. It can be expected that such information could be retrieved from the inventory, for instance by means of [I-D.cwbgp-green-inventory-energy-management].

However, the real-time and the average values are dynamic and require periodic advertisement for a timely characterization of the energy consumed by the network. This document builds on top of such metrics, proposing new structures for the TLV fields in IGP protocols so to extend FlexAlgo capabilities to provide energy-aware topologies.

The following section proposes additional information ot be considered as part of the reported metrics to build energy-aware routing topologies with FlexAlgo and to allow path computation based on that information.

2.1. Additional information to be reported

Additional information for the energy-related metrics is described next.

2.1.1. Adjustment factors

To ensure flexibility, network providers can apply configurable weighting factors to adjust power values during path computation.

These values can account for external environmental factors, such as temperature that may impact energy consumption, or for any other administrative consideration of interest for the network provider. Moreover, that factor could reflect aspects such as the ratio of renewable enrgy sources of applicable to the network element (e.g., by extrapolating the ratio of renewable energy feeding the Point of Presence hosting the network element) at the time of reporting the metric. This can be useful in situations where nodes operate in diverse physical conditions.

The following configurable parameters

* Node-Specific Adjustment Factor: a node-level factor allows operators to increase or decrease the advertised power value for specific nodes.

* Link-Specific Adjustment Factor: Similarly, a per-link adjustment factor modifies the reported power values for individual links.

The adjustment factor is defined in the range of 0 to 1, with the aim of multiplying the reported metric by such factor. The adjustment factor can be assumed to be populated time to time by an external network management entity (e.g., the network controller) based on policies or external administrative information (that could be the case of updating the percentage of renewable energy). The calculation of these adjustment factors will be detailed in next versions of the document.

2.1.2. Configurable default value

Energy-aware routing must address cases where not all network elements provide power metrics, as could be teh case of legacy nodes. These factors help balance energy-aware routing decisions when data is incomplete or inconsistent.

For network elements that do not advertise energy metrics a default static value is configured.

* Node-Specific static default value: A node-level static value.

* Link-Specific static default value: Similarly, a per-link static value is defined.

Nodes reporting static node or per-link values can be taking into account as a criterion for the path calculation providing the option to avoid use of paths with static default values, or prefer paths not containing static default values .

2.1.3. Power modes

The metrics being reported can be collected under different power state modes, thus being important to provide such context to the measurement for optimization purposes.

IANA [IANA-power] maintains different registries for enumeration of power states. Thus indication of the registry used plus the power state in place is an important information.

3. Proposed TLV formats

Note. This version describes only TLVs for IS-IS. Next versions will nclude TLVs for OSPF.

The proposed TLV fields follow a generic structure as shown in the next figure.

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |      Type     |     Length    |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                             Value                             |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 1: Generic format for the TLV

The Type octect serves for the purpose of identifying the type of the TLV for its proper interpretation. The Lenght octect serves for indicating the length of the entire TLV structure. The Value field follows the encoding described in the next sections.

3.1. Value field for reporting measured metrics

When reporting an absolute power consumption metric, the Value field has the following structure:

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |A|   AF Value  | Rsv |D|P| Reg.|          Power Mode           |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |   Power Value (entire part)   |    Power Value (fraction)     |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: Value field for absolute power consumption metrics

The meaning of the fields is as follows:

* A (1 bit): this bit declares the presence of the adjustment factor. The value 0 means that no adjustment factor applies (then implying that the field AF Value is ignored), while the value 1 means that an adjustment factor must be applied.

* AF VAlue (7 bits): this field reports the value of the adjustment factor to be applied to the reported metric. It represents a decimal value between 0 and 100, that multiplies the power metric (as expressed by the field Power Value).

* Rsv (3 bits): this bits are reserved for future use.

* D (1 bit): this bit indicates if the power metric (as expressed by the field Power Value) is a default value for backward compatibility with nodes not capable of reporting power metrics

directly. A value of 1 means that the reported metric is a default metric, while the value 0 means that the metric reported is result of a mesurement by sensors in the network element.

* P (1 bit): when set to 1 this bit means that a power state is associated to the reported metric.

* Registry (3 bit): this field indicates the IANA registry used [IANA-power] for interpreting the identifyer of the power state reported in the field Power Mode.

* Power Mode (8 bits): this field identify the power state which applies to the reported metric.

* Power Value (32 bits) provide the actual power measurement which is divided in an entire part (the first 16 bits) and a decimal part (the last 16 bits) expressed in watts.

3.2. Value field for reporting average metrics

When reporting a relative power consumption metric (e.g., watts per unit of traffic), the Value field has the following structure:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                                +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                |         Time Interval         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Power value (entire part)   |    Power value (fraction)     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: Value field for relative power consumption metrics

The meaning of the fields is as follows:

* TIme Interval (16 bits): this fileds indicates the interval in which the calculation of the average for the relative power consumption metric, expressed in seconds.

* Power Value (32 bits) provide the average metric of power per unit of traffic which is divided in an entire part (the first 16 bits) and a decimal part (the last 16 bits), being expressed in watts per transmitted Gb.

4. Integration into Path Computation

The enhanced power metric framework introduces an energy-aware approach to path computation by incorporating both node-level and link-level power attributes. This framework enables more efficient and sustainable network operations by quantifying energy usage across all elements contributing to a forwarding path.

At the node level, a Node Score is calculated by aggregating power efficiency ratings from all hardware components within a network device (e.g., forwarding engines, line cards, and control plane modules). At the link level, a Link Score is determined by summing the power consumption values of each link participating in the candidate path. These two components are combined by the Path Selection Algorithm, which computes the total Path Score as the sum of the Node Score and the Link Scores along the path. This score represents the cumulative energy cost of routing traffic through that path.

During path computation, Segment Routing (SR) controllers leverage this framework to prioritize paths with the lowest cumulative power scores, effectively favoring energy-efficient routes. The algorithm may also operate in hybrid mode, where power metrics are considered alongside traditional routing metrics such as latency, bandwidth, or reliability, enabling flexible, policy-driven optimization strategies tailored to the network’s operational goals.

5. Open points for further discussion

A number of open points require further discussion.

* Check whether the adjustment factor and power mode should be applied to the relative W/Gb measurement.

* If power mode is applied to the relative measurement, determine whether to indicate if the usage was total or partial during the measurement period.

* Define how frequently to report the measurement (in the case of relative metrics, it may make sense to report at the same cadence as traffic monitoring).

* Analyze how to report the power consumption of line cards.

* Determine how to anticipate the impact of adding a new flow on the node or link power consumption.

* Assess how to handle scenarios where computation leads to a change in power mode.

* Consider defining custom power modes.

6. Security Considerations

To be completed.

7. IANA Considerations

To be provided.

8. Acknowledgements

TBC

9. References

9.1. Normative References

[RFC8491] Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg, "Signaling Maximum SID Depth (MSD) Using IS-IS", RFC 8491, DOI 10.17487/RFC8491, November 2018, <https://www.rfc-editor.org/info/rfc8491>.

[RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <https://www.rfc-editor.org/info/rfc8665>.

[RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <https://www.rfc-editor.org/info/rfc8667>.

[RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <https://www.rfc-editor.org/info/rfc9256>.

[RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <https://www.rfc-editor.org/info/rfc9350>.

9.2. Informative References

[I-D.bclp-green-terminology]
Liu, P. C., Boucadair, M., Wu, Q., Contreras, L. M., and M. P. Palmero, "Terminology for Energy Efficiency Network Management", Work in Progress, Internet-Draft, draft-bclp-green-terminology-02, 14 June 2025, <https://datatracker.ietf.org/doc/html/draft-bclp-green-terminology-02>.

[I-D.belmq-green-framework]
Claise, B., Contreras, L. M., Lindblad, J., Palmero, M. P., Stephan, E., and Q. Wu, "Framework for Energy Efficiency Management", Work in Progress, Internet-Draft, draft-belmq-green-framework-03, 13 June 2025, <https://datatracker.ietf.org/doc/html/draft-belmq-green-framework-03>.

[I-D.cwbgp-green-common-energy-management]
Ma, Q., Wu, Q., Stephan, E., de Dios, O. G., Pignataro, C., and S. Han, "A Common YANG Data Model for Energy Efficiency Network Management", Work in Progress, Internet-Draft, draft-cwbgp-green-common-energy-management-00, 2 March 2025, <https://datatracker.ietf.org/doc/html/draft-cwbgp-green-common-energy-management-00>.

[I-D.cwbgp-green-inventory-energy-management]
Chen, G., Wu, Q., Stephan, E., de Dios, O. G., Pignataro, C., and S. Han, "A Network Inventory Data Model for Energy Efficiency Management", Work in Progress, Internet-Draft, draft-cwbgp-green-inventory-energy-management-00, 2 March 2025, <https://datatracker.ietf.org/doc/html/draft-cwbgp-green-inventory-energy-management-00>.

[I-D.cx-green-green-metrics]
Clemm, A., Pignataro, C., Schooler, E., Ciavaglia, L., Rezaki, A., Mirsky, G., and J. Tantsura, "Green Networking Metrics for Environmentally Sustainable Networking", Work in Progress, Internet-Draft, draft-cx-green-green-metrics-00, 21 October 2024, <https://datatracker.ietf.org/doc/html/draft-cx-green-green-metrics-00>.

[I-D.ietf-idr-sr-policy-metric]
KaZhang, Dong, J., and K. Talaulikar, "BGP SR Policy Extensions for metric", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-metric-02, 29 December 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-metric-02>.

[I-D.li-lsr-flex-algo-energy-efficiency]
Li, J. and C. Lin, "Flexible Algorithms for Energy Efficiency", Work in Progress, Internet-Draft, draft-li-lsr-flex-algo-energy-efficiency-00, 3 March 2025, <https://datatracker.ietf.org/doc/html/draft-li-lsr-flex-algo-energy-efficiency-00>.

[I-D.liu-spring-sr-policy-energy-efficiency]
Liu, Y. and C. Lin, "Computing Energy Consumption Path in Segment Routing Networks", Work in Progress, Internet-Draft, draft-liu-spring-sr-policy-energy-efficiency-00, 3 March 2025, <https://datatracker.ietf.org/doc/html/draft-liu-spring-sr-policy-energy-efficiency-00>.

[IANA-power]
"IANA Power State Sets (https://www.iana.org/assignments/power-state-sets/power-state-sets.xhtml)", n.d..

[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/info/rfc2119>.

Authors’ Addresses

Raul Arco
Nokia
Email: raul.arco@nokia.com

Luis M. Contreras
Telefonica
Email: luismiguel.contrerasmurillo@telefonica.com