SPRING Y. Liu
Internet Draft China Mobile
Intended status: Standards Track C. Lin
Expires: September 03, 2025 New H3C Technologies
March 03, 2025

Computing Energy Consumption Path in Segment Routing Networks
draft-liu-spring-sr-policy-energy-efficiency-00

Abstract

This document describes a method for computing energy consumption paths in Segment Routing (SR) networks, aiming to optimize network traffic routing for energy efficiency, including procedures for energy consumption data collection, path calculation, and issuance, as well as considerations for data plane implementation in both MPLS SR and SRv6 networks.

Table of Contents

1. Introduction

The importance of energy consumption 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] outlines a variety of metrics that can be utilized to assess energy consumption. However, the intricate details of these metrics extend beyond the scope of this document.

[RFC9252] defines the fundamental architecture and operational principles of Segment Routing (SR) and describes the SR network programming model, which enables flexible network path control through the definition of Segment Identifiers (SIDs). This document focuses on path computation based on energy consumption information and utilizes SR to implement energy-aware path control.

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. Terminology

TBD.

3. Background

In the modern digital era, network energy consumption has become a critical focus, driven by the growing demand for sustainable practices and the need to reduce operational costs. Networks consume substantial energy, leading to carbon emissions and environmental degradation. Optimizing energy usage helps reduce their carbon footprint and supports global efforts to combat climate change. Energy is a major operational expense for network operators, and improving efficiency directly lowers electricity costs, especially in large-scale networks, resulting in significant financial savings. As network traffic grows exponentially, energy-efficient designs ensure sustainable scalability without proportional increases in energy consumption, which is essential for supporting future technologies such as 5G, IoT, and cloud computing.

The source routing characteristics of SR make it a flexible, scalable, and efficient networking technology. By simplifying network control, enabling explicit path definition, and ensuring compatibility with existing technologies, SR meets the demands of modern networks for traffic engineering, fault recovery, and scalability while reducing complexity and overhead. Additionally, SR networks support network slicing, allowing the creation of independent paths for different service types.

SR networks can be utilized for energy-efficient path optimization in large-scale networks and seamlessly integrate with existing

IPv4/IPv6 infrastructures. By collecting energy consumption data from each node and link, SR enables the planning of energy-efficient paths based on routing policies, thereby achieving the goal of reducing overall network energy consumption.

The motivations for addressing energy consumption in SR networks include, but are not limited to:

Reducing energy consumption in network communications by selecting energy-efficient paths and leveraging energy-related information associated with SR paths and policies.

Allowing the source node or controller/PCE to use energy consumption metrics as constraints and optimization criteria for path computation, thereby optimizing the routing of network communications.

4. Energy consumption

Based on the scope of energy consumption measurement, it includes overall device energy consumption, board-level energy consumption, and interface-level energy consumption. Since routing protocols typically use node-level or interface-level energy consumption information for path selection, energy consumption measurements can be conducted at the overall device or board level. However, when advertising the information, board-level energy consumption can be converted into corresponding interface-level information for dissemination.

Energy consumption metrics, measured in watts per gigabyte (W/GB), indicate the energy consumed for every gigabyte of data transmitted. Based on the measurement objectives, these metrics can be classified into the following types: maximum energy consumption, real-time energy consumption.

1) Maximum Energy Consumption: The energy consumed per unit of traffic when the device operates at maximum load.

2) Real-Time Energy Consumption: The energy consumed per unit of traffic under current operating conditions.

The first metric is a static parameter of the device, while the second one is a dynamic parameter that requires real-time measurement and dissemination.

When the device is not currently forwarding traffic, the real-time energy consumption is meaningless. In such cases, maximum energy consumption can be used to calculate the path.

Differentiating by the scope of energy consumption testing, it includes overall energy consumption and interface energy consumption.

1) overall energy consumption: Measuring energy consumption of the device as a whole.

2) Energy consumption by interface: Measuring energy consumption at the granularity of interfaces. Generally, measuring energy consumption by interface is challenging to implement on devices, so a rough measurement can be conducted on the entire board and then averaged for each interface.

5. Mechanism

The framework of computing energy consumption path in SR networks: The controller centrally collects energy consumption information from all nodes within the SR network domain, computes the most energy-efficient path uniformly, and distributes the optimized path as SR-policy to head end.

```
                          +------------------+
                 +--------|Network Controller| Energy Consumption
                 |        +--------/|\-------+
                 |                  |
             SR-Policy    Energy Consumption Collection
                 |                  |
              +-\|/-+   +-----------|-----------+   +-----+
    Handling  |Head |---|    Segment Routing    |---|End  |
    behaviors |Point|   |    Network Domain     |   |Point|
              |     |   |  PE ----- P ------ PE |   |     |
              +-----+   +-----------------------+   +-----+
```

Figure 1. Framework of Computing Energy Consumption path in SR network

5.1. Energy Consumption Collection

Energy consumption information is disseminated and collected within a SR network domain through IGP protocol extensions. In inter-domain scenarios, it can be propagated and collected using BGP protocol extensions by BGP-LS extensions.

Energy consumption information is collected between the SR network domain and the Network Controller using methods such as YANG, NETCONF, and SNMP.

The specific information classification of Energy consumption is detailed in section 4.

5.2. Path Calculation Based on Energy Consumption

The network controller selects network paths based on the collected energy consumption information and calculates the paths according to the specified policy. During the calculation, both node energy consumption and link energy consumption are considered. If the device advertises link energy consumption, it is prioritized; otherwise, node energy consumption is used.

These consumption metrics may include maximum energy consumption, real-time energy consumption. When selecting energy consumption parameters, if the device is currently forwarding traffic, the real-time energy consumption is used as the energy consumption parameter for path selection. When the device is not forwarding traffic, the maximum energy consumption is used.

During the calculation process, nodes and links that do not meet the energy consumption criteria are excluded, and the path with the lowest energy consumption is prioritized for selection.

5.3. Issuance of Path

The network controller distributes path to the head end. This distribution can be performed using YANG, BGP or PCEP. The head end then conducts network forwarding based on the distributed SR-Policy. When using YANG, BGP and PCEP, necessary expansions for the energy consumption metric should be made.

6. Procedures

6.1. Energy Consumption Collection

Energy consumption information can be integrated into network topology as attributes of nodes and links, serving as criteria for routing calculations.

Energy consumption information can be directly reported to the controller by each node through the NETCONF reporting mechanism.

Alternatively, energy consumption information can be propagated within the domain through IGP flooding and then reported to the controller via BGP-LS at a designated point. To prevent frequent changes in energy consumption information from causing excessive updates to IGP LSPs, a refresh interval must be established, during which the energy consumption information in the LSP remains unchanged. If the refresh interval is too long, the energy consumption information may become outdated; if it is too short, it could lead to frequent LSP flooding.

6.2. Path calculation based on Energy Consumption

When performing routing calculations, the controller can adopt various strategies based on energy consumption metrics. It may exclude nodes and links with excessively high maximum energy consumption, filter out those with high traffic-related energy consumption, or eliminate nodes and links with significant energy consumption fluctuation rates. The specific strategy can rely on a single energy consumption parameter or a combination of multiple parameters for decision-making.

When planning paths, the network can be divided into different topologies using Flex-Algo and Multi-Topo technologies to accommodate varying energy consumption requirements.

To prevent traffic oscillation, the controller must set a threshold when calculating paths based on energy consumption information. Traffic will only be switched to a new path if the calculated energy consumption change exceeds this threshold.

6.3. Data Planes

After the controller performs routing calculations and generates the path, it can deliver the path to the headend via PCEP or NETCONF. Depending on the data plane, the generated path can be implemented as an SR-Policy or SRv6-Policy.

For an MPLS SR network, during route calculation, energy consumption information is combined with node label and adjacent label information. By specifying node labels and adjacent labels, nodes and links can be selected while excluding those with high energy consumption.

For SRv6 networks, during route calculation, energy consumption information is combined with node SIDs and adjacent End.X SIDs. By specifying node SIDs and adjacent End.X SIDs, nodes and links can be selected while excluding those with high energy consumption.

7. Use Case

7.1. Path Calculation Based on Maximum Energy Consumption

```
                             +------------------+
                    +--------|Network Controller|
                    |        +------------------+
                    |
                    |
                    |                 Max:10
                 +-\|/-+     +---------P1-------+    +-----+
                 |     |   10|                  |10  |     |
                 |Head |--- PE1                PE2---|End  |
                 |Point|     |      Max:15      |    |Point|
                 +-----+     +---------P2-------+    +-----+
```

When calculating the energy consumption path based on maximum energy consumption, the path computation is performed using the node energy consumption, interface energy consumption, or board energy consumption information published by each node.

The maximum energy consumption is a fixed value and does not require dynamic measurement or updates.

During the energy consumption path calculation, both primary and backup paths are supported.

In the diagram, the maximum energy consumption of P1 is 10, while that of P2 is 15. Therefore, the computed primary path is PE1 -> P1 -> PE2, and the backup path is PE1 -> P2 -> PE2.

In the event of a failure on the primary path, such as a fault in P1, traffic is quickly switched to the backup path PE1 -> P2 -> PE2 for forwarding.

7.2. Path Calculation Based on Average Energy Consumption

```
                         +------------------+
                +--------|Network Controller|
                |        +------------------+
                |
                |                Real:10
                |                Real:20
                |                Real:30
             +-\|/-+     +---------P1-------+    +-----+
             |     | 10  |                  |10  |     |
             |Head |--- PE1                PE2---|End  |
             |Point|     |      Real:15     |    |Point|
             +-----+     +---------P2-------+    +-----+
```

When calculating the energy consumption path based on the real-time energy consumption, the path computation is performed using the node energy consumption, interface energy consumption, or board energy consumption information published by each node.

During the energy consumption path calculation, both primary and backup paths are supported. In the event of a failure on the primary path, traffic can be quickly switched to the backup path.

In the diagram, the real-time energy consumption of P2 is 15, while that of other devices is 10. The primary path is calculated as PE1 -> P1 -> PE2, with a path energy consumption of 30, and the backup path is PE1 -> P2 -> PE2, with a path energy consumption of 35.

If the primary path fails, such as a fault in P1, traffic is quickly switched to the backup path PE1 -> P2 -> PE2 for forwarding.

When the path energy consumption changes, the path energy consumption must be recalculated. To avoid traffic oscillation, switching to a new path is only triggered when the energy consumption of the new path falls below a threshold relative to the original path.

For example, we set the threshold to 80%, meaning the new path is only adopted if its energy consumption is less than 80% of the original path.

For instance, when the energy consumption of P1 increases to 20, the energy consumption of PE1 -> P1 -> PE2 becomes 40, while that of PE1 -> P2 -> PE2 remains 35. At this point, the backup path energy consumption is 87.5% of the original path, which does not meet the 80% threshold, so no path switch occurs.

However, when the energy consumption of P1 further increases to 30, the original path energy consumption becomes 50, and the backup path energy consumption is now 70% of the original path. This triggers a path switch, and the controller updates the primary path to PE1 -> P2 -> PE2 and the backup path to PE1 -> P1 -> PE2.

8. IANA Considerations

This document does not have any IANA requests.

9. Security Considerations

TBD.

10. References

10.1. Normative References

[draft-cx-green-green-metrics] A. Clemm, Ed., Santa Clara University, C. Pignataro, Ed., NC State University, E. Schooler, University of Oxford, L. Ciavaglia, A. Rezaki, Nokia, G. Mirsky, Ericsson, J. Tantsura, Nvidia, "Green Networking Metrics for Environmentally Sustainable Networking", draft-cx-green-green-metrics, DOI 10.17487/draft-cx-green-green-metrics, October 2024, <https://www.rfc-editor.org/info/draft-cx-green-green-metrics>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <https://www.rfc-editor.org/info/rfc8402>.

[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>.

10.2. 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