

Network Working Group  
Internet Draft  
Intended status: Standards Track  
Expires: August 21, 2024

L. Gong  
W. Cheng  
China Mobile  
C. Lin  
New H3C Technologies  
A. Lindem  
LabN Consulting LLC  
R. Chen  
ZTE Corporation  
February 23, 2024

**Advertising Unreachable Links in OSPF  
draft-gong-lsr-ospf-unreachable-link-05**

Abstract

This document proposes the method to advertise links as unreachable in OSPF. In some scenarios, there are requirements to advertise unreachable links in OSPF for purposes other than building the normal Shortest Path Tree.

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), 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/1id-abstracts.txt>

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

This Internet-Draft will expire on August 21, 2024.



Copyright Notice

Copyright (c) 2023 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 (<http://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 Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction.....2](#)
- [1.1. Requirements Language.....3](#)
- [2. Use Case.....3](#)
- [2.1. Case 1: Traffic Engineering.....3](#)
- [2.2. Case 2: Flexible Algorithm.....3](#)
- [3. Solution based on MaxLinkMetric.....4](#)
- [4. Backward Compatibility.....5](#)
- [4.1. Stub Router Advertisement Backward Compatibility.....6](#)
- [5. Management Considerations.....6](#)
- [6. Security Considerations.....6](#)
- [7. IANA Considerations.....6](#)
- [8. References.....7](#)
- [8.1. Normative References.....7](#)
- [8.2. Informative References.....7](#)
- [Contributors.....8](#)
- [Authors' Addresses.....8](#)

**1. Introduction**

In some scenarios, there are requirements to advertise unreachable links in OSPF for purposes other than building the normal Shortest Path Tree. One example is a link that is available for Traffic Engineering (TE), but not for hop-by-hop routing. Another example is that specific links with dedicated resources for network slicing are included in Flexible Algorithm (Flex-Algorithm), but should be excluded in the default topology.

This document proposes the method to advertise unreachable links in OSPF.



### 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. Use Case

### 2.1. Case 1: Traffic Engineering

A network topology is shown in Figure 1. There is a link only available for Traffic Engineering between Node A and E. If that link is reachable in the SPF computation, undesired flows of best-effort traffic service may utilize the link.




Figure 1: Network Topology

### 2.2. Case 2: Flexible Algorithm

A network topology is shown in Figure 2. Nodes A, B, C, and D have an extra link between each other. These links have an Extended Administrative Group (EAG) [[RFC7308](#)] attribute specifying the "red" color.




Figure 2: Network Topology

Flex-Algorithm 128 is enabled on Nodes A, B, C, and D, with an EAG rule of including "red". Flex-Algorithm allows an IGP to compute the



paths along the constrained topology. The topology used by Flex-Algorithm 128 is shown in Figure 3.



Figure 3: Topology of Flex-Algorithm 128

Flex-Algorithm 128 is used to transmit particular flows, such as for a network slice. The "red" links used by Flex-Algorithm 128 are sub-interfaces with dedicated queues for bandwidth guarantee. So, it is expected that only the particular flows are transmitted on these links using Flex-Algorithm 128. However, these links are also contained in the default topology used by normal SPF calculation, and unexpected flows of best-effort service may be steered onto these links. Therefore, it is a problem that the dedicated links for Flex-Algorithm are still reachable in normal SPF calculation.

If all the "red" links are advertised as unreachable, the default topology used in normal SPF calculation will be as Figure 4. This allows only the network slice traffic will be steered into the "red" links by Flex-Algorithm 128.




Figure 4: SPF Topology after Excluding Unreachable Links

### 3. Solution based on MaxLinkMetric

This document specifies that if a link is advertised with the MaxLinkMetric (0xffff), it MUST NOT be considered during the normal SPF computation.

In OSPF protocol, there are some inconsistencies when a link is advertised with the MaxLinkMetric (0xffff). [RFC1247] specified that, if the cost of the link is 0xffff, the link should not be used for data traffic. However, this was changed in [RFC1583] and subsequent OSPF versions to not treat links with the cost 0xffff as unreachable.





However, such inconsistency may lead to routing loops. For example, in the network shown as Figure 5, link D-F is advertised with MaxLinkMetric (65535/0xffff). Router A supports MaxLinkMetric, but router B does not. Router A sees link D-F as reachable, and the shortest path to F is A->B->D->F. Router B sees link D-F as unreachable, and the shortest path to F is B->A->C->E->F. As a result, A forwards the packets to B, but B returns them to A, which causes routing loops.




Figure 5: Inconsistency of MaxLinkMetric Causing Loops

To improve backward compatibility, this document defines that all routers supporting MaxLinkMetric must advertise a Router Information (RI) LSA with a Router Information Capabilities TLV [RFC7770] including the following Router Information Capability Bit:

Bit	Capabilities
TBD	MaxLinkMetric support

Upon detecting the presence of a reachable Router-LSA without a companion RI LSA that has the bit set, all routers in the area MUST recalculate routes without considering MaxLinkMetric.

MaxLinkMetric is applicable for the following TLVs/LSAs:

- o The Router-LSA [RFC2328]
- o The OSPFv2 Extended Link TLV of OSPFv2 Extended Link Opaque LSA [RFC7684]
- o The Router-Link TLV of OSPFv3 E-Router-LSA [RFC8362]

#### 4. Backward Compatibility

To avoid topology inconsistency and achieve backward compatibility, routers MUST advertise the corresponding capability as described in [Section 3](#).



Upon detecting the absence of that capability from any router in the same area, all routers MUST recalculate routes without considering MaxLinkMetric.

#### **4.1. Stub Router Advertisement Backward Compatibility**

Stub Router Advertisement [[RFC6987](#)] also uses MaxLinkMetric (0xffff) to indicate a router-LSA link should not be used for transit traffic.

When an OSPFv2 router supports [[RFC6987](#)] and the MaxLinkMetric capability defined in this document, it MUST also support [[RFC8770](#)]. When announcing itself as a stub router, it MUST set the H-bit in the router-LSA and advertise all its non-stub links with a link cost of MaxLinkMetric - 1 (0xffff). Since MaxLinkMetric will not be used to indicate a link is unreachable unless all OSPFv2 routers support this specification as specified in [section 3](#), all routers will also support the H-bit and the usage of MaxLinkMetric - 1 to indicate a link should not be used for transit traffic.

An OSPFv3 router can simply use the R-bit [[RFC5340](#)] for stub router advertisement.

### **5. Management Considerations**

Support of the MaxLinkMetric capability SHOULD be configurable.

In some networks, the operator may still want links with maximum metric to be treated as reachable. For example, the auto-costing of links is used and there is a mix of low-speed and high-speed links. In such cases, the updated routers can disable the MaxLinkMetric capability and still treat links with maximum metric as reachable.

It is also RECOMMENDED that implementations supporting this document and auto-costing limit the maximum cost to MaxLinkMetric - 1 (0xffff).

### **6. Security Considerations**

The document does not introduce any new security issues into the OSPF protocol.

### **7. IANA Considerations**

This document defines a new bit in the registry "OSPF Router Informational Capability Bits":



Bit Number	Capability Name	Reference
TBA	MaxLinkMetric support	This document

## 8. References

### 8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.
- [RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", [RFC 7684](#), DOI 10.17487/RFC7684, November 2015, <<https://www.rfc-editor.org/info/rfc7684>>.
- [RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", [RFC 7770](#), DOI 10.17487/RFC7770, February 2016, <<http://www.rfc-editor.org/info/rfc7770>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), May 2017.
- [RFC8362] Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", DOI 10.17487/RFC8362, [RFC 8362](#), April 2018, <<https://www.rfc-editor.org/info/rfc8362>>.

### 8.2. Informative References

- [RFC1247] Moy, J., "OSPF Version 2", [RFC 1247](#), July 1991.
- [RFC1583] Moy, J., "OSPF Version 2", [RFC 1583](#), March 1994.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", [RFC 5340](#), DOI 10.17487/RFC5340, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.
- [RFC6987] Retana, A., Nguyen, L., Zinin, A., White, R., and D. McPherson, "OSPF Stub Router Advertisement", [RFC 6987](#), DOI 10.17487/RFC6987, September 2013, <<http://www.rfc-editor.org/info/rfc6987>>.



[RFC7308] Osborne, E., "Extended Administrative Groups in MPLS Traffic Engineering (MPLS-TE)", [RFC 7308](#), DOI 10.17487/RFC7308, July 2014, <<https://www.rfc-editor.org/info/rfc7308>>.

[RFC8770] Patel, K., Pillay-Esnault, P., Bhardwaj, M., and S. Bayraktar, "Host Router Support for OSPFv2", [RFC 7308](#), DOI 10.17487/RFC8770, April 2020, <<https://www.rfc-editor.org/info/rfc8770>>.

#### Contributors

Mengxiao Chen  
New H3C Technologies  
China  
Email: chen.mengxiao@h3c.com

Yanrong Liang  
Ruijie Networks Co., Ltd.  
China  
Email: liangyanrong@ruijie.com.cn

#### Authors' Addresses

Liyan Gong  
China Mobile  
China  
Email: gongliyan@chinamobile.com

Weiqiang Cheng  
China Mobile  
China  
Email: chengweiqiang@chinamobile.com

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

Acee Lindem  
LabN Consulting LLC  
United States of America  
Email: [acee.ietf@gmail.com](mailto:acee.ietf@gmail.com)

Ran Chen  
ZTE Corporation  
China  
Email: [chen.ran@zte.com.cn](mailto:chen.ran@zte.com.cn)