

Address-Prefix-Based Outbound Route Filter for BGP-4

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document defines a new Outbound Router Filter (ORF) type for BGP, termed "Address Prefix Outbound Route Filter", that can be used to perform address-prefix-based route filtering. This ORF-type supports prefix-length- or range-based matching, wild-card-based address prefix matching, as well as the exact address prefix matching for address families.

1. Introduction

The Outbound Route Filtering Capability defined in [BGP-ORF] provides a mechanism for a BGP speaker to send to its BGP peer a set of Outbound Route Filters (ORFs) that can be used by its peer to filter its outbound routing updates to the speaker.

This documents defines a new ORF-type for BGP, termed "Address Prefix Outbound Route Filter (Address Prefix ORF)", that can be used to perform address-prefix-based route filtering. The Address Prefix ORF supports prefix-length- or range-based matching, wild-card-based address prefix matching, as well as the exact address prefix matching for address families [BGP-MP].

2. Address Prefix ORF-Type

The Address Prefix ORF-Type allows one to express ORFs in terms of address prefixes. That is, it provides address-prefix-based route filtering, including prefix-length- or range-based matching, as well as wild-card address prefix matching.

Conceptually, an Address Prefix ORF entry consists of the fields <Sequence, Match, Length, Prefix, Minlen, Maxlen>.

The "Sequence" field specifies the relative ordering of the entry among all the Address Prefix ORF entries.

The "Match" field specifies whether this entry is "PERMIT" (value 0) or "DENY" (value 1).

The "Length" field indicates the length (in bits) of the address prefix. A length of zero indicates a prefix that matches all (as specified by the address family) addresses (with the prefix itself of zero octets).

The "Prefix" field contains an address prefix of an address family.

The "Minlen" field indicates the minimum prefix length (in bits) that is required for "matching". The field is considered unspecified with a value of 0.

The "Maxlen" field indicates the maximum prefix length (in bits) that is required for "matching". The field is considered unspecified with a value of 0.

The fields "Sequence", "Length", "Minlen", and "Maxlen" are all unsigned integers.

This document imposes the following requirement on the values of these fields:

$$0 \leq \text{Length} < \text{Minlen} \leq \text{Maxlen}$$

However, tests related to the "Minlen" or "Maxlen" value should be omitted when the "Minlen" or "Maxlen" field (respectively) is unspecified.

In addition, the "Maxlen" value must be no more than the maximum length (in bits) of a host address for a given address family [BGP-MP].

3. Address Prefix ORF Encoding

The value of the ORF-Type for the Address Prefix ORF-Type is 64.

An Address Prefix ORF entry is encoded as follows. The "Match" field of the entry is encoded in the "Match" field of the common part [BGP-ORF], and the remaining fields of the entry are encoded in the "Type specific part", as shown in Figure 1.




Figure 1: Address Prefix ORF Encoding

Note that the "Prefix" field contains the address prefix followed by enough trailing bits to make the end of the field fall on an octet boundary. The value of the trailing bits is irrelevant.

4. Address Prefix ORF Matching

In addition to the general matching rules defined in [BGP-ORF], several Address-Prefix-ORF-specific matching rules are defined as follows.

Consider an Address Prefix ORF entry, and a route maintained by a BGP speaker with Network Layer Reachability Information (NLRI) in the form of <Prefix, Length>.

The route is considered as "no match" to the ORF entry if the NLRI is neither more specific than, nor equal to, the <Prefix, Length> fields of the ORF entry.

When the NLRI is either more specific than, or equal to, the <Prefix, Length> fields of the ORF entry, the route is considered as a match to the ORF entry only if the NLRI match condition as listed in Table 1 is satisfied.

ORF Entry		NLRI
Minlen	Maxlen	Match Condition
un-spec.	un-spec.	NLRI.length == ORF.length
specified	un-spec.	NLRI.length >= ORF.Minlen
un-spec.	specified	NLRI.length <= ORF.Maxlen
specified	specified	NLRI.length >= ORF.Minlen AND NLRI.length <= ORF.Maxlen

Table 1: Address Prefix ORF Matching

When more than one Address Prefix ORF entry matches the NLRI of the route, the "first-match" rule applies. That is, the ORF entry with the smallest sequence number (among all the matching ORF entries) is considered as the sole match, and it would determine whether the route should be advertised.

The assignment of the sequence numbers is a local matter for the BGP speaker that sends the Address Prefix ORF entries.

5. IANA Considerations

This document specifies a new Outbound Route Filtering (ORF) type, Address Prefix ORF. The value of the ORF-type is 64.

6. Security Considerations

This extension to BGP does not change the underlying security issues in [BGP-4].

7. Normative References

- [BGP-4] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [BGP-MP] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.
- [BGP-ORF] Chen, E., and Y. Rekhter, "Outbound Route Filtering Capability for BGP-4", RFC 5291, August 2008.

Authors' Addresses

Enke Chen
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134

EMail: enkechen@cisco.com

Srihari R. Sangli
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134

EMail: rsrihari@cisco.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

