

GROW  
Internet-Draft  
Updates: [7854](#) (if approved)  
Intended status: Standards Track  
Expires: 9 November 2024

M. Srivastava  
Juniper Networks  
Y. Liu  
China Mobile  
C. Lin  
New H3C Technologies  
J. Li  
China Mobile  
8 May 2024

**Definition For New BMP Statistics Type  
draft-ietf-grow-bmp-bgp-rib-stats-03**

Abstract

[RFC 7854](#) defined different BMP statistics messages types to observe interesting events that occur on the router. This document updates [RFC 7854](#) by adding new statistics type to monitor BMP rib-in and rib-out Ribs.

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 9 November 2024.

Copyright Notice

Copyright (c) 2024 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 (<https://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 Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1.](#) Introduction . . . . . [2](#)
- [1.1.](#) Requirements Language . . . . . [2](#)
- [2.](#) Statistics Definition . . . . . [2](#)
- [2.1.](#) RIB-IN Statistics Definition . . . . . [3](#)
- [2.2.](#) RIB-OUT Statistics Definition . . . . . [7](#)
- [3.](#) IANA Considerations . . . . . [8](#)
- [4.](#) Security Considerations . . . . . [8](#)
- [5.](#) Acknowledgements . . . . . [8](#)
- [6.](#) Normative References . . . . . [8](#)
- Authors' Addresses . . . . . [9](#)

**[1.](#) Introduction**

[RFC7854] defines a number of different BMP statistics types to observe interesting events that occur on the router. Stats are either counters or gauges. A 32-bit Counter is a non-negative integer that monotonically increases until it reaches a maximum value, when it wraps around and starts increasing again from 0. A 64-bit Gauge is a non-negative integer that may increase or decrease, but shall never exceed a maximum value, nor fall below a minimum one.

This document defines new gauges for BMP statistics message. The format of the BMP statistics message remains same as defined in [\[RFC7854\]](#).

**[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.](#) Statistics Definition**

This section defines different statistics type for RIB-IN and RIB-OUT monitoring type.



## **2.1. RIB-IN Statistics Definition**

- \* Type = TBD1: (64-bit Gauge) Current number of routes in Adj-RIBs-In-Pre-Policy. The value can increase or decrease based upon ongoing configuration change. Note that this counter updates stats type 7 defined in [\[RFC7854\]](#) and makes it explicit for Adj-RIBs-In-Pre-Policy.
- \* Type = TBD2: (64-bit Gauge) Current number of routes in per-AFI/SAFI Adj-RIBs-In-Pre-Policy. The value can increase or decrease based upon ongoing configuration change. Note that this counter is similar from stats type 9 defined in [\[RFC7854\]](#) and makes it explicit for Adj-RIBs-In-Pre-Policy. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.
- \* Type = TBD3: (64-bit Gauge) Current number of routes in Adj-RIBs-In-Post-Policy. The value can increase or decrease based upon ongoing configuration change.
- \* Type = TBD4: (64-bit Gauge) Current number of routes in per-AFI/SAFI Adj-RIBs-In-Post-Policy. The value can increase or decrease based upon ongoing configuration change. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.
- \* Type = TBD5: (64-bit Gauge) Current number of routes in per-AFI/SAFI rejected by inbound policy. The value can increase or decrease based upon ongoing configuration change. Note that this counter is different from stats type 0 defined in [\[RFC7854\]](#). The stats type 0 in [\[RFC7854\]](#) is the a 32 counter which is monotonically increasing number and doesn't represents the current number of routes rejected by inbound policy due to ongoing configuration changes. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.
- \* Type = TBD6: (64-bit Gauge) Number of routes in per-AFI/SAFI accepted by inbound policy. The value can increase or decrease based upon ongoing configuration change or network events. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge. Some implementations, or configurations in implementations, may discard routes that do not match policy and thus the accepted count and the rib-in counts will be identical in such cases.



- \* Type = TBD7: (64-bit Gauge) Number of routes in per-AFI/SAFI selected as primary route. The value can increase or decrease based upon ongoing configuration change or network events. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge. A primary path is a recursive or non-recursive path whose nexthop resolution ends with an adjacency [draft-ietf-rtgwg-bgp-pic](#) [[I-D.ietf-rtgwg-bgp-pic](#)]. A prefix can have more than one primary path if multipath is configured [draft-lapukhov-bgp-ecmp-considerations](#) [[I-D.lapukhov-bgp-ecmp-considerations](#)]. A best-path is also considered as a primary path.
- \* Type = TBD8: (64-bit Gauge) Number of routes in per-AFI/SAFI selected as backup route. The value can increase or decrease based upon ongoing configuration change or network events. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge. A backup path is also installed in the RIB, but it is not used until some or all primary paths become unreachable. Backup paths are used for fast convergence in the event of failures.
- \* Type = TBD9: (64-bit Gauge) Number of routes in per-AFI/SAFI suppressed by configured route damping policy. The value can increase or decrease based upon configuration change or network events. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge. Suppressed refers to a path which has been declared suppressed by the BGP Route Flap Damping mechanism as described in [Section 2.2 of \[RFC2439\]](#).
- \* Type = TBD10: (64-bit Gauge) Number of routes in per-AFI/SAFI marked as stale by any configuration. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge. Stale refers to a path which has been declared stale by the BGP Graceful Restart mechanism as described in [Section 4.1 of \[RFC4724\]](#).
- \* Type = TBD11: (64-bit Gauge) Number of routes in per-AFI/SAFI marked as stale by LLGR. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge. Stale refers to a path which has been declared stale by the BGP Long-Lived Graceful Restart mechanism as described in [Section 4.3 of \[RFC9494\]](#). This is the route that are marked stale as part of LLGR process.



- \* Type = TBD12: (64-bit Gauge) Number of routes left until reaching the received route threshold as defined in [Section 6.7 of \[RFC4271\]](#). This value can increase or decrease base on ongoing configuration changes of the peer's received route threshold.
- \* Type = TBD13: (64-bit Gauge) Number of routes in per-AFI/SAFI left until reaching the received route threshold as defined in [Section 6.7 of \[RFC4271\]](#). This value can increase or decrease base on ongoing configuration changes of the peer's received route threshold. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.
- \* Type = TBD14: (64-bit Gauge) Number of routes left until reaching the license-customized route threshold. This value is affected by whether a customized license exists for the relevant address family, and when the customized license is installed. This value can increase or decrease based on current customized information of license.

Since the license-customized route threshold is vendor specific, this type value (TBD14) SHOULD use enterprise-specific TLV encoding as described in Section 3.3 of [\[I-D.ietf-grow-bmp-tlv-ebit\]](#).

For the enterprise-specific TLV of TBD14, E-bit SHOULD set to 1, and TBD14 is 15 bits of TLV Type. As enterprise-specific Type value, TBD14 is determined by the Enterprise number, and does not have to be defined by INNA.

- \* Type = TBD15: (64-bit Gauge) Number of routes in per-AFI/SAFI left until reaching the license-customized route threshold. This value is affected by whether a customized license exists for the relevant address family, and when the customized license is installed. This value can increase or decrease based on current customized information of license. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.

It's also since the license-customized route threshold is vendor specific, this type value (TBD15) SHOULD also use enterprise-specific TLV encoding as described in Section 3.3 of [\[I-D.ietf-grow-bmp-tlv-ebit\]](#).

For the enterprise-specific TLV of TBD15, E-bit SHOULD set to 1, and TBD15 is 15 bits of TLV Type. As enterprise-specific Type value, TBD15 is determined by the Enterprise number, and also does not have to be defined by INNA.





- \* Type = TBD16: (64-bit Gauge) Current Number of routes rejected by exceeding the length threshold of as-path. This value can increase or decrease base on ongoing configuration changes of the length threshold of as-path. For as-path attribute information, please refer to [[RFC4271](#)].
- \* Type = TBD17: (64-bit Gauge) Current Number of routes in per-AFI/SAFI rejected by exceeding the length threshold of as-path. This value can increase or decrease base on ongoing configuration changes of the length threshold of as-path. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.
- \* Type = TBD18: (64-bit Gauge) Current Number of routes in per-AFI/SAFI Adj-RIBs-In-Post-Policy invalidated by verifying route origin Autonomous System (AS) number through the Route Origin Authorization (ROA) of Resource Public Key Infrastructure (RPKI) [[RFC6811](#)]. This value can increase or decrease base on ongoing ROA changes of RPKI. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.
- \* Type = TBD19: (64-bit Gauge) Current Number of routes in per-AFI/SAFI Adj-RIBs-In-Post-Policy validated by verifying route origin Autonomous System (AS) number through the Route Origin Authorization (ROA) of Resource Public Key Infrastructure (RPKI) [[RFC6811](#)]. This value can increase or decrease base on ongoing ROA changes of RPKI. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.
- \* Type = TBD20: (64-bit Gauge) Current Number of routes in per-AFI/SAFI Adj-RIBs-In-Post-Policy not found by verifying route origin Autonomous System (AS) number through the Route Origin Authorization (ROA) of Resource Public Key Infrastructure (RPKI) [[RFC6811](#)]. This value can increase or decrease base on ongoing ROA changes of RPKI. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.



## **2.2. RIB-OUT Statistics Definition**

- \* Type = TBD21: (64-bit Gauge) Current number of routes in per-AFI/SAFI rejected by outbound policy. These routes are active routes which should otherwise would have been advertised in absense of outbound policy which rejected them. The value can increase or decrease based upon ongoing configuration change. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge. This counter only considers routes distributed from loc-rib into the adj-ribs-out and does not include cases like BGP add-paths [[RFC7911](#)].
- \* Type = TBD22: (64-bit Gauge) Current Number of routes refused to be sent by exceeding the length threshold of as-path. This value can increase or decrease base on ongoing configuration changes of the length threshold of as-path. For as-path attribute information, please refer to [[RFC4271](#)].
- \* Type = TBD23: (64-bit Gauge) Current Number of routes in per-AFI/SAFI refused to be sent by exceeding the length threshold of as-path. This value can increase or decrease base on ongoing configuration changes of the length threshold of as-path. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.
- \* Type = TBD24: (64-bit Gauge) Current Number of routes in per-AFI/SAFI Adj-RIBs-Out-Post-Policy invalidated by verifying route origin Autonomous System (AS) number through the Route Origin Authorization (ROA) of Resource Public Key Infrastructure (RPKI) [[RFC6811](#)]. This value can increase or decrease base on ongoing ROA changes of RPKI. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.
- \* Type = TBD25: (64-bit Gauge) Current Number of routes in per-AFI/SAFI Adj-RIBs-Out-Post-Policy validated by verifying route origin Autonomous System (AS) number through the Route Origin Authorization (ROA) of Resource Public Key Infrastructure (RPKI) [[RFC6811](#)]. This value can increase or decrease base on ongoing ROA changes of RPKI. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.
- \* Type = TBD26: (64-bit Gauge) Current Number of routes in per-AFI/SAFI Adj-RIBs-Out-Post-Policy not found by verifying route origin Autonomous System (AS) number through the Route Origin



Authorization (ROA) of Resource Public Key Infrastructure (RPKI) [RFC6811]. This value can increase or decrease base on ongoing ROA changes of RPKI. The value is structured as: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), followed by a 64-bit Gauge.

### 3. IANA Considerations

This document requests that IANA assign the following new parameters (Except TBD14 and TBD15) to the BMP parameters name space (<https://www.iana.org/assignments/bmp-parameters/bmp-parameters.xhtml>).

### 4. Security Considerations

The considerations in [Section 11 of \[RFC7854\]](#) apply to this document. It is also believed that this document does not add any additional security considerations.

### 5. Acknowledgements

The author would like to thank Jeff Haas for his valuable input.

### 6. Normative References

[I-D.ietf-grow-bmp-tlv-ebit]

Lucente, P. and Y. Gu, "Support for Enterprise-specific TLVs in the BGP Monitoring Protocol", Work in Progress, Internet-Draft, [draft-ietf-grow-bmp-tlv-ebit-05](#), 18 March 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-grow-bmp-tlv-ebit-05>>.

[I-D.ietf-rtgwg-bgp-pic]

Bashandy, A., Filsfils, C., and P. Mohapatra, "BGP Prefix Independent Convergence", Work in Progress, Internet-Draft, [draft-ietf-rtgwg-bgp-pic-20](#), 1 October 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-bgp-pic-20>>.

[I-D.lapukhov-bgp-ecmp-considerations]

Lapukhov, P. and J. Tantsura, "Equal-Cost Multipath Considerations for BGP", Work in Progress, Internet-Draft, [draft-lapukhov-bgp-ecmp-considerations-12](#), 28 December 2023, <<https://datatracker.ietf.org/doc/html/draft-lapukhov-bgp-ecmp-considerations-12>>.



- [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>>.
- [RFC2439] Villamizar, C., Chandra, R., and R. Govindan, "BGP Route Flap Damping", [RFC 2439](#), DOI 10.17487/RFC2439, November 1998, <<https://www.rfc-editor.org/info/rfc2439>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4724] Sangli, S., Chen, E., Fernando, R., Scudder, J., and Y. Rekhter, "Graceful Restart Mechanism for BGP", [RFC 4724](#), DOI 10.17487/RFC4724, January 2007, <<https://www.rfc-editor.org/info/rfc4724>>.
- [RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", [RFC 6811](#), DOI 10.17487/RFC6811, January 2013, <<https://www.rfc-editor.org/info/rfc6811>>.
- [RFC7854] Scudder, J., Ed., Fernando, R., and S. Stuart, "BGP Monitoring Protocol (BMP)", [RFC 7854](#), DOI 10.17487/RFC7854, June 2016, <<https://www.rfc-editor.org/info/rfc7854>>.
- [RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", [RFC 7911](#), DOI 10.17487/RFC7911, July 2016, <<https://www.rfc-editor.org/info/rfc7911>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC9494] Uttaro, J., Chen, E., Decraene, B., and J. Scudder, "Long-Lived Graceful Restart for BGP", [RFC 9494](#), DOI 10.17487/RFC9494, November 2023, <<https://www.rfc-editor.org/info/rfc9494>>.

Authors' Addresses





Mukul Srivastava  
Juniper Networks  
10 Technology Park Dr  
Westford, MA 01886  
United States of America  
Email: msri@juniper.net

Yisong Liu  
China Mobile  
32 Xuanwumen West Street  
Beijing  
Xicheng District, 100053  
China  
Email: liuyisong@chinamobile.com

Changwang Lin  
New H3C Technologies  
8 Yongjia North Road  
Beijing  
Haidian District, 100094  
China  
Email: linchangwang.04414@h3c.com

Jinming Li  
China Mobile  
32 Xuanwumen West Street  
Beijing  
Xicheng District, 100053  
China  
Email: lijnming@chinamobile.com

