GROW Working Group
Internet-Draft
Intended status: Standards Track
Expires: 24 October 2026

T. Wu
S. Zhuang
N. Geng
Huawei

22 April 2026

BMP Extension for Monitoring Multiple BGP Instances
draft-wu-grow-bmp-multi-instance-00

Abstract

The BGP Monitoring Protocol (BMP), as defined in RFC 7854, provides a means for monitoring BGP sessions. In deployments where a single device runs multiple BGP instances simultaneously - typically a base instance together with one or more additional instances, each identified by an Instance Name and potentially running in a separate process with an independent Router-ID and Autonomous System (AS) number - the existing BMP message format does not allow a Monitoring Station to distinguish peers or routes belonging to different BGP instances on the same monitored router.

This document defines a new BMP Type-Length-Value (TLV) element, the BGP Instance Name TLV, that carries the Instance Name of the BGP instance associated with a given peer or route. Following the TLV extension framework defined in [I-D.ietf-grow-bmp-tlv], the new TLV MAY appear in Peer Up Notification, Peer Down Notification, Route Monitoring, Statistics Report, and Route Mirroring messages. The TLV enables a Monitoring Station to unambiguously attribute monitored BGP information to the correct BGP instance on the monitored router.

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 24 October 2026.

Table of Contents

1. Introduction

The BGP Monitoring Protocol (BMP) [RFC7854] defines a mechanism by which a monitored router sends BGP-related state to a BMP Monitoring Station. Each monitored BGP peer is identified on the wire by a Per-Peer Header that carries fields such as Peer Type, Peer Flags, Peer Distinguisher, Peer Address, Peer AS, Peer BGP ID, and a timestamp.

On modern routing platforms it is increasingly common to run multiple, independent BGP instances on the same physical device. A typical deployment comprises:

* One BGP base instance (the default instance traditionally implied by [RFC4271]).

* One or more additional BGP instances, each identified by a locally significant Instance Name.

Instances on the same device are strictly isolated: they MAY use identical or different AS numbers, MAY have independent Router-IDs, and are typically implemented in distinct operating-system processes. Each instance independently maintains its own set of BGP peers, address families, and routing tables. BGP instances and VRFs are orthogonal concepts: a given BGP instance may itself host multiple VRFs.

When such a device is monitored through a single BMP session (or even through multiple BMP sessions multiplexed to the same Monitoring Station), the BMP messages produced by the different BGP instances become indistinguishable. Two peers located in different BGP instances may share the same Peer Address, Peer AS, and Peer BGP ID, and a given prefix may legitimately appear in more than one instance with a different best path. As a result, the Monitoring Station cannot correctly attribute state to the originating BGP instance.

This document addresses this gap by defining a new TLV element - the BGP Instance Name TLV - that carries the Instance Name of the BGP instance associated with the BMP message. The TLV is defined within the TLV extension framework of [I-D.ietf-grow-bmp-tlv], which extends RFC 7854 to support TLV-encoded optional data across all BMP message types.

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

This document uses the following terms:

BGP Instance: An independent instantiation of the BGP protocol on a monitored router, characterized by its own BGP state machine, peer table, and RIBs. Multiple BGP instances MAY coexist on the same device and are typically implemented in separate processes.

BGP Base Instance: The default BGP instance on a monitored router;

i.e., the instance historically assumed by implementations of [RFC4271].

Instance Name: A locally significant, administratively assigned string that uniquely identifies a BGP instance on a monitored router. The Instance Name is encoded as a UTF-8 string [RFC3629].

Monitored Router: A BMP-speaking device that runs one or more BGP instances and sends BMP messages to a Monitoring Station, as defined in [RFC7854].

Monitoring Station: A BMP-receiving entity, as defined in [RFC7854].

TLV: Type-Length-Value element, as defined in [I-D.ietf-grow-bmp-tlv].

3. Problem Statement

A monitored router running multiple BGP instances produces BMP messages whose Per-Peer Headers are, taken in isolation, insufficient to identify the originating BGP instance. Specifically:

* Two BGP instances on the same router MAY use the same AS number. The Peer AS field of the Per-Peer Header therefore cannot disambiguate instances.

* Two BGP instances MAY establish peerings with peers that share the same Peer Address (for example, when the instances have overlapping address spaces, or simply by administrative coincidence). The Peer Address field therefore cannot disambiguate instances.

* The Peer Distinguisher field defined in [RFC7854] is used to express a Route Distinguisher [RFC4364] or is zero for the global instance, and is orthogonal to the concept of a BGP instance as defined in this document.

* The Peer Type field enumerates BGP peer kinds (Global, RD, Local, Loc-RIB, and extensions such as those in [RFC9069]) but does not convey which BGP instance a peer belongs to.

Consequently, when a Monitoring Station receives a Route Monitoring or Peer Up message from such a router, it cannot reliably attribute the reported information to a specific BGP instance. This is problematic for operations such as per-instance route-table reconstruction, per-instance policy auditing, and per-instance fault isolation.

This document resolves the ambiguity by tagging BMP messages with the Instance Name of the originating BGP instance.

4. BGP Instance Name TLV

This document defines a new TLV, the BGP Instance Name TLV, using the TLV encoding and semantics of [I-D.ietf-grow-bmp-tlv]. The TLV carries the Instance Name of the BGP instance with which the enclosing BMP message is associated.

4.1. TLV Format

The BGP Instance Name TLV is encoded as follows:

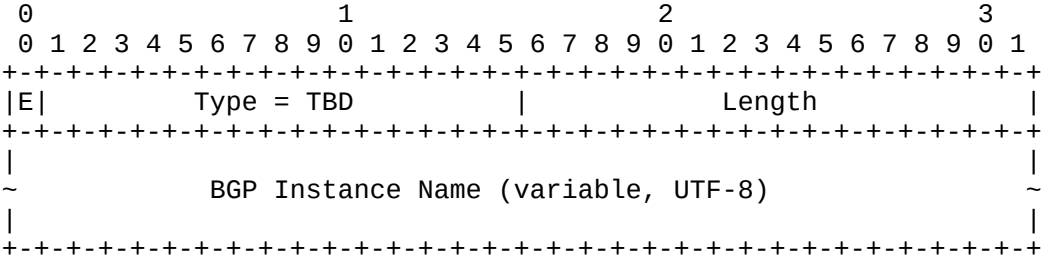

Figure 1: BGP Instance Name TLV

E (1 bit): Enterprise Bit, as defined in [I-D.ietf-grow-bmp-tlv]. MUST be set to 0 when this IANA-registered Type value is used.

Type (15 bits): Set to TBD (to be assigned by IANA, see Section 8). This value identifies the TLV as a BGP Instance Name TLV.

Length (2 octets): Length, in octets, of the BGP Instance Name field. The Length MUST be non-zero; a Length of 0 is invalid and the TLV MUST be ignored by the receiver.

BGP Instance Name (variable): The Instance Name of the BGP instance associated with the enclosing BMP message, encoded as a UTF-8 string [RFC3629]. There is no requirement to terminate the string with a null or any other character. The Instance Name is locally significant on the monitored router; no semantics beyond uniqueness on that router are implied by this document.

4.2. Scope and Cardinality

The BGP Instance Name TLV describes a property of the BGP peer (and therefore of the BGP instance) that originates the enclosing BMP message. The TLV is therefore defined to apply at the scope of the whole BMP message rather than at the scope of an individual NLRI within a Route Monitoring message. In the terminology of [I-D.ietf-grow-bmp-tlv], it is a global TLV ("G-Type").

A BMP message SHOULD contain at most one BGP Instance Name TLV. If a receiver encounters more than one BGP Instance Name TLV in the same BMP message, it MUST use the first occurrence and MUST ignore subsequent occurrences.

When this document is implemented, the monitored router SHOULD include the BGP Instance Name TLV in every BMP message emitted on behalf of a non-default BGP instance. The monitored router MAY include the TLV for messages originated by the BGP base instance; when the TLV is omitted, the Monitoring Station SHOULD treat the message as originating from the BGP base instance.

4.3. Applicability to BMP Message Types

The BGP Instance Name TLV MAY be carried in the following BMP message types:

* Peer Up Notification (Section 4.10 of [RFC7854], using the BMP Peer Up TLVs registry as updated by [RFC9736]).

* Peer Down Notification (Section 4.9 of [RFC7854], using the TLV mechanism defined in [I-D.ietf-grow-bmp-tlv]).

* Route Monitoring (Section 4.6 of [RFC7854], using the TLV mechanism defined in [I-D.ietf-grow-bmp-tlv]).

* Statistics Report (Section 4.8 of [RFC7854], using the TLV mechanism defined in [I-D.ietf-grow-bmp-tlv]).

* Route Mirroring (Section 4.7 of [RFC7854]).

This document does not define the carriage of the BGP Instance Name TLV in Initiation or Termination messages. Initiation and Termination describe properties of the BMP session itself rather than of an individual BGP instance. Future documents MAY define such usage if required by new use cases.

For Route Monitoring, Statistics Report, Peer Down, and Route Mirroring messages, the BGP Instance Name TLV is encoded in the optional TLV area defined by [I-D.ietf-grow-bmp-tlv]. For Peer Up messages, the TLV is carried in the Peer Up Information TLVs area, using the registry defined by [RFC9736].

The BGP Instance Name conveyed in Peer Up for a given peer MUST remain constant for the lifetime of the BMP-monitored peering session. A Peer Down Notification for that peer SHOULD carry the same BGP Instance Name TLV, to allow the Monitoring Station to correlate the teardown with the corresponding Peer Up even in the presence of message reordering or partial state loss.

5. Operational Considerations

Because a BGP instance is a property of the peer rather than of each individual route, the Instance Name is, strictly speaking, redundant once the Monitoring Station has received the Peer Up message for a peer. Nevertheless, including the BGP Instance Name TLV in Route Monitoring, Statistics Report, and Route Mirroring messages provides the following operational benefits:

* Self-contained messages: each BMP message can be independently interpreted without relying on prior state, which simplifies Monitoring Station implementations that do not maintain a complete peer cache.

* Robustness to state loss: a Monitoring Station that restarts, or that connects to the monitored router mid-session, can still correctly attribute messages to their BGP instance.

* Correlation across sessions: when a single Monitoring Station aggregates data from multiple BMP sessions or multiple routers, per-message instance tagging simplifies downstream processing.

Implementations SHOULD make the inclusion of the BGP Instance Name TLV in high-volume message types (notably Route Monitoring) configurable, so that operators may trade off per-message self-containment against the additional bandwidth and processing overhead imposed by repeated Instance Names.

The length of the Instance Name is bounded only by the BMP message framing. Operators SHOULD choose short, stable Instance Names to limit per-message overhead.

6. Backward Compatibility

A Monitoring Station that does not implement this extension will treat the BGP Instance Name TLV as an unknown TLV. Per [I-D.ietf-grow-bmp-tlv], unknown TLVs are silently ignored; legacy Monitoring Stations therefore remain interoperable with monitored routers that emit the new TLV.

In mixed deployments where legacy Monitoring Stations are present, operators SHOULD ensure that BGP instance disambiguation is achieved by some other means (for example, by using separate BMP sessions per BGP instance, or by arranging that peer addresses are unique across instances).

A Monitoring Station that implements this extension but receives BMP messages from a legacy monitored router - i.e., a router that does not emit the BGP Instance Name TLV - SHOULD treat such messages as originating from the BGP base instance.

7. Security Considerations

This document defines a new TLV for use within BMP messages and does not alter the underlying BMP transport, authentication, or authorization model. The security considerations of [RFC7854] and [I-D.ietf-grow-bmp-tlv] therefore apply unchanged.

The BGP Instance Name is administratively assigned on the monitored router and may reveal information about the internal structure of the operator’s network (for example, the existence and naming of tenant-specific or service-specific BGP instances). Operators who regard Instance Names as sensitive SHOULD ensure that BMP sessions are carried over a confidential and integrity-protected transport, as recommended by [RFC7854].

A misbehaving or compromised monitored router could emit forged or misleading BGP Instance Name TLVs. Monitoring Stations SHOULD NOT take policy or control actions based solely on the value of the BGP Instance Name without additional authentication of the monitored router.

8. IANA Considerations

This document requests that IANA allocate one new code point in each of the following registries under the "BGP Monitoring Protocol (BMP) Parameters" registry group, referring to this document.

8.1. BMP Peer Up and Peer Down TLVs Registry

IANA is requested to allocate a new code point in the "BMP Peer Up and Peer Down TLVs" registry (as defined by [I-D.ietf-grow-bmp-tlv], updating [RFC9736]):

Value	Description	Reference
TBD	BGP Instance Name	This document

Table 1: BMP Peer Up and Peer Down TLVs Allocation

8.2. BMP Route Monitoring TLVs Registry

IANA is requested to allocate a new code point in the "BMP Route Monitoring TLVs" registry (as defined by [I-D.ietf-grow-bmp-tlv]):

Value	Description	Reference
TBD	BGP Instance Name	This document

Table 2: BMP Route Monitoring TLVs Allocation

8.3. BMP Statistics TLVs Registry

IANA is requested to allocate a new code point in the "BMP Statistics TLVs" registry (as defined by [I-D.ietf-grow-bmp-tlv]):

Value	Description	Reference
TBD	BGP Instance Name	This document

Table 3: BMP Statistics TLVs Allocation

8.4. BMP Route Mirroring TLVs Registry

IANA is requested to allocate a new code point in the "BMP Route Mirroring TLVs" registry [RFC7854]:

Value	Description	Reference
TBD	BGP Instance Name	This document

Table 4: BMP Route Mirroring TLVs Allocation

It is RECOMMENDED that IANA assign the same numeric code point across all of the above registries, for consistency of implementation. The Value field for each allocation is defined in Section 4 of this document.

9. References

9.1. Normative References

[I-D.ietf-grow-bmp-tlv] Lucente, P. and T. Graf, "BMP v4: Extended TLV Support for BGP Monitoring Protocol (BMP)", Work in Progress, Internet-Draft, draft-ietf-grow-bmp-tlv-20, March 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-grow-bmp-tlv-20>.

[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>.

[RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <https://www.rfc-editor.org/info/rfc3629>.

[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>.

[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>.

[RFC9736] Scudder, J. and P. Lucente, "The BGP Monitoring Protocol (BMP) Peer Up Message Namespace", RFC 9736, DOI 10.17487/RFC9736, March 2025, <https://www.rfc-editor.org/info/rfc9736>.

9.2. Informative References

[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>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <https://www.rfc-editor.org/info/rfc4364>.

[RFC9069] Evens, T., Bayraktar, S., Lucente, P., Mi, P., and S. Zhuang, "Support for Local RIB in the BGP Monitoring Protocol (BMP)", RFC 9069, DOI 10.17487/RFC9069, February 2022, <https://www.rfc-editor.org/info/rfc9069>.

Authors' Addresses

Tianhao Wu
Huawei
Beijing
China
Email: wutianhao10@huawei.com

Shunwan Zhuang
Huawei
Beijing
China
Email: zhuangshunwan@huawei.com

Nan Geng
Huawei
Beijing
China
Email: gengnan@huawei.com