

Delay-Tolerant Networking Working Group
Internet Draft
Intended status: Standards Track
Expires: November 2018

S. Burleigh
JPL, Calif. Inst. Of Technology
May 20, 2018

Bundle-in-Bundle Encapsulation
draft-burleigh-dtn-bibect-01.txt

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 November 21, 2018.

Copyright Notice

Copyright (c) 2018 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.

Abstract

This document describes Bundle-in-Bundle Encapsulation (BIBE), a Delay-Tolerant Networking (DTN) Bundle Protocol (BP) "convergence layer" protocol that tunnels BP "bundles" through encapsulating bundles. The services provided by the BIBE convergence-layer protocol adapter encapsulate an outbound BP "bundle" in a BIBE convergence-layer protocol data unit for transmission as the payload of a bundle. Security measures applied to the encapsulating bundle may augment those applied to the encapsulated bundle. The protocol includes a mechanism for recovery from loss of an encapsulating bundle, called "custody transfer". This mechanism is adapted from the custody transfer procedures described in the experimental Bundle Protocol specification developed by the Delay-Tolerant Networking Research group of the Internet Research Task Force and documented in [RFC 5050](#).

Table of Contents

1.	Introduction.....	2
2.	Conventions used in this document.....	4
3.	BIBE Design Elements.....	4
3.1.	BIBE Endpoints.....	4
3.2.	BIBE Protocol Data Units.....	4
3.3.	Custody Signals.....	6
3.4.	Custody Transfer Status Reports.....	7
4.	BIBE Procedures.....	8
4.1.	BPDU Transmission.....	8
4.2.	BPDU Reception.....	8
4.3.	Retransmission Timer Expiration.....	10
4.4.	Custody Signal Reception.....	10
5.	Security Considerations.....	11
6.	IANA Considerations.....	11
7.	References.....	11
7.1.	Normative References.....	11
7.2.	Informative References.....	11
8.	Acknowledgments.....	11
Appendix A.	For More Information.....	13
Appendix B.	CDDL expression.....	14

[1.](#) Introduction

This document describes Bundle-in-Bundle Encapsulation (BIBE), a Delay-Tolerant Networking (DTN) Bundle Protocol (BP) [[RFC5050](#)] "convergence layer" protocol that tunnels BP "bundles" through encapsulating bundles.

Conformance to the bundle-in-bundle encapsulation (BIBE) specification is OPTIONAL for BP nodes. Each BP node that conforms to the BIBE specification provides a BIBE convergence-layer adapter (CLA) that is implemented within the administrative element of the BP node's application agent. Like any convergence-layer adapter, the BIBE CLA provides:

- . A transmission service that sends an outbound bundle (from the bundle protocol agent) to a peer CLA. In the case of BIBE, the sending CLA and receiving peer CLA are both BP nodes.
- . A reception service that delivers to the bundle protocol agent an inbound bundle that was sent by a peer CLA (itself a BP node) via the BIBE convergence layer protocol.

The BIBE CLA performs these services by:

- . Encapsulating outbound bundles in BIBE protocol data units, which take the form of Bundle Protocol administrative records as described later.
- . Requesting that the bundle protocol agent transmit bundles whose payloads are BIBE protocol data units.
- . Taking delivery of BIBE protocol data units that are the payloads of bundles received by the bundle protocol agent.
- . Delivering to the bundle protocol agent the bundles that are encapsulated in delivered BIBE protocol data units.

Bundle-in-bundle encapsulation may have broad utility, but the principal motivating use case is the deployment of "cross domain solutions" in secure communications. Under some circumstances a bundle may arrive at a node that is on the frontier of a sector of network topology in which augmented security is required, from which the bundle must egress at some other designated node. In that case, the bundle may be encapsulated within a bundle to which the requisite additional BP Security (BPSEC) [bpsec] extension block(s) can be attached, whose source is the point of entry into the insecure region (the "security source") and whose destination is the point of egress from the insecure region (the "security destination").

Note that:

- . If the payload of the encapsulating bundle is protected by a Bundle Confidentiality Block (BCB), then the source and destination of the encapsulated bundle are encrypted, providing defense against traffic analysis that BPSEC alone cannot offer.
- . Bundles whose payloads are BIBE protocol data units may themselves be forwarded via a BIBE convergence-layer adapter,

- enabling nested bundle encapsulation to arbitrary depth as required by security policy.
- . Moreover, in the event that no single point of egress from an insecure region of network topology can be determined at the moment a bundle is to be encapsulated, multiple copies of the bundle may be encapsulated individually and forwarded to all candidate points of egress.

The protocol includes a mechanism for recovery from loss of an encapsulating bundle, called "custody transfer". This mechanism is adapted from the custody transfer procedures described in the experimental Bundle Protocol specification developed by the Delay-Tolerant Networking Research group of the Internet Research Task Force and documented in [RFC 5050](#). Custody transfer is a convention by which the loss or corruption of BIBE encapsulating bundles can be mitigated by the exchange of other bundles, which are termed "custody signals".

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-2119](#) [[RFC2119](#)].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying [RFC-2119](#) significance.

3. BIBE Design Elements

3.1. BIBE Endpoints

BIBE convergence-layer protocol endpoints, also known as BIBE convergence-layer adapters (BCLAs), are the Administrative Elements of Bundle Protocol nodes that conform to the BIBE protocol specification. The node of which a given BCLA is one component is termed the BCLA's "local node".

3.2. BIBE Protocol Data Units

Notionally, a BCLA is assumed to implement in some way, for each neighboring node to which the local node issues Bundle Protocol Data Units (BPDUs), the following two data resources:

1. A "custodial transmission count" (CTC). A CTC is a monotonically increasing integer indicating the number of "custodial" BPDUs - that is, BPDUs for which custody transfer

was requested - that have been issued to the neighboring node by the local node since instantiation of the local node.

2. A "custodial transmission database" (CTDB), a notional array of "custodial transmission items" (CTIs). The CTDB contains one CTI for each custodial BPDU issued to the neighboring node, by the local node, for which (a) no custody disposition has yet been received in any custody signal (as discussed later) and (b) the bundle encapsulated in that BPDU has not yet been destroyed due to, e.g., time-to-live expiration. Each CTI notionally contains:
 - a. A reference to the bundle encapsulated in the corresponding BPDU.
 - b. The "transmission ID" of the corresponding BPDU, as discussed below.
 - c. A "retransmission time" indicating the time by which custody disposition for the corresponding BPDU is expected.

A BIBE protocol data unit is a Bundle Protocol administrative record whose record type code is 3 (i.e., bit pattern 0011), constructed as follows.

Each BPDU SHALL be represented as a CBOR array. The number of elements in the array SHALL be 3.

The first item of the BPDU array SHALL be the "transmission ID" for the BPDU, represented as a CBOR unsigned integer. The transmission ID for a BPDU for which custody transfer is NOT requested SHALL be zero. The transmission ID for a BPDU for which custody transfer IS requested SHALL be the current value of the local node's custodial transmission count, plus 1.

The second item of the BPDU array SHALL be the BPDU's retransmission time (i.e., the time by which custody disposition for this BPDU is expected), represented as a CBOR unsigned integer. Retransmission time for a BPDU for which custody transfer is NOT requested SHALL be zero. Retransmission time for a BPDU for which custody transfer IS requested SHALL take the form of a "DTN Time" as defined in the Bundle Protocol specification; determination of the value of retransmission time is an implementation matter that is beyond the scope of this specification and may be dynamically responsive to changes in connectivity.

The third item of the BPDU array SHALL be a single BP bundle, termed the "encapsulated bundle", represented as a CBOR byte string.

3.3. Custody Signals

A "custody signal" is defined as a Bundle Protocol administrative record whose record type code is 4 (i.e., bit pattern 0100) and whose content is constructed as follows.

The content of each custody signal SHALL be represented as a CBOR array. The number of elements in the array SHALL be 2.

The first item of the custody signal content array SHALL be a disposition code represented as a CBOR unsigned integer. Valid disposition codes are defined as follows:

+-----+-----+-----+-----+-----+-----+			
Value		Meaning	
+=====+=====+=====+=====+=====+=====+			
0		Custody accepted.	
+-----+-----+-----+-----+-----+-----+			
1		No further information.	
+-----+-----+-----+-----+-----+-----+			
2		Reserved for future use.	
+-----+-----+-----+-----+-----+-----+			
3		Redundant (reception by a node that	
		already has a copy of this bundle).	
+-----+-----+-----+-----+-----+-----+			
4		Depleted storage.	
+-----+-----+-----+-----+-----+-----+			
5		Destination endpoint ID unintelligible.	
+-----+-----+-----+-----+-----+-----+			
6		No known route destination from here.	

+-----+-----+-----+-----+-----+-----+-----+			
	7	No timely contact with next node on route.	
+-----+-----+-----+-----+-----+-----+-----+			
	8	Block unintelligible.	
+-----+-----+-----+-----+-----+-----+-----+			
	(other)	Reserved for future use.	
+-----+-----+-----+-----+-----+-----+-----+			

Figure 1: Disposition Codes

The second item of the custody signal content array SHALL be a "disposition scope report", represented as a CBOR indefinite-length array. Each item of the disposition scope report array SHALL be a "disposition scope sequence", represented as a CBOR array of two elements. The first element of each disposition scope sequence array SHALL be the first transmission ID in a sequence of 1 or more consecutive transmission IDs corresponding to BPDUs to which the custody signal's disposition is declared to apply; the second element of each disposition scope sequence array SHALL be the number of transmission IDs in that sequence. Both are represented as CBOR unsigned integers.

A custody signal constitutes an assertion by the source of that administrative bundle that the indicated disposition code applies to all BPDUs identified by the transmission IDs enumerated in the custody signal's disposition scope report. If the disposition code is zero, then the source of the custody signal has accepted custody of all bundles that were encapsulated in the indicated BPDUs. Otherwise the source of the custody signal has refused custody of all bundles that were encapsulated in the indicated BPDUs, for the indicated reason.

3.4. Custody Transfer Status Reports

A "custody transfer status report" is a bundle status report with the "reporting node attempted custody transfer" flag set to 1.

4. BIBE Procedures

4.1. BPDUs Transmission

When a BCLA is requested by the bundle protocol agent to send a bundle to the peer BCLA(s) included in the BP endpoint identified by a specified BP endpoint ID:

- . The BCLA SHALL generate, as defined in [Section 6.2](#) of the Bundle Protocol specification (a work in progress), a BPDU for which the third element of the content array is the bundle that is to be transmitted. The destination of the bundle whose payload is the BPDU (termed the "encapsulating bundle") SHALL be the specified BP endpoint. Selection of the values of the parameters governing the forwarding of the encapsulating bundle, other than the destination endpoint ID, is an implementation matter. The parameter values governing the forwarding of the BPDU's encapsulated bundle MAY be consulted for this purpose.
- . Note that any transmission request presented to a BCLA MAY request that the transmission be subject to Custody Transfer, provided that the destination EID of the request identifies a singleton endpoint.
- . If Custody Transfer is requested:
 - o The first element of the BPDU's content array MUST be the BPDU's transmission ID, which SHALL be 1 more than the current value of the BCLA's CTC for the node that is the sole occupant of the BPDU's destination endpoint.
 - o The second element of the BPDU's content array MUST be the BPDU's retransmission time as discussed in 3.2 above.
 - o The bundle protocol agent MUST add the retention constraint "Custody accepted" to the encapsulated bundle.
 - o The BCLA MAY establish a retransmission timer for the encapsulated bundle. If a retransmission timer is established, it MUST be set to expire at the BPDU's retransmission time.
- . Otherwise, the first two elements of the BPDU's content array MUST both be zero.

Note that the custody transfer retransmission timer mechanism provides a means of recovering from loss of an encapsulating bundle as indicated by non-arrival of a responding custody signal.

4.2. BPDUs Reception

When a BCLA receives a BPDU from the bundle protocol agent (that is, upon delivery of the payload of an encapsulating bundle):

- . If Custody Transfer was requested for this BPDU (as indicated by a non-zero value of transmission ID):
 - o If the encapsulated bundle has the same source node ID, creation timestamp, and (if that bundle is a fragment) fragment offset and payload length as another bundle that is currently retained at the BCLA's local node, then custody transfer redundancy MUST be handled as follows:
 - . The BCLA SHALL add the BPDU's transmission ID to the disposition scope report of a pending outbound custody signal, destined for the node that was the source of the encapsulating bundle, whose disposition is the reason code from Figure 1 for "Redundant reception".,
 - o Otherwise, if the BCLA determines that its local node can neither deliver nor forward the encapsulated bundle for any of the reasons listed in Figure 1, then custody transfer has failed. Custody transfer failure SHALL be handled as follows:
 - . The BCLA SHALL add the BPDU's transmission ID to the disposition scope report of a pending outbound custody signal, destined for the node that was the source of the encapsulating bundle, whose disposition is the reason code from Figure 2 that indicates the reason for the custody transfer failure.
 - o Otherwise, custody transfer has succeeded:
 - . The BCLA SHALL add the BPDU's transmission ID to the disposition scope report of a pending outbound custody signal, destined for the node that was the source of the encapsulating bundle, whose disposition is zero (indicating that custody was accepted).
 - o In each of these three cases:
 - . The pending outbound custody signal MAY then be issued immediately, but alternatively it MAY be issued at some time in the future, possibly enabling additional BPDUs' transmission IDs to be added to the same disposition scope report.
 - . If the "request reporting of custody transfer attempted" flag in the encapsulating bundle's status report request field is set to 1, and status reporting is enabled, a custody transfer status report whose reason code is the same as the pending outbound custody signal's disposition SHOULD be generated, destined for the report-to endpoint of the encapsulating bundle.
- . If Custody Transfer was NOT requested for this BPDU, or if Custody Transfer was requested for this BPDU and custody transfer succeeded, then the encapsulated bundle SHALL be

delivered from the convergence layer adapter to the bundle protocol agent, whereupon bundle reception SHALL be performed as defined in [section 5.6](#) of the Bundle Protocol specification (a work in progress) as usual: the encapsulated bundle may be forwarded, delivered, etc.

Note that the manner in which pending outbound custody signals are managed, disposition scope reports are aggregated, and custody signal transmission is initiated is an implementation matter that is beyond the scope of this specification. Note, however, that failure to deliver a custody signal prior to the earliest value of retransmission time among all BPDUs enumerated in the custody signal's disposition scope report may result in unnecessary retransmission of one or more BPDUs.

[4.3.](#) Retransmission Timer Expiration

Upon expiration of a retransmission timer, the BCLA SHOULD remove the corresponding CTI from the CTDB (destroying the associated retransmission timer, if any) and notify the bundle protocol agent that custodial transmission of the indicated bundle failed. This notification may cause the indicated bundle to be re-forwarded (possibly on a different route).

[4.4.](#) Custody Signal Reception

When a BCLA receives a custody signal from the bundle protocol agent (that is, upon delivery of the payload of a custody-signal-bearing bundle):

- . If the custody signal's disposition is 0 (custody acceptance), then for each transmission ID in the custody signal's disposition scope report:
 - o The bundle protocol agent MUST remove the retention constraint "Custody accepted" on the bundle referenced by the corresponding CTI.
 - o The corresponding CTI MUST be removed from the CTDB (destroying the associated retransmission timer, if any).
- . Otherwise (custody refusal), for each transmission ID in the custody signal's disposition scope report:
 - o The corresponding CTI MUST be removed from the CTDB (destroying the associated retransmission timer, if any).
 - o Any further action taken by the BCLA is implementation-specific and may depend on the reason code cited for the refusal. For example, if the custody signal's reason code was "Depleted storage", the BCLA might choose to notify the bundle protocol agent that custodial transmission of

the indicated bundle failed. If the reason code was "Redundant reception", on the other hand, this might cause the BCLA simply to instruct the bundle protocol agent to remove the retention constraint "Custody accepted" on the bundle referenced by the corresponding CTI and to revise its algorithm for computing retransmission time.

5. Security Considerations

An adversary on a DTN-based network that can delete bundles could delete a BIBE custody signal in transit. This could result in unnecessary custodial retransmission, degrading network performance.

Alternatively, an adversary on a DTN-based network that can reorder bundles could cause bundles to be delivered to a BCLA in an order that complicates the efficient construction of disposition scope reports in pending outbound custody signals. This could result in inefficient custody transfer communications, again degrading network performance.

Custody transfer in BIBE may be contraindicated in environments characterized by such attacks.

6. IANA Considerations

The BIBE specification requires IANA registration of the new BIBE administrative records (type codes 3 and 4) defined above.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

7.2. Informative References

[RFC5050] Scott, M. and S. Burleigh, "Bundle Protocol Specification", [RFC 5050](#), November 2007.

8. Acknowledgments

This work is freely adapted from [[RFC5050](#)], which was an effort of the Delay Tolerant Networking Research Group. The following DTNRG participants contributed significant technical material and/or inputs to that document: Dr. Vinton Cerf of Google, Scott Burleigh, Adrian Hooke, and Leigh Torgerson of the Jet Propulsion Laboratory,

Michael Demmer of the University of California at Berkeley, Robert Durst, Keith Scott, and Susan Symington of The MITRE Corporation, Kevin Fall of Carnegie Mellon University, Stephen Farrell of Trinity College Dublin, Peter Lovell of SPARTA, Inc., Manikantan Ramadas of Ohio University, and Howard Weiss of SPARTA, Inc.

The custody transfer procedures defined in this specification are adapted from the Aggregate Custody Signals draft specification authored in 2010-2012 by Sebastian Kuzminsky and Andrew Jenkins, then of the University of Colorado at Boulder.

Although the BIBE specification diverges in some ways from the original Bundle-in-Bundle Encapsulation Internet Draft authored by Susan Symington, Bob Durst, and Keith Scott of The MITRE Corporation ([draft-irtf-dtnrg-bundle-encapsulation-06](#), 2009), the influence of that earlier document is gratefully acknowledged.

This document was prepared using 2-Word-v2.0.template.dot.

Appendix A.**For More Information**

Please refer comments to dtn@ietf.org. The Delay Tolerant Networking Research Group (DTNRG) Web site is located at <http://www.dtnrg.org>.

Copyright (c) 2018 IETF Trust and the persons identified as authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in [Section 4.c](#) of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

Appendix B.**CDDL expression**

For informational purposes, Carsten Bormann has kindly provided an expression of the Bundle Protocol specification in the CBOR Data Definition Language (CDDL). Portions of CDDL expression that bear on the custody transfer extension are presented below, somewhat edited by the authors. Note that wherever the CDDL expression is in disagreement with the textual representation of the BP specification presented in the earlier sections of this document, the textual representation rules.

```
admin-record-choice /= BIBE-PDU
```

```
BIBE-PDU = [3, [transmission-ID: uint,  
                retransmission-time: uint,  
                encapsulated-bundle: bytes,  
                admin-common]]
```

```
admin-record-choice /= custody-signal
```

```
custody-signal = [4, [disposition-code: uint,  
                     disposition-scope-report,  
                     admin-common]]
```

```
disposition-scope-report = *disposition-scope-sequence
```

```
disposition-scope-sequence = [first-transmission-ID: uint,  
                              number-of-transmission-IDs: uint]
```

Authors' Address

Scott Burleigh
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109-8099
US
Phone: +1 818 393 3353
Email: Scott.Burleigh@jpl.nasa.gov