DTN Research Group
Internet-Draft
Intended status: Experimental
Expires: April 11, 2010

S. Symington
The MITRE Corporation
October 8, 2009

Delay-Tolerant Networking Retransmission Block
draft-irtf-dtnrg-bundle-retrans-block-06

Status of this Memo

Copyright Notice

Abstract

This document defines an optional extension block, called a Retransmission Block (RB), that may be used with the Bundle Protocol [refs.DTNBP] within the context of a Delay-Tolerant Network architecture [refs.DTNarch]. The Retransmission Block (RB) is designed to be used within a DTN that, as a matter of policy, deletes all replayed bundles from the network. It is designed to be used in a network that permits duplicate bundles to be forwarded if those bundles have been retransmitted by a custodian, that may (if possible) permit duplicate bundles to be forwarded if those bundles are in intentional or unintentional routing loops (contingent on the availability of mechanisms to distinguish looping bundles from other bundles), but that will consider all other duplicate bundles to be maliciously replayed bundles and delete them as such. The Retransmission Block is designed to be inserted into a bundle by a custodian when the custodian is retransmitting that bundle. The purpose of the RB is to mark the bundle as a custody-based retransmission so that it can be distinguished from other types of duplicate bundles and thereby be spared from deletion. This document defines the format and processing of this new Retransmission Block.

Table of Contents

1. Introduction

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 [refs.RFC2119].

The DTN bundle protocol [refs.DTNBP] defines the bundle as its protocol data unit. As discussed in the DTN Security Overview [refs.DTNsecOver], due to the resource-scarcity that characterizes DTNs, unauthorized access and use of DTNs is a serious concern. As described in the Bundle Security Protocol [refs.DTNBPsec], use of the Bundle Authentication Block (BAB) at every node in the DTN can be used to thwart an attacker wanting to launch a denial of service attack by injecting bogus or modified bundles into the network. Use of the BAB enables bogus or modified bundles to be detected and deleted at the very first node at which they are received. Use of the BAB, however, does not enable maliciously replayed bundles to be detected, because such replayed bundles will contain valid authenticators. Replayed bundles will only be deleted from the network when they expire. Given the high latency typical in some DTNs, bundles may be valid for days or weeks. For those networks in which waiting for replayed bundles to expire is not an adequate form of protection against the unauthorized use of the network posed by replayed bundles, additional measures will be required to actively detect and delete maliciously replayed bundles.

The detection and deletion of maliciously replayed bundles at any given node is not simply a matter of configuring the node to maintain a record of every bundle it receives, comparing each new bundle received against the list of already-received bundles, and deleting all duplicates. While such an approach would result in the deletion of maliciously replayed bundles, it could also result in the deletion of other types of duplicate bundles that are not replays and that shouldn't be deleted.

Indiscriminate deletion of all duplicates could result in the deletion of bundles that are in intentional or unintentional routing loops. Such bundles may legitimately be found in DTNs and their deletion may not be desirable. Unfortunately, there are currently no mechanisms available to enable a DTN node to distinguish bundles that are in loops from maliciously replayed bundles; suppression of replays currently requires suppression of looping bundles. If replay suppression is needed in a DTN that employs routing strategies that result in routing loops, the routing protocols will need to provide a mechanism for distinguishing looping bundles from maliciously replayed bundles. If such a mechanism is provided, it can be used in conjunction with the Retransmission Block as defined in this

document.

Indiscriminate deletion of all duplicates could also result in the deletion of bundles that have been retransmitted by a custodian as part of the custody transfer process. The Bundle Protocol includes a custody-based retransmission mechanism that may result in a custodian retransmitting a stored bundle when the bundle's retransmission timer expires or when the custodian receives a "failed" custody signal for the bundle. Such retransmitted bundles are duplicates of previously forwarded bundles. A DTN node that is configured to simply delete all duplicates received would delete such custodial retransmissions, thereby rendering custody transfer ineffective. In order to be able to enable custody transfer to operate correctly while also detecting and deleting malicious replays, DTN nodes need a way to determine whether or not a duplicate bundle received is a custodial retransmission so that custodial retransmissions can be spared from deletion.

This document defines an optional bundle block called a Retransmission Block (RB) that is intended to be used in DTNs that, as a matter of policy, delete all replayed bundles from the network. Such a DTN may be configured to permit duplicate bundles to be forwarded only if those duplicate bundles are bundles that have been retransmitted by a custodian. In this case, the DTN would be one in which not only replayed bundles, but also bundles resulting from intentional or unintentional routing loops would also be deleted. If, on the other hand, the routing protocols being used in the DTN enable bundles that are in loops to be distinguished from replayed bundles, then the DTN could be configured such that only those duplicate bundles that are replayed bundles are deleted.

In either case, the RB is designed to be inserted into a bundle by a custodian when the custodian is retransmitting that bundle in response to a custody transfer failure or retransmission timer expiration. The intent of the RB is to mark a custodially retransmitted bundle as such, so that when the bundle is received at downstream nodes that detect it to be a duplicate of a previously-received bundle, those nodes can understand it to be a custody-based retransmission that should be preserved rather than another type of duplicate that may (according to network policy) be deleted.

The RB is intended to enable custodially retransmitted bundles to be distinguished from other duplicates. Other mechanisms would need to be defined in order to be able to distinguish looping bundles from other duplicates. If the RB is used with duplicate suppression in the absence of these other mechanisms, then it will result in the deletion of all looping bundles in addition to all replays.

This document defines the format and processing of this new Retransmission Block. The capabilities described in this document are OPTIONAL for deployment with the Bundle Protocol. Bundle Protocol implementations claiming to support Retransmission Blocks MUST be capable of:

-Generating a Retransmission Block and inserting it into a bundle,

-Logging the relevant fields of all bundles received until those bundles expire,

-Calculating a checksum or digest (as determined by local policy) of the payload block of a bundle,

-Receiving bundles containing a Retransmission Block and using the information contained in this Retransmission Block (in conjunction with information from logged bundles and with a mechanism, if available, for determining whether a bundle is in a routing loop) to make duplicate deletion decisions, and

-Deleting a Retransmission Block from a bundle

as defined in this document.

2. Applicability Statement

The objective of the Retransmission Block (RB) is to make custodially-retransmitted bundles distinguishable from other duplicate bundles. As such, the RB is designed to be used within a DTN that does not, as a matter of policy, permit replayed bundles to be forwarded and that is willing to enforce the detection and deletion of replayed bundles by having every node

-Authenticate all bundles

-Keep track of bundles that have been received to identify which newly-received bundles are duplicates

-Implement and use the RB as a way to enable nodes to determine which duplicate bundles are custodial retransmissions

-Spare from deletion those duplicate bundles that are custodial retransmissions

2.1. Bundle Authentication Requirement

Use of the RB to distinguish custodial retransmissions from replayed bundles requires that bundles be authenticated in order to ensure that the RB was in fact inserted by a legitimate node and that the RB has not been modified since its insertion. There is no point in using the RB within a DTN that does not perform bundle authentication because DTNs that do not perform bundle authentication are susceptible to denial of service attacks caused by all types of bundles that can be modified or inserted by an adversary, not just replays; in such an environment, detecting and deleting replays does little to protect against denial of service attacks.

2.2. Deletion of All Replays, including Bundles in Routing Loops

Note that the RB is intended to be used only in DTNs that do not permit the forwarding of replayed bundles. If a DTN does permit replayed bundles to be forwarded, there is no point in using the RB. Use of the RB makes it possible to distinguish which duplicate bundles are custodial retransmissions so that they can be spared from deletion. The RB may be used in DTNs that are configured to delete all duplicates that are not custodial retransmissions. In this case, duplicates that are the result of intentional or unintentional routing loops will also be deleted along with replayed bundles. In order to be able to suppress replays in a DTN that employs routing strategies that result in routing loops, a mechanism for distinguishing a bundle that is in a routing loop from a replayed bundle will need to be provided for use in conjunction with the RB.

Currently, no such mechanisms are known to exist, so replay deletion will also result in the deletion of bundles that are in routing loops.

2.3. Universal Support for Replay Suppression

One component of a DTN node's security policy should be whether or not replays are allowed to be forwarded by that node. A node that is not allowed to forward replays should delete all duplicate bundles that are not custodial retransmissions and that cannot be determined to be the result of routing loops. In order to be applied consistently, correctly, and effectively, replay detection and deletion is something that should be enforced at all nodes in the DTN rather than something that is enforced on a node-by-node basis. If not every node supports replay detection and deletion then some replays will be allowed to remain in the network. The RB is designed to be used within a DTN in which every node is configured to detect and delete replays.

2.4. Universal Support for the Retransmission Block

While implementation of and support for the RB is optional, the RB is designed to be used within a DTN in which every node supports the RB. Failure to support the RB at one or more nodes in a DTN that, as a matter of policy, deletes replays may result in the erroneous deletion of custodially retransmitted bundles. Section 5 further discusses the ramifications of non-uniform support of the RB.

3. Retransmission Block Format

The Retransmission Block (RB) MAY be included in a bundle. A RB uses the Canonical Bundle Block Format as defined in the Bundle Protocol [refs.DTNBP]. That is, it is comprised of the following elements:

-Block-type code (one byte) - defined as in all bundle protocol blocks except the primary bundle block. The block type code for the Retransmission Block is 0x07.

-Block processing control flags (SDNV) - defined as in all bundle protocol blocks except the primary bundle block. The following block processing control flag MUST NOT be set:

-Block must be replicated in every fragment

The following block processing control flag MUST be set:

-Block contains an EID-reference field

- Block EID reference count and EID reference - composite field containing a count of EID references with a value of 1 (expressed as an SDNV) followed by a single EID reference (expressed as a pair of SDNVs). Presence of this field is indicated by the setting of the "block contains an EID reference field" flag in the block processing control flags to 1. The EID referenced MUST be that of the retransmitting custodian that inserted this block.

-Block data length (SDNV) - defined as in all bundle protocol blocks except the primary bundle block.

-Block-type-specific data field as follows:

- Retransmission sequence number (SDNV) - An unsigned integer indicating the number of times this bundle has been retransmitted by this custodian.

The Retransmission Block format is as follows:

```
+------+--------------+-----------------------------+--------------+
|type  |flags (SDNV)  |EID ref count and list (comp) |length (SDNV) |
+------+--------------+-----------------------------+--------------+
|   Retransmission Sequence Number (SDNV)                          |
+------------------------------------------------------------------+
```

Figure 1

4. Retransmission Block Processing

The following are the processing steps that a bundle node must take relative to generation, reception, and processing of Retransmission Blocks, assuming that the node is configured to detect and discard replays.

4.1. Bundle Reception

According to the Bundle Protocol, if a node receives a bundle that it currently has in custody as custodian, the received bundle will be discarded.

Upon receipt of any other type of bundle, the node SHALL delete the bundle's Retransmission Block if the custodian EID referenced in the RB is not the same as the custodian EID referenced in the Primary Bundle Block.

4.2. Detecting Duplicates and Determining which ones are Custodial Retransmissions

We define a duplicate to be a bundle that has the same source endpoint ID, creation timestamp, fragment offset and payload length (if the bundle is a fragment), and checksum or digest of the payload block as another bundle. (Whether to use a checksum or a digest of the payload block is determined by local policy.)

If a bundle is received that is a duplicate of a previously received bundle, then

-If the received bundle does not include a Retransmission Block, the bundle is not a custodial retransmission.

-If the received bundle does include a Retransmission Block and the RB's EID reference and retransmission sequence number values are the same as those in the Retransmission Block (if any) of the previously-received, duplicate bundle, the bundle is not a custodial retransmission.

-Otherwise, the received bundle is a custodial retransmission.

The receiving node MUST delete this bundle if it is a duplicate, if it is not a custodial retransmission, and if it cannot be determined (by some mechanism that may be defined elsewhere) to be a result of a routing loop.

4.3. Keeping Track of Bundles Received

If the bundle is not deleted as a replay, the node must store at least the following information from or about the bundle for comparison with future received bundles: source EID; creation timestamp; fragment offset and payload length (if the bundle is a fragment); checksum or digest of the payload block; custodian EID (if the bundle does not include a Retransmission Block) and Retransmission Block EID reference and retransmission sequence number (if the bundle does include a Retransmission Block).

4.4. Purging stored bundle information

The stored information for all bundles whose creation timestamp + lifetime is less than the current time MAY be deleted.

4.5. Bundle Forwarding

As part of the custody acceptance procedures, the accepting node MUST delete the bundle's Retransmission Block (if it has one).

4.6. Custodial Retransmission

Upon deciding to re-forward a bundle as a result of custody transfer failure, the re-forwarding custodian MUST:

- insert a RB with a retransmission sequence number value of 0 into the bundle if the bundle does not already include a RB, or

- increment the retransmission sequence number value in the Retransmission Block if the bundle does already include a RB.

- Store the inserted/modified retransmission block values along with the other information from the bundle as part of its custody storage procedures.

The EID reference in the Retransmission Block MUST refer to the re-forwarding custodian.

If a custodian decides to re-forward only a fragment of a bundle that it had previously forwarded, the re-forwarded fragment will not be a duplicate of any bundle that had previously been transmitted by this custodian. Therefore, the re-forwarded fragment SHALL NOT include a Retransmission Block.

5. Non-Uniform Support for the Retransmission Block

Failure to support the RB at one or more nodes in a DTN in which, as a matter of policy, all nodes are configured to delete replayed bundles may result in the erroneous deletion of custodially retransmitted bundles in the following cases:

> A node that does not support the RB but that is configured to delete replays could delete duplicate bundles even if they include RBs that mark them as being custodial retransmissions.
>
> A custodial node that does not support the RB but that retransmits a bundle would not include a RB to mark the bundle as a custodial retransmission, so that when the bundle is received at a downstream node that is configured to suppress replays, the bundle would be deleted by that downstream node (even if that downstream node supports the RB).

Consequently, the RB SHOULD be supported at all nodes in a DTN that, as a matter of policy, deletes replayed bundles. If not all nodes in the DTN support the RB, then to preserve support for custodial retransmission while maximizing replay suppression, the security policies of the nodes and the Block Processing Flags in the RB should be configured as follows:

- The "Discard bundle if block can't be processed" Block Processing Flag SHOULD NOT be set,

- The "Discard block if it can't be processed" Block Processing Flag SHOULD NOT be set,

- Nodes that support the RB should be configured to delete duplicates that are not custodial retransmissions,

- Nodes that do not support the RB should be configured to forward duplicates (so that they don't inadvertently delete custodial retransmissions), and

- Nodes that do not support the RB should be configured not to take custody of bundles (to ensure that custodial retransmissions will always include RBs).

The above configuration ensures that custodial retransmissions will not be erroneously deleted, and that all duplicate that are received at nodes that support the RB will be deleted. Only duplicates that are received at nodes that do not support the RB will be forwarded and allowed to remain in the network. If these are forwarded to a node that supports the RB, however, they will be deleted at that

node. Therefore, a network configured in this way is vulnerable to a denial-of-service attack only from duplicate bundles that circulate exclusively among nodes that do not support the RB.

6. Security Considerations

As mentioned in the Applicability Statement Section, it does not make sense to detect and suppress replayed bundles without first authenticating that those bundles have not been modified. Without authentication that a bundle has been forwarded intact, a network is vulnerable to denial of service attacks launched merely by the injection of any spurious bundles into the network or the modification of any authentic bundles. There seems little value in protecting against denial-of-service attacks resulting from replayed bundles if denial-of-service attacks resulting from such modified or spurious bundles will be permitted. Therefore, in determining the security policy of a node, nodes that support the RB and that are configured to suppress replays should also be required to authenticate bundles. Furthermore, all nodes in the DTN should be configured in the same way, to ensure that replays will be suppressed consistently without also resulting in the erroneous deletion of custodial retransmissions.

If the integrity of the RB is not protected, an adversary could inject many replayed bundles into the network yet include an RB in each that makes these bundles appear to be legitimate retransmissions. Integrity protection for the entire bundle, including the RB, MUST be provided by using the BAB with a ciphersuite, such as the BAB-HMAC ciphersuite defined in the Bundle Security Protocol, that uses a strict canonicalisation algorithm to protect the entire bundle between one hop and the next. Because of the hop-by-hop nature of the protection provided by the BAB, every node in the network would need to require all bundles to be protected with the BAB in order to ensure bundle authentication across the network. If, instead, integrity protection were to be provided using the PIB, with a ciphersuite that uses mutable canonicalization, the DTN would still be vulnerable to a replay attack in which an adversary modifies the fragment offset information of a previously-transmitted, valid bundle, and injects this modified bundle into the network. Such a bundle would not be deleted as a replay because its offset information is unique, but it would authenticate using the PIB because ciphersuites using mutable canonicalization do not calculate their security results over the fragment offset information (due to the fact that this information may change as the bundle traverses the network).

If a node or BAB key is compromised, authentication provided through use of the BAB does not help to protect against replays, but in this case the network's vulnerability to denial-of-service attacks is much larger than just a vulnerability to replays. If a node is compromised, any bundle could be created and injected into the network.

If a node or key is compromised, however, payload content must be taken into consideration in order to protect the DTN from insertion attacks that may be possible as a result of the RB being used. In some cases it might be possible for an adversary to know or guess that a specific source will emit a bundle at a specific time. In this case, the adversary could send out its own bundle that purports to be from that source and that contains correctly-guessed timestamp information, with the intent that this bundle be received at a forwarding node before the authentic bundle from the actual source. If the adversary that is injecting the spurious bundle is in possession of a compromised BAB key, this spurious bundle would appear to be valid when received by a forwarding node. If the forwarding node were to use only bundle source, creation timestamp, and fragment information (and not a checksum or digest of the payload block, as is required) to identify duplicates, then when the forwarding node receives the second bundle, it would delete this bundle as a duplicate even though this second bundle is actually the authentic bundle from the actual source. If payload block content is being used to identify duplicates, on the other hand, then the two bundles would not appear to be duplicates and the second one would not be deleted. The fact that the bundle source, creation timestamp, and fragment information of the bundles match whereas the payloads do not, however, would serve as a red flag that something is amiss and needs to be investigated. If an adversary launching this kind of insertion attack is in possession of a compromised BAB key, then insertion of a PIB into the bundle by the bundle's source would enable the forwarding node to determine which of the bundles is legitimate and which is not (assuming the forwarding node is in possession of the keying material necessary to authenticate the PIB security result). If the adversary launching this attack is in possession of the source's PIB key, however, then determining which bundle is legitimate would be impossible. Still, the presence of two bundles with identical bundle source, creation timestamp, and fragment information but different payloads would serve as a red flag.

7. IANA Considerations

If the bundle protocol becomes a standards track protocol, then we may want to consider having IANA establish a register of block types, of which the Retransmission Block would be one.

8. References

8.1. Normative References

[refs.RFC2119]
Bradner, S. and J. Reynolds, "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, October 1997.

[refs.DTNBP]
Scott, K. and S. Burleigh, "Bundle Protocol Specification", RFC 5050, November 2007.

[refs.DTNBPsec]
Symington, S., Farrell, S., Weiss, H., and P. Lovell, "Bundle Security Protocol Specification", draft-irtf-dtnrg-bundle-security-08.txt, work-in-progress, March 2009.

8.2. Informative References

[refs.DTNarch]
Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., and H. Weiss, "Delay-Tolerant Network Architecture", RFC 4838, April 2007.

[refs.DTNsecOver]
Farrell, S., Symington, S., Weiss, H., and P. Lovell, "Delay-Tolerant Network Security Overview", draft-irtf-dtnrg-sec-overview-06.txt, work-in-progress, March 2009.

Author's Address

Susan Flynn Symington
The MITRE Corporation
7515 Colshire Drive
McLean, VA 22102
US

Phone: +1 (703) 983-7209
Email: susan@mitre.org
URI: http://mitre.org/