

Network Transport Circuit Breakers

draft-ietf-tsvwg-circuit-breaker

Most recent version -08 (uploaded for this meeting).

Editor: Gorry Fairhurst



Status

Shepherd Review: David Black

TSVWG Last Call: Completed, 13 Sep 2015.

Submitted to IESG for Publication: Rev-04, Intended status BCP

AD Review: Rev-05, Spencer Dawkins (complete)

Transport Review: Bob Briscoe (issued to be resolved).

APPs-Area Review: Elliot Lear (issued to be resolved).

IETF Last Call: Not yet (Awaiting AD go-ahead with new Rev.)

Normative REF: RFC50405.bis (still to complete WGLC)

Examples of current references to this draft by other IETF drafts

[draft-ietf-avtcore-rtp-circuit-breakers](#)

Multimedia Congestion Control: Circuit Breakers for Unicast RTP Sessions

[draft-ietf-trill-over-ip](#)

Transparent Interconnection of Lots of Links (TRILL) over IP

[draft-ietf-tsvwg-ecn-encap-guidelines](#)

Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP

[draft-ietf-tsvwg-gre-in-udp-encap](#)

GRE-in-UDP Encapsulation

[draft-ietf-tsvwg-rfc5405bis](#)

UDP Usage Guidelines

[draft-ietf-tsvwg-tunnel-congestion-feedback](#)

Tunnel Congestion Feedback

[RFC 7510](#)

Encapsulating MPLS in UDP

[draft-ietf-pals-congcons](#)

Pseudowire Congestion Considerations

Updates since end of WGLC

Rev-05 updated after AD comments

Clarifications to wording, addition to abstract and security considerations.

Rev-06 updated after TSV Review (Bob Briscoe, David Black)

Clarifications to say that reaction time is of the order of 10's secs to minutes; More detail on what is intended by Fast-Trip CBs and the differences to other CBs; Updated feedback requirements for managed CBs

Rev-07 updated after TSV Review (David Black)

Control communication used by managed CBs may be in-band or out-of-band

Rev-08 updated after Apps Area Review (Elliot Lear: 21 October 2015)

Improved introduction and description of use; Clarification that a trigger is abnormal; Minor comments.

Review Question:

Should the RTP-CB still be included?

Pros: The technique is called a CB; It has many similarities with other CB techniques; It is well-known within the RTP community.

Con: The technique operates on single flows, typically between end hosts; and can include multiple types of measurements. It operates on much shorter timescales.

Recommendation by editor: Keep description, but additional text has been proposed in -08 to clarify the above differences.

Review Question:

What protection is needed for control communication?

Current Text:

Simple protection can be provided by using a randomized source port, or equivalent field in the packet header (such as the RTP SSRC value and the RTP sequence number) expected not to be known to an off-path attacker.

Review Comment:

- > I think the draft should recommend that for most scenarios, randomized
- > ports will be insufficient protection for CB control messages, which
- > should be properly cryptographically authenticated. Otherwise, a
- > CB-controlled aggregate is too vulnerable to these off-path attacks.

Recommendation by editor: Do not currently understand/agree with this comment, so seeking opinion of the group on what to recommend.

New Text:

2 control communication methods for the Managed CB

In-Band (Recommended, provides fate-sharing):

- Repeated loss of control messages SHOULD trigger Circuit Breaker.

Out-of-Band (Sometimes needed):

- Loss of control messages SHOULD NOT trigger Circuit Breaker (Avoids failure amplification/propagation)
- SHOULD be regarded as abnormal network event and be logged
- Additional operator action might be appropriate (Depends on network and traffic)

Check the text in rev > -07, if you have questions...

Next Steps:

There have been quite a lot of changes to written text and structure of English.

In the editor's opinion, all significantly improved the correctness and clarity of what is presented.

The intention of the draft has not changed, and apart from the control feedback case, it is not thought that any new concepts introduced.

So, now request to check updates have been correctly applied and request AD to continue publication process.

Necessary Acknowledgment

This work was part-funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700).