Network Working Group	L. Eggert
Internet-Draft	Nokia
Intended status: Experimental	June 23, 2010
Expires: December 25, 2010	

Congestion Control for the Constrained Application Protocol (CoAP) draft-eggert-core-congestion-control-00

Abstract

The Constrained Application Protocol (CoAP) is a simple, low-overhead, UDP-based protocol for use with resource-constrained IP networks and nodes. CoAP defines a simple technique to individually retransmit lost messages, but has no other congestion control mechanisms. This document motivates the need for additional congestion control mechanisms, and defines some simple strawman proposals. The goal is to encourage experimentation with these and other proposals, in order to determine which mechanisms are feasible to implement on resource-constrained nodes and are effective real deployments.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. This document may not be modified, and derivative works of it may not be created, except to format it for publication as an RFC or to translate it into languages other than English.

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 http://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 December 25, 2010.

Copyright Notice

Copyright (c) 2010 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/licenseinfo) 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.

1. Introduction

The Constrained Application Protocol (CoAP) [I-D.ietf-core-coap] (Shelby, Z., Frank, B., and D. Sturek, "Constrained Application Protocol (CoAP)," June 2010.) is a simple, low-overhead, UDP-based protocol for use with resource-constrained IP networks and nodes. CoAP defines two kinds of interactions between end-points:

- a client/server interaction model, where request or notify messages initiate a transaction with a server, which may send a response to the client with a matching transaction ID
- an asynchronous subscribe/notify interaction model, where a server can send notify messages to a client about a resource which the client has subscribed to

CoAP uses the User Datagram Protocol (UDP) [RFC0768] (Postel, J., "User Datagram Protocol," August 1980.) to transmit these messages and defines a simple mechanism to individually retransmit lost messages using an exponentially backed-off timer. This document argues that although this retransmission mechanism is a required first step to implement congestion control for CoAP, it alone is not sufficient to alleviate network overload in all conditions. Section 2 (Discussion of Internet Congestion Control Principles) gives a short summary of Internet congestion control principles, and Section 3 (CoAP Congestion Control) presents some simple strawman proposals that attempt to complement the current message retransmission mechanism in CoAP.

2. Discussion of Internet Congestion Control Principles

[RFC2914] (Floyd, S., "Congestion Control Principles," September 2000.) describes the best current practices for congestion control in the Internet, and requires that Internet communication employ congestion control mechanisms. Because UDP itself provides no congestion control mechanisms, it is up to the applications and application-layer protocols that use UDP for Internet communication to employ suitable mechanisms to prevent congestion collapse and establish a degree of fairness. CoAP is one such application-layer protocol. [RFC2914] (Floyd, S., "Congestion Control Principles," September 2000.) identifies two major reasons why congestion control mechanisms are critical for the stable operation of the Internet:

- 1. The prevention of congestion collapse, i.e., a state where an increase in network load results in a decrease in useful work done by the network.
- The establishment of a degree of fairness, i.e., allowing multiple flows to share the capacity of a path reasonably equitably.

The overwhelming majority of the bytes on the Internet are caused by bulk transfers, and the traditional congestion control mechanisms are engineered to saturate the network without driving it into congestive collapse. Fairness between flows is an important consideration when the network operates around the saturation point, so that new flows are not disadvantaged compared to established flows, and can obtain a reasonable share of the capacity quickly.

The environments that CoAP targets are IP networks, although more resource-constrained ones than the "big-I" Internet. This does not eliminate the need for end-point-based congestion control. If anything, the environments that CoAP will be deployed in have fewer capabilities for network provisioning, traffic engineering and capacity allocation, which are among the techniques that can sometimes offset the need for end-to-end congestion control to some degree.

However, the environments that CoAP targets are sufficiently different from the "big-I" Internet so that the motivations for congestion control from [RFC2914] (Floyd, S., "Congestion Control Principles," September 2000.) should probably be weighted differently. CoAP networks will not be used for bulk data transfers and CoAP nodes will not need to use a significant fraction of the capacity of a path to provide a useful service. (In fact, they are often too resource-constrained to do so in the first place.) Under normal operation, a CoAP network will be mostly idle, which means that fairness between the transmissions of different CoAP nodes is not a large issue. A CoAP congestion control mechanism can hence focus on preventing congestion collapse, which is a much more tractable problem given the specific conditions of CoAP environments.

The current IETF congestion control mechanisms, such as TCP [RFC5681] (Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control," September 2009.) or TFRC [RFC5348] (Floyd, S., Handley, M., Padhye, J., and J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification," September 2008.), all focus on determining a "safe" sending rate for a bulk transfer, i.e., for a single flow of many packets between a sender and destination where many packets are in flight at any given time. They measure the path characteristics, such as round-trip time (RTT) and packet loss rate, by monitoring the ongoing transfer and use this information to adjust the sending rate of the flow during the transmission.

This approach is not feasible for CoAP. The infrequent request/response interaction that CoAP supports does not generate sufficient data about the path characteristics to drive a traditional congestion control loop, even if the notion of "a flow" to a destination is extended from "one CoAP transaction" to "a sequence of CoAP transactions". This approach is also not applicable to multicast transmissions, which CoAP offers.

[RFC5405] (Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers," November 2008.) documents the IETF's current best practices for using UDP for unicast communication in the Internet. It provides guidance on topics such as message sizes, reliability, checksums, middlebox traversal and congestion control. Section 3.1.2 of [RFC5405] (Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers," November 2008.), which focuses on congestion control for low data-volume applications, is especially relevant to CoAP.

Section 3.1.2 of <u>[RFC5405] (Eggert, L. and G. Fairhurst, "Unicast UDP</u> <u>Usage Guidelines for Application Designers," November 2008.</u>) acknowledges that the traditional IETF congestion control mechanisms are not applicable for low data-volume application protocols such as COAP. Instead, it recommends that such application protocols:

*maintain an estimate of the RTT for any destination with which they communicate, or assume a conservative fixed value of 3 seconds when no RTT estimate can be obtained (e.g., unidirectional communication)

*control their transmission behavior by not sending on average more than one UDP datagram per RTT to a destination

*detect packet loss and exponentially back their retransmission timer off when a loss event occurs

*employ congestion control for both directions of a bi-directional communication

CoAP follows some of these guidelines already. It uses a fixed value of 1 second for its retransmission timer for both requests and responses, which although shorter than the recommended value in <u>[RFC5405] (Eggert,</u> L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application <u>Designers," November 2008.</u>) is likely appropriate for many of its deployment scenarios. CoAP also uses exponential back-off for its retransmission timer.

This alone, however, does not result in a complete congestion control mechanism for CoAP. <u>Section 3 (CoAP Congestion Control)</u> defines an experimental complement to the current CoAP mechanism described in

[I-D.ietf-core-coap] (Shelby, Z., Frank, B., and D. Sturek, "Constrained Application Protocol (CoAP)," June 2010.).

3. CoAP Congestion Control

This section proposes several congestion control techniques for CoAP that are intended to improve its ability to prevent congestion collapse. At the moment, these techniques are described with the intent of encouraging experimentation with such proposals in CoAP simulations and testbed deployments. Of particular interest are mechanism requiring little computation and state, i.e., mechanisms that can be implemented in resource-constrained nodes without much overhead.

3.1. Retransmissions

CoAP already defines a simple retransmission scheme with exponential back-off, where messages that have not been responded to in RESPONSE_TIMEOUT are retransmitted, followed by doubling RESPONSE_TIMEOUT. Up to MAX_RETRANSMIT retransmission attempts are made. (At the moment, [I-D.ietf-core-coap] (Shelby, Z., Frank, B., and D. Sturek, "Constrained Application Protocol (CoAP)," June 2010.) defines RESPONSE_TIMEOUT to be 1 second and MAX_RETRANSMIT to be five attempts.) As stated above, although RESPONSE_TIMEOUT is shorter than what [RFC5405] (Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers," November 2008.) recommends, the shorter value is likely to not cause large issues in many deployments that CoAP targets.

However, using a fixed value for RESPONSE_TIMEOUT instead of basing it on the measured RTT to a destination has some minor drawbacks. CoAP may be used in deployments where the path RTTs can approach the currently defined RESPONSE_TIMEOUT of 1 second, such as Internet deployments involving GSM or 3G links, or cases where preparing a response can involve significant computation or where it otherwise incurs delays, such as long sleep cycles at the receiver. Fixed timeouts that are too short can cause spurious retransmissions, i.e., unnecessary retransmissions in cases where either the request or the response are still in transit. Spurious retransmissions, especially persistent ones, waste resources.

This section therefore proposes that CoAP deployments experiment with maintaining an estimate of the RTT for any destination with which they communicate. Specifically, it is suggested that deployments experiment with the algorithm specified in [RFC2988] (Paxson, V. and M. Allman, "Computing TCP's Retransmission Timer," November 2000.) to compute a smoothed RTT (SRTT) estimate, and compute RESPONSE_TIMEOUT in the same

way [RFC2988] (Paxson, V. and M. Allman, "Computing TCP's
Retransmission Timer," November 2000.) computes RTO.
A second suggestion is to experiment with a longer RESPONSE_TIMEOUT,
such as 3 seconds, which is what [RFC5405] (Eggert, L. and G.
Fairhurst, "Unicast UDP Usage Guidelines for Application Designers,"
November 2008.) recommends, in order to determine if there are
significant drawbacks or whether this value could be lengthened.

3.2. Aggregate Congestion Control

Traditional Internet congestion control algorithms control the sending rate of a single flow. When a node establishes multiple, parallel flows, their congestion control loops run (mostly) independently of one another. Interactions between the control loops of parallel flows is (mostly) indirect, e.g., a rate increase of one flow may cause packet loss and a rate decrease to another.

CoAP "flows", i.e., sequences of infrequent CoAP transactions between the same two nodes, do not require much more per-flow congestion control than a retransmission scheme that reduces the rate (increases the back-off) of a flow under loss, and a (low) cap on the number of allowed outstanding requests to a destination. ([RFC5405] (Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers," November 2008.) recommends "on average not more than one"

outstanding transaction to a given destination.)

On the other hand, CoAP applications may potentially want to initiate many transactions with different nodes at the same time. Allowing CoAP applications to initiate an unlimited number of parallel transactions gives them the means for causing overload, and depends on applicationlevel measures to detect and correctly mitigate this failure. Because each transaction only consumes a very limited amount of resources, it is arguably more important to control the total outstanding number of transactions, compared to controlling the rate at which each individual one is being (re)transmitted. The CoAP spec [I-D.ietf-core-coap] (Shelby, Z., Frank, B., and D. Sturek, "Constrained Application

<u>Protocol (CoAP), " June 2010.</u>) does currently not impose any limit on how many parallel transactions to different nodes an end-point may have outstanding.

Given the importance of preventing congestion collapse, this document argues that the CoAP protocol should specify a common mechanism for congestion controlling the aggregate traffic a CoAP node sends into the network. In other words, the CoAP stack should locally drop application-generated messages under overload situations, rather than attempting to send them into the network, irrespective of the destination.

One proposal is to implement a simple windowing algorithm. In this mechanism, a CoAP node has a certain number of "transmission credits"

available during a time interval. Sending one CoAP message consumes one transmission credit, independent of which destination it is being sent to. If all transmission credits have been used up during a time interval, the CoAP node drops any additional messages that the applications attempt to send during the remainder of the time interval. At the end of a time interval, the CoAP node determines whether responses have been received for all requests it has issued within the time interval. If this is the case, the CoAP node increases the number of send credits by one for the following time interval. If responses fail to arrive for some of the requests issued during the time interval, the number of permitted CoAP requests is cut in half for the next interval.

The description above leaves several questions unanswered. These include the length of the time interval and whether it is fixed or adapted over time, whether an increase by one and a reduction by half are the correct parameters for the proposed AIMD (additive increase, multiplicative decrease) scheme, whether the decrease should be proportional to the loss rate, and others.

This document does at the moment not attempt to answer these questions. Instead, it encourages simulations and implementations to explore the design space, and also consider other non-windowing approaches.

3.3. Explicit Congestion Notification

Explicit Congestion Notification (ECN) [RFC3168] (Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP," September 2001.) is an extension to IP that allows routers to inform end nodes when they approach congestion by setting a bit in the IP header. The receiver of a message echoes this bit to the sender, which reacts as if a packet loss had occurred for the flow.

Deployment of ECN can reduce overall packet loss, because senders can react to congestion early, i.e., before packet loss occurs. This is especially attractive in resource-constrained environments, because retransmissions can be avoided.

If CoAP uses an aggregate congestion control mechanism such as described in Section <u>Section 3.2 (Aggregate Congestion Control)</u>, it will reduce the amount of transmission credits for the next time interval when some of the responses received had the ECN bit set. (Other reactions to ECN markings may be possible.) Whether ECN support is possible in CoAP deployments remains to be investigated, because ECN usage requires a negotiation handshake (can potentially be avoided if support is made mandatory for CoAP deployments) and because routers need to support ECN marking. At this point, simulations attempting to quantify the benefits may therefore be easiest to obtain.

3.4. Multicast Considerations

CoAP requests may be multicast, and result in several replies from different end-points, potentially consuming much more resource capacity for the request and response transmissions than a single unicast transaction. It can therefore be argued that sending multicast requests should be more conservatively controlled than the sending of unicast requests.

CoAP already acknowledges this to some degree by not retransmitting multicast requests at the CoAP-level. Unfortunately, CoAP currently has no means for preventing an application from doing application-level retransmissions of multicast requests. Given that the prevention of congestion collapse is important, such a mechanism should be added. The aggregate congestion control proposal in Section <u>Section 3.2</u> (Aggregate Congestion Control) puts a cap on the number of transmissions allowed during a time interval, including multicast requests. It is currently unclear whether additional means are required for CoAP deployments that make heavy use of multicast. As before, experimentation is encouraged to understand the problem space.

4. IANA Considerations

This document requests no actions from IANA. [Note to the RFC Editor: Please remove this section upon publication.]

5. Security Considerations

This document has no known security implications. [Note to the RFC Editor: Please remove this section upon publication.]

6. Acknowledgments

Lars Eggert is partly funded by <u>[TRILOGY] (, "Trilogy Project," .)</u>, a research project supported by the European Commission under its Seventh Framework Program.

7. References

7.1. Normative References

[I-D.ietf- core-coap]	Shelby, Z., Frank, B., and D. Sturek, <u>"Constrained</u> <u>Application Protocol (CoAP)</u> ," draft-ietf-core-coap-00 (work in progress), June 2010 (<u>TXT</u>).
[RFC0768]	Postel, J., " <u>User Datagram Protocol</u> ," STD 6, RFC 768, August 1980 (<u>TXT</u>).
[RFC2914]	Floyd, S., " <u>Congestion Control Principles</u> ," BCP 41, RFC 2914, September 2000 (<u>TXT</u>).
[RFC2988]	Paxson, V. and M. Allman, " <u>Computing TCP's</u> <u>Retransmission Timer</u> ," RFC 2988, November 2000 (<u>TXT</u>).
[RFC3168]	Ramakrishnan, K., Floyd, S., and D. Black, " <u>The</u> <u>Addition of Explicit Congestion Notification (ECN) to</u> <u>IP</u> ," RFC 3168, September 2001 (<u>TXT</u>).
[RFC5405]	Eggert, L. and G. Fairhurst, " <u>Unicast UDP Usage</u> <u>Guidelines for Application Designers</u> ," BCP 145, RFC 5405, November 2008 (<u>TXT</u>).

7.2. Informative References

[RFC5348]	Floyd, S., Handley, M., Padhye, J., and J. Widmer, " <u>TCP</u> <u>Friendly Rate Control (TFRC): Protocol Specification</u> ," RFC 5348, September 2008 (<u>TXT</u>).
[RFC5681]	Allman, M., Paxson, V., and E. Blanton, " <u>TCP Congestion</u> <u>Control</u> ," RFC 5681, September 2009 (<u>TXT</u>).
[TRILOGY]	"Trilogy Project," http://www.trilogy-project.org/.

Author's Address

	Lars Eggert
	Nokia Research Center
	P.O. Box 407
	Nokia Group 00045
	Finland
Phone:	+358 50 48 24461
Email:	<u>lars.eggert@nokia.com</u>
URI:	http://research.nokia.com/people/lars_eggert