

Network Working Group  
Internet-Draft  
Intended status: Experimental  
Expires: 2 August 2026

P. Balasubramanian  
Confluent  
O. Ertugay  
D. Havey  
Microsoft  
M. Bagnulo  
Universidad Carlos III de Madrid  
29 January 2026

LEDBAT++: Congestion Control for Background Traffic  
draft-irtf-iccr-ledbat-plus-plus-06

Abstract

This memo describes LEDBAT++, a set of enhancements to the LEDBAT (Low Extra Delay Background Transport) congestion control algorithm for background traffic. The LEDBAT congestion control algorithm has several shortcomings that prevent it from working effectively in practice. LEDBAT++ extends LEDBAT by adding a set of improvements, including reduced congestion window gain, modified slow-start, multiplicative decrease and periodic slowdowns. This set of improvement mitigates the known issues with the LEDBAT algorithm, such as latency drift, latecomer advantage and inter-LEDBAT fairness. LEDBAT++ has been implemented as a TCP congestion control algorithm in the Windows operating system. LEDBAT++ has been deployed in production at scale on a variety of networks and been experimentally verified to achieve the original stated goals of LEDBAT. This document is a product of the Internet Congestion Control Research Group (ICCRG) of the Internet Research Task Force (IRTF).

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 2 August 2026.

## Copyright Notice

Copyright (c) 2026 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

## Table of Contents

1. Introduction . . . . .	2
2. Terminology . . . . .	3
3. LEDBAT Issues . . . . .	3
3.1. Latecomer advantage . . . . .	4
3.2. Inter-LEDBAT fairness . . . . .	4
3.3. Latency drift . . . . .	4
3.4. Low latency competition . . . . .	5
3.5. Dependency on one-way delay measurements . . . . .	5
4. LEDBAT++ Mechanisms . . . . .	5
4.1. Slower than Reno increase . . . . .	5
4.2. Multiplicative decrease . . . . .	6
4.3. Modified slow start . . . . .	7
4.4. Initial and periodic slowdown . . . . .	7
4.5. Use of Round Trip Time instead of one way delay . . . . .	8
5. Experiment Considerations . . . . .	9
5.1. Status of the experiment at the time of this writing. . . . .	9
6. Security Considerations . . . . .	10
7. IANA Considerations . . . . .	10
8. Acknowledgements . . . . .	10
9. References . . . . .	10
9.1. Normative References . . . . .	10
9.2. Informative References . . . . .	11
Authors' Addresses . . . . .	12

## 1. Introduction

Operating systems and applications use background connections for a variety of tasks, such as software updates, large media downloads, telemetry, or error reporting. These connections should operate without affecting the general usability of the system. Usability is measured in terms of available network bandwidth and network latency. LEDBAT [RFC6817] is designed to minimize the impact of lower than

best effort connections on the latency and bandwidth of other connections. To achieve that, each LEDBAT connection monitors the delay experienced by the packets it sends across an Internet path, and compares them to the minimum delay observed on the connection. The difference between the transmission delay and the minimum delay is used as an estimate of the queuing delay. If the queuing delay is above a target, LEDBAT directs the connection to reduce the bandwidth used by the flow. If the queuing delay is below the target, the connection is allowed to increase its transmission rate. The increase or decrease in the used bandwidth are proportional to the difference between the observed values and the target. LEDBAT reacts to packet losses and other congestion signals in the same way as standard TCP.

However, there are a few issues that plague LEDBAT, some previously documented, and some discovered by experiments (see Section 3). LEDBAT++ specifies additional mechanisms (and in some cases deviates from) LEDBAT to overcome these problems. The remaining sections describe the problems and the mechanisms in detail.

LEDBAT++ is defined for use with TCP and its use with other transports is out-of-scope of this specification.

The consensus of the Internet Congestion Control Research Group (ICCRG) is to publish this document to encourage further experimentation and review of LEDBAT++. The objective of this RFC is to document LEDBAT++ modifications of the base LEDBAT implementation and encourage its use so the algorithm can be further verified and improved.

This document is not an IETF product and is not a standard. The status of this document is experimental. In section 5 titled Experiment Considerations, we describe the purpose of the experiment and its current status.

## 2. Terminology

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 [RFC2119] and [RFC8174].

## 3. LEDBAT Issues

This section lists some known LEDBAT issues from existing literature and also list some new problems observed as a result of experimentation with an implementation of [RFC6817].

### 3.1. Latecomer advantage

Delay based congestion control protocols like LEDBAT are known to suffer from a latecomer advantage. When the newcomer establishes a connection, the transmission delay that it encounters includes any queuing delay caused by existing connections that share common network queues. The newcomer considers this delay, that is larger than the base delay, the minimum, and thereby increases its transmission rate while other LEDBAT connections slow down. Eventually, the latecomer will result in using the entire bandwidth available to the connection. Standard TCP congestion control as described in [RFC9293] and [RFC5681], causes some queuing, the LEDBAT delay measurements incorporate that queuing, and the base delay as measured by the connection is thus set to a larger value than the actual minimum. As a result, the queues remain mostly full. In some cases, this queuing persists even after the closing of the competing TCP connection. This phenomenon was already known during the design of LEDBAT, but there is no mitigation in the LEDBAT design. The designers of the protocol relied instead on the inherent burstiness of network traffic. Small gaps in transmission schedules would allow the latecomer to measure the true delay of the connection. This reasoning is not satisfactory because workloads can upload large amount of data, and would not always see such gaps.

### 3.2. Inter-LEDBAT fairness

The latecomer advantage is caused by the improper evaluation of the base delay, with the latecomer using a larger value than the preexisting connections. However, even when all competing connections have a correct evaluation of the base delay, some of them will receive a larger share of resources. The reason for that persistent unfairness is explained in [RethinkLEDBAT]. LEDBAT specifies proportional feedback based on a ratio between the measured queuing delay and a target. Proportional feedback uses both additive increases and additive decreases. This does stabilize the queue sizes, but it does not guarantee fair sharing between the competing connections.

### 3.3. Latency drift

LEDBAT estimates the base delay of a connection as the minimum of all observed transmission delays over a 10-minute interval. It uses an interval rather than a measurement over the whole duration of the connection, because network conditions may change over time. For example, an existing connection may be transparently rerouted over a longer path, with a longer transmission delay. Keeping the old estimate would then cause LEDBAT to unnecessarily reduce the connection throughput. However experiments show that this causes a

ratcheting effect when LEDBAT connections are allowed to operate for a long time. The delay feedback in LEDBAT causes the queuing delay to stabilize just below the target. After an initial interval, all new measurements are equal to the initial transmission delay plus a fraction of the target. Every 10 minutes, the measured base delay increases by that fraction of the target queuing delay, leading to potentially large values over time.

### 3.4. Low latency competition

LEDBAT compares the observed queuing delays to a fixed target. The target value cannot be set too low, because that would cause poor operation on slow networks (where the link transmission rate add a significant serialisation delay). In practice, it is set to 60ms, a value that allows proper operation of latency sensitive applications like VoIP. But if the bottleneck buffer is small such that the queuing delay will never reach the target, then the LEDBAT connection behaves just like an ordinary connection. It competes aggressively, and obtains the same share of the bandwidth as regular TCP connections. On high speed links the problem is exacerbated.

### 3.5. Dependency on one-way delay measurements

The LEDBAT algorithm requires use of one-way delay measurements. This makes it harder to use with transport protocols like TCP that have no reliable way to obtain one way delay measurements. TCP timestamps do not standardize clock frequency, and the endpoints will need to rely on heuristics to guess the clock frequency of the remote peer to detect and correct for clock skew. TCP timestamps do not include clock synchronization, and would need some non-standard invention to compensate for clock skew. Any such mechanism is very fragile.

## 4. LEDBAT++ Mechanisms

### 4.1. Slower than Reno increase

When the measured delays are below the target delay, LEDBAT behaves like standard TCP [RFC9293]. LEDBAT introduces a GAIN parameter which can be set between 0 and 1. In order to solve the low latency competition problem, LEDBAT++ makes the GAIN parameter dynamic. When standard and LEDBAT connections share the same bottleneck, they usually experience the same packet drop rate (i.e. when drop tail queues are used). The GAIN value ensures that the throughput of the LEDBAT++ connection will be a fraction ( $1/\text{SQRT}(1/\text{GAIN})$ ) of the throughput of the regular connections. Small values of GAIN work well when the base delay is small, and ensure that the LEDBAT++ connection will yield to regular connections in these networks.

However, small values of GAIN do not work well on paths with long delays. In the absence of competing traffic, combining large base delays with small GAIN values causes the throughput used by the connection to remain well under capacity for a long time. In LEDBAT++, GAIN is a function of the ratio between the base delay and the target delay:

$$\text{GAIN} = 1 / (\min(16, \text{CEIL}(2 * \text{TARGET} / \text{base})))$$

where  $\text{CEIL}(X)$  is defined as the smallest integer larger than  $X$ . Implementations MAY experiment with the constant value  $1/16$  as a tradeoff between responsiveness and performance.

#### 4.2. Multiplicative decrease

[RethinkLEDBAT] suggests combining additive increases and multiplicative decreases in order to solve the Inter-LEDBAT fairness problem. It proposes to change the way LEDBAT increases and decreases the congestion window based on the ratio between the observed delay and the target. Assume that the congestion window is changed once per roundtrip measurement. In standard LEDBAT, the per RTT window when delay is less than target is:

$$W += \text{GAIN} * (1 - \text{delay} / \text{target})$$

In LEDBAT++, with multiplicative decrease, the per RTT window when delay is less than target is:

$$W += \text{GAIN}$$

Similarly in standard LEDBAT, the per RTT window when the delay is higher than target is:

$$W -= \text{GAIN} * (\text{delay} / \text{target} - 1)$$

In LEDBAT++, with multiplicative decrease, the per RTT window delay is higher than target is:

$$W += \max((\text{GAIN} - \text{Constant} * W * (\text{delay} / \text{target} - 1)), -W/2)$$

It is RECOMMENDED that the Constant be set to 1. Implementations MAY experiment with this value.

If a connection has various estimates of the base delay, the multiplicative decrease MUST be capped to at most  $W/2$ . Otherwise, spikes in delay can cause the window to immediately drop to its minimal value.

LEDBAT++ sender SHOULD also ensure that the congestion window never decreases below 2 packets. This lower bound is intended to prevent starvation of the LEDBAT++ connection in the presence of transient congestion or delayed feedback. However, on severely constrained paths where the fair-share congestion window of competing flows may legitimately fall below two packets, enforcing such a bound could induce unnecessary queueing or congestion to itself and other flows.

#### 4.3. Modified slow start

Traditional initial slow start can cause spikes in bandwidth usage. However skipping exponential congestion window increase results in significantly reduced performance over long Internet Paths. LEDBAT++ applies the dynamic GAIN parameter to the congestion window increases. In standard TCP operation, the congestion window increases for every ACK by exactly the amount of bytes acknowledged. A LEDBAT++ sender increases the congestion window by that number multiplied by the dynamic GAIN value. In low latency links, this ensures that LEDBAT++ connections ramp up slower than regular connections. LEDBAT++ sender limits the initial window to 2 packets. LEDBAT++ sender monitors the transmission delays during the slow start period. If the queuing delay is larger than 3/4ths of the target delay, exit slow start and immediately move to the congestion avoidance phase. After initial slow start, the increase of congestion window is bounded by the SSTHRESH estimate acquired during congestion avoidance, and the risk of creating congestion spikes is very low. Exiting slow start on excessive delay SHOULD be applied only during the initial slow start.

#### 4.4. Initial and periodic slowdown

The LEDBAT specification assumes that there will be natural gaps in the aggregated traffic through the bottleneck of the path, and that during those gaps the observed delay corresponds to a state where the queues are empty. However, there are workloads where the traffic is sustained for long periods. This causes base delay estimates to be inaccurate and is one of the major reasons behind latency drift as well as the lack of inter-LEDBAT fairness. To ensure stability, LEDBAT++ forces these gaps, or slow down periods. A slowdown is an interval during which the LEDBAT++ connection voluntarily reduces its traffic, allowing queues to drain and transmission delay measurements to converge to the base delay. The slowdown works as follows:

- \* Upon entering slowdown, set SSTHRESH to the current version of the congestion window CWND, and then reduce CWND to 2 packets.
- \* Keep CWND frozen at 2 packets for 2 RTT.

- \* After 2 RTT, ramp up the congestion window according to the slow start algorithm, until the congestion window reaches SSTHRESH.

Keeping the CWND frozen at 2 packets for 2 RTT allows the queues to drain, and is key to obtaining accurate delay measurements. The initial slowdown starts shortly after the connection completes the initial slow start phase; 2 RTT after the initial slow start completes. After the initial slowdown, LEDBAT++ sender performs periodic slowdowns. The interval between slowdown is computed so that slowdown does not cause more than a 10% drop in the utilization of the bottleneck. LEDBAT++ sender measures the duration of the slowdown, from the time of entry to the time at which the congestion window regrows to the previous SSTHRESH value. The next slowdown is then scheduled to occur at 9 times this duration after the exit point. The combination of initial and periodic slowdowns allows competing LEDBAT connections to obtain good estimates of the base delay, and when combined with multiplicative decrease solves both the latecomer advantage and the Inter-LEDBAT fairness problems.

#### 4.5. Use of Round Trip Time instead of one way delay

LEDBAT++ uses Round Trip Time measurements instead of one way delay. One possible shortcoming of round trip delay measurements is that they incorporate queuing delays in both directions. This can lead to unnecessary slowdowns, such as slowing down an upload connection because a download is saturating the downlink but in practice this seems to benefit the workloads because the bottleneck link can carry ACK traffic in the other direction for the competing flows. Round trip measurements also include the delay at the receiver between receiving a packet and sending the corresponding acknowledgement. These delays are normally quite small, except when the delayed acknowledgement logic kicks in. Effect of delayed ACKs can be particularly acute when the congestion window only includes a few packets, for example at the beginning of the connection.

The problems of using one way delay are mitigated through a set of implementation choices. First, LEDBAT++ sender enables the TCP Timestamp option, in order to obtain RTT samples with each acknowledgement. A LEDBAT++ sender SHOULD filter the round trip measurements by using the minimum of the 4 most recent delay samples, as suggested in the LEDBAT specification. Finally, the queueing delay target is set larger than the typical TCP maximum acknowledgement delay. This avoids over reacting to a single delayed ACK measurement. The LEDBAT++ default delay target of 60ms is different from the 100ms value recommended in [RFC6817].

## 5. Experiment Considerations

The status of this document is Experimental. The general purpose of the proposed experiment is to gain more experience running LEDBAT++ over different network paths to see if the proposed LEDBAT++ parameters perform well in different situations. Specifically, it is desirable to understand the following aspects of the LEDBAT++ mechanism:

- The impact of transparent proxies which prevent measurement of end-to-end delay and might interfere with the effective operation of LEDBAT++.
- Interaction between LEDBAT++ and Active Queue Management techniques such as Codel [RFC8289], PIE [RFC8003] and L4S [RFC9330].
- How the LEDBAT++ should resume after a period during which there was no incoming traffic and the information about the rLEDBAT state information is potentially dated.

[RFC9743] provides guidelines for the IETF to use when evaluating a proposed congestion control algorithm that differs from the general congestion control principles outlined in [RFC2914]. The guidance and evaluation criteria is intended to be useful to authors proposing congestion control algorithms and for the IETF community when evaluating whether a proposal is appropriate for publication in the RFC Series and for deployment in the Internet.

### 5.1. Status of the experiment at the time of this writing.

LEDBAT++ is available in Microsoft's Windows 11 22H2 since October 2023 [Windows11] and in Windows Server 2022 since September 2022 [WindowsServer].

In addition, LEDBAT++ has been deployed by Microsoft in wide scale in the following services:

- BITS (Background Intelligent Transfer Service)
- DO (Delivery Optimization) service
- Windows update # using DO
- Windows Store # using DO
- OneDrive

- Windows Error Reporting # wermgr.exe; werfault.exe
- System Center Configuration Manager (SCCM)
- Windows Media Player
- Microsoft Office
- Xbox (download games) # using DO

An experimental evaluation of the LEDBAT++ algorithm is presented in [COMNET1]. Experiments involving the interaction of LEDBAT++ and BBR are presented in [COMNET2]

## 6. Security Considerations

LEDBAT++ enhances LEDBAT and inherits the general security considerations discussed in [RFC6817].

LEDBAT++ uses the RTT measurements to modulate the rate of the sender. An attacker wishing to starve a flow can introduce an artificial delay to the packets either by actually delaying the packets. This would cause the LEDBAT++ receiver to believe that a queue is building up and reduce the congestion window. Note that an attacker to do that must be on path, so if that is the case, it is probably more direct to simply drop the packets and achieve even a larger window reduction.

## 7. IANA Considerations

This document has no actions for IANA.

## 8. Acknowledgements

The LEDBAT++ algorithm was designed and implemented by Osman Ertugay, Christian Huitema, Praveen Balasubramanian, and Daniel Havey.

We would like to thank Reese Enghardt, Brian Trammell, Gorrry Fairhurst and Ari Keranen for the review and comments on earlier versions of this document.

This work was supported by the EU through the StandICT project RXQ.

## 9. References

### 9.1. Normative References

- [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>>.
- [RFC6817] Shalunov, S., Hazel, G., Iyengar, J., and M. Kuehlewind, "Low Extra Delay Background Transport (LEDBAT)", RFC 6817, DOI 10.17487/RFC6817, December 2012, <<https://www.rfc-editor.org/info/rfc6817>>.
- [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>>.

## 9.2. Informative References

- [COMNET1] Bagnulo, M.B. and A.G. Garcia-Martinez, "An experimental evaluation of LEDBAT++", Computer Networks Volume 212, 2022.
- [COMNET2] Bagnulo, M.B. and A.G. Garcia-Martinez, "When less is more: BBR versus LEDBAT++", Computer Networks Volume 219, 2022.
- [RethinkLEDBAT] Carofiglios, G., Muscariello, L., Rossi, D., Testa, C., and S. Valenti, "Rethinking the Low Extra Delay Background Transport (LEDBAT) Protocol", Computer Networks, Volume 57, Issue 8, 4 June 2013, Pages 18381852, 2013, <<http://perso.telecom-paristech.fr/~drossi/paper/rossil3comnet.pdf>>.
- [RFC2018] Mathis, M., Mahdavi, J., Floyd, S., and A. Romanow, "TCP Selective Acknowledgment Options", RFC 2018, DOI 10.17487/RFC2018, October 1996, <<https://www.rfc-editor.org/info/rfc2018>>.
- [RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, DOI 10.17487/RFC2914, September 2000, <<https://www.rfc-editor.org/info/rfc2914>>.
- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", RFC 5681, DOI 10.17487/RFC5681, September 2009, <<https://www.rfc-editor.org/info/rfc5681>>.
- [RFC8003] Laganier, J. and L. Eggert, "Host Identity Protocol (HIP) Registration Extension", RFC 8003, DOI 10.17487/RFC8003, October 2016, <<https://www.rfc-editor.org/info/rfc8003>>.

- [RFC8289] Nichols, K., Jacobson, V., McGregor, A., Ed., and J. Iyengar, Ed., "Controlled Delay Active Queue Management", RFC 8289, DOI 10.17487/RFC8289, January 2018, <<https://www.rfc-editor.org/info/rfc8289>>.
- [RFC9293] Eddy, W., Ed., "Transmission Control Protocol (TCP)", STD 7, RFC 9293, DOI 10.17487/RFC9293, August 2022, <<https://www.rfc-editor.org/info/rfc9293>>.
- [RFC9330] Briscoe, B., Ed., De Schepper, K., Bagnulo, M., and G. White, "Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture", RFC 9330, DOI 10.17487/RFC9330, January 2023, <<https://www.rfc-editor.org/info/rfc9330>>.
- [RFC9743] Duke, M., Ed. and G. Fairhurst, Ed., "Specifying New Congestion Control Algorithms", BCP 133, RFC 9743, DOI 10.17487/RFC9743, March 2025, <<https://www.rfc-editor.org/info/rfc9743>>.
- [Windows11] Forsmann, C.F., "What's new in Delivery Optimization", Microsoft Documentation <https://learn.microsoft.com/en-us/windows/deployment/do/whats-new-do>, 2023.
- [WindowsServer] Havey, D.H., "LEDBAT Background Data Transfer for Windows", Microsoft Blog <https://techcommunity.microsoft.com/t5/networking-blog/ledbat-background-data-transfer-for-windows/ba-p/3639278>, 2022.

#### Authors' Addresses

Praveen Balasubramanian  
Confluent  
Email: [pravb.ietf@gmail.com](mailto:pravb.ietf@gmail.com)

Osman Ertugay  
Microsoft  
Phone: +1 425 706 2684  
Email: [osmaner@microsoft.com](mailto:osmaner@microsoft.com)

Daniel Havey  
Microsoft  
Email: [dhavey@gmail.com](mailto:dhavey@gmail.com)

Marcelo Bagnulo  
Universidad Carlos III de Madrid  
Email: marcelo@it.uc3m.es