Network Working Group
Internet-Draft
Intended status: Informational
Expires: 27 November 2026

C. Giese
RtBrick
R. Patterson
Sky UK

26 May 2026

DHCP Explicit Rate Signaling
draft-giese-dhcp-rate-signaling-00

Abstract

This document defines new Dynamic Host Configuration Protocol (DHCP) options for both DHCPv4 and DHCPv6 to explicitly signal available upstream and downstream data rates. In many broadband access networks, Customer Premises Equipment (CPE) and intermediate nodes lack visibility into the subscriber’s provisioned service tier. By communicating these capacities natively via DHCP, clients, relay agents, and snooping switches can dynamically configure localized traffic shaping and queuing. This explicit signaling improves overall network performance by reducing the reliance on indiscriminate packet dropping and policing at the service edge. Additionally, it provides the necessary capacity awareness to enable effective Active Queue Management (AQM) and the Low Latency, Low Loss, and Scalable Throughput (L4S) architecture.

About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at https://datatracker.ietf.org/doc/draft-giese-dhcp-rate-signaling/.

Source for this draft and an issue tracker can be found at https://github.com/GIC-de/draft-dhcp-rate-signaling.

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/.

Copyright Notice

Table of Contents

1. Introduction

1.1. Problem Statement

In typical broadband access networks, the Customer Premises Equipment (CPE) is often unaware of the actual available data rates. This lack of visibility often occurs when an external modem or Optical Network Terminal (ONT) [G.984.1] connects to the CPE at a physical link speed that significantly exceeds the subscriber's provisioned service rate. Furthermore, operators commonly deploy unified access profiles where the available rate is artificially limited at the service edge, such as the Broadband Network Gateway (BNG) [TR101], to match the subscriber's purchased service tier.

When the network bottleneck resides between the BNG and the CPE, intermediate devices are typically not well equipped to provide deep buffering, priority scheduling, or Active Queue Management (AQM) [RFC7567]. Relying on indiscriminate packet dropping or policing at the service edge severely degrades the user experience. Conversely, network performance improves significantly by performing intelligent shaping and prioritization. When combined with AQM and the Low Latency, Low Loss, and Scalable Throughput (L4S) architecture [RFC9330], these localized traffic management benefits are further amplified.

1.2. Architectural Context

In many IP over Ethernet (IPoE) [TR101] architectures, the Broadband Network Gateway (BNG) operates strictly as a DHCP relay agent. While per-subscriber traffic management policies, such as queues, shapers, and policers, are typically provisioned out-of-band via Authentication, Authorization, and Accounting (AAA) protocols like RADIUS [RFC2865], deployments lacking direct AAA integration at the service edge require an in-band signaling alternative.

Transporting available data rates natively within DHCP options addresses this architectural constraint. This mechanism allows the BNG to dynamically instantiate downstream shapers and upstream policers based on the DHCP payload, while concurrently equipping the Customer Premises Equipment (CPE) with the parameters required to manage upstream traffic locally.

Furthermore, intermediate Layer 2 access nodes performing DHCP snooping can passively extract these explicitly signaled rate parameters to optimize their local queues and shapers. By distributing capacity awareness across the entire forwarding path, this in-band signaling mitigates buffer congestion within the access segment and significantly improves end-to-end transport performance.

1.3. Applicability and Benefits

While auto-configuration protocols such as TR-069 [TR069] can provision rate information, they are not universally deployed by all service providers. Furthermore, the increasing prevalence of customer-owned, unmanaged CPE devices, including devices running open-source firmware or custom projects, limits the effectiveness of operator-managed configuration servers. A standardized DHCP option addresses this gap by providing a universal mechanism to explicitly signal available data rates directly from the DHCP server, across BNGs and access nodes, down to the CPE. This localized approach is particularly advantageous because it serves the entire path, whereas auto-configuration servers exclusively target the end device. This method also integrates seamlessly with architectures where RADIUS servers inject DHCP options.

Although primarily designed for IPoE deployments, this mechanism is equally applicable to Point-to-Point Protocol over Ethernet (PPPoE) [RFC2516] environments. Since DHCP is frequently utilized over PPPoE, most notably for IPv6 Prefix Delegation, this option provides a standardized method for rate signaling. Consequently, this approach can supersede the fragmented, proprietary methods currently in use, such as embedding rate limits within PPP [RFC1661] authentication reply messages.

Conveying rates via DHCP natively supports dynamic updates through regular lease renewals or triggered reconfiguration requests. As an auxiliary benefit, this explicit rate information can be exposed in the CPE user interface to satisfy customer expectations regarding bandwidth visibility. Ultimately, introducing a DHCP option to signal available data rates represents a simple, standardized enhancement that yields widespread improvements in internet service delivery.

Operational requirements necessitate the definition of this option for both DHCPv4 and DHCPv6. This ensures coverage for networks lacking IPv6 support and prevents configuration gaps in dual-stack scenarios where DHCPv4 is established prior to DHCPv6.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

This document makes use of the terms defined in [RFC2131] and [RFC8415]. The following additional terms are used:

DHCP The abbreviation DHCP is used throughout this document to refer to both DHCPv4 and DHCPv6 protocols.

DHCPv4 Dynamic Host Configuration Protocol [RFC2131]

DHCPv6 Dynamic Host Configuration Protocol for IPv6 [RFC8415]

Subscriber
: The individual, organization, or entity that maintains a contractual relationship with a Broadband Service Provider for network access services. Within the network infrastructure, a subscriber is typically represented by an authenticated logical session (e.g., IPoE or PPPoE) and an associated policy profile that dictates service attributes, including provisioned upstream and downstream data rates.

3. DHCP Rate Option

The DHCP Rate Option specified in this document employs a unified sub-option structure for both DHCPv4 and DHCPv6, utilizing the format explicitly known from DHCPv4 Option 82 (the Relay Agent Information Option) [RFC3046]. The top-level option encapsulation strictly conforms to the requirements of each base protocol. To simplify cross-protocol implementation, this document proposes the uniform assignment of OPTION_RATE with the option code TBD1 (value to be assigned by IANA) for both IP versions. Specifically, DHCPv4 utilizes an 8-bit option code set to TBD1 alongside an 8-bit length field, whereas DHCPv6 utilizes a 16-bit option code set to TBD1 alongside a 16-bit length field. Despite these differences in outer header sizing, the internal payload remains completely identical. The encapsulated sub-options maintain consistent 8-bit sub-option code and 8-bit length fields across both protocol versions, ensuring

common parsing and processing logic regardless of the underlying IP version.

3.1. Sub-Options Format

The rate information field consists of a sequence of SubOpt/Length/Value tuples for each sub-option, encoded in the following manner:

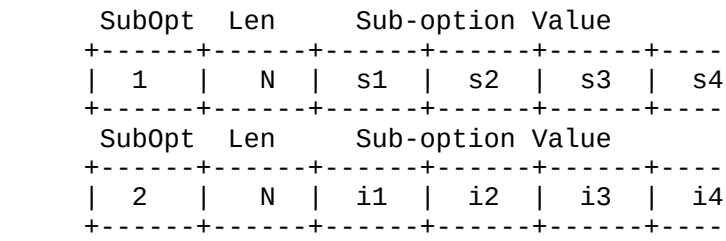

No "pad" sub-option is defined, and the rate information field SHALL NOT be terminated with a 255 sub-option. The length of OPTION_RATE MUST include all bytes of the sub-option code/length/value tuples. The length N of the sub-options MUST be the number of octets in only that sub-option’s value field. A sub-option length MAY be zero. The sub-options need not appear in sub-option code order.

The initial assignment of DHCP Rate Sub-options is as follows:

Sub-option Code	Length	Description
1	8	Available Rate Upstream in bits per second (bps)
2	8	Available Rate Downstream in bits per second (bps)
3	1	Rate Type (L2 or L3)

Table 1

3.2. Sub-Options

3.2.1. Available Rate Upstream

The sub-option Available Rate Upstream defines the rate in bits per second (bps) available from the DHCP client towards the DHCP server direction. The rate format is a 64-bit unsigned integer in network byte order.

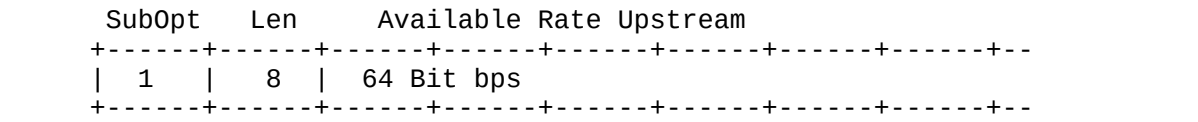

A value of 0 in this context signifies an unrestricted rate. It can be interpreted as a request to remove a previously set rate, thereby resetting to the device default configuration.

3.2.2. Available Rate Downstream

The available rate downstream defines the rate in bits per second (bps) available from the DHCP server towards the DHCP client direction. The rate format is a 64-bit unsigned integer in network byte order.

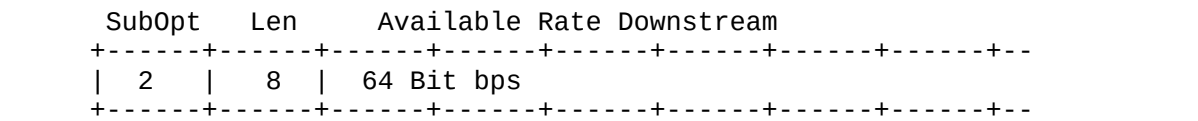

A value of 0 in this context signifies an unrestricted rate. It can be interpreted as a request to remove a previously set rate, thereby resetting to the device default configuration.

3.2.3. Rate Type

The rate type defines the networking layer to which the stated rates apply. The default value of 2 is defined as the Layer 2 rate, which signifies that the rate encompasses the entire Ethernet frame. Implementations SHOULD calculate this rate using the Ethernet header and all payload, excluding the Ethernet Frame Check Sequence (FCS) and Inter-Packet Gap (IPG).

If the rate type is set to 0, the explicitly signaled rates are informational only. Devices receiving this rate type MUST NOT apply the specified rate limits to their physical interfaces, traffic shapers, policers, or Active Queue Management (AQM) parameters. This value accommodates deployments where the network exposes the provisioned service tier to the Customer Premises Equipment (CPE) solely to populate user interfaces or for telemetry purposes, without altering the device’s localized forwarding behavior.

If the rate type is set to 3, the rate applies to Layer 3, encompassing only the IP header and its payload. This rate calculation is frequently utilized by end-user speed test applications and is often regarded as the marketable "product bandwidth". Its primary advantage is its independence from the variable overhead introduced by differing numbers of VLAN tags or tunnel encapsulations.

In the absence of the Rate Type sub-option, the client MUST assume that the signaled values are defined as the Layer 2 rate.

The values 1 and 4-255 are currently unassigned and reserved for future use. If a client, server, or relay agent receives a Rate Type sub-option containing an unrecognized or reserved value, it MUST ignore the entire OPTION_RATE. Applying explicitly signaled rate limits without understanding the intended networking layer could result in incorrect localized traffic management. Therefore, to fail safely, the receiving device MUST discard the option entirely and rely on its default rate configuration.

A client MAY include the Rate Type sub-option within its initial requests to serve as a hint to the server regarding its preferred calculation method (e.g., requesting a Layer 3 rate instead of a Layer 2 rate). Sending this hint is OPTIONAL for the client, and honoring the hint is OPTIONAL for the server.

4. DHCPv4

4.1. DHCPv4 Rate Option

The DHCPv4 OPTION_RATE code is TBD1.

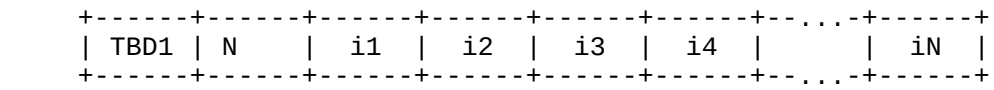

The length (N) gives the total number of octets in the Rate Information Field, which is either zero or longer than 2 bytes, which is the sub-options header length.

4.2. DHCPv4 Client Behavior

DHCPv4 clients that support the DHCP Rate Option SHOULD include the corresponding OPTION_RATE code in the Parameter Request List (PRL). This inclusion explicitly signals client support, enabling the DHCP server to determine whether to include the rate parameters in its response. Furthermore, this signaling provides network operators with visibility into client capabilities, which aids in troubleshooting and resolving customer service quality complaints.

A client MAY include the DHCP Rate Option directly within its DHCPREQUEST messages to serve as a hint to the server proposing their maximum data rates or preferred rate type (L2 or L3 rates). For example, when a CPE device is connected via a 1 Gbps WAN interface to an external Optical Network Terminal (ONT) or modem capable of exceeding 1 Gbps on the WAN side, the CPE can explicitly signal this physical limitation to the service provider. This enables the network to align its shaping parameters directly with the device's actual capacity, ensuring traffic does not exceed the physical limit.

However, providing this hint is OPTIONAL. The manner in which a DHCP server processes or utilizes these client-provided hints is implementation-specific and outside the scope of this document. Because the server remains authoritative, the client MUST accept and apply the rate type ultimately provided by the server in the DHCPACK message, regardless of the hint it originally sent.

Clients MUST ignore the OPTION_RATE when received within a message other than DHCPACK. If the OPTION_RATE is present in a DHCPOFFER message, the client MUST NOT apply the specified rate limits to its interfaces. However, a client MAY evaluate the rate information provided in a DHCPOFFER as a selection criterion to prefer one server's offer over another.

4.3. DHCPv4 Server Behavior

When a DHCPv4 server is configured to support the DHCP Rate Option and receives a client request (e.g., DHCPDISCOVER or DHCPREQUEST) that includes the OPTION_RATE within the Parameter Request List (PRL), the server MUST include the OPTION_RATE in the resulting DHCPOFFER and DHCPACK messages.

The server MAY derive the specific upstream and downstream rates and rate type from local configuration profiles, centralized Authentication, Authorization, and Accounting (AAA) systems such as RADIUS, or external policy servers. If no non-zero values are configured or signaled to be used, the server MAY return rate values of 0.

A server MAY include the OPTION_RATE in its responses even if the client did not explicitly request it via the Parameter Request List (PRL), provided the operator has explicitly configured the server to forcefully inject the option to provision intermediate nodes, such as DHCP relay agents or Layer 2 snooping switches, which MAY drop these options before forwarding the message to the client.

4.4. DHCPv4 Relay Agent Behavior

DHCPv4 Relay Agents, including L2 DHCPv4 Relay Agents [TR101], MAY extract the OPTION_RATE from DHCPACK messages traversing the network. Relay agents that perform localized traffic management MAY utilize these extracted values to dynamically instantiate shapers and policers on their subscriber-facing interfaces.

Furthermore, a relay agent MAY add, modify, or remove the OPTION_RATE before forwarding the DHCP message to the client. This accommodates deployments where the relay agent (e.g., a BNG) is responsible for policy enforcement and populates or overrides the OPTION_RATE based on subscriber attributes retrieved directly from an external Authentication, Authorization, and Accounting (AAA) server, such as RADIUS.

5. DHCPv6

5.1. DHCPv6 Rate Option

The DHCPv6 OPTION_RATE code is TBD1.

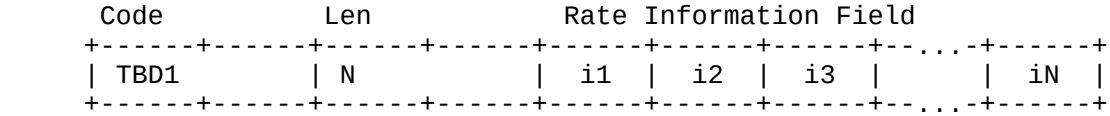

The length (N) gives the total number of octets in the Rate Information Field, which is either zero or longer than 2 bytes, which is the sub-options header length.

5.2. DHCPv6 Client Behavior

DHCPv6 clients that support the DHCP Rate Option SHOULD include the corresponding OPTION_RATE code in the Option Request Option (ORO) [RFC8415]. This inclusion explicitly signals client support, enabling the DHCPv6 server to determine whether to include the rate parameters in its response. Furthermore, this signaling provides network operators with visibility into client capabilities, which aids in troubleshooting and resolving customer service quality complaints.

A client MAY include the OPTION_RATE with any or all sub-options within its DHCPv6 Request messages to serve as a hint to the server proposing their maximum data rates or preferred rate type (L2 or L3 rates).

Providing this hint is OPTIONAL. The manner in which a DHCPv6 server processes these client-provided hints is implementation-specific. Because the server remains authoritative, the client MUST accept and apply the rate type ultimately provided by the server in the REPLY message, regardless of the hint it originally sent.

Clients MUST ignore the OPTION_RATE when received within a message other than REPLY. If the OPTION_RATE is present in an ADVERTISE message, the client MUST NOT apply the specified rate limits to its interfaces. However, the client MAY evaluate this early rate visibility as a selection criterion to prefer one server’s advertisement over another.

5.3. DHCPv6 Server Behavior

A DHCPv6 server MAY embed the OPTION_RATE directly within a REPLY message encapsulated inside RELAY-REPL messages to explicitly provision the end client. The server MAY include the option within the RELAY-REPL message to target the corresponding relay agent, instructing it to apply rate limits locally. The nested relay header architecture of DHCPv6 empowers the server to explicitly address each relay agent in the path, in addition to the end client, ensuring precise targeting of signaling parameters.

If network policy dictates localized traffic management at both the Customer Premises Equipment (CPE) and the relay node, the server MAY include the OPTION_RATE at all encapsulation levels simultaneously. When provisioning multiple levels, the server MAY supply different rate values to each respective node. For example, an operator might configure a relay agent’s upstream policer with a slightly higher rate limit than the CPE’s upstream shaper. This operational delta accommodates minor traffic burstiness from the CPE and prevents premature packet drops at the intermediate access node. This capability to independently target different nodes along the forwarding path is unique to the nested relay header architecture of DHCPv6, as DHCPv4 lacks a comparable mechanism for addressing multiple relay agents distinctly.

Furthermore, to dynamically update a client's rate limits mid-lease, the server MAY utilize RECONFIGURE messages to apply updates before the T1 timer expires. By triggering the client to initiate a Renew or Information-request transaction, this mechanism allows the server to push newly modified rate parameters without waiting for timer expiration.

5.4. DHCPv6 Relay Agent Behavior

DHCPv6 Relay Agents, including Lightweight DHCPv6 Relay Agents (LDRA) [RFC6221], MUST extract and consume the OPTION_RATE from their corresponding Relay-Reply header. Because the DHCPv6 architecture provides a dedicated signaling channel for intermediate nodes, relay agents MUST NOT passively inspect the encapsulated client-facing REPLY payload to extract rate information. Relay agents that perform localized traffic management MAY utilize these explicitly targeted values to dynamically instantiate shapers, policers, or Active Queue Management (AQM) disciplines on their subscriber-facing interfaces.

In many architectures, the Broadband Network Gateway (BNG) or similar intermediary device serves as the authoritative policy enforcement point rather than a transparent relay. In such deployments, the intermediary device MAY populate, modify, or remove the OPTION_RATE destined for the client. This allows the network edge to inject dynamic rate parameters based on subscriber attributes retrieved directly from an external Authentication, Authorization, and Accounting (AAA) server, such as RADIUS, before the message reaches the downstream client.

6. DHCP Snooping

DHCP snooping switches are typically deployed as intermediate Layer 2 devices that passively monitor DHCP message exchanges to enforce security policies and build binding databases, all without modifying the DHCP payloads. These devices MAY passively inspect the DHCP Rate Option within messages destined for clients. By extracting these explicit rate parameters, snooping devices can dynamically provision appropriate traffic shapers, policers, or hardware queues on the corresponding downstream, client-facing ports.

7. PPPoE

In Point-to-Point Protocol over Ethernet (PPPoE) [RFC2516] architectures, the Customer Premises Equipment (CPE) typically employs DHCPv6 over the PPP [RFC1661] link to request an IPv6 Delegated Prefix (IA_PD) [RFC8415]. This encapsulated DHCPv6 exchange provides a standardized transport mechanism for the explicit DHCP Rate Option. While less prevalent in modern deployments, DHCPv4

transactions operating within a PPPoE session MAY similarly convey these rate options.

The foundational processing rules and client behavior for rate options received over PPPoE are identical to those defined for IP over Ethernet (IPoE) environments.

If a client receives rate limits embedded within the PPP authentication reply message and concurrently receives the DHCP OPTION_RATE, the explicitly signaled DHCP OPTION_RATE MUST take priority. This precedence ensures that the standardized, dynamic DHCP signaling supersedes fragmented or proprietary rate limits previously negotiated during the PPP authentication phase.

Because the PPP session dictates the primary logical link state, the applied rate MUST revert to the device’s default configuration under two specific conditions. First, the client MUST reset the applied limits if it receives a valid DHCP message explicitly signaling rate removal with a sub-option containing a rate value of zero. Second, the rate MUST be implicitly revoked if the underlying PPPoE session itself is terminated.

8. Interaction with AQM and L4S

Active Queue Management (AQM) mechanisms, as recommended in [RFC7567], are most effective when they operate at or near the true bottleneck rate for a given service. In many broadband deployments today, Customer Premises Equipment (CPE) and intermediate access nodes configure shaping and AQM parameters against the physical port speed rather than the subscriber’s provisioned rate, which can lead either to persistent queues and excess latency when configured too high, or to underutilization when configured too low.

By explicitly signaling per-subscriber upstream and downstream rates via the DHCP Rate Option, this document enables CPE devices, relay agents, and snooping switches to instantiate shapers and AQM instances that closely track the actual bottleneck capacity for each subscriber. Placing the bottleneck queue under control of an AQM that follows the recommendations in [RFC7567] allows operators to limit queue growth and reduce queuing delay while still efficiently utilizing the contracted bandwidth. The Low Latency, Low Loss, and Scalable Throughput (L4S) architecture [RFC9330] relies on shallow queues under an L4S-compatible AQM and on congestion controllers that react promptly to Explicit Congestion Notification (ECN) signals. Providing accurate rate information to devices at or near the bottleneck link allows those devices to configure L4S-capable AQMs at the appropriate shaping rate, so that L4S flows can achieve consistently low queuing delay while still fully utilizing the

subscriber’s provisioned service tier. The DHCP Rate Option defined in this document is therefore an enabler for deploying L4S and other modern AQM schemes in access networks, even though the detailed design of AQM and congestion control algorithms remains outside the scope of this specification.

9. Errors and Conflicts

Clients receiving conflicting rate information across DHCPv4 and DHCPv6 protocols SHOULD apply the most recently received value.

To ensure forward compatibility, clients, servers, and relay agents MUST ignore unrecognized sub-option codes and continue processing the remainder of the Rate Option.

Conversely, if a device receives a recognized sub-option containing an unrecognized or reserved value that dictates the fundamental interpretation of the rate parameters (such as an unassigned Rate Type), it MUST discard the entire OPTION_RATE. Applying explicitly signaled rates without understanding the intended networking layer could result in incorrect localized traffic management. In such cases, the device MUST rely on its default rate configuration.

If multiple instances of the same sub-option code are present, the last instance MUST be processed.

A client may receive an OPTION_RATE indicating an Available Rate that exceeds the maximum physical link speed of its upstream or downstream interfaces (e.g., signaling 2 Gbps to a CPE with a 1 Gbps WAN port). In such scenarios, the client SHOULD cap the applied traffic shaping, policing, or Active Queue Management (AQM) parameters to the maximum capacity of the physical link.

10. Operational or Manageability Considerations

Deploying explicit rate signaling via DHCP introduces several operational benefits and deployment considerations for network management. When combined with appropriately configured AQM and, where deployed, L4S-compatible queue management, these per-subscriber rate parameters help to concentrate congestion control at a well-defined bottleneck and minimize queuing delay in the access segment.

Implementations SHOULD expose the explicitly signaled, DHCP-learned rate parameters within Customer Premises Equipment (CPE) management interfaces, such as the local Web User Interface (UI) or remote management protocols. Providing end-users and operators with immediate visibility into the locally provisioned service tier significantly reduces support inquiries related to perceived bandwidth issues and improves overall user satisfaction.

As noted in Errors and Conflicts, a client may receive an OPTION_RATE indicating an available rate that exceeds the maximum physical link speed of its interfaces. In such scenarios, management interfaces SHOULD expose both the originally signaled rate value and the effective, capped rate value applied by the device. Additionally, implementations SHOULD log this discrepancy if logging facilities are enabled. Capturing and exposing these specific events provides critical telemetry for network operators, as they frequently indicate a mismatch between the subscriber’s provisioned service tier and their installed physical equipment.

11. Customer-Owned and Open Source CPE

A key motivation for this option is to support customer-owned or subscriber-managed CPE, including retail routers and devices running open-source firmware (for example, OpenWrt) that are not integrated with operator auto-configuration systems such as TR-069. In these environments, the access network can expose the provisioned upstream and downstream rates via DHCP, and CPE implementations that understand this option MAY use the learned values to configure local shaping, policing, queue management, or simple rate indicators in their user interfaces.

Because the option format is intentionally simple and identical for DHCPv4 and DHCPv6, it is straightforward for open-source projects and custom CPE implementations to add support without requiring coupling to any specific vendor management system. Even if no advanced AQM features are present, aligning any local rate limits with the signaled values helps avoid misconfiguration and reduces the likelihood of bufferbloat in the customer-owned equipment.

12. Security Considerations

DHCP messages are typically transmitted as plaintext and are unauthenticated. Consequently, the DHCP Rate Options defined in this document are vulnerable to interception, modification, or spoofing by on-path attackers. A malicious actor or a successfully deployed rogue DHCP server could inject artificially low rate limits to severely throttle a client’s connection, resulting in a localized Denial of Service (DoS).

The severity of this specific risk is generally no greater than standard DHCP threat vectors, such as rogue default gateway assignment, DNS hijacking, or IP pool exhaustion, which typically yield a much more critical and immediate loss of service.

To mitigate the impact of malicious or malformed options, clients MAY implement basic sanity checks and threshold validations before applying rate parameters. For example, clients MAY ignore downstream or upstream rates that fall below a basic operational minimum (e.g., 1000 bps) to prevent complete session starvation. Furthermore, as mandated in the client behavior specification, if a maliciously injected rate is impractically high, the client implicitly mitigates this by capping the applied rate to the physical link capacity.

13. IANA Considerations

This document requests that IANA assign the same numeric value in both registries for DHCPv4 and DHCPv6, if feasible.

13.1. DHCPv4 Option

IANA is requested to assign a new DHCP Option Code (TBD1) in the "BOOTP Vendor Extensions and DHCP Options" registry for OPTION_RATE.

13.2. DHCPv6 Option

IANA is requested to assign a new DHCPv6 Option Code (TBD1) in the "Option Codes" registry for OPTION_RATE.

13.3. DHCP Rate Sub-Options Registry

IANA is requested to create a new registry titled "DHCP Rate Sub-Options".

Value	Description	Reference
0	Unassigned	
1	Available Rate Upstream in bits per second (bps)	RFC TBD2
2	Available Rate Downstream in bits per second (bps)	RFC TBD2
3	Rate Type (L2, L3 or informational)	RFC TBD2
4-255	Unassigned	

Table 2

14. References

14.1. 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/rfc/rfc2119>.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <https://www.rfc-editor.org/rfc/rfc2131>.

[RFC3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, DOI 10.17487/RFC3046, January 2001, <https://www.rfc-editor.org/rfc/rfc3046>.

[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/rfc/rfc8174>.

[RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <https://www.rfc-editor.org/rfc/rfc8415>.

14.2. Informative References

[G.984.1] ITU-T, "Gigabit-capable passive optical networks (GPON): General characteristics", ITU-T Recommendation G.984.1, March 2008.

[RFC1661] Simpson, W., Ed., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, DOI 10.17487/RFC1661, July 1994, <https://www.rfc-editor.org/rfc/rfc1661>.

[RFC2516] Mamakos, L., Lidl, K., Evarts, J., Carrel, D., Simone, D., and R. Wheeler, "A Method for Transmitting PPP Over Ethernet (PPPoE)", RFC 2516, DOI 10.17487/RFC2516, February 1999, <https://www.rfc-editor.org/rfc/rfc2516>.

[RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, DOI 10.17487/RFC2865, June 2000, <https://www.rfc-editor.org/rfc/rfc2865>.

[RFC6221] Miles, D., Ed., Ooghe, S., Dec, W., Krishnan, S., and A. Kavanagh, "Lightweight DHCPv6 Relay Agent", RFC 6221, DOI 10.17487/RFC6221, May 2011, <https://www.rfc-editor.org/rfc/rfc6221>.

[RFC7567] Baker, F., Ed. and G. Fairhurst, Ed., "IETF Recommendations Regarding Active Queue Management", BCP 197, RFC 7567, DOI 10.17487/RFC7567, July 2015, <https://www.rfc-editor.org/rfc/rfc7567>.

[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/rfc/rfc9330>.

[TR069] Broadband Forum, "CPE WAN Management Protocol", TR 069 Issue 6, June 2020.

[TR101] Broadband Forum, "Migration to Ethernet-Based Broadband Aggregation", TR 101 Issue 2, July 2011.

Acknowledgments

The authors would like to thank Glenn Deen and Jason Livingood for their valuable review comments and discussion, which helped to significantly improve the clarity, applicability, and operational guidance of this document.

Authors' Addresses

Christian Giese
RtBrick
Email: christian@rtbrick.com

Richard Patterson
Sky UK
Email: Richard.Patterson@sky.uk