



Consumer  
Technology  
Association™



# CTA Specification

Common Media Server Data

## CTA-5006

November 2022

## NOTICE

Consumer Technology Association (CTA)<sup>™</sup> Standards, Bulletins and other technical publications are designed to serve the public interest through eliminating misunderstandings between manufacturers and purchasers, facilitating interchangeability and improvement of products, and assisting the purchaser in selecting and obtaining with minimum delay the proper product for his particular need. Existence of such Standards, Bulletins and other technical publications shall not in any respect preclude any member or nonmember of the Consumer Technology Association from manufacturing or selling products not conforming to such Standards, Bulletins or other technical publications, nor shall the existence of such Standards, Bulletins and other technical publications preclude their voluntary use by those other than Consumer Technology Association members, whether the standard is to be used either domestically or internationally.

Standards, Bulletins and other technical publications are adopted by the Consumer Technology Association in accordance with the American National Standards Institute (ANSI) patent policy. By such action, the Consumer Technology Association does not assume any liability to any patent owner, nor does it assume any obligation whatever to parties adopting the Standard, Bulletin or other technical publication.

This document does not purport to address all safety problems associated with its use or all applicable regulatory requirements. It is the responsibility of the user of this document to establish appropriate safety and health practices and to determine the applicability of regulatory limitations before its use.

Copyright © 2022 by the Consumer Technology Association (CTA)<sup>™</sup>. All rights reserved. This document may not be reproduced, in whole or part, without written permission. Federal copyright law prohibits unauthorized reproduction of this document by any means. Organizations may obtain permission to reproduce a limited number of copies by entering into a license agreement. Requests to reproduce text, data, charts, figures or other material should be made to the Consumer Technology Association (CTA)<sup>™</sup>.

(Formulated under the cognizance of the CTA **WAVE Project**; for information please see [cta.tech/WAVE](http://cta.tech/WAVE).)

Published by  
CONSUMER TECHNOLOGY ASSOCIATION  
Technology & Standards Department  
[www.cta.tech](http://www.cta.tech)

All rights reserved

## Table of Contents

<b>1</b>	<b>Introduction</b> .....	<b>1</b>
<b>2</b>	<b>Network Nomenclature</b> .....	<b>1</b>
<b>3</b>	<b>Data Transmission Mode</b> .....	<b>2</b>
<b>4</b>	<b>Payload Definition</b> .....	<b>3</b>
<b>5</b>	<b>Inheritance</b> .....	<b>5</b>
<b>6</b>	<b>Caching</b> .....	<b>5</b>
6.1	Caching of Dynamic Fields .....	5
6.2	Supporting Conditional Responses (HTTP 304) .....	6
6.3	Computation of ETag .....	6
<b>7</b>	<b>Reserved Keys</b> .....	<b>6</b>
<b>8</b>	<b>Client Processing Requirements</b> .....	<b>10</b>
8.1	Common Client Processing Requirements .....	10
8.2	Additional Intermediary Processing Requirements .....	11
8.3	Additional Player Processing Requirements .....	11
<b>9</b>	<b>Security and Privacy</b> .....	<b>11</b>
9.1	Threat Environment .....	11
9.2	Threats to the Intermediary and Media Clients .....	11
9.3	Specific Mitigations .....	12
<b>10</b>	<b>References</b> .....	<b>12</b>
<b>Annex A.</b>	<b>CMSD Header Examples (Informative)</b> .....	<b>14</b>
<b>Annex B.</b>	<b>Informative Use-Case Definitions (Informative)</b> .....	<b>16</b>

## Figures

---

Figure 1: Response flows from an origin to a user-agent .....	2
---	---

## Tables

---

Table 1: Reserved keys .....	6
------------------------------	---

## FOREWORD

---

This document was developed by the Web Application Video Ecosystem (WAVE) Project of the Consumer Technology Association<sup>1</sup>. The WAVE Project is a broad industry initiative of content, technology, infrastructure and device companies, all working together towards commercial Internet video interoperability based on industry standards.

---

<sup>1</sup> See <https://cta.tech/WAVE>

# Common Media Server Data

## 1 Introduction

---

Adaptive streaming of segmented media is enabled by media players requesting media objects from servers. These servers are arranged in a hierarchy starting with the origin server, which holds the authoritative copy of the content requested by user agents and other servers. Outbound [RFC9110] responses traverse a series of mid-tier and edge intermediaries, known collectively as Content Distribution Networks (CDNs). These CDNs may themselves be stacked. The edge servers are the outermost servers. Edge servers are the first intermediaries to receive user-agent requests, in a given request/response flow, and the last intermediary to forward a response when communicating with media players. The origin servers know information about the media object which the CDN servers do not. For example, they may know the format, the duration and the encoded bitrate of a media object. In the case of live streams, they may know for how long the object has been available and the likely next object in the sequence. The edge servers in turn know information unavailable to the origin or players. For example, they may know the throughput available in the next network hop, or the cache status of the various objects or the accumulated history of the media object as it was moved from origin to edge server.

The purpose of the Common Media Server Data (CMSD) specification is to define a standard means by which every media server (intermediate and origin) can communicate data with each media object response and have it received and processed consistently by every intermediary and player, for the purpose of improving the efficiency and performance of distribution and ultimately the quality of experience enjoyed by the users.

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.

## 2 Network Nomenclature

---

CMSD defines structure for data transmitted in the response to a request from a media player for an HTTP adaptive streaming media object. The response usually originates at an origin server and is then propagated through a series of intermediaries to the player. Figure 1 illustrates various combinations of intermediaries possible in response data flows from an origin to a user agent.

## Web Application Video Ecosystem – Common Media Server Data

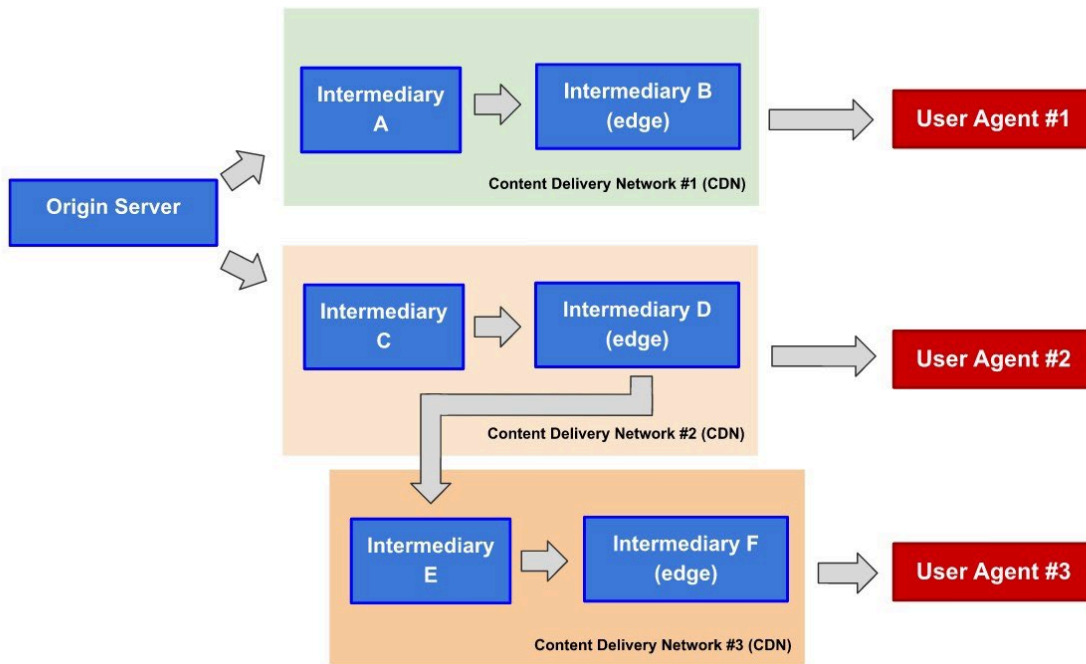


Figure 1: Response flows from an origin to a user-agent

The definitions of these terms are given in [RFC9110] and are repeated here for convenience:

**Server:** An HTTP server is an entity that accepts connections in order to service HTTP requests by sending HTTP responses.

**Client:** An HTTP client is an entity that establishes a connection to a server for the purpose of sending one or more HTTP requests.

**Intermediary:** HTTP enables the use of intermediaries to satisfy requests through a chain of connections. There are three common forms of intermediaries: proxy, gateway and tunnel.

**Edge Server:** An intermediary that communicates directly with a requesting user agent.

**Origin:** A program that can originate authoritative responses for a given target resource. In the context of this document, this typically refers to an HTTP service provided by a collection of servers.

**User Agent:** Any of the various client programs that initiate a request and that run on an entity that is not an intermediary. In the context of media delivery, the user agent is commonly referred to as the 'player'.

### 3 Data Transmission Mode

Common Media Service Data is transferred via a set of custom HTTP response headers:

## Web Application Video Ecosystem – Common Media Server Data

- CMSD-Dynamic: Keys whose values apply only to the next transmission hop. Typically a new CMSD-Dynamic header instance will be added by each intermediary participating in the delivery.
- CMSD-Static: Keys whose values persist over multiple requests for the object.

Examples of values that vary with each response would be estimated throughput and round-trip-time, while values that persist over multiple requests would be format type, media duration, etc.

### 4 Payload Definition

---

The data payload for the CMSD response headers consists of a series of key/value pairs constructed according to [RFC8941]. Specifically, the CMSD-Static header is of type *sf-dictionary* and CMSD-Dynamic is of type *sf-list*. Here, we provide an informational summary:

1. All information in the payload is represented as <key>=<value> pairs. In the CMSD-Static header, this is represented as dictionary members; in the CMSD-Dynamic header, this is represented as parameters of a list member.
2. The key and value are separated by an equals sign (Unicode 0x3D). If the value type is BOOLEAN and the value is TRUE, then the equals sign and the value are omitted.
3. Successive key/value pairs are delimited by a comma (Unicode 0x2C) unless they are carried in the CMSD-Dynamic header in which case they must be separated by a semicolon (Unicode 0x3B). This approach is referred to as parameterization of an item as per [RFC8941].
4. Any value of type String must be enclosed by opening and closing double quotes (Unicode 0x22). Double quotes and backslashes must be escaped using a backslash "\" (Unicode 0x5C) character. Any value of type Token does not require quoting. Unicode is not directly supported in Strings.
5. Inner Lists are denoted by surrounding parentheses (Unicode 0x28 and Unicode 0x29), and their values are delimited by one or more spaces (Unicode 0x20).
6. Data payloads transmitted via headers must not be URL-encoded. The definition for URL-encoding can be found in Section 5 of [WHATWG].

[RFC8941] remains normative and specifies many more formatting and structural requirements. The following rules are additionally applied by this specification:

1. The key names described in this specification are reserved. Custom key names may be used, but they MUST carry a hyphenated prefix to ensure that there will not be a namespace collision with future revisions to this specification. Servers SHOULD use a

## Web Application Video Ecosystem – Common Media Server Data

reverse-DNS syntax when defining their own prefix. Custom keys MUST only utilize the String value type. A compliant example custom key would be `com.example-mykey="23"`.

2. All key names are case-sensitive and lower case.
3. All keys are OPTIONAL.
4. Keys in the CMSD-Dynamic header are associated with the server from which the parameters originated. These parameters are separated from one another by a semicolon (Unicode 0x3B).
5. The order of CMSD-Dynamic header entries is significant and critical to the interpretation by the player. Per [RFC9110], a proxy MUST NOT change the order of the field lines with the same header name. Therefore, a proxy server may add CMSD-Dynamic keys in one of two ways:
  - a. By adding a new CMSD-Dynamic header. In this case, it MUST NOT change the order of the existing CMSD-Dynamic headers and the new CMSD-Dynamic header MUST be added as the last member.
  - b. By appending key/value pairs to the end of an existing CMSD-Dynamic header entry. In this case, values MUST be delimited using a comma (Unicode 0x2C).
6. Transport Layer Security (TLS) SHOULD be used to protect all transmission of CMSD data.

CMSD-Dynamic list elements are added by and associated with a single server that handles the response. Each server adds its identifier and associated parameters to the end of the list according to the rules above. This way, the first list element is associated with the server nearest to the origin or the origin itself, and the last element is associated with the most recent server to handle the response.

The list member is a String that identifies the server. The value SHOULD identify both the organization and the intermediary that is writing the key. Identifiers SHOULD be as concise as possible to reduce log file size, while remaining unique.

The identifier used for the CMSD-Dynamic list member SHOULD be identical to the identifier used for a Proxy-Status, Cache-Status, or any other similar agglutinative header. It MUST be identical to the identifier used for the Origin Identifier (n) of the CMSD-Static header if the same server sets both headers. Each intermediary SHOULD implement Proxy-Status [RFC9209] and Cache-Status [RFC9211].

If an intermediary (such as an Origin Shield) has additional information about an object and it belongs to the same entity which has already written CMSD-Static data, then it is permissible to update or extend the CMSD-Static data without adding a new entity key.

Below are two example sets of headers as received by the client. Either is acceptable and they are semantically equivalent. They represent the journey of the object from an origin, through four intermediaries across two entities, to a player:



## Web Application Video Ecosystem – Common Media Server Data

### Example 1:

CMSD-Static:ot=v,sf=h,st=v,d=5000,br=2000,n="OriginProviderA"

CMSD-Dynamic: "CDNB-3ak1";etp=96;rtt=8

CMSD-Dynamic: "CDNB-w35k";etp=76;rtt=32

CMSD-Dynamic: "CDNA-987.343";etp=48;rtt=30

CMSD-Dynamic: "CDNA-312.663";etp=115;rtt=16;mb=5000

### Example 2:

CMSD-Dynamic: "CDNB-3ak1";etp=96;rtt=8,"CDNB-w35k";etp=76;rtt=32,"CDNA-987.343";etp=48;rtt=30,"CDNA-312.663";etp=115;rtt=16;mb=5000

CMSD-Static:ot=v,sf=h,st=v,d=5000,br=2000,n="OriginProviderA"

More comprehensive examples are available in Annex A.

## 5 Inheritance

---

An intermediary SHOULD persist any CMSD key/value pairs that it receives from an upstream server. An intermediary MAY choose to concatenate multiple prior CMSD-Dynamic headers into a single CMSD-Dynamic entry per [RFC9110], Section 5.3. If this is done, then the order MUST be preserved and the entries MUST be separated by a comma (Unicode 0x2C).

An intermediary MAY add key/value pairs to the header and MAY choose to modify or remove certain keys whose propagation would lead to deleterious behavior.

Since all headers are optional, an intermediary that does not have any information to add to a CMSD-Dynamic header MAY signal such by adding a member with no parameters. Furthermore, if it does not wish to add any information at all, it can use an empty string as its identifier. An intermediary MUST NOT respond with a CMSD-Dynamic header that it has not modified.

## 6 Caching

---

### 6.1 Caching of Dynamic Fields

The CMSD-Dynamic header fields can be cached by intermediaries. This means that a client could falsely attribute the CMSD-Dynamic header data to the current response, when in fact it applies to a previous response. It is not guaranteed to receive the latest data which may relate to a prior connection. To mitigate this risk, content distributors who control and operate players can require that CDN service providers, who operate edge servers, correctly add (or, if support for CMSD is unavailable, remove) CMSD-Dynamic headers at the edge. This requirement allows players to trust that the CMSD-Dynamic data added does in fact refer to the last-mile connection.

# Web Application Video Ecosystem – Common Media Server Data

## 6.2 Supporting Conditional Responses (HTTP 304)

Change of CMSD key values SHOULD NOT cause the HTTP body to be retransmitted, akin to non-conditional request/response.

## 6.3 Computation of ETag

Origin servers need to consider the impact of these static and dynamic fields when computing ETag. ETag MUST be computed on the body and MUST exclude the CMSD-Dynamic headers.

## 7 Reserved Keys

---

The reserved key names are defined in Table 1 below.

**Table 1: Reserved keys**

Description	Key Name	Header Name	Type & Unit	Value definition	Use-case reference
Availability time	at	CMSD-Static	Integer Milliseconds	The wallclock time at which the first byte of this object became available at the origin for successful request. The time is expressed as integer milliseconds since the Unix Epoch, i.e., the number of milliseconds that have elapsed since January 1, 1970 (midnight UTC/GMT), not counting leap seconds (in ISO 8601: 1970-01-01T00:00:00Z).	12
Duress	du	CMSD-Dynamic	Boolean	Key is included without a value if the server is under duress, due to cpu, memory, disk IO, network IO or other reasons. The thresholds for signaling duress are left to the discretion of the server operator. The intent is that the client will use this signal to move away to an alternate server if possible. This key MUST NOT be sent if it is false.	3
Encoded bitrate	br	CMSD-Static	Integer Kbps	The encoded bitrate of the audio or video object being requested. If the instantaneous bitrate varies over the duration of the object, the average value over the duration of the object SHOULD be communicated. This key should only accompany objects that have an implicit bitrate.	6

## Web Application Video Ecosystem – Common Media Server Data

Description	Key Name	Header Name	Type & Unit	Value definition	Use-case reference
Estimated Throughput	etp	CMSD-Dynamic	Integer Kbps	<p>The throughput between the server and the client over the currently negotiated transport as estimated by the server at the start of the response. The value is expressed in units of kilobits per second and rounded to the nearest integer. The time window for this estimate is expected to be the duration of the current response at most. The throughput may vary during the response and the client SHOULD use this data as an adjunct to its own throughput estimates. As an informative example, this estimate could be derived in this way:</p> $\text{etp} = 8 * \text{send\_window} / (\text{rtt})$ <p>where send_window = min (cwnd * mss, rwnd) with Congestion Window (cwnd) measured in packets, Maximum Segment Size (mss) in bytes, Receiver Window (rwnd) in bytes and rtt in milliseconds. Note that multiple client processes adjacent to the media player may pool their requests into the same connection to the server. In this case the server estimate of throughput will be against the entirety of the connection, not all of which will be accessible to the media player.</p>	2
Held time	ht	CMSD-Static	Integer Milliseconds	The number of milliseconds that this response was held back by an origin before returning. This is applicable to blocking responses under LL-HLS [HLSbis].	12
Intermediary identifier	n	CMSD-Static	String	An identifier for the processing server. The value SHOULD identify both the organization and the intermediary that is writing the key. Identifiers SHOULD be as concise as possible to reduce log file and transferred size, while still remaining unique.	1
Max suggested bitrate	mb	CMSD-Dynamic	Integer Kbps	The maximum bitrate value that the player SHOULD play in its Adaptive Bit Rate (ABR) ladder. If the player is playing a bitrate higher than this value, it SHOULD immediately switch to a bitrate lower than or equal to this value.	4,10

## Web Application Video Ecosystem – Common Media Server Data

Description	Key Name	Header Name	Type & Unit	Value definition	Use-case reference
Next Object Response	nor	CMSD-Static	Vertical line [Unicode 0x7C] delimited string	The URL-encoded relative path to one or more objects which can reasonably be expected to be requested by a media client consuming the current response. This key will typically be added by the origin. An intermediate server MAY use this key to perform a prefetch action. In the case of redirects, this path is relative to the final request. Each object SHOULD be fetched in its entirety unless a matching 'nrr' entry exists for that list element. Special care must be taken to percent-encode the " " character if it appears in the path.	5,7,11
Next Range Response	nrr	CMSD-Static	Vertical line [Unicode 0x7C] delimited string of ranges in the form "<range-start>-<range-end>"	If the next response will be a partial object response, then this string denotes the byte range that will be returned. If the 'nor' field is not set, then the object is assumed to match the object currently being served. Formatting is similar to the HTTP Range header, except that the unit MUST be 'byte', the 'Range:' prefix is NOT required, specifying multiple ranges is NOT allowed and the only valid form is "<range-start>-<range-end>".	5,7,11
Object duration	d	CMSD-Static	Integer in milliseconds	The playback duration in milliseconds of the object. If the value of playback duration is not known accurately, this parameter MUST be omitted. This key MUST NOT be used in responses to range requests against objects.	6
Object type	ot	CMSD-Static	Token - one of [m,a,v,av,i,c,tt,k,o]	<p>The media role of the current object being returned:</p> <ul style="list-style-type: none"> <li>m = text file, such as a manifest or playlist</li> <li>a = audio only</li> <li>v = video only</li> <li>av = muxed audio and video</li> <li>i = init segment</li> <li>c = caption or subtitle</li> <li>tt = ISOBMFF timed text track</li> <li>k = cryptographic key, license or certificate.</li> <li>o = other</li> </ul> <p>It is assumed that the server adding this key knows the object role. If not, then this key MUST NOT be used.</p>	6

## Web Application Video Ecosystem – Common Media Server Data

Description	Key Name	Header Name	Type & Unit	Value definition	Use-case reference
Response delay	rd	CMSD-Dynamic	Integer milliseconds	The time elapsed between the receipt of the request and when the first byte of the body becomes available to send to the client. The intention is for receivers to use this value to more accurately calculate the throughput of the connection [MHV22].	3
Round Trip Time	rtt	CMSD-Dynamic	Integer milliseconds	Estimated round trip time between client and server. This estimate may be derived from the transport handshake. For subsequent requests over the same connection, the value can be refined to be an exponentially weighted moving average of prior instantaneous values. An informative example algorithm for this averaging is provided by [18].	3
Startup	su	CMSD-Static	Boolean	If used, Key MUST be included without a value if the object is needed for startup of the stream. This key MUST NOT be sent if it is FALSE. The threshold of startup is left to the determination of the origin. It should approximate the starting buffer of the intended players.	8
Stream type	st	CMSD-Static	Token - one of [v,l]	v = all segments are available – e.g., VoD. l = segments become available over time – e.g., live. This state information SHOULD be trusted for no longer than the cache time of the object.	6
Streaming format	sf	CMSD-Static	Inner list of tokens - ([d h s o])	The streaming format that defines the current response.  d = MPEG DASH h = HTTP Live Streaming (HLS) s = Smooth Streaming o = other If the streaming format being returned is unknown, then this key MUST NOT be used. If the object is serving multiple streaming formats (such as a CMAF container for HLS and DASH), then the inner-list SHOULD contain both target formats – e.g., (d h).	6

## Web Application Video Ecosystem – Common Media Server Data

Description	Key Name	Header Name	Type & Unit	Value definition	Use-case reference
Version	v	CMSD-Static	Integer	The version of this specification used for interpreting the defined key names and values. If this key is omitted, any recipients MUST interpret the values as being defined by version 1. A server SHOULD omit this field if the version is 1.	9

## 8 Client Processing Requirements

---

There are two types of clients for CMSD: intermediaries and players. Some processing requirements are common to all clients, while others are particular to either intermediaries or players.

### 8.1 Common Client Processing Requirements

1. The client MUST only process these requirements when data is received via a valid CMSD header. A valid CMSD header is one which is formatted according to the rules defined in sections [3 – Data Transmission Mode] and [4 – Payload Definition].
2. A client, upon receiving common media server data, MUST interpret the keys according to their definition in this document.
3. Unknown keys, which the client does not understand, MUST be ignored. The presence of unknown keys does not make a CMSD header invalid.
4. Values that do not meet the structured data definition (such as an invalid token, or a string when an integer is expected) MUST be ignored.
5. Since there is no guarantee that keys are included, a client MUST be robust against the absence of individual keys on any given response.
6. The client MUST be able to correctly process the key-value pairs irrespective of the order in which they are defined.
7. Clients SHOULD be aware that malicious origins or proxies may send false key data with the objective of either attacking the client or gaining an unfair delivery advantage. The client SHOULD validate incoming key data before any performance impacting behaviors are executed.
8. Clients MUST ignore the entire data set if the signaled version is greater than they understand, as they cannot know which fields have been modified or deprecated. Clients SHOULD log the version incompatibility so that there is a record of why the data is not getting processed.

### 8.2 Additional Intermediary Processing Requirements

1. The intermediary, upon receiving the nor 'next object request' or nrr 'next range request' attributes, MAY optionally decide not to implement any pre-fetch action against that data.
2. Intermediaries SHOULD NOT include the CMSD-Dynamic header in the Vary header (Section 12.5.5 of [RFC9110]).
3. Edge servers SHOULD provide the necessary CORS responses to allow browser-based clients to receive the custom headers, specifically by setting the Access-Control-Expose-Headers value to include the names 'CMSD-Dynamic' and 'CMSD-Static'.
4. As specified in Section 5.3 of [RFC9110] an intermediary MUST NOT change the order of the CMSD header field values when forwarding a message.

### 8.3 Additional Player Processing Requirements

1. The player, upon receiving the maximum suggested bitrate (mb) attribute SHOULD evaluate its available bitrate tiers and switch to a combination of video and audio tiers such that the aggregate bitrate signaled via the playlist or manifest is less than or equal to the signaled maximum suggested bitrate.

## 9 Security and Privacy

---

### 9.1 Threat Environment

The data transmitted as defined by this specification is passed between a server and a client as described in Section 3, Data Transmission Mode. Data is exchanged over HTTP/HTTPS during the regular process of a server responding to a client request. The HTTP Response Headers utilized for data transmission are mature and established technologies for data transmission over the web. Transport Layer Security (TLS) can provide confidentiality and integrity for data during transmission.

### 9.2 Threats to the Intermediary and Media Clients

A malicious origin may inject false data. This tactic may be part of replay, message insertion, or modification attacks. If the server-client communication is delivered over HTTP, then man-in-the-middle attacks are feasible. Use of HTTPS for communications mitigates these attacks. All client behaviors are optional, which aids in client-based protection strategies. The client is not required to take any action upon receiving CMSD data. Some limited denial-of-service amplification opportunity exists for malicious origins utilizing the next-object-request key. Requiring the next object to be a relative path to the current response, along with the prefetch operation being optional for the edge server, helps mitigate this amplification.

When a client utilizes a connection that may be shared, such as when connections are coalesced to the same origin which is possible with HTTP/2 and HTTP/3, certain metrics

## Web Application Video Ecosystem – Common Media Server Data

contained in the CMSD-Dynamic header could be considered as a side channel if they pertain to their common transport connection.

### 9.3 Specific Mitigations

As discussed above, this specification does not expose any security issues that are not already exposed to an intermediary receiving a response from an origin or a client receiving a response from an edge server. A number of steps have been taken to mitigate specific security and privacy concerns:

- The 'nor' key values must be relative paths to the current request. This makes it harder to inject false requests to arbitrary objects. While multiple items may be added to the 'nor' field, leading to an amplification attack, these items would need to be inserted by an origin itself and they could only point back at that same origin, thus removing the opportunity to target third parties.
- All requests to intermediaries and clients are optionally executed by those entities, meaning that an entity can ignore them for security concerns (such as a rate-based threshold being exceeded) and still be compliant with the specification.
- Personally Identifiable Information fields, such as IP address, cookie information and location data, are intentionally not carried by the specification. Any use of custom keys defined outside this specification must take care to avoid transmitting values containing personally identifiable information.
- To address potential fingerprinting and side channel issues, the 'etp' and 'rtt' metrics may apply an appropriate level of quantization and/or noise to the values to a level that provides privacy whilst still allowing for their utility.

## 10 References

---

- [RFC2119] IETF RFC 2119, *Key words for use in RFCs to Indicate Requirement Levels*, <https://tools.ietf.org/html/rfc2119>.
- [RFC8174] IETF RFC 8174, *Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words*, <https://tools.ietf.org/html/rfc8174>.
- [WHATWG] WHATWG *URL Living Standard*. Accessed 7 October 2022. <https://url.spec.whatwg.org/#application/x-www-form-urlencoded>.
- [RFC8941] IETF RFC 8941, *Structured Field Values for HTTP*, <https://tools.ietf.org/html/rfc8941>.



## Web Application Video Ecosystem – Common Media Server Data

- [RFC9209] IETF RFC 9209, *The Proxy-Status HTTP Response Header Field*, <https://tools.ietf.org/html/rfc9209>.
- [RFC9211] IETF RFC 9211, *The Cache-Status HTTP Response Header Field*, <https://tools.ietf.org/html/rfc9211>.
- [HLSbis] IETF Draft, *HTTP Live Streaming 2nd Edition*. Accessed on 7 October 2022. <https://datatracker.ietf.org/doc/draft-pantos-hls-rfc8216bis/>.
- [RFC9110] IETF RFC 9110, *HTTP Semantics*, <https://tools.ietf.org/html/rfc9110>.
- [MHV22] May Lim, Mehmet N. Akcay, Abdelhak Bentaleb, Ali C. Begen and Roger Zimmermann, "The benefits of server hinting when DASHing or HLSing," in Proc. ACM Mile-High Video (MHV), Denver, CO, Mar. 2022 ([DOI: 10.1145/3510450.3517317](https://doi.org/10.1145/3510450.3517317))

## Annex A. CMSD Header Examples (Informative)


1. Origin-only VOD DASH manifest
  - a. CMSD-Static:ot=m,sf=d,st=v,su,n="OriginProviderA"
2. Origin-only VOD DASH manifest, with prefetching of all init segments and range limited media segments
  - a. CMSD-Static:ot=m,sf=d,st=v,su,n="OriginProviderA",nor="/1080/init.mp4|/720/init.mp4|/1080/track.mp4|/720/track.mp4|/480/init.mp4",nrr="|1048-3658|898-1974|"
3. Origin-only prefetching of segments whose path and filenames contain the vertical line delimiter
  - a. CMSD-Static:nor="/video-segments/a%7Cb/a%7Cb-365.mp4|/video-segments/a%7Cb/a%7Cb-366.mp4"
4. Origin-only prefetching of a segment whose relative path lies in a parent folder
  - a. CMSD-Static:nor=".././1080/segment34.mp4"
5. Origin-only VOD HLS primary playlist, with prefetching for all variant playlists, init segments and first segments
  - a. CMSD-Static:ot=m,sf=h,st=v,su,n="OriginProviderA",nor="/video/1080/1080p-playlist.m3u8|/video/720/720p-playlist.m3u8|/video/480/480p-playlist.m3u8|/audio/audio-playlist.m3u8|/video/1080/init.mp4|/video/720/init.mp4|/video/480/init.mp4|/audio/init.mp4|/video/1080/seg1.mp4|/video/720/seg1.mp4|/video/480/seg1.mp4"
6. Origin-only HLS live video segment
  - a. CMSD-Static:ot=v,sf=h,st=l,d=2002,br=5400,tl=12000,n="OriginProviderA"
7. HLS VOD segment moving from an origin through two different CDNs to the player
  - a. CMSD-Static:ot=v,sf=h,st=v,d=6006,br=1450,n="OriginProviderA"
  - b. CMSD-Dynamic: "CDNB-3ak1";etp=115000;rtt=8
  - c. CMSD-Dynamic: "CDNB-w35k";etp=93000;rtt=32
  - d. CMSD-Dynamic: "CDNA-987.343";etp=140000;rtt=30
  - e. CMSD-Dynamic: "CDNA-312.663";etp=32000;rtt=56;
8. Origin-only chunk-transferred LL-DASH live video segment
  - a. CMSD-Static:ot=v,sf=d,st=l,d=4004,br=2500,tl=3000,at=1656703938470,n="OriginProviderA"
9. VOD HLS audio segment with prefetch of next segment and dynamic edge-applied data

## Web Application Video Ecosystem – Common Media Server Data

- a. CMSD-Static:ot=a,sf=h,st=v,d=6006,br=128,n="OriginProviderA",nor="segment-3256.m4a"
  - b. CMSD-Dynamic: "CDNA-312.663";etp=12000;rtt=28;
10. Edge server signaling that is under duress when delivering an HLS primary playlist
- a. CMSD-Static:ot=m,sf=h,su,st=v,n="OriginProviderA",nor="/video/1080/1080p-playlist.m3u8|/video/720/720p-playlist.m3u8"
  - b. CMSD-Dynamic: "CDNA-312.663";etp=12;rtt=28;du
11. Origin-only VOD HLS playlist with custom key
- a. CMSD-Static:ot=m,sf=h,st=v,su,n="OriginProviderA",com.example-city="245"

### Annex B. Informative Use-Case Definitions (Informative)

1. Identifying intermediaries in the media distribution chain, for the purposes of investigating delivery and resolving availability issues.
2. Provide an estimate of the throughput available between an edge server and a player, such that the player can use that information to enrich its decision about which bitrate to select, especially the initial bitrate decision made at the start of playback.
3. Provide link and server characteristics to a client for the purposes of the client making load balancing decisions between multiple potential servers.
4. Provide a means for a CDN to instruct all connected players to limit their upper bitrate, in response to an ISP request to reduce congestion on the last-mile network.
5. Provide information which an intermediary could use to prefetch the next item in a sequence of media objects. Prefetching increases performance by moving the object closer to the player ahead of the player's request for that object.
6. Provide information about the media characteristics of a binary object, for forensic, logging, or delivery optimization purposes.
7. Allow an intermediary to prefetch init segments ahead of a player request.
8. Identify media segments located at the start of playback when delivery performance is most critical due to the player's need to establish a buffer. By identifying these segments, intermediaries can take special measures to optimize their delivery performance.
9. Allow for an extensibility mechanism for CMSD so that future upgrades can be made and clients and servers can both agree on the generation and interpretation of the keys and values.
10. Allow a CDN to signal to a player a suggested playback bitrate in order to optimize collective QoE.
11. Allow a CDN edge server to prefetch a range of a track file that will be guaranteed to include the next media segment range requested by the client.
12. Provide information about the timing of media segments to assist players with low-latency live streaming.



### **Consumer Technology Association Document Improvement Proposal**

If in the review or use of this document a potential change is made evident for safety, health or technical reasons, please email your reason/rationale for the recommended change to [standards@CTA.tech](mailto:standards@CTA.tech).

Consumer Technology Association  
Technology & Standards Department  
1919 S Eads Street, Arlington, VA 22202  
FAX: (703) 907-7693 [standards@CTA.tech](mailto:standards@CTA.tech)

**Consumer  
Technology  
Association™**