Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 18 September 2026

M. Dierken
paywalls.net

17 March 2026

HTTP Usage Reporting for Cached Resources
draft-dierken-usage-log-http-00

Abstract

This document defines a mechanism by which an HTTP origin can advertise an endpoint for downstream usage reporting and by which an authenticated client operator can report usage of cached representations that were used without issuing additional requests to the origin.

The mechanism complements access-time response metadata by allowing a client operator to report downstream usage associated with a prior origin-served response context. The protocol defines discovery of a usage reporting endpoint, submission semantics, and a minimal record format suitable for reconciliation and billing.

This specification does not define payment, settlement, or enforcement mechanisms. It defines a good-faith accounting protocol for reporting downstream usage.

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 18 September 2026.

Copyright Notice

Table of Contents

1. 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).

2. Introduction

Some HTTP resources may be served from cache or otherwise reused without an additional request to the origin. In such cases, the origin may not directly observe every downstream use of the representation.

This document defines a mechanism for reporting such usage back to the origin.

The goals of this specification are to:

* allow an origin to advertise a usage reporting endpoint,

* allow authenticated client operators to report downstream usage,

* support both per-event and aggregated usage reporting, and

* provide sufficient structure for reconciliation and billing.

This document does not define commercial models, payment systems, or enforcement mechanisms for dishonest reporting.

The trust model assumed by this protocol is that authenticated client operators submit usage reports in good faith and that the origin performs reconciliation and optional deduplication.

3. Terminology

Origin The server that served the response or its delegated reporting service.

Client Operator The authenticated entity responsible for requesting resources and reporting downstream usage.

Cached Representation A representation used without issuing a new request to the origin.

Usage Report A batch submission containing one or more usage records.

Usage Record A single assertion of downstream usage.

Response Identifier An identifier associated with a prior origin-served response context.

4. Usage Key

The accounting identity for reported usage in this specification is the tuple:

(resource, response_id)

where resource identifies the URI of the representation used and response_id identifies the origin-served response context associated with that representation.

The response_id ties the reported usage to the representation version, access event, or other origin-defined response context previously recorded by the origin.

All usage records reported by this protocol refer to exactly one usage key.

A response_id value is assigned by the origin or its delegated service. Clients MUST NOT invent substitute values when none was provided.

5. Protocol Overview

The mechanism consists of two parts:

1. discovery of a usage reporting endpoint, and

2. submission of usage reports.

A typical interaction proceeds as follows:

1. The origin serves a response and advertises a usage reporting endpoint.

2. The client operator later reuses the cached representation without contacting the origin.

3. The client operator periodically submits usage reports describing downstream usage.

Usage reports are submitted using HTTP POST and may contain either individual usage events or aggregated usage records.

6. Discovery

An origin MAY advertise a usage reporting endpoint using the HTTP Link response header with relation type usage-log.

Example:

```
Link: <https://example.com/usage-log>; rel="usage-log"
```

The target URI identifies the endpoint that accepts usage reports.

An origin MAY include this header on any response relevant to future cached use.

6.1. Authoritative Discovery

Resource-associated discovery via the Link header is authoritative.

If multiple discovery mechanisms exist, the usage-log link associated with the resource response takes precedence.

6.2. Optional Default Discovery

Deployments MAY expose a default usage reporting endpoint via a well-known resource or equivalent mechanism.

Such mechanisms are secondary to resource-associated discovery and are outside the scope of this specification.

7. Submission

Usage reports are submitted via HTTP POST to the advertised usage reporting endpoint.

Clients MUST authenticate using an origin-supported HTTP authentication mechanism.

Example:

```
POST /usage-log HTTP/1.1
Host: example.com
Authorization: Bearer agt_XYZ
Content-Type: application/usage-report+jsonl
```

The request body contains one or more usage records encoded as JSON Lines.

Each line is an independent JSON object.

8. Media Type

This specification defines the media type application/usage-report+jsonl.

Encoding is UTF-8 line-delimited JSON objects. Each line represents one usage record.

9. Usage Record Format

Two reporting forms are defined.

9.1. Event Form

In event form, each record represents one usage event.

Fields:

* resource -- string; URI of the resource used
* response_id -- string; identifier of the associated origin-served response context
* used_at -- timestamp in RFC 3339 format

Example:

{"resource":"https://example.com/news/123","response_id":"resp_82fd","used_at":"2026-03-06T18:05:00Z"}

9.2. Aggregate Form

In aggregate form, a record summarizes multiple usage events for the same usage key over a reporting window.

Fields:

* resource -- string
* response_id -- string
* window_start -- timestamp in RFC 3339 format
* window_end -- timestamp in RFC 3339 format
* count -- integer

Example:

{"resource":"https://example.com/news/123","response_id":"resp_82fd","window_start":"2026-03-06T00:00:00Z","window_end":"2026-03-07T00:00:00Z","count":148}

Aggregation of identical usage keys is optional but RECOMMENDED.

10. Client Requirements

A client operator submitting usage reports:

* MUST authenticate to the usage reporting endpoint,
* MUST ensure that submitted records represent good-faith assertions of usage,
* SHOULD batch records when practical,
* SHOULD aggregate usage where appropriate, and
* SHOULD retain internal records sufficient for reconciliation.

Reports MAY be submitted asynchronously and out of chronological order.

11. Origin Requirements

An origin receiving usage reports:

* MUST authenticate the sender,
* MUST treat accepted reports as operator assertions of usage,
* SHOULD tolerate accidental duplicate submission,
* MAY deduplicate records using deployment-specific methods, and
* MAY reject malformed reports.

This specification does not require exactly-once delivery semantics.

12. Response Semantics

Typical responses include:

202 report accepted for processing

200 report accepted and processed

400 malformed report

401 or 403 authentication or authorization failure

413 report batch too large

415 unsupported media type

429 sender should slow down

5xx server error; client MAY retry

202 Accepted is RECOMMENDED for asynchronous ingestion systems.

13. Relationship to Access-Time Response Metadata

This specification complements, but does not replace, access-time response metadata.

Access-time metadata describes the terms or context under which a response was served by the origin.

Usage reporting describes downstream usage that occurred after access without a new origin request.

How reported usage maps to pricing, licensing, or billing terms is outside the scope of this specification.

14. Duplicate Submission

Duplicate submission is considered an operational concern rather than a core protocol threat.

Usage reports create obligations for the reporting operator; therefore strict anti-replay protocols are not required.

Origins SHOULD implement reconciliation or duplicate suppression sufficient for normal operational reliability.

15. Security Considerations

Usage reports may create accounting or billing consequences. Origins MUST authenticate reporting clients.

Implementations SHOULD use HTTPS transport.

Implementations MAY apply additional integrity protections such as HTTP Message Signatures if required for dispute resolution.

16. Privacy Considerations

Usage reports may expose resource access patterns.

Deployments SHOULD avoid including user-level identifiers unless strictly necessary.

Usage reporting SHOULD identify resources and operator-level usage without exposing individual user data.

17. Link Relation Registration

Relation Name: usage-log

Description: identifies an endpoint that accepts usage reports for downstream usage of origin resources or cached representations.

Reference: this document.

18. IANA Considerations

This document requests registration of:

* the usage-log link relation

* the media type application/usage-report+jsonl

* Response-Id

19. Example

Origin response:

```
HTTP/1.1 200 OK
Response-Id: resp_82fd
Link: <https://example.com/usage-log>; rel="usage-log"
Cache-Control: max-age=86400
```

Later, the client submits a daily aggregate usage report:

```
POST /usage-log HTTP/1.1
Host: example.com
Authorization: Bearer agt_XYZ
Content-Type: application/usage-report+jsonl

{"resource":"https://example.com/news/123","response_id":"resp_82fd","window_start":"2
026-03-06T00:00:00Z","window_end":"2026-03-07T00:00:00Z","count":148}
```

The origin acknowledges receipt:

```
HTTP/1.1 202 Accepted
```