

Internet-Draft
Intended status: Informational
Expires: 3 June 2026

Y. Motiwala
mesibo
December 2025

Consolidated Parameters Part for Multipart/Form-Data
draft-motiwala-consolidated-multipart-params-00

Abstract

This document describes a deployment pattern for reducing overhead in multipart/form-data requests that include both file uploads and numerous text parameters. The approach consolidates multiple text fields into one or more application/x-www-form-urlencoded parts inside the multipart body.

This technique requires no changes to existing HTTP, MIME, or form specifications and is fully compatible with RFC 7578. It uses only existing, standardized Content-Type values in their already-defined manner, making it immediately deployable with current infrastructure.

The pattern guarantees backward compatibility: servers that don't recognize it still receive all parameters as standard form data. RFC 7578 defines multipart/form-data but does not specify optimization strategies; this document fills that gap.

This document records the pattern, processing rules, and deployment considerations, and reports implementation experience from production systems processing millions of uploads daily.

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 3 June 2026.

Copyright Notice

Copyright (c) 2025 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.

Table of Contents

1. Introduction	3
1.1. Requirements Language	4

1.2.	Terminology	4
1.3.	Relationship to Existing Standards	5
2.	Motivation	6
2.1.	Current Inefficiency	6
2.2.	Applicability	7
2.3.	Standards Gap	8
3.	Specification	9
3.1.	Format Definition	9
3.2.	Standards Compliance	10
3.3.	Requirements	11
3.3.1.	Client Requirements	11
3.3.2.	Server Requirements	12
3.4.	Field Name Conflicts	12
4.	Examples	13
4.1.	Standard Multipart (Before)	13
4.2.	Consolidated Parameters (After)	14
4.3.	Multiple Files with Parameters	15
5.	Implementation Experience	16
5.1.	Efficiency Measurements	16
5.2.	Compatibility Validation	16
5.3.	Backward Compatibility	17
6.	Security Considerations	17
7.	IANA Considerations	18
8.	References	18
8.1.	Normative References	18
	Acknowledgments	19
	Author's Address	20

1. Introduction

HTML forms with file upload capabilities must use the multipart/form-data encoding type [RFC7578]. This encoding wraps each form field, including simple text fields, in a MIME part with headers and boundary delimiters. For forms containing numerous text parameters alongside file uploads, this creates significant overhead.

Consider a file upload form with 15 text fields (user_id, session_token, file_type, permissions, metadata fields, etc.) and one file. Using standard multipart/form-data, each text field requires approximately 100 bytes of overhead to transmit perhaps 20 bytes of actual data, representing a 4-6x inflation factor for parameter transmission.

RFC 7578 [RFC7578] defines multipart/form-data but does not specify how implementations should balance efficiency with compatibility. This has led to divergent practices: browsers generate one part per field, some APIs embed JSON metadata, and custom implementations use various ad-hoc approaches.

Browsers generate one MIME part per form field because the HTML form submission algorithm requires each form control to become one entry in the form data set, and each entry becomes one multipart part. This behavior originates from the HTML Standard, not from multipart/form-data itself. Nothing in HTTP, MIME multipart semantics, or RFC 7578 prevents other user agents from using more efficient encoding strategies that result in fewer multipart parts and reduced overhead.

This document describes Consolidated Parameters, a general-purpose multipart/form-data encoding pattern in which multiple text parameters are combined into one or more application/x-www-form-urlencoded parts inside the multipart body. The technique requires no extensions to HTTP, no additions to MIME, and no deviations from RFC 7578. It is fully standards-compliant and may be used by any HTTP user agent, reducing parameter transmission overhead while remaining fully backward compatible with existing parsers.

The pattern applies to all HTTP clients, and any client may adopt it while remaining fully compliant with existing standards.

The pattern guarantees backward compatibility because servers that do not recognize it still receive all parameters as standard form-data fields, which can be manually parsed using existing form handling logic. This enables safe deployment without requiring coordination between clients and servers.

This approach has been deployed in production systems processing millions of uploads daily, demonstrating its effectiveness and compatibility in real-world use.

1.1. Requirements Language

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 uppercase, as shown here.

As an Informational specification, the use of normative language describes implementation recommendations for interoperability rather than mandatory protocol requirements.

1.2. Terminology

This document uses the following terms:

Standard Multipart: A multipart/form-data request where each form field is encoded as a separate MIME part, as traditionally generated by web browsers.

Consolidated Parameters: An approach where multiple text parameters are encoded together in a single multipart part using application/x-www-form-urlencoded format. This term is used purely as a descriptive label and has no normative meaning.

Parameters Part: A MIME part within multipart/form-data that has Content-Type: application/x-www-form-urlencoded and contains multiple name=value pairs in URL-encoded format.

1.3. Relationship to Existing Standards

This document does not modify, extend, or add anything to existing standards. It describes an application-level pattern using only existing, standardized Content-Type values:

- o multipart/form-data as defined in [RFC7578]
- o application/x-www-form-urlencoded as defined in [HTML5]

Both specifications already permit arbitrary Content-Type values in multipart parts. This document simply describes consolidating parameters into a single part rather than using separate parts for each parameter - a choice already available within current specifications.

The inefficiency of one-part-per-field encoding arises from the HTML form submission algorithm, which currently requires browsers to generate one multipart part per form control. However, this is a requirement of the HTML Standard, not of HTTP, MIME multipart semantics, or RFC 7578. Nothing in these protocol specifications mandates one MIME part per field. As a result, HTTP clients have full flexibility to use more efficient encodings while remaining fully compatible with existing parsers. Future browser implementations may adopt this optimization

if the HTML Standard is updated to permit more efficient multipart encoding strategies.

Since this pattern uses only existing, standardized content types within their defined semantics, all standard HTTP and multipart processing rules continue to apply. Security validation, duplicate field handling, streaming uploads, chunked transfer encoding, and other standard behaviors remain unchanged. No special-case handling is required beyond what is already defined in the referenced specifications.

2. Motivation

2.1. Current Inefficiency

The multipart/form-data format transmits each form field as an individual MIME part with its own Content-Disposition header, optional Content-Type, and boundary delimiters. For small forms this overhead is negligible, but it becomes significant when many short text parameters accompany one or more files.

For example, transmitting the value "12345" for a user_id field requires:

```
-----WebKitFormBoundary7MA4YWxkTrZu0gW
Content-Disposition: form-data; name="user_id"

12345
```

This represents approximately 105 bytes total to transmit 12 bytes of data (field name "user_id" plus value "12345").

Each text field adds significant structural overhead. When many short text parameters accompany files, cumulative overhead becomes disproportionately large relative to the actual content. This increases bandwidth and CPU cost when forms include many parameters and upload volume is high.

Per-field overhead in a typical multipart form:

- o Boundary delimiter: ~40 bytes (--boundary + CRLF)
- o Content-Disposition header: ~45-60 bytes
- o Empty line separator: 2 bytes (CRLF)
- o Trailing CRLF: 2 bytes
- o Total: ~90-105 bytes per field
- o Overhead-to-payload ratio: approximately 9:1

(Note: These calculations assume HTTP transport where binary data is transmitted directly. Email contexts requiring base64 encoding would have additional overhead not shown here.)

For a form with 15 such fields (each ~20 bytes of data):

- o Standard multipart: ~1575 bytes overhead + 300 bytes data = 1875
- o URL-encoded consolidated: ~105 bytes overhead + 300 bytes = 405
- o Savings: 1470 bytes (78% reduction in total request size)

The problem is further amplified in deployments that process large volumes of uploads. Across millions of requests per day, excessive multipart overhead directly increases operational bandwidth, and processing costs. A social platform processing 100 million photo uploads per day with 15 metadata fields per upload saves multiple terabytes (TB) monthly in bandwidth costs alone, not including reduced latency benefits for users on mobile networks.

2.2. Applicability

This approach applies universally to any HTTP-based file upload that

includes multiple text parameters. Typical examples include:

- o User-generated content (photos, videos, documents) with metadata (user ID, timestamps, privacy settings, descriptions)
- o Application data synchronization with context (device ID, sync tokens, version information)
- o API requests combining binary payloads with authentication, configuration, or processing parameters

The benefit scales with the number of text parameters. Forms with 5 or more parameters show measurable improvement; forms with 10+ parameters achieve significant bandwidth reduction.

Virtually all modern applications perform file uploads with metadata, making this approach widely applicable across web, mobile, IoT, and API contexts. The specific application domain (social media, e-commerce, healthcare, etc.) does not affect applicability or benefit.

2.3. Standards Gap

RFC 7578 [RFC7578] defines multipart/form-data but does not specify efficient encoding strategies for forms with numerous text parameters. The specification describes the traditional one-part-per-field approach used by browsers but remains silent on efficiency considerations.

This omission has resulted in implementation fragmentation. Browsers follow the HTML form submission algorithm (one part per field), while API implementations have adopted various approaches: some embed all parameters in JSON objects, others use custom encodings, and many simply accept the inefficiency of standard multipart encoding.

This document addresses this gap by providing a standardized efficient encoding pattern that:

- o Uses only standardized Content-Type values (multipart/form-data and application/x-www-form-urlencoded)
- o Maintains full backward compatibility with RFC 7578 parsers
- o Provides measurable efficiency improvements (Section 2.1)
- o Has been validated in production deployments (Section 5)

3. Specification

3.1. Format Definition

An optimized multipart/form-data request MAY include one or more parts with Content-Type: application/x-www-form-urlencoded. Such parts consolidate multiple text parameters using the URL-encoded format defined in [HTML5] and [RFC3986].

Example structure:

```
POST /upload HTTP/1.1
Host: api.example.com
Content-Type: multipart/form-data; boundary=----Boundary123

-----Boundary123
Content-Disposition: form-data; name="params"
Content-Type: application/x-www-form-urlencoded

user_id=12345&session=abc&type=document&permissions=read
-----Boundary123
Content-Disposition: form-data; name="file"; filename="doc.pdf"
Content-Type: application/pdf
```

[binary PDF data]
-----Boundary123--

The part name (e.g., "params") is arbitrary and has no special meaning. Servers MUST identify the parameters part by its Content-Type header, not by its name.

URL-encoded content MUST follow the application/x-www-form-urlencoded specification:

- o Field names and values are URL-encoded per [RFC3986]
- o Name=value pairs are separated by "&" (ampersand)

The parameters part MAY appear at any position within the multipart sequence. However, placing it first is RECOMMENDED for efficient server processing, as authentication tokens and request metadata can be validated before processing file data. This is particularly beneficial for streaming uploads, where early validation can reject unauthorized requests before receiving large file payloads.

3.2. Standards Compliance

This approach operates entirely within existing specifications:

RFC 7578 [RFC7578] (multipart/form-data):

Defines multipart/form-data and specifies that each part has an optional Content-Type (Section 4.4). It does not prescribe which Content-Type values are permitted, nor does it mandate that each form field must be a separate part.

HTML5 [HTML5] (application/x-www-form-urlencoded):

Defines the URL-encoded format for form data. This document uses this format within a multipart part rather than as the entire request body.

RFC 2046 [RFC2046] (MIME multipart):

The foundation for multipart formats. Section 5.1 establishes that parts can have any Content-Type.

3.3. Requirements

3.3.1. Client Requirements

Clients MUST properly URL-encode all parameter names and values in consolidated parameters according to [RFC3986].

Clients may send consolidated parameters in all requests or only when servers support them. Both approaches provide the same compatibility guarantees:

- Always: Send consolidated parameters in all requests. Even servers that don't recognize consolidated parameters receive all data as a single URL-encoded field, which can be manually parsed using existing form parsing logic, ensuring backward compatibility.
- Conditionally: Detect server support through implementation-specific means (API documentation, configuration, version negotiation) and send consolidated parameters only when supported.

Clients MAY omit the consolidated parameters part entirely if there are no text fields to send. A multipart request containing only file parts is valid and requires no special handling.

3.3.2. Server Requirements

Servers processing multipart requests with consolidated parameters:

1. MUST parse the multipart/form-data structure according to [RFC7578] without modification.
2. When encountering a part with Content-Type: application/x-www-form-urlencoded, MUST parse that part's body using URL-decoding rules.
3. MUST extract name=value pairs from the URL-encoded part and make them available to the application as form parameters.
4. MUST process other parts (files, other content types) using standard multipart parsing.

Example server-side processing logic:

```
for (part in multipartRequest.parts) {
    if (part.contentType == "application/x-www-form-urlencoded") {
        params.merge(urlDecode(part.body));
    } else {
        // Existing processing for files and other parts
        processStandardPart(part);
    }
}
```

Graceful Degradation:

Servers that do not recognize this optimization will treat the URL-encoded part as a single text field, which provides graceful degradation. The URL-encoded content is preserved as the field value and can be manually parsed if needed.

3.4. Field Name Conflicts

The handling of duplicate field names in multipart/form-data is already defined by RFC 7578 and existing server implementations. This pattern does not change how duplicate field names are handled; servers apply their existing precedence rules.

Servers typically handle duplicates using last-wins semantics, first-wins semantics, or by collecting all values into an array. The same rules apply here. Clients SHOULD avoid sending duplicate field names across parts to prevent ambiguity.

4. Examples

4.1. Standard Multipart (Before)

Traditional browser-generated multipart for a file upload with metadata:

```
POST /api/upload HTTP/1.1
Host: api.example.com
Content-Type: multipart/form-data; boundary=----WebKitBoundary

-----WebKitBoundary
Content-Disposition: form-data; name="user_id"

12345
-----WebKitBoundary
Content-Disposition: form-data; name="session_token"

abc123xyz789
-----WebKitBoundary
Content-Disposition: form-data; name="file_type"

document
```

```
-----WebKitBoundary
Content-Disposition: form-data; name="permissions"

read
-----WebKitBoundary
Content-Disposition: form-data; name="file"; filename="report.pdf"
Content-Type: application/pdf

[binary PDF data - 50KB]
-----WebKitBoundary--

Size analysis:
o Parameter overhead: ~400 bytes
o Parameter data: ~50 bytes
o File overhead: ~120 bytes
o File data: ~50,000 bytes
o Total: ~50,570 bytes
```

4.2. Consolidated Parameters (After)

The same upload using consolidated parameters:

```
POST /api/upload HTTP/1.1
Host: api.example.com
Content-Type: multipart/form-data; boundary=----WebKitBoundary

-----WebKitBoundary
Content-Disposition: form-data; name="params"
Content-Type: application/x-www-form-urlencoded

user_id=12345&session_token=abc123xyz789&file_type=document&permissions=read
-----WebKitBoundary
Content-Disposition: form-data; name="file"; filename="report.pdf"
Content-Type: application/pdf

[binary PDF data - 50KB]
-----WebKitBoundary--

Size analysis:
o Parameter overhead: ~120 bytes (one part)
o Parameter data: ~85 bytes (includes "&" separators)
o File overhead: ~120 bytes
o File data: ~50,000 bytes
o Total: ~50,325 bytes
```

Savings: 245 bytes (0.48% of total, but 61% of parameter overhead)

The total savings are small when files are large, but the reduction in parameter overhead is substantial. For forms with more parameters or smaller files, the percentage savings increases substantially. For a 5KB file with 20 parameters, savings can exceed 15% of total request size.

4.3. Multiple Files with Parameters

Uploading multiple files with shared metadata:

```
POST /api/batch-upload HTTP/1.1
Host: api.example.com
Content-Type: multipart/form-data; boundary=----Boundary

-----Boundary
Content-Disposition: form-data; name="params"
Content-Type: application/x-www-form-urlencoded

user_id=12345&batch_id=B789&timestamp=1234567890&location=office
```



```
-----Boundary
Content-Disposition: form-data; name="file1"; filename="scan1.jpg"
Content-Type: image/jpeg
```

```
[JPEG data]
-----Boundary
Content-Disposition: form-data; name="file2"; filename="scan2.jpg"
Content-Type: image/jpeg
```

```
[JPEG data]
-----Boundary
Content-Disposition: form-data; name="file3"; filename="scan3.jpg"
Content-Type: image/jpeg
```

```
[JPEG data]
-----Boundary--
```

This pattern is particularly efficient when the same metadata applies to multiple files, avoiding repetition across parts.

5. Implementation Experience

5.1. Efficiency Measurements

Production deployment shows this pattern reduces parameter transmission overhead by 45-60%. For forms with 10+ text parameters, total request size decreases by 8-15%.

No performance degradation was observed in server parsing or client encoding. At scale, the approach has saved terabytes (TB) of bandwidth monthly across millions of daily uploads in chat and messaging applications.

5.2. Compatibility Validation

Testing confirms standard multipart parsers correctly handle consolidated parameter parts as single text fields containing URL-encoded data, validating the graceful degradation property.

5.3. Backward Compatibility

Production testing confirms this consolidation does not introduce breakage: all tested multipart parsers correctly process consolidated parameter parts as single text fields containing URL-encoded data.

Since servers already handle URL-encoded parameters (from standard HTML forms or application/x-www-form-urlencoded requests), they can extract individual parameters from the consolidated field using existing parsing logic. No server modifications are required to maintain functionality; only optimized servers need additional logic to automatically extract parameters.

6. Security Considerations

Consolidated parameters do not introduce new security risks beyond those already present in the underlying formats. Parameter-handling risks remain identical to those of standard application/x-www-form-urlencoded encoding, and the multipart encapsulation risks remain identical to those of standard multipart/form-data encoding.

This pattern combines existing, well-understood formats. Implementers should consult the security considerations in the specifications for these formats:

- o RFC 7578 Section 8 - Security Considerations for

- multipart/form-data
- o HTML5 Section 4.10.22.6 - Security Considerations for application/x-www-form-urlencoded
- o RFC 3986 Section 7 - Security Considerations for URL encoding

Servers should apply the same validation, size limits, and sanitization to consolidated parameters as they would to standard form fields. The security model remains unchanged: parameters arrive in a different encoding but do not change the threat model or the required security controls.

7. IANA Considerations

This document has no IANA actions. It does not define new media types, require registrations, or modify existing registries. The optimization uses existing, well-defined media types:

- o multipart/form-data [RFC7578]
- o application/x-www-form-urlencoded [HTML5]

8. References

8.1. 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>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC7578] Masinter, L., "Returning Values from Forms: multipart/form-data", RFC 7578, DOI 10.17487/RFC7578, July 2015, <<https://www.rfc-editor.org/info/rfc7578>>.
- [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>>.
- [HTML5] WHATWG, "HTML Living Standard - Forms", <<https://html.spec.whatwg.org/multipage/forms.html>>.
- [RFC2046] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", RFC 2046, DOI 10.17487/RFC2046, November 1996, <<https://www.rfc-editor.org/info/rfc2046>>.

Acknowledgments

The author would like to thank the mesibo engineering team for implementing and testing this optimization in production environments.

Author's Address

Yusuf Motiwala
mesibo
San Francisco, CA
United States of America

Email: yusuf@mesibo.com

URI: <https://mesibo.com/>