

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: April 28, 2011

L. Masinter
Adobe
October 25, 2010

MIME and the Web
draft-masinter-mime-web-info-01

Abstract

This document describes some of the ways in which parts of the MIME system, originally designed for electronic mail, have been used in the Web, and some of the ways in which those uses have resulted in difficulties. Given this background and justification, this document then goes on to outline requirements for changes to MIME registries and practices for their use within W3C and IETF, in order to address those difficulties. Within IETF, a companion Best Current Practice document will be developed to specifically make some changes to the Internet Media Types and Charset registries.

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 <http://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 April 28, 2011.

Copyright Notice

Copyright (c) 2010 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. History	3
2.1. Origins of MIME	3
2.2. Introducing MIME into the Web	4
2.3. Distributed Extensibility	5
3. Problems with application to the Web	5
3.1. Lack of clarity	5
3.2. Differences between email and Web delivery	6
3.3. The Rules Weren't Quite Followed	7
3.4. Consequences	7
3.5. The Down Side of Extensibility	8
4. Additional considerations	8
4.1. There are related problems with charsets	8
4.2. Embedded, downloaded, launch independent application	9
4.3. Additional Use Cases: Polyglot and Multiview	9
4.4. Evolution, Versioning, Forking	9
4.5. Content Negotiation	10
4.6. Fragment identifiers	11
5. Recommendations	11
5.1. Internet Media Type registration	12
5.1.1. MIME registry magic numbers for sniffing	12
5.1.2. Scripting and scriptable content safety	12
5.1.3. Fragment identifiers	12
5.1.4. Application info	12
5.1.5. File extensions in registry	12
5.2. Sniffing	13
5.2.1. Sniffing uses Media Type magic number	13
5.2.2. Sniffing when there are multiple different definitions	13
5.2.3. Sniffing charsets	13
5.2.4. Sniffing security uses scriptability info	13
5.3. Changes to IANA processes for MIME registries	13
5.4. FTP specification	13
5.5. Update some URI definitions	14
5.6. Changes to W3C findings, processes	14
6. Acknowledgements	14
7. IANA Considerations	14
8. Security Considerations	14
9. Informative References	14
Author's Address	15

1. Introduction

This document was initially prompted by a set of discussions about Web architecture and the difficulties surrounding evolution of the Web, Internet Media types, multiple specifications for a single media type, and related discussions.

The document gives some of the history of MIME and its introduction and use in the web Section 2. It then describes some of the current difficulties with the use of MIME in the web context Section 3. This background and context is then followed by a description of changes which would reduce some of those difficulties; the changes involve specifications, practices, and registries within IETF and W3C Section 5. In particular, changes to the registry and maintenance procedures for MIME-related registries maintained by IANA are describes.

Currently, discussion of this document is suggested on the mailing list www-tag@w3c.org (mailing list open for subscription to all), archives at <http://lists.w3.org/Archives/Public/www-tag/>.

NOTE: This document is still quite rough; some of the facts need to be checked, many sections still need expansion. Any help with references and such appreciated.

2. History

2.1. Origins of MIME

MIME ("Multipurpose Internet Mail Extensions") was invented originally for email, based on general principles of "messaging" (a foundational architecture framework). The role of MIME was to extend Internet email messaging from ASCII-only plain text, to include other character sets, images, rich documents, etc.) [RFC1521], [RFC1522]. The basic architecture of complex content messaging is:

- o Message sent from A to B.
- o Message includes some data. Sender A includes standard 'headers' telling recipient B enough information that recipient B knows how sender A intends the message to be interpreted.
- o Recipient B gets the message, interprets the headers for the data and uses it as information on how to interpret the data.

MIME is a "tagging and bagging" specification:

tagging: How to label content so the intent of how the content should be interpreted is known.

bagging: How to wrap the content so the label is clear, or, if there are multiple parts to a single message, how to combine them.

"MIME types" (renamed "Internet Media Types" in later specs [RFC2046]) are part of the "tagging" -- a way to describe the content of a message so that it could be used to initiate interpretation of a message. The "Internet Media Type registry" (MIME type registry) is where someone can tell the world what a particular label means, as far as the sender's intent of how recipients should process a message of that type, and the description of a recipients capability and ability for senders.

2.2. Introducing MIME into the Web

The original World Wide Web (the 0.9 version of HTTP, see [RFC1945]) didn't have "tagging and bagging" -- everything sent via HTTP was assumed to be HTML. However, at the time (early 1990's) other distributed information access systems, including Gopher (distributed menu system) and WAIS (remote access to document databases) were adding capabilities for accessing many things other text and hypertext and the WWW folks were considering type tagging. It was agreed that HTTP should use MIME as the vocabulary for talking about file types and character sets. The result was that HTTP 1.0 added the "content-type" header, following (more or less) MIME. Later, for content negotiation, additional uses of this technology (in 'Accept' headers) were also added.

The differences between the use of Internet Media Types between email and HTTP have minor:

- o default charset: HTTP specified ISO-8859-1 as the default character set, not US-ASCII
- o requirement for CRLF in plain text: in practice, web clients didn't restrict content to use CRLF in text/* MIME bodies.

These minor differences have caused a lot of trouble.

2.3. Distributed Extensibility

The real advantage of using Internet Media Types to label content meant that the Web was no longer restricted to a single format. This one addition meant expanding from Global Hypertext to Global Hypermedia (as suggested in a 1992 email [connolly92])

```
+-----+
| The Internet currently serves as the backbone for a global |
| hypertext.  FTP and email provided a good start, and the gopher, |
| WWW, or WAIS clients and servers make wide area information |
| browsing simple.  These systems even interoperate, with email |
| servers talking to FTP servers, WWW clients talking to gopher |
| servers, on and on. |
| This currently works quite well for text.  But what should WWW |
| clients do as Gopher and WAIS servers begin to serve up pictures, |
| sounds, movies, spreadsheet templates, postscript files, etc.? |
| It would be a shame for each to adopt its own multimedia typing |
| system. |
| If they all adopt the MIME typing system (and as many other |
| features from MIME as are appropriate), we can step from global |
| hypertext to global hypermedia that much easier. |
+-----+
```

The fact that HTTP could reliably transport images of different formats, for example, allowed NCSA to add to HTML. MIME allowed other document formats (Word, PDF, Postscript) and other kinds of hypermedia, as well as other applications, to be part of the Web. MIME was arguably the most important extensibility mechanism in the Web.

3. Problems with application to the Web

Unfortunately, while the use of Internet Media Types for the Web added incredible power, a number of problems have arisen.

3.1. Lack of clarity

Many people are confused about the purpose of MIME in the Web, its uses, the meaning of Internet Media Types. Many W3C specifications TAG findings and Internet Media Type registrations make what are incorrect assumptions about the meaning and purposes of a Internet Media Type registration.

3.2. Differences between email and Web delivery

Some of the differences between the application contexts of email and Web delivery determine different requirements:

- o In the Web, the transfer of data is initiated differently than in email: the "messages" with labeled content are usually HTTP responses to a specific (GET) request (although the request is itself a message, GET has no content). In the most common case, then, the receiver knows more about the data before it has been sent.
- o Clients would like to know more about the content before they retrieve it. The "tagging" is often not sufficient to know, for example, "can I interpret this if I retrieve it", because of versioning, capabilities, or dependencies on things like screen size or interaction capabilities of the recipient.
- o Some content isn't delivered over the HTTP (files on local file system), or there is no opportunity for tagging (data delivered over FTP) and in those cases, some other ways are needed for determining file type.

Operating systems use (and continued to evolve) different systems to determine the 'type' of something, different from the MIME tagging and bagging:

- o 'magic numbers': in many contexts, file types could be guessed pretty reliably by looking for headers.
- o Originally MAC OS had a 4 character 'file type' and another 4 character 'creator code' for file types.
- o Windows evolved to use the "file extension" -- 3 letters (and then more) at the end of the file name -- as the initial determination of the overall type of a file. This practice has now extended to other systems.

Information about these other ways of determining type (rather than by the content-type label) were gathered for the Internet Media Type registry; those registering types are encouraged to also describe 'magic numbers', Mac file type, common file extensions. However, since there was no formal use of that information, the quality of that information in the registry is haphazard.

Finally, there was the fact that tagging and bagging might be OK for unilaterally initiated (one-way) messaging, you might want to know whether you could handle the data before reading it in and

interpreting it, but the Internet Media Types weren't enough to tell.

3.3. The Rules Weren't Quite Followed

The behavior of the community when the Internet Media Type registry was designed hasn't matched expectations:

- o Lots of file types aren't registered (no entry in IANA for file types).
- o Those that are, the registration is incomplete or incorrect (people doing registration didn't understand 'magic number' or other fields).
- o The actual content deployed or created by deployed software doesn't match the registration.

In particular, Web implementations of Internet Media Types diverged from expected behavior:

- o Browser implementors would be liberal in what they accepted, and use what looked like a file extension in the URL and/or magic number or other 'sniffing' techniques to decide file type, without assuming content-label was authoritative. This was necessary anyway for files that weren't delivered by HTTP.
- o HTTP server implementors and administrators didn't supply ways of easily associating the 'intended' file type label with the file, resulting in files frequently being delivered with a label other than the one they would have chosen if they'd thought about it, and if browsers *had* assumed content-type was authoritative. Some popular servers had default configuration files that treated any unknown type as "text/plain" (plain ext in ASCII). Since it didn't matter (the browsers worked anyway), it was hard to get this fixed.

Incorrect senders coupled with liberal readers wind up feeding a negative feedback loop based on the robustness principle ([WikiRobust], [RFC3117]).

3.4. Consequences

The result, alas, is that increased unreliability, in that

- o servers sending responses to browsers don't have a good guarantee that the browser won't "sniff" the content and decide to do something other than treat it as it is labeled

- o browsers receiving content don't have a good guarantee that the content isn't mis-labeled
- o intermediaries (gateways, proxies, caches, and other pieces of the Web infrastructure) don't have a good way of telling what the conversation means.

This ambiguity and 'sniffing' also applies to packaged content in webapps ('bagging' but using ZIP rather than MIME multipart). (NOTE: NEEDS EXPANSION, REFERENCE TO WEBAPPS)

3.5. The Down Side of Extensibility

Extensibility adds great power, and allows the Web to evolve without committee approval of every extension. For some (those who want to extend and their clients who want those extensions), this is power! For others (those who are building Web components or infrastructure), extensibility is a drawback -- it adds to the unreliability and difference of the Web experience. When senders use extensions recipients aren't aware of, implement incorrectly or incompletely, then communication often fails. With messaging, this is a serious problem, although most 'rich text' documents are still delivered in multiple forms (using multipart/alternative).

If your job is to support users of a popular browser, however, where each user has installed a different configuration of file handlers and extensibility mechanisms, MIME may appear to add unnecessary complexity and variable experience for users of all but the most popular types.

4. Additional considerations

This section notes some additional considerations.

4.1. There are related problems with charsets

MIME includes provisions not only for file 'types', but also, importantly the "character encoding" used by text types: for example, simple US ASCII, Western European ISO-8859-1, Unicode UTF8. A similar vicious cycle also happened with character set labels: mislabeled content happily processed correctly by liberal browsers encouraged more and more sites to proliferate text with mis-labeled character sets, to the point where browsers feel they *have* to guess the wrong label. (NEEDS EXPANSION)

There are sites that intentionally label content as iso-2022-jp or euc-jp when it is in fact one of the Microsoft extension charsets

(e.g., for access to circled digits. This is an intentional misuse of the definitions of the charsets themselves -- definitions which originated at the national standards body level.

4.2. Embedded, downloaded, launch independent application

The type of a document might be determined not only for entire documents "HTML" vs "Word" vs "PDF", but also to embedded components of documents, "JPEG image" vs. "PNG image". However, the use cases, requirements and likely operational impact of MIME handling is likely different for those use cases.

4.3. Additional Use Cases: Polyglot and Multiview

There are some interesting additional use cases which add to the design requirements:

- o "Polyglot" documents: A 'polyglot' document is one which is some data which can be treated as two different Internet Media Types, in the case where the meaning of the data is the same. This is part of a transition strategy to allow content providers (senders) to manage, produce, store, deliver the same data, but with two different labels, and have it work equivalently with two different kinds of receivers (one of which knows one Internet Media Type, and another which knows a second one.) This use case was part of the transition strategy from HTML to an XML-based XHTML, and also as a way of a single service offering both HTML-based and XML-based processing (e.g., same content useful for news articles and Web pages.
- o "Multiview" documents: This use case seems similar but it's quite different. In this case, the same data has very different meaning when served as two different content-types, but that difference is intentional; for example, the same data served as text/html is a document, and served as an RDFa type is some specific data.

4.4. Evolution, Versioning, Forking

The subject of format/language/type evolution is complex; this section is a little terse.

Formats and their specifications evolve over time. There are several reasons for the evolution: innovation, compatibility with other implementations, attempts to gain control.

Some times new evolutions are "compatible", although compatibility has several variations. It is part of the responsibility of the designer of a new version of a file type to try to insure both

forward and backward compatibility: new documents work reasonably (with some fallback) with old viewers and that old documents work reasonably with new viewers. In some cases this is accomplished, others not; in some cases, "works reasonably" is softened to "either works reasonably or gives clear warning about nature of problem (version mismatch)."

In MIME, the 'tag', the Internet Media Type, corresponds to the versioned series. Internet Media Types do not identify a particular version of a file format. Rather, the general idea is that the Internet Media Type identifies the family, and also how you're supposed to otherwise find version information on a per-format basis. Many (most) file formats have an internal version indicator, with the idea that you only need a new Internet Media Type to designate a completely incompatible format. The notion of an "Internet Media Type" is very coarse-grained. The general approach to this has been that the actual Media Type includes provisions for version indicator(s) embedded in the content itself to determine more precisely the nature of how the data is to be interpreted. That is, the message itself contains further information.

Unfortunately, lots has gone wrong in this scenario as well -- processors ignoring version indicators encouraging content creators to not be careful to supply correct version indicators, leading to lots of content with wrong version indicators.

Those updating an existing Internet Media Type registration to account for new versions are admonished to not make previously conforming documents non-conforming. This is harder to enforce than would seem, because the previous specifications are not always accurate to what the Internet Media Type was used for in practice.

(NOTE: MULTIPLE INCOMPATIBLE AUTHORITATIVE SPECS)

4.5. Content Negotiation

The general idea of content negotiation is when party A communicates to party B, and the message can be delivered in more than one format (or version, or configuration), there can be some way of allowing some negotiation, some way for A to communication to B the available options, and for B to be able to accept or indicate preferences.

Content negotiation happens all over. When one fax machine twirps to another when initially connecting, they are negotiating resolution, compression methods and so forth. In Internet mail, which is a one-way communication, the "negotiation" consists of the sender preparing and sending multiple versions of the message, one in text/html, one in text/plain, for example, in sender-preference order. The

recipient then chooses the first version it can understand.

HTTP added "Accept" and "Accept-language" to allow content negotiation in HTTP GET, based on Internet Media Types, and there are other methods explained in the HTTP spec.

4.6. Fragment identifiers

The Web added the notion of being able to address part of a content and not the whole content by adding a 'fragment identifier' to the URL that addressed the data. Of course, this originally made sense for the original Web with just HTML, but how would it apply to other content. The URL spec glibly noted that "the definition of the fragment identifier meaning depends on the Internet Media Type", but unfortunately, few of the Internet Media Type definitions included this information, and practices diverged greatly.

If the interpretation of fragment identifiers depends on the MIME type, though, this really crimps the style of using fragment identifiers differently if content negotiation is wanted.

5. Recommendations

This section outlines the kinds of changes needed to bring the MIME system in line with current practice and to address the problems outlined above. The purpose of this text is not to specify the exact details of how changes can be accomplished, but rather to find broad agreement.

We need a clear direction on how to make the Web more reliable, not less. We need a realistic transition plan from the unreliable Web to the more reliable one. Part of this is to encourage senders (Web servers) to mean what they say, and encourage recipients (browsers) to give preference to what the senders are sending.

We should try to create specifications for protocols and best practices that will lead the Web to more reliable and secure communication. To this end, we give an overall architectural approach to use of MIME, and then specific specifications, for HTTP clients and servers, Web Browsers in general, proxies and intermediaries, which encourage behavior which, on the one hand, continues to work with the already deployed infrastructure (of servers, browsers, and intermediaries), but which advice, if followed, also improves the operability, reliability and security of the Web.

This section outlines requirements for standards and practices

intended to address some of the difficulties. This is an early version, which mainly contains "strawman" proposals for changes. It is intended to stimulate discussion -- however, the hope is that we can get agreement about the nature of the problems and current situation before focusing in detail about possible solutions. However, having at least strawman proposals seems to be helpful. For some problems, additional changes to IETF and W3C specifications are also be advisable; the expectations are briefly outlined here.

5.1. Internet Media Type registration

Update the Internet Media Type registry and registration process.

5.1.1. MIME registry magic numbers for sniffing

Be clearer about relationship of 'magic numbers' to sniffing; review Internet Media Types already registered and update.

5.1.2. Scripting and scriptable content safety

Be clearer about requiring Security Considerations to address risks of sniffing

5.1.3. Fragment identifiers

Problem: MIME type definitions don't talk about fragment identifiers.

require definition of fragment identifier applicability; add fragID semantics

5.1.4. Application info

Problem: ((hasn't been expanded))

Could the 'applications that use this type' section to be clearer about whether the file type is frequently for embedding (plug-in) or as a separate document with auto-launch (MIME handler), or should always be downloaded? Is there a separate issue for 'auto-play on download' vs. 'ask user for permission'?

5.1.5. File extensions in registry

Problem: Sniffing needs to use file extensions too; signify which file extensions are useful for sniffing.

Be clearer about file extension use and relationship of file extensions to MIME handlers

5.2. Sniffing

Various new specifications discuss, promote or mandate the use of 'sniffing' -- using the content of the data to supplement or even override the declared content-type or charset. Update these specifications.

5.2.1. Sniffing uses Media Type magic number

Update the proposed Media Type sniffing document so that sniffing uses MIME registry for 'magic numbers'.

5.2.2. Sniffing when there are multiple different definitions

Address issue of sniffing when there are multiple independent definitions of the same MIME type.

5.2.3. Sniffing charsets

Update sniffing of charsets to use charset reference info.

5.2.4. Sniffing security uses scriptability info

If the Internet Media Type registry is more explicit about which kinds of content contain what kind of scriptability access, then the specifications for sniffing can reference the Internet Media Type registry to determine what kinds of sniffing constitute a 'privilege upgrade'.

Note that all sniffing can be a privilege upgrade, if there is a buggy recipient, although bugs can be fixed, but spec violations are a problem.

5.3. Changes to IANA processes for MIME registries

Problem: Internet Media Type registries are hard to update, and there can be different definitions of the same MIME type.

STRAWMAN: Allow commenting or easier update; not all Internet Media Type owners need or have all the information the internet needs. Wiki for Internet Media Types as well as formal registry? Ability to add comments about deployed senders, deployed content, deployed receivers.

5.4. FTP specification

Do FTP clients also change rules about guessing file types based on OS of FTP server?

5.5. Update some URI definitions

ftp, file, need sniffing, http sometimes does; data defaults to text/plain rather than sniffing. Should this info be in the URI definitions.

5.6. Changes to W3C findings, processes

Update Tag finding on authoritative metadata: is it possible to remove 'authority'?

new: MIME and Internet Media Type section to WebArch, referencing this memo

New: Add a W3C Web architecture material on MIME in HTML to W3C web site, referencing this memo

Reconsider other extensibility mechanisms (namespaces, for example): should they use MIME or something like it?

6. Acknowledgements

This document is the result of discussions among many individuals in the IETF and W3C. Special thanks to Noah Mendelsohn.

7. IANA Considerations

This document includes no specific changes to IANA registries or processes. However, it outlines several considerations for future explicit recommendations to IANA, to change Internet Media Type and Charset registries and the processes around their maintenance. IANA evaluation of the feasibility of these changed processes is required.

8. Security Considerations

This document discusses some of the security issues resulting from use (and mis-use) of MIME content types in the Web.

9. Informative References

[RFC1521] Borenstein, N. and N. Freed, "MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies", RFC 1521, <<http://tools.ietf.org/html/rfc1521>>.

- [RFC1522] Moore, K., "MIME (Multipurpose Internet Mail Extensions) Part Two: Message Header Extensions for Non-ASCII Text", RFC 1522, September 1993, <<http://tools.ietf.org/html/rfc1522>>.
- [RFC1945] Berners-Lee, T., Fielding, R., and H. Nielsen, "Hypertext Transfer Protocol -- HTTP/1.0", RFC 1945, May 1996, <<http://tools.ietf.org/rfc/rfc1945>>.
- [RFC2046] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", RFC 2046, November 1996, <<http://tools.ietf.org/html/rfc2046>>.
- [RFC3117] Rose, M., "On the Design of Application Protocols", RFC 3117, November 2001, <<http://tools.ietf.org/html/rfc3117>>.
- [WikiRobust] "Robustness principle", 2010, <http://en.wikipedia.org/wiki/Robustness_principle>.
- [connolly92] Connolly, D., "Global Hypermedia", Oct 1992, <<http://lists.w3.org/Archives/Public/www-talk/1992SepOct/0024.html>>.
- [mime-sniff] Barth, A. and I. Hickson, "Media Type Sniffing", May 2010, <<http://tools.ietf.org/html/draft-abarth-mime-sniff>>.

Author's Address

Larry Masinter
Adobe
345 Park Ave.
San Jose, 95110
USA

Phone: +1 408 536 3024
Email: masinter@adobe.com
URI: <http://larry.masinter.net>

