

Workgroup: JMAP
Internet-Draft: draft-ietf-jmap-blob-13
Updates: [8620](#) (if approved)
Published: 14 September 2022
Intended Status: Standards Track
Expires: 18 March 2023
Authors: B. Gondwana, Ed.
Fastmail

JMAP Blob management extension

Abstract

The JMAP base protocol (RFC8620) provides the ability to upload and download arbitrary binary data via HTTP POST and GET on defined endpoint. This binary data is called a "blob".

This extension adds additional ways to create and access blobs, by making inline method calls within a standard JMAP request.

This extension also adds a reverse lookup mechanism to discover where blobs are referenced within other data types.

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 March 2023.

Copyright Notice

Copyright (c) 2022 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. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Conventions Used In This Document](#)
- [3. Addition to the Capabilities Object](#)
 - [3.1. urn:ietf:params:jmap:blob](#)
 - [3.1.1. Capability Example](#)
- [4. Blob Methods](#)
 - [4.1. Blob/upload](#)
 - [4.1.1. Blob/upload simple example](#)
 - [4.1.2. Blob/upload complex example](#)
 - [4.2. Blob/get](#)
 - [4.2.1. Blob/get simple example](#)
 - [4.2.2. Blob/get example with range and encoding errors](#)
 - [4.3. Blob/lookup](#)
 - [4.3.1. Blob/lookup example](#)
- [5. Security considerations](#)
- [6. IANA considerations](#)
 - [6.1. JMAP Capability registration for "blob"](#)
 - [6.2. JMAP Error Codes Registration for "unknownDataType"](#)
 - [6.3. Creation of "JMAP Data Types" Registry](#)
- [7. Changes](#)
- [8. Acknowledgements](#)
- [9. Normative References](#)
- [10. Informative References](#)
- [Author's Address](#)

1. Introduction

Sometimes JMAP ([RFC8620]) interactions require creating a blob and then referencing it. In the same way that IMAP Literals were extended by [RFC7888], embedding small blobs directly into the JMAP method calls array can be an option for reducing roundtrips.

Likewise, when fetching an object, it can be useful to also fetch the raw content of that object without a separate roundtrip.

Since raw blobs may contain arbitrary binary data, this document defines a use of the base64 coding specified in [RFC4648] for both creating and fetching blob data.

Where JMAP is being proxied through a system which applies additional access restrictions, it can be useful to know which

objects reference any particular blob, and this document defines a way to discover those references.

2. Conventions Used In This Document

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.

The definitions of JSON keys and datatypes in the document follow the conventions described in the core JMAP specification [[RFC8620](#)].

3. Addition to the Capabilities Object

The capabilities object is returned as part of the JMAP Session object; see [[RFC8620](#)], Section 2.

This document defines an additional capability URI.

3.1. urn:ietf:params:jmap:blob

The capability urn:ietf:params:jmap:blob being present in the "accountCapabilities" property of an account represents support for additional API methods on the Blob datatype. Servers that include the capability in one or more "accountCapabilities" properties MUST also include the property in the "capabilities" property.

The value of this property in the JMAP session "capabilities" property MUST be an empty object.

The value of this property in an account's "accountCapabilities" property is an object that MAY contain the following information on server capabilities and permissions for that account:

*maxSizeBlobSet: UnsignedInt|null

If present, this is the maximum size of blob (in octets) that the server will allow to be created (including blobs created by concatenating multiple data sources together).

Clients MUST NOT attempt to create blobs larger than this size.

If this value is not present or null, then clients are not required to limit the size of blob they try to create, though servers can always reject creation of blobs regardless of size; e.g. due to lack of disk space, or per-user rate limits.

*maxDataSources: UnsignedInt|null

If present, gives the maximum number of of DataSourceObjects allowed per creation in a Blob/upload. Servers MUST allow at least 64 items.

If this value is not present or null, then clients SHOULD assume a limit of 64 items.

*supportedTypeNames: String[]|null

An array of data type names that are supported for Blob/lookup. If the server does not support lookups then this could be the empty list, not present or null.

*supportedDigestAlgorithms String[]|null

An array of supported digest algorithms that are supported for Blob/get. If the server does not support calculating blob digests, then this could be the empty list, not present, or null. Algorithms on this list *MUST* be present in the HTTP Digest Algorithms registry defined by [[RFC3230](#)], and are always lowercased.

3.1.1. Capability Example

```
{
  "capabilities": {
    ...,
    "urn:ietf:params:jmap:blob": {}
  },
  "accounts": {
    "A13842": {
      ...
      "accountCapabilities": {
        "urn:ietf:params:jmap:blob": {
          "maxSizeBlobSet": 50000000,
          "maxDataSources": 100,
          "supportedTypeNames" : [
            "Mailbox",
            "Thread",
            "Email"
          ],
          "supportedDigestAlgorithms" : [
            "sha",
            "sha-256"
          ]
        }
      }
    }
  }
}
```

4. Blob Methods

A blob is a sequence of zero or more octets.

The JMAP base spec [[RFC8620](#)] defines the Blob/copy method, which is unchanged by this specification, and is selected by the urn:ietf:params:jmap:core capability.

The following JMAP Methods are selected by the urn:ietf:params:jmap:blob capability.

4.1. Blob/upload

This is similar to a Foo/set from [[RFC8620](#)] in some ways, however blobs can't be updated or deleted, so only create is allowed in the method call, and blobs don't have state, so there is no state field present in the method response.

Parameters

*accountId: Id

The id of the account in which the blobs will be created.

`*create: Id[UploadObject]`

A map of creation id to UploadObjects.

Result

The result is the same as for Foo/set in RFC8620, with created and notCreated objects mapping from the creationId.

The created objects contain:

`*id: Id`

the blobId which was created

`*type: String|null`

the media type as given in the creation (if any); or detected from content; or null

`*size: UnsignedInt`

as per RFC8620 - the size of the created blob in octets

Plus any other properties identical to those that would be returned in the JSON response of the RFC8620 upload endpoint (which may be extended in the future - this document anticipates that implementations will extend both the upload endpoint and the Blob/upload responses in the same way)

Or if there is a problem with a creation, then the server will return a notCreated response with a map from the failed creationId to a SetError object.

For each successful upload, servers MUST add an entry to the creationIds map for the request. This allows the blob id to be used via back-reference in subsequent method calls.

UploadObject

`*data: DataSourceObject[]`

an array of zero or more octet sources in order (zero to create an empty blob). The result of each of these sources is concatenated together in order to create the blob.

`*type: String|null (default: null)`

hint for media type of the data

DataSourceObject

Exactly one of:

*data:asText: String|null (raw octets, must be UTF-8)

*data:asBase64: String|null (base64 representation of octets)

or a blobId source:

*blobId: Id

*offset: UnsignedInt|null (MAY be zero)

*length: UnsignedInt|null (MAY be zero)

If null then offset is assumed to be zero.

If null then length is the remaining octets in the blob.

If the range can not be fully satisfied (i.e. begins or extends past the end of the data in the blob) then the DataSourceObject is invalid and results in a notCreated response for this creation id.

If the data properties have any invalid references or invalid data contained in them, the server MUST NOT guess as to the user's intent, and MUST reject the creation and return a notCreated response for that creation id.

Likewise, invalid characters in the base64 of data:asBase64, or invalid UTF-8 in data:asText MUST result in a nonCreated response.

It is envisaged that the definition for DataSourceObject might be extended in future, for example to fetch external content.

A server MUST accept at least 64 DataSourceObjects per create, as described in Section 3.1 of this document.

4.1.1. Blob/upload simple example

The data:asBase64 field is set over multiple lines for ease of publication here, however all data:asBase64 would be sent as a continuous string with no whitespace on the wire.

Method Call:

```
[
  "Blob/upload",
  {
    "accountId": "account1",
    "create": {
      "1": {
        "data" : [
          {
            "data:asBase64": "iVBORw0KGgoAAAANSUgAAAAEAAAABAQMAAAAl21
              AAAA1BMVEX/AAAZ4gk3AAAAAXRSTlN/gFy0ywAAAAp
              FUeJxjYgAAAAYAAzY3fKgAAAAASUVORK5CYII=",
          }
        ],
        "type": "image/png"
      },
    },
    "R1"
  ]
```

Response:

```
[
  "Blob/upload",
  {
    "accountId" : "account1",
    "created" : {
      "1": {
        "id" : "G4c6751edf9dd6903ff54b792e432fba781271beb",
        "type" : "image/png",
        "size" : 95
      },
    },
    "R1"
  ]
```


4.1.2. Blob/upload complex example

Method Calls:

```
[
  [
    "Blob/upload",
    {
      "create": {
        "b4": {
          "data": [
            {
              "data:asText": "The quick brown fox jumped over the lazy d
            }
          ]
        }
      }
    },
    "S4"
  ],
  [
    "Blob/upload",
    {
      "create": {
        "cat": {
          "data": [
            {
              "data:asText": "How"
            },
            {
              "blobId": "#b4",
              "length": 7,
              "offset": 3
            },
            {
              "data:asText": "was t"
            },
            {
              "blobId": "#b4",
              "length": 1,
              "offset": 1
            },
            {
              "data:asBase64": "YXQ/"
            }
          ]
        }
      }
    },
    "CAT"
  ],
]
```

```
[
  "Blob/get",
  {
    "properties": [
      "data:asText",
      "size"
    ],
    "ids": [
      "#cat"
    ]
  },
  "G4"
]
]
```

Responses:

```
[
  [
    "Blob/upload",
    {
      "oldState": null,
      "created": {
        "b4": {
          "id": "Gc0854fb9fb03c41cce3802cb0d220529e6eef94e",
          "size": 45,
          "type": "application/octet-stream"
        }
      },
      "notCreated": null,
      "accountId": "account1"
    },
    "S4"
  ],
  [
    "Blob/upload",
    {
      "oldState": null,
      "created": {
        "cat": {
          "id": "Gcc60576f036321ae6e8037ffc56bdee589bd3e23",
          "size": 19,
          "type": "application/octet-stream"
        }
      },
      "notCreated": null,
      "accountId": "account1"
    },
    "CAT"
  ]
]
```

```
],  
[  
  "Blob/get",  
  {  
    "list": [  
      {  
        "id": "Gcc60576f036321ae6e8037ffc56bdee589bd3e23",  
        "data:asText": "How quick was that?",  
        "size": 19  
      }  
    ],  
    "notFound": [],  
    "accountId": "account1"  
  },  
  "G4"  
]  
]
```

4.2. Blob/get

A standard JMAP get, with two additional optional parameters:

`*offset: UnsignedInt|null`

start this many octets into the blob data. If null or unspecified, this defaults to zero.

`*length: UnsignedInt|null`

return at most this many octets of the blob data. If null or unspecified, then all remaining octets in the blob are returned. This can be considered equivalent to an infinitely large length value, except that the `isTruncated` warning is not given unless the start offset is past the end of the blob.

Request Properties:

Any of

`*data:asText`

`*data:asBase64`

`*data` (returns `data:asText` if the selected octets are valid UTF-8, or `data:asBase64`)

`*digest:<algorithm>` (where `<algorithm>` is one of the named algorithms in the `supportedDigestAlgorithms` capability)

`*size`

If not given, properties defaults to data and size.

Result Properties:

`*data:asText: String|null`

the raw octets of the selected range if they are valid UTF-8, otherwise null

`*data:asBase64: String`

the base64 encoding of the octets in the selected range

`*digest:<algorithm> String`

the base64 encoding of the digest of the octets in the selected range, calculated using the named algorithm

`*isEncodingProblem: Boolean (default: false)`

`*isTruncated: Boolean (default: false)`

*size: UnsignedInt

the number of octets in the entire blob

The size value MUST always be the number of octets in the underlying blob, regardless of offset and length.

The data fields contain a representation of the octets within the selected range that are present in the blob. If the octets selected are not valid UTF-8 (including truncating in the middle of a multi-octet sequence) and data or data:asText was requested, then the key isEncodingProblem MUST be set to true and the data:asText response value MUST be null. In the case where data was requested and the data is not valid UTF-8, then data:asBase64 MUST be returned.

If the selected range requests data outside the blob (i.e. the offset+length is larger than the blob) then the result is either just the octets from the offset to the end of the blob, or an empty string if the offset is past the end of the blob. Either way, the isTruncated property in the result MUST be set to true to tell the client that the requested range could not be fully satisfied. If digest was requested, any digest is calculated on the octets that would be returned for a data field.

Servers SHOULD store the size for blobs in a format which is efficient to read, and clients SHOULD limit their request to just the size parameter if that is all they need, as fetching blob content could be significantly more expensive and slower for the server.

4.2.1. Blob/get simple example

Where a blob containing the string "The quick brown fox jumped over the lazy dog." has blobId Gc0854fb9fb03c41cce3802cb0d220529e6eef94e.

The first method call requests just the size for multiple blobs, and the second requests both size and a short range of the data for one of the blobs.

Method Calls:

```
[
  [
    "Blob/get",
    {
      "ids" : [
        "Gc0854fb9fb03c41cce3802cb0d220529e6eef94e"
        "not-a-blob"
      ],
      "properties" : [
        "data:asText",
        "digest:sha",
        "size"
      ]
    },
    "R1"
  ],
  [
    "Blob/get",
    {
      "accountId" : "account1",
      "ids" : [
        "Gc0854fb9fb03c41cce3802cb0d220529e6eef94e"
      ],
      "properties" : [
        "data:asText",
        "data:asBase64",
        "digest:sha",
        "digest:sha-256",
        "size"
      ],
      "offset" : 4,
      "length" : 9
    },
    "R2"
  ]
]
```

Responses:

```
[
  [
    "Blob/get",
    {
      "accountId": "account1",
      "list": [
        {
          "id": "Gc0854fb9fb03c41cce3802cb0d220529e6eef94e",
```

```
        "data:asText": "The quick brown fox jumped over the lazy dog."
        "digest:sha": "wIVPufsDxBz00ALLDSIFKebu+U4=",
        "size": 45
    }
],
    "notFound": [
        "not-a-blob"
    ]
},
    "R1"
],
[
    "Blob/get",
    {
        "accountId": "account1",
        "list": [
            {
                "id": "Gc0854fb9fb03c41cce3802cb0d220529e6eef94e",
                "data:asText": "quick bro",
                "digest:sha": "QiRAPtfyX8K6tm1i0AtZ87Xj3Ww=",
                "digest:sha-256": "gdg9INW7lwhK60Q9u0dwDz2ZY/gubi0En0x1FpKt00A",
                "size": 45
            }
        ]
    },
    "R2"
]
]
```


4.2.2. Blob/get example with range and encoding errors

The b1 value is the text: "The quick brown fox jumped over the \x81\x81 fox" which contains an invalid utf8 sequence.

The results have the following interesting properties:

- *G1: defaults to data and size - so b1 returns isEncodingProblem and a base64 value.
- *G2: since data:asText was explicitly selected, does not attempt to return a value for the data, just isEncodingProblem for b1.
- *G3: since only data:asBase64 was requested, there is no encoding problem and both values are returned.
- *G4: since the requested range could be satisfied as text, both blobs are returned as data:asText and there is no encoding problem.
- *G5: both blobs cannot satisfy the requested range, so isTruncated is true for both.

Note: some values have been wrapped for line length - there would be no whitespace in the data:asBase64 values on the wire

Method calls:

```
[
  "Blob/upload",
  {
    "create": {
      "b1": {
        "data": [
          {
            "data:asBase64": "VGhlIHFlYWNRIGJyb3duIGZveCBqdW1wZWwgc3ZlcjB0aGUggYEgZG9nLg=="
          }
        ]
      },
      "b2": {
        "data": [
          {
            "data:asText": "hello world"
          }
        ],
        "type" : "text/plain"
      }
    }
  },
  "S1"
],
[
  "Blob/get",
  {
    "ids": [
      "#b1",
      "#b2"
    ]
  },
  "G1"
],
[
  "Blob/get",
  {
    "ids": [
      "#b1",
      "#b2"
    ],
    "properties": [
      "data:asText",
      "size"
    ]
  },
  "G1"
],
  "G1"
]
```

```

    "G2"
  ],
  [
    "Blob/get",
    {
      "ids": [
        "#b1",
        "#b2"
      ],
      "properties": [
        "data:asBase64",
        "size"
      ]
    },
    "G3"
  ],
  [
    "Blob/get",
    {
      "offset": 0,
      "length": 5,
      "ids": [
        "#b1",
        "#b2"
      ]
    },
    "G4"
  ],
  [
    "Blob/get",
    {
      "offset": 20,
      "length": 100,
      "ids": [
        "#b1",
        "#b2"
      ]
    },
    "G5"
  ]
]

```

Responses:

```

[
  [
    "Blob/upload",
    {
      "oldState": null,

```

```

    "created": {
      "b2": {
        "id": "G2aae6c35c94fcfb415dbe95f408b9ce91ee846ed",
        "size": 11,
        "type": "application/octet-stream"
      },
      "b1": {
        "id": "G72cfa4804194563685d9a4b695f7ba20e7739576",
        "size": 43,
        "type": "text/plain"
      }
    },
    "updated": null,
    "destroyed": null,
    "notCreated": null,
    "notUpdated": null,
    "notDestroyed": null,
    "accountId": "account1"
  },
  "S1"
],
[
  "Blob/get",
  {
    "list": [
      {
        "id": "G72cfa4804194563685d9a4b695f7ba20e7739576",
        "isEncodingProblem": true,
        "data:asBase64": "VGhlIHF1aWNrIGJyb3duIGZveCBqdW1wZWQgb3ZlciB0aGUggYEgZG9nLg==",
        "size": 43
      },
      {
        "id": "G2aae6c35c94fcfb415dbe95f408b9ce91ee846ed",
        "data:asText": "hello world",
        "size": 11
      }
    ],
    "notFound": [],
    "accountId": "account1"
  },
  "G1"
],
[
  "Blob/get",
  {
    "list": [
      {
        "id": "G72cfa4804194563685d9a4b695f7ba20e7739576",

```

```

        "isEncodingProblem": true,
        "size": 43
    },
    {
        "id": "G2aae6c35c94fcfb415dbe95f408b9ce91ee846ed",
        "data:asText": "hello world",
        "size": 11
    }
],
"notFound": [],
"accountId": "account1"
},
"G2"
],
[
    "Blob/get",
    {
        "list": [
            {
                "id": "G72cfa4804194563685d9a4b695f7ba20e7739576",
                "data:asBase64": "VGhlIHF1aWNrIGJyb3duIGZveCBqdW1wZWQgb3ZlciB0aGUggYEgZG9nLg==",
                "size": 43
            },
            {
                "id": "G2aae6c35c94fcfb415dbe95f408b9ce91ee846ed",
                "data:asBase64": "aGVsbG8gd29ybGQ=",
                "size": 11
            }
        ],
        "notFound": [],
        "accountId": "account1"
    },
    "G3"
],
[
    "Blob/get",
    {
        "list": [
            {
                "id": "G72cfa4804194563685d9a4b695f7ba20e7739576",
                "data:asText": "The q",
                "size": 43
            },
            {
                "id": "G2aae6c35c94fcfb415dbe95f408b9ce91ee846ed",
                "data:asText": "hello",
                "size": 11
            }
        ]
    }
]

```

```
    ],
    "notFound": [],
    "accountId": "account1"
  },
  "G4"
],
[
  "Blob/get",
  {
    "list": [
      {
        "id": "G72cfa4804194563685d9a4b695f7ba20e7739576",
        "isTruncated": true,
        "isEncodingProblem": true,
        "data:asBase64": "anVtcGVkIG92ZXIgdGhlIIIGBIGRvZy4=",
        "size": 43
      },
      {
        "id": "G2aae6c35c94fcfb415dbe95f408b9ce91ee846ed",
        "isTruncated": true,
        "data:asText": "",
        "size": 11
      }
    ],
    "notFound": [],
    "accountId": "account1"
  },
  "G5"
]
]
```

4.3. Blob/lookup

Given a list of blobIds, this method does a reverse lookup in each of the provided type names to find the list of Ids within that data type which reference the provided blob.

The definition of reference is somewhat loosely defined, but roughly means "you could discover this blobId by looking inside this object", for example if a Mailbox contains an Email which references the blobId, then it references that blobId. Likewise for a Thread.

Parameters

*accountId: Id

The id of the account used for the call.

*typeNames: String[]

A list of names from the "JMAP Data Types" registry. Only names for which "Can reference blobs" is true may be specified, and the capability which defines each type must also be used by the overall JMAP request in which this method is called.

If a type name is not known by the server, or the associated capability has not been requested, then the server returns an "unknownDataType" error.

*ids: Id[]

A list of blobId values to be looked for.

Response

*list: BlobInfo[]

A list of BlobInfo objects.

BlobInfo Object

*id: Id

The Blob Identifier.

*matchedIds: String[Id[]]

A map from type name to list of Ids of that data type (e.g. the name "Email" maps to a list of emailIds)

If a blob is not visible to a user at all, then the server SHOULD return that blobId in the notFound array, however it may also return an empty list for each type name, as it may not be able to know if other data types do reference that blob.

4.3.1. Blob/lookup example

Method call:

```
[
  "Blob/lookup",
  {
    "typeNames": [
      "Mailbox",
      "Thread",
      "Email"
    ],
    "ids": [
      "Gd2f81008cf07d2425418f7f02a3ca63a8bc82003",
      "not-a-blob"
    ]
  },
  "R1"
]
```

Response:

```
[
  "Blob/lookup",
  {
    "list": [
      {
        "id": "Gd2f81008cf07d2425418f7f02a3ca63a8bc82003",
        "matchedIds": {
          "Mailbox": [
            "M54e97373",
            "Mcbe6b662"
          ],
          "Thread": [
            "T1530616e"
          ],
          "Email": [
            "E16e70a73eb4",
            "E84b0930cf16"
          ]
        }
      }
    ],
    "notFound": [
      "not-a-blob"
    ]
  },
  "R1"
]
```

5. Security considerations

JSON parsers are not all consistent in handling non-UTF-8 data. JMAP requires that all JSON data be UTF-8 encoded, so servers **MUST** only return a null value if data:asText is requested for a range of octets which is not valid UTF-8, and set isEncodingProblem: true.

Servers **MUST** apply any access controls, such that if the authenticated user would be unable to discover the blobId by making queries, then this fact can't be discovered via a Blob/lookup. For example, if an Email exists in a Mailbox which the authenticated user does not have access to see, then that emailId **MUST** not be returned in a lookup for a blob which is referenced by that email.

If a server might sometimes return all names empty rather than putting a blobId in the notFound response to a Blob/get, then the server **SHOULD** always return the same type of response, regardless of whether a blob exists but the user can't access it, or doesn't exist at all. This avoids leaking information about the existence of the blob.

The server **MUST NOT** trust that the data given to a Blob/upload is a well formed instance of the specified media type, and if the server attempts to parse the given blob, only hardened parsers designed to deal with arbitrary untrusted data should be used. The server **SHOULD NOT** reject data on the grounds that it is not a valid specimen of the stated type.

Blob/upload with carefully chosen data sources can be used to recreate dangerous content on the far side of security scanners (anti-virus or exfiltration scanners for example) which may be watching the upload endpoint. Server implementations **SHOULD** provide a hook to allow security scanners to check the resulting blob after concatenating the data sources in the same way that they do for the upload endpoint.

6. IANA considerations

6.1. JMAP Capability registration for "blob"

IANA is requested to register the "blob" JMAP Capability as follows:

Capability Name: urn:ietf:params:jmap:blob

Specification document: this document

Intended use: common

Change Controller: IETF

Security and privacy considerations: this document, Section XXX

6.2. JMAP Error Codes Registration for "unknownDataType"

IANA is requested to register the "unknownDataType" JMAP Error Code as follows:

JMAP Error Code: unknownDataType

Intended use: common

Change Controller: IETF

Reference: this document

Description: The server does not recognise this data type, or the capability to enable it was not present.

6.3. Creation of "JMAP Data Types" Registry

IANA is requested to create a new registry "JMAP Data Types" with the initial content:

| Type Name | Can reference blobs | Can use for state change | Capability | Reference |
|------------------|---------------------|--------------------------|---------------------------------------|-----------------------------|
| Core | No | No | urn:ietf:params:jmap:core | [RFC8620] |
| PushSubscription | No | No | urn:ietf:params:jmap:core | [RFC8620] |
| Mailbox | Yes | Yes | urn:ietf:params:jmap:mail | [RFC8621] |
| Thread | Yes | Yes | urn:ietf:params:jmap:mail | [RFC8621] |
| Email | Yes | Yes | urn:ietf:params:jmap:mail | [RFC8621] |
| EmailDelivery | No | Yes | urn:ietf:params:jmap:mail | [RFC8621] |
| SearchSnippet | No | No | urn:ietf:params:jmap:mail | [RFC8621] |
| Identity | No | Yes | urn:ietf:params:jmap:submission | [RFC8621] |
| EmailSubmission | No | Yes | urn:ietf:params:jmap:submission | [RFC8621] |
| VacationResponse | No | Yes | urn:ietf:params:jmap:vacationresponse | [RFC8621] |
| MDN | No | No | urn:ietf:params:jmap:mdn | [RFC9007] |

Table 1

This policy for this registry is "Specification required", either an RFC or a similarly stable reference document which defines a JMAP Data Type and associated capability.

7. Changes

EDITOR: please remove this section before publication.

The source of this document exists on github at: <https://github.com/brong/draft-gondwana-jmap-blob/>

draft-ietf-jmap-blob-13

- *added examples of digest responses

draft-ietf-jmap-blob-12

- *updates based on Neil Jenkins' feedback:

- fixed [] positions for type specs
- documented delta between /upload and /set better
- allowed zero-length blobId sources
- fixed examples with /set leftovers
- documented datatypes registry policy

- *added optional "digest" support

draft-ietf-jmap-blob-11:

- *updates based on IETF113 feedback:

- added wording to suggest the a Blob/get of just size might be faster
- added an example with just the size field being selected

draft-ietf-jmap-blob-10:

- *removed remaining references to catenate.

draft-ietf-jmap-blob-09:

- *tidied up introduction text
- *replaced Blob/set with Blob/upload
- *made all upload creates take an array of sources to normalise behaviour at the cost of a slightly more complex default case.

draft-ietf-jmap-blob-08:

- *Fixed spelling of Neil's name in acknowledgements
- *Last call review (thanks Jim Fenton)

- fixed mmark silliness causing RFC8620 to be non-normative in the references
- clarified the capability object and accountCapability object requirements
- made capability keys much more tightly defined, with mandatory minimum catenate limit and default values.
- increased use of normative language generally
- lowercased 'blob' anywhere it wasn't explicitly the object

-lowercased titles of the columns in the registry

draft-ietf-jmap-blob-07:

- *more examples to cover the interactions of offset, length and encoding checks.

draft-ietf-jmap-blob-06:

- *removed asHex - we only need base64 and text
- *added reference to where base64 is defined
- *made 'destroy' not be allowed
- *expanded JSON examples for readability
- *removed 'expires' from examples

draft-ietf-jmap-blob-05:

- *discovered I hadn't actually included typeName and matchedIds anywhere except the updates section, oops!
- *added a catenate example
- *tightened up some text

draft-ietf-jmap-blob-04:

- *added security considerations for scanning catenate results

draft-ietf-jmap-blob-03:

- *added capabilities object
- *renamed types to typeName and matchedIds
- *added details of how to handle non-UTF8 data and truncation in Blob/get
- *added isTruncated and isEncodingProblem to Blob/get to tell the client if the request wasn't entirely satisfied.

draft-ietf-jmap-blob-02:

- *fixed incorrect RFC number in reference and HTTP PUT -> POST, thanks Ken.
- *added acknowledgements section
- *removed all 'datatype' text and changed to 'data type' or 'type name' as appropriate (issue #1 proposal)
- *expanded security considerations section and moved optional Blob/lookup empty case into Blob/lookup section

draft-ietf-jmap-blob-01:

- *renamed 'datatypes' to 'types' to align with PushSubscription from RFC8620.
- *added example for Blob/get

*specified offset and length precisely

draft-ietf-jmap-blob-00:

*initial adoption as an IETF document, otherwise identical to
draft-gondwana-jmap-blob-02

draft-gondwana-jmap-blob-02

*renamed 'objects' to 'datatypes'
*specified Blob/lookup
*added IANA registry for datatypes

draft-gondwana-jmap-blob-01

*added an example

draft-gondwana-jmap-blob-00

*initial proposal

8. Acknowledgements

Joris Baum, Jim Fenton, Neil Jenkins, Alexey Melnikov, Ken Murchison, Robert Stepanek and the JMAP working group at the IETF.

9. 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>>.
- [RFC3230] Mogul, J. and A. Van Hoff, "Instance Digests in HTTP", RFC 3230, DOI 10.17487/RFC3230, January 2002, <<https://www.rfc-editor.org/info/rfc3230>>.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/info/rfc4648>>.
- [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>>.
- [RFC8620] Jenkins, N. and C. Newman, "The JSON Meta Application Protocol (JMAP)", RFC 8620, DOI 10.17487/RFC8620, July 2019, <<https://www.rfc-editor.org/info/rfc8620>>.

10. Informative References

[RFC7888]

Melnikov, A., Ed., "IMAP4 Non-synchronizing Literals", RFC 7888, DOI 10.17487/RFC7888, May 2016, <<https://www.rfc-editor.org/info/rfc7888>>.

[RFC8621]

Jenkins, N. and C. Newman, "The JSON Meta Application Protocol (JMAP) for Mail", RFC 8621, DOI 10.17487/RFC8621, August 2019, <<https://www.rfc-editor.org/info/rfc8621>>.

Author's Address

Bron Gondwana (editor)
Fastmail
Level 2, 114 William St
Melbourne VIC 3000
Australia

Email: brong@fastmailteam.com

URI: <https://www.fastmail.com>