

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: August 29, 2013

H. VandeSompel
Los Alamos National Laboratory
M. Nelson
Old Dominion University
R. Sanderson
Los Alamos National Laboratory
February 25, 2013

**HTTP framework for time-based access to resource states -- Memento
draft-vandesompel-memento-06**

Abstract

The HTTP-based Memento framework bridges the present and past Web. It facilitates obtaining representations of prior states of a resource by leveraging the resource's URI and a preferred datetime. To this end, the framework introduces datetime negotiation, a variation on content negotiation. It also facilitates recognizing a resource that encapsulates a frozen prior state of another resource by expressing that state's datetime and by providing an appropriately typed link to that other resource.

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 August 29, 2013.

Copyright Notice

Copyright (c) 2013 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	4
1.1.	Terminology	4
1.2.	Notational Conventions	4
1.3.	Purpose	5
2.	HTTP headers, Relation Types	6
2.1.	HTTP Headers	7
2.1.1.	Accept-Datetime, Memento-Datetime	7
2.1.2.	Vary	8
2.1.3.	Link	8
2.2.	Relation Types	8
2.2.1.	Relation Type "original"	9
2.2.2.	Relation Type "timegate"	9
2.2.3.	Relation Type "timemap"	9
2.2.4.	Relation Type "memento"	10
2.3.	Memento response headers: Summary	10
3.	General considerations regarding Datetime Negotiation	11
4.	HTTP Interactions for Datetime Negotiation	12
4.1.	Pattern 1 - The Original Resource acts as its own TimeGate	14
4.1.1.	Pattern 1.1 - URI-R=URI-G ; 302-style negotiation ; distinct URI-M for Mementos	15
4.1.2.	Pattern 1.2 - URI-R=URI-G ; 200-style negotiation ; distinct URI-M for Mementos	17
4.1.3.	Pattern 1.3 - URI-R=URI-G ; 200-style negotiation ; no distinct URI-M for Mementos	18
4.2.	Pattern 2 - A remote resource acts as a TimeGate for the Original Resource	19
4.2.1.	Pattern 2.1 - URI-R<>URI-G ; 302-style negotiation ; distinct URI-M for Mementos	21
4.2.2.	Pattern 2.2 - URI-R<>URI-G ; 200-style negotiation ; distinct URI-M for Mementos	23
4.2.3.	Pattern 2.3 - URI-R<>URI-G ; 200-style negotiation ; no distinct URI-M for Mementos	24
4.3.	Pattern 3 - The Original Resource is a Fixed Resource	25
4.4.	Pattern 4 - Mementos without a TimeGate	26
4.5.	Special Cases	28
4.5.1.	Original Resource provides no "timegate" link	28
4.5.2.	Server exists but Original Resource no longer does	28

4.5.3.	Issues with Accept-Datetime	28
4.5.4.	Memento of a 3XX response	29
4.5.5.	Memento of responses with 4XX or 5XX HTTP status codes	31
4.5.6.	Sticky "Memento-Datetime" value for Mementos	32
4.5.7.	Intermediate Resources	32
5.	TimeMaps	33
5.1.	Index and Paging TimeMaps	35
5.2.	Mementos for TimeMaps	37
6.	IANA Considerations	37
7.	Security Considerations	38
8.	Changelog	38
9.	Acknowledgements	40
10.	References	41
10.1.	Normative References	41
10.2.	Informative References	41
	Authors' Addresses	42

1. Introduction

1.1. Terminology

This specification uses the terms "resource", "request", "response", "entity-body", "content negotiation", "user agent", "server" as described in [[RFC2616](#)], and it uses the terms "representation" and "resource state" as described in [[W3C.REC-aww-20041215](#)].

In addition, the following terms specific to the Memento framework are introduced:

- o Original Resource: An Original Resource is a resource that exists or used to exist, and for which access to one of its prior states may be required.
- o Memento: A Memento for an Original Resource is a resource that encapsulates a prior state of the Original Resource. A Memento for an Original Resource as it existed at time T is a resource that encapsulates the state the Original Resource had at time T.
- o TimeGate: A TimeGate for an Original Resource is a resource that is capable of datetime negotiation to support access to prior states of the Original Resource.
- o TimeMap: A TimeMap for an Original Resource is a resource from which a list of URIs of Mementos of the Original Resource is available.

1.2. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

When needed for extra clarity, the following conventions are used:

- o URI-R is used to denote the URI of an Original Resource.
- o URI-G is used to denote the URI of a TimeGate.
- o URI-M is used to denote the URI of a Memento.
- o URI-T is used to denote the URI of a TimeMap.

1.3. Purpose

The state of an Original Resource may change over time. Dereferencing its URI at any specific moment yields a response that reflects the resource's state at that moment: a representation of the resource's state (e.g. "200 OK" HTTP status code), an indication of its non-existence (e.g. "404 Not Found" HTTP status code), a relation to another resource (e.g. "302 Found" HTTP status code), etc. However, responses may also exist that reflect prior states of an Original Resource: a representation of a prior state of the Original Resource, an indication that the Original Resource did not exist at some time in the past, a relation that the Original Resource had to another resource at some time in the past, etc. Mementos that provide such responses exist in web archives, content management systems, or revision control systems, among others. For any given Original Resource several Mementos may exist, each one reflecting a frozen prior state of the Original Resource.

Examples are:

Mementos for Original Resource <http://www.ietf.org/> :

- o <http://web.archive.org/web/19970107171109/http://www.ietf.org/>
- o <http://webarchive.nationalarchives.gov.uk/20080906200044/http://www.ietf.org/>

Mementos for Original Resource http://en.wikipedia.org/wiki/Hypertext_Transfer_Protocol :

- o http://en.wikipedia.org/w/index.php?title=Hypertext_Transfer_Protocol&oldid=366806574
- o http://en.wikipedia.org/w/index.php?title=Hypertext_Transfer_Protocol&oldid=33912
- o http://web.archive.org/web/20071011153017/http://en.wikipedia.org/wiki/Hypertext_Transfer_Protocol

Mementos for Original Resource <http://www.w3.org/TR/webarch/> :

- o <http://www.w3.org/TR/2004/PR-webarch-20041105/>
- o <http://www.w3.org/TR/2002/WD-webarch-20020830/>
- o <http://webarchive.nationalarchives.gov.uk/20100304163140/http://www.w3.org/TR/webarch/>

In the abstract, the Memento framework introduces a mechanism to access versions of web resources that:

- o Is fully distributed in the sense that resource versions may reside on multiple servers, and that any such server is likely only aware of the versions it holds;
- o Uses the global notion of datetime as a resource version indicator and access key;
- o Leverages the following primitives of [\[W3C.REC-aww-20041215\]](#): resource, resource state, representation, content negotiation, and link.

The core components of Memento's mechanism to access resource versions are:

1. The abstract notion of the state of an Original Resource (URI-R) as it existed at datetime T. Note the relationship with the ability to identify the state of a resource at datetime T by means of a URI as intended by the proposed Dated URI scheme [\[I-D.masinter-dated-uri\]](#).
2. A "bridge" from the present to the past, consisting of:
 - o The existence of a TimeGate (URI-G), which is aware of (at least part of the) version history of the Original Resource (URI-R);
 - o The ability to negotiate in the datetime dimension with that TimeGate (URI-G), as a means to access the state that the Original Resource (URI-R) had at datetime T.
3. A "bridge" from the past to the present, consisting of an appropriately typed link from a Memento (URI-M), which encapsulates the state the Original Resource (URI-R) had at datetime T, to the Original Resource (URI-R).

This document is concerned with specifying an instantiation of these abstractions for resources that are identified by HTTP(S) URIs.

2. HTTP headers, Relation Types

The Memento framework is concerned with HEAD and GET interactions with Original Resources, TimeGates, Mementos, and TimeMaps that are identified by HTTP or HTTPS URIs. Details are only provided for resources identified by HTTP URIs but apply similarly to those with HTTPS URIs.

2.1. HTTP Headers

The Memento framework operates at the level of HTTP request and response headers. It introduces two new headers ("Accept-Datetime", "Memento-Datetime") and introduces new values for two existing headers ("Vary", "Link"). Other HTTP headers are present or absent in Memento response/request cycles as specified by [\[RFC2616\]](#).

2.1.1. Accept-Datetime, Memento-Datetime

The "Accept-Datetime" request header is transmitted by a user agent to indicate it wants to access a past state of an Original Resource. To that end, the "Accept-Datetime" header is conveyed in an HTTP request issued against a TimeGate for an Original Resource, and its value indicates the datetime of the desired past state of the Original Resource.

Example of an "Accept-Datetime" request header:

```
Accept-Datetime: Thu, 31 May 2007 20:35:00 GMT
```

The "Memento-Datetime" response header is used by a server to indicate that a response reflects a prior state of an Original Resource. Its value expresses the datetime of that state. The URI of the Original Resource for which the response reflects a prior state is provided as the Target IRI of a link provided in the HTTP "Link" header that has a Relation Type of "original" (see [Section 2.2](#)).

The presence of a "Memento-Datetime" header and associated value for a given response constitutes a promise that the resource state reflected in the response will no longer change (see [Section 4.5.6](#)).

Example of a "Memento-Datetime" response header:

```
Memento-Datetime: Wed, 30 May 2007 18:47:52 GMT
```

Values for the "Accept-Datetime" and "Memento-Datetime" headers consist of a MANDATORY datetime expressed according to the [\[RFC1123\]](#) format, which is formalized by the [rfc1123-date](#) construction rule of the BNF in Figure 1. The datetime **MUST** be represented in Greenwich Mean Time (GMT).


```
accept-dt-value = rfc1123-date *SP
rfc1123-date = wkday "," SP date1 SP time SP "GMT"
date1         = 2DIGIT SP month SP 4DIGIT
               ; day month year (e.g., 20 Mar 1957)
time          = 2DIGIT ":" 2DIGIT ":" 2DIGIT
               ; 00:00:00 - 23:59:59 (e.g., 14:33:22)
wkday         = "Mon" | "Tue" | "Wed" | "Thu" | "Fri" | "Sat" |
               "Sun"
month         = "Jan" | "Feb" | "Mar" | "Apr" | "May" | "Jun" |
               "Jul" | "Aug" | "Sep" | "Oct" | "Nov" | "Dec"
```

Figure 1: BNF for the datetime format

[2.1.2.](#) Vary

Generally, the "Vary" header is used in HTTP responses to indicate the dimensions in which content negotiation is possible. In the Memento framework, a TimeGate uses the "Vary" header with a value that includes "accept-datetime" to convey that datetime negotiation is possible.

For example, this use of the "Vary" header indicates that datetime is the only dimension in which negotiation is possible:

```
Vary: accept-datetime
```

The use of the "Vary" header in this example shows that both datetime negotiation, and media type content negotiation are possible:

```
Vary: accept-datetime, accept
```

[2.1.3.](#) Link

The Memento framework defines the "original", "timegate", "timemap", and "memento" Relation Types to convey typed links among Original Resources, TimeGates, Mementos, and TimeMaps. They are defined in [Section 2.2](#), below. In addition, existing Relation Types may be used, for example, to support navigating among Mementos. Examples are "first", "last", "prev", "next", "predecessor-version", "successor-version" as detailed in [[RFC5988](#)] and [[RFC5829](#)].

[2.2.](#) Relation Types

This section introduces the Relation Types used in the Memento framework. They are defined in a general way and their use in HTTP "Link" Headers [[RFC5988](#)] is described in detail. The use of these Relation Types in TimeMaps is described in [Section 5](#).

2.2.1. Relation Type "original"

"original" -- A link with an "original" Relation Type is used to point from a TimeGate or a Memento to its associated Original Resource.

Use in HTTP "Link" headers: A TimeGate and a Memento MUST include exactly one link with an "original" Relation Type in their HTTP "Link" header.

2.2.2. Relation Type "timegate"

"timegate" -- A link with a "timegate" Relation Type is used to point from the Original Resource, as well as from a Memento associated with the Original Resource, to a TimeGate for the Original Resource.

Use in HTTP "Link" headers: An Original Resource and a Memento SHOULD include a link with a "timegate" Relation Type in their HTTP "Link" header. Since multiple TimeGates can exist for any Original Resource, multiple "timegate" links MAY occur, each with a distinct Target IRI.

2.2.3. Relation Type "timemap"

"timemap" -- A link with a "timemap" Relation Type is used to point from a TimeGate or a Memento associated with an Original Resource, as well as from the Original Resource itself, to a TimeMap for the Original Resource.

Attributes: A link with a "timemap" Relation Type SHOULD use the "type" attribute to convey the mime type of the TimeMap serialization. The "from" and "until" attributes may be used to express the start and end of the temporal interval covered by Mementos listed in the TimeMap. That is, the linked TimeMap will not contain Mementos with archival datetimes outside of the expressed temporal interval. Attempts SHOULD be made to convey this interval as accurately as possible. The value for the these attributes MUST be a datetime expressed according to the [rfc1123](#)-date construction rule of the BNF in Figure 1 and it MUST be represented in Greenwich Mean Time (GMT).

Use in HTTP "Link" headers: An Original Resource, a TimeGate and a Memento SHOULD include a link with a "timemap" Relation Type in their HTTP "Link" header. Multiple such links, each with a distinct Target IRI, MAY be expressed as a means to point to different TimeMaps or to different serializations of the same TimeMap. In all cases, use of the "from" and "until" attributes is OPTIONAL.

2.2.4. Relation Type "memento"

"memento" -- A link with a "memento" Relation Type is used to point from a TimeGate or a Memento for an Original Resource, as well as from the Original Resource itself, to a Memento for the Original Resource.

Attributes: A link with a "memento" Relation Type MUST include a "datetime" attribute with a value that matches the "Memento-Datetime" of the Memento that is the target of the link; that is, the value of the "Memento-Datetime" header that is returned when the URI of the linked Memento is dereferenced. The value for the "datetime" attribute MUST be a datetime expressed according to the [rfc1123](#)-date construction rule of the BNF in Figure 1 and it MUST be represented in Greenwich Mean Time (GMT). This link MAY include a "license" attribute to associate a license with the Memento; the value for the "license" attribute SHOULD be a URI.

Use in HTTP "Link" headers: An Original Resource, a TimeGate and a Memento MAY include links in their HTTP "Link" headers with a "memento" Relation Type. For responses in which a Memento is selected, the provision of navigational links that lead to Mementos other than the selected one can be beneficial to the user agent. Of special importance are links that lead to the temporally first and last Memento known to the responding server, as well as links leading to Mementos that are temporally adjacent to the selected one.

2.3. Memento response headers: Summary

The below table summarizes the use of the Vary, Memento-Datetime, and Link headers in responses from Original Resources, TimeGates, and Mementos. The value of the Vary header MUST contain the "accept-datetime" value, whereas the value of the Memento-Datetime header MUST convey the datetime of the state of the Original Resource that is encapsulated in the Memento selected by the response. From now on, this datetime will be referred to as the archival datetime of the Memento.

The Link header can contain links with the "original", "timegate", "timemap", and "memento" Relation Types and their use for the various resource types is also detailed in the table.

The inclusion of Memento headers and links in responses from the respective resources is as shown in the table irrespective of whether an HTTP HEAD/GET request against their URI includes an "Accept-Datetime" request header: requests with and without an "Accept-Datetime" header yield the inclusion of the same response headers. Values for those response headers depend on the request.

The use of response headers and links for specific scenarios can be derived from this table. For example:

- o URI-R<>URI-G<>URI-M - If the TimeGate is distinct from the associated Original Resource and Memento, then their use is as shown in the respective columns.
- o URI-R=URI-G - If the TimeGate coincides with the Original Resource, then their use is obtained by taking the union of the headers for a TimeGate and for an Original Resource and by selecting REQUIRED over RECOMMENDED and REQUIRED over PROHIBITED. For example, this case must have a "Vary" header that includes an "accept-datetime" value.
- o URI-R=URI-M - If the Original Resource and its Memento coincide (the Original Resource is a FixedResource as per [\[W3C.gen-ont-20090420\]](http://www.w3.org/TR/2009/REC-gen-ont-20090420/)), their use is obtained by taking the union of the headers for an Original Resource and for a Memento and by selecting REQUIRED over RECOMMENDED and REQUIRED over PROHIBITED. For example, this case must have a "Memento-Datetime" header.

Response Header	Original Resource	TimeGate	Memento
Vary:	PROHIBITED	REQUIRED, 1	PROHIBITED
accept-datetime			
Memento-Datetime	PROHIBITED	PROHIBITED	REQUIRED, 1
Link	RECOMMENDED, 0 or 1	REQUIRED, 1	REQUIRED, 1
"original" link	NA	REQUIRED, 1	REQUIRED, 1
"timegate" link	RECOMMENDED, 0 or more	NA	RECOMMENDED, 0 or more
"timemap" link	RECOMMENDED, 0 or more	RECOMMENDED, 0 or more	RECOMMENDED, 0 or more
"memento" link	OPTIONAL, 0 or more	OPTIONAL, 0 or more	OPTIONAL, 0 or more

Table 1: Memento Headers

3. General considerations regarding Datetime Negotiation

In order to respond to a datetime negotiation request, the server uses an internal algorithm to select the Memento that best meets the user agent's datetime preference. The exact nature of the selection algorithm is at the server's discretion but SHOULD be consistent, for

example, always selecting the Memento that is nearest in time relative to the requested datetime, always selecting the Memento that is nearest in the past relative to the requested datetime, etc.

Due to the sparseness of Mementos in most systems, the value of the "Memento-Datetime" header returned by a server may differ (significantly) from the value conveyed by the user agent in "Accept-Datetime".

Although a Memento encapsulates a prior state of an Original Resource, the entity-body returned in response to an HTTP GET request issued against a Memento may very well not be byte-to-byte the same as an entity-body that was previously returned by that Original Resource. Various reasons exist why there are significant chances these would be different yet do convey substantially the same information. These include format migrations as part of a digital preservation strategy, URI-rewriting as applied by some web archives, and the addition of banners as a means to brand web archives.

When negotiating in the datetime dimension, the regular content negotiation dimensions (media type, character encoding, language, and compression) remain available. It is the TimeGate server's responsibility to honor (or not) such content negotiation, and in doing so it MUST always first select a Memento that meets the user agent's datetime preference, and then consider honoring regular content negotiation for it. As a result of this approach, the returned Memento will not necessarily meet the user agent's regular content negotiation preferences. Therefore, it is RECOMMENDED that the server provides "memento" links in the HTTP "Link" header pointing at Mementos that do meet the user agent's regular content negotiation requests and that have a value for the "Memento-Datetime" header in the temporal vicinity of the user agent's preferred datetime value.

A user agent that engages in datetime negotiation with a resource typically starts by issuing an HTTP HEAD, not GET, request with an "Accept-Datetime" header in order to determine how to proceed. This strategy is related to the existence of various server implementation patterns as will become clear in the below.

In the following sections, the HTTP status code of the responses with an entity-body is shown as "200 OK", but a series of "206 Partial Content" responses could be substituted.

4. HTTP Interactions for Datetime Negotiation

This section describes the HTTP interactions of the Memento framework

for a variety of scenarios. First, Figure 2 provides a schematic overview of a successful request/response chain that involves datetime negotiation. Dashed lines depict HTTP transactions between user agent and server. The interactions are for a scenario where the Original Resource resides on one server, whereas both its TimeGate and Mementos reside on another. Scenarios also exist in which all these resources are on the same server (for example, content management systems) or all are on different servers (for example, an aggregator of TimeGates).

```

1: UA --- HTTP HEAD/GET; Accept-Datetime: T -----> URI-R
2: UA <-- HTTP 200; Link: URI-G ----- URI-R
3: UA --- HTTP HEAD/GET; Accept-Datetime: T -----> URI-G
4: UA <-- HTTP 302; Location: URI-M; Vary; Link:
    URI-R,URI-T -----> URI-G
5: UA --- HTTP GET URI-M; Accept-Datetime: T -----> URI-M
6: UA <-- HTTP 200; Memento-Datetime: T; Link:
    URI-R,URI-T,URI-G ----- URI-M

```

Figure 2: A datetime negotiation request/response chain

- o Step 1: The user agent that wants to access a prior state of the Original Resource issues an HTTP HEAD/GET against URI-R that has an "Accept-Datetime" HTTP header with a value of the datetime of the desired state.
- o Step 2: The response from URI-R includes an HTTP "Link" header with a Relation Type of "timegate" pointing at a TimeGate (URI-G) for the Original Resource.
- o Step 3: The user agent starts the datetime negotiation process with the TimeGate by issuing an HTTP GET request against URI-G that has an "Accept-Datetime" HTTP header with a value of the datetime of the desired prior state of the Original Resource.
- o Step 4: The response from URI-G includes a "Location" header pointing at a Memento (URI-M) for the Original Resource. In addition, the response contains an HTTP "Link" header with a Relation Type of "original" pointing at the Original Resource (URI-R), and an HTTP "Link" header with a Relation Type of "timemap" pointing at a TimeMap (URI-T).
- o Step 5: The user agent issues an HTTP GET request against URI-M.
- o Step 6: The response from URI-M includes a "Memento-Datetime" HTTP header with a value of the archival datetime of the Memento. It also contains an HTTP "Link" header with a Relation Type of "original" pointing at the Original Resource (URI-R), with a

Relation Type of "timegate" pointing at a TimeGate (URI-G) for the Original Resource, and with a Relation Type of "timemap" pointing at a TimeMap (URI-T) for the Original Resource. The state that is expressed by the response is the state the Original Resource had at the archival datetime expressed in the "Memento-Datetime" header.

Figure 2 depicts a specific pattern to implement the Memento framework. Multiple patterns exist and they can be grouped as follows:

- o Pattern 1 ([Section 4.1](#)) - The Original Resource acts as its own TimeGate
- o Pattern 2 ([Section 4.2](#)) - A remote resource acts as a TimeGate for the Original Resource
- o Pattern 3 ([Section 4.3](#)) - The Original Resource is a Fixed Resource
- o Pattern 4 ([Section 4.4](#)) - Mementos without a TimeGate

Details of the HTTP interactions for common cases for each of those patterns are provided in [Section 4.1](#) through [Section 4.4](#). Special cases are described in [Section 4.5](#).

Figure 3 shows a user agent that attempts to datetime negotiate with the Original Resource `http://a.example.org/` by including an "Accept-Datetime" header in its HTTP HEAD request. This initiating request is the same for Pattern 1 ([Section 4.1](#)) through Pattern 3 ([Section 4.3](#)).

```
HEAD / HTTP/1.1
Host: a.example.org
Accept-Datetime: Tue, 11 Sep 2001 20:35:00 GMT
Connection: close
```

Figure 3: User Agent Attempts Datetime Negotiation With Original Resource

[4.1](#). Pattern 1 - The Original Resource acts as its own TimeGate

In this implementation pattern, the Original Resource acts as its own TimeGate, which means that URI-R and URI-G coincide. Content management systems and revision control systems can support datetime negotiation in this way as they are commonly aware of the version history of their own resources.

The response to this request when datetime negotiation for this resource is supported depends on the negotiation style it uses (200-style or 302-style) and on the existence or absence of a URI-M for Mementos that is distinct from the URI-R of the associated Original Resource. The various cases are summarized in the below table and the server responses for each are detailed in the remainder of this section.

Pattern	Original Resource	TimeGate	Memento	Negotiation Style
Pattern 1.1 (Section 4.1.1)	URI-R	URI-R	URI-M	302
Pattern 1.2 (Section 4.1.2)	URI-R	URI-R	URI-M	200
Pattern 1.3 (Section 4.1.3)	URI-R	URI-R	URI-R	200

Table 2: Pattern 1

[4.1.1.1](#). Pattern 1.1 - URI-R=URI-G ; 302-style negotiation ; distinct URI-M for Mementos

In this case, the response to the user agent's request of Figure 3 has a "302 Found" HTTP status code, and the "Location" header conveys the URI-M of the selected Memento. As per [Section 2.3](#), the use of Memento response headers and links in the response from URI-R=URI-G is as follows:

- o The "Vary" header MUST be provided and it MUST include the "accept-datetime" value.
- o The response MUST NOT contain a "Memento-Datetime" header.
- o The "Link" header MUST be provided and it MUST contain at least a link with the "original" Relation Type that has the URI-R of the Original Resource as Target IRI. The provision of other links is encouraged and is subject to the considerations described in [Section 2.2](#).

The server's response to the request of Figure 3 is shown in Figure 4. Note the inclusion of the recommended link to the TimeGate that, in this case, has a Target IRI that is the URI-R of the Original Resource.


```
HTTP/1.1 302 Found
Date: Thu, 21 Jan 2010 00:06:50 GMT
Server: Apache
Vary: accept-datetime
Location:
  http://a.example.org/?version=20010911203610
Link: <http://a.example.org/>; rel="original timegate"
Content-Length: 0
Content-Type: text/plain; charset=UTF-8
Connection: close
```

Figure 4: Response from URI-R=URI-G for Pattern 1.1

In a subsequent request, shown in Figure 5, the user agent can obtain the selected Memento by issuing an HTTP GET request against the URI-M that was provided in the "Location" header. The inclusion of the "Accept-Datetime" header in this request is not needed but will typically occur as the user agent is in datetime negotiation mode.

```
GET /?version=20010911203610 HTTP/1.1
Host: a.example.org
Accept-Datetime: Tue, 11 Sep 2001 20:35:00 GMT
Connection: close
```

Figure 5: User Agent Requests Selected Memento

The response has a "200 OK" HTTP status code and the entity-body of the response contains the representation of the selected Memento. As per [Section 2.3](#), the use of Memento response headers and links in the response from URI-M is as follows:

- o A "Vary" header that includes an "accept-datetime" value MUST NOT be provided.
- o The response MUST include a "Memento-Datetime" header. Its value expresses the archival datetime of the Memento.
- o The "Link" header MUST be provided and it MUST contain at least a link with the "original" Relation Type that has the URI-R of the Original Resource as Target IRI. The provision of other links is encouraged and is subject to the considerations described in [Section 2.2](#).

The server's response to the request of Figure 5 is shown in Figure 6. Note the provision of the required "original", and the recommended "timegate" and "timemap" links. The former two point to the Original Resource, which acts as its own TimeGate. The latter has "from" and "until" attributes to indicate the temporal interval

covered by Mementos listed in the linked TimeMap.

```
HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:06:51 GMT
Server: Apache-Coyote/1.1
Memento-Datetime: Tue, 11 Sep 2001 20:36:10 GMT
Link: <http://a.example.org/>; rel="original timegate",
      <http://a.example.org/?version=all&style=timemap>
      ; rel="timemap"; type="application/link-format"
      ; from="Tue, 15 Sep 2000 11:28:26 GMT"
      ; until="Wed, 20 Jan 2010 09:34:33 GMT"
Content-Length: 23364
Content-Type: text/html; charset=utf-8
Connection: close
```

Figure 6: Response from URI-M for Pattern 1.1

4.1.2. Pattern 1.2 - URI-R=URI-G ; 200-style negotiation ; distinct URI-M for Mementos

In this case, the response to the user agent's request of Figure 3 has a "200 OK" HTTP status code, and the "Content-Location" header conveys the URI-M of the selected Memento. As per [Section 2.3](#), the use of Memento response headers and links in the response from URI-R=URI-G is as follows:

- o The "Vary" header MUST be provided and it MUST include the "accept-datetime" value.
- o The response MUST include a "Memento-Datetime" header. Its value expresses the archival datetime of the selected Memento.
- o The "Link" header MUST be provided and it MUST contain at least a link with the "original" Relation Type that has the URI-R of the Original Resource as Target IRI. The provision of other links is encouraged and is subject to the considerations described in [Section 2.2](#).

The server's response to the request of Figure 3 is shown in Figure 7. Note the provision of optional "memento" links pointing at the oldest and most recent Memento for the Original Resource known to the responding server.


```
HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:06:50 GMT
Server: Apache
Vary: accept-datetime
Content-Location:
  http://a.example.org/?version=20010911203610
Memento-Datetime: Tue, 11 Sep 2001 20:36:10 GMT
Link: <http://a.example.org/>; rel="original timegate",
  <http://a.example.org/?version=20000915112826>
  ; rel="memento first"; datetime="Tue, 15 Sep 2000 11:28:26 GMT",
  <http://a.example.org/?version=20100120093433>
  ; rel="memento last"; datetime="Wed, 20 Jan 2010 09:34:33 GMT"
Content-Length: 2312
Content-Type: text/html; charset=UTF-8
Connection: close
```

Figure 7: Response from URI-R=URI-G for Pattern 1.2

In a subsequent request, which is the same as Figure 3 but with HTTP GET instead of HEAD, the user agent can obtain the representation of the selected Memento. It will be provided as the entity-body of a response that has the same Memento headers as in Figure 7.

4.1.3. Pattern 1.3 - URI-R=URI-G ; 200-style negotiation ; no distinct URI-M for Mementos

In this case, the response to the user agent's request of Figure 3 has a "200 OK" HTTP status code, and it does not contain a "Content-Location" nor "Location" header as there is no URI-M of the selected Memento to convey. As per [Section 2.3](#), the use of Memento response headers and links in the response from URI-R=URI-G is as follows:

- o The "Vary" header MUST be provided and it MUST include the "accept-datetime" value.
- o The response MUST include a "Memento-Datetime" header. Its value expresses the archival datetime of the selected Memento.
- o The "Link" header MUST be provided and it MUST contain at least a link with the "original" Relation Type that has the URI-R of the Original Resource as Target IRI. The provision of other links is encouraged and is subject to the considerations described in [Section 2.2](#).

The server's response to the request of Figure 3 is shown in Figure 8. The recommended "timemap" and "timegate" links are included in addition to the mandatory "original" link.


```

HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:06:50 GMT
Server: Apache
Vary: accept-datetime
Memento-Datetime: Tue, 11 Sep 2001 20:36:10 GMT
Link: <http://a.example.org/>; rel="original timegate",
      <http://a.example.org/?version=all&style=timemap>
      ; rel="timemap"; type="application/link-format"
Content-Length: 0
Content-Type: text/plain; charset=UTF-8
Connection: close

```

Figure 8: Response from URI-R=URI-G for Pattern 1.3

In a subsequent request, which is the same as Figure 3 but with HTTP GET instead of HEAD, the user agent can obtain the representation of the selected Memento. It will be provided as the entity-body of a response that has the same Memento headers as in Figure 8.

4.2. Pattern 2 - A remote resource acts as a TimeGate for the Original Resource

In this implementation pattern, the Original Resource does not act as its own TimeGate, which means that URI-R and URI-G are different. This pattern is typically implemented by servers for which the history of their resources is recorded in remote systems such as web archives and transactional archives [[Fitch](#)]. But servers that maintain their own history, such as content management systems and Version Control Systems, may also implement this pattern, for example, to distribute the load involved in responding to requests for current and prior representations of resources between different servers.

This pattern is summarized in the below table and is detailed in the remainder of this section. Three cases exist that differ regarding the negotiation style that is used by the remote TimeGate and regarding the existence of a URI-M for Mementos that is distinct from the URI-G of the TimeGate.

Pattern	Original Resource	TimeGate	Memento	Negotiation Style
Pattern 2.1 (Section 4.2.1)	URI-R	URI-G	URI-M	302
Pattern 2.2 (Section 4.2.2)	URI-R	URI-G	URI-M	200

	Pattern 2.3		URI-R		URI-G		URI-G		200	
	(Section 4.2.3)									
+-----+-----+-----+-----+-----+										

Table 3: Pattern 2

The response by the Original Resource to the request shown in Figure 3 is the same for all three cases. As per [Section 2.3](#), the use of headers and links in the response from URI-R is as follows:

- o A "Vary" header that includes an "accept-datetime" value MUST NOT be provided.
- o The response MUST NOT contain a "Memento-Datetime" header.
- o The "Link" header SHOULD be provided. It MUST NOT include a link with an "original" Relation Type. It SHOULD include a link with a "timegate" Relation Type that has the URI-G of the TimeGate as Target IRI. It SHOULD include a link with a "timemap" Relation Type that has the URI-T of the TimeGate as Target IRI. Multiple "timegate" and "timemap" links can be provided to accommodate situations in which the server is aware of multiple TimeGates or Timemaps for the Original Resource.

Figure 9 shows such a response. Note the absence of an "original" link as the responding resource is neither a TimeGate or a Memento.

```
HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:02:12 GMT
Server: Apache
Link: <http://arxiv.example.net/timegate/http://a.example.org/>
    ; rel="timegate"
Content-Length: 255
Connection: close
Content-Type: text/html; charset=iso-8859-1
```

Figure 9: Response from URI-R<>URI-G for Pattern 2

Once a user agent has obtained the URI-G of a remote TimeGate for the Original Resource it can engage in datetime negotiation with that TimeGate. Figure 10 shows the request issued against the TimeGate whereas [Section 4.2.1](#) through [Section 4.2.3](#) detail the responses for various TimeGate implementation patterns.


```
HEAD /timegate/http://a.example.org/ HTTP/1.1
Host: arxiv.example.net
Accept-Datetime: Tue, 11 Sep 2001 20:35:00 GMT
Connection: close
```

Figure 10: User Agent Engages in Datetime Negotiation With Remote TimeGate

[4.2.1](#). Pattern 2.1 - URI-R<>URI-G ; 302-style negotiation ; distinct URI-M for Mementos

In case the TimeGate uses a 302 negotiation style, the response to the user agent's request of Figure 10 has a "302 Found" HTTP status code, and the "Location" header conveys the URI-M of the selected Memento. As per [Section 2.3](#), the use of Memento response headers and links in the response from URI-G is as follows:

- o The "Vary" header MUST be provided and it MUST include the "accept-datetime" value.
- o The response MUST NOT contain a "Memento-Datetime" header.
- o The "Link" header MUST be provided and it MUST contain at least a link with the "original" Relation Type that has the URI-R of the Original Resource as Target IRI. The provision of other links is encouraged and is subject to the considerations described in [Section 2.2](#).

The server's response to the request of Figure 10 is shown in Figure 11. It contains the mandatory "original" link that points back to the Original Resource associated with this TimeGate and it shows the recommended "timemap" link that includes "from" and "until" attributes.


```
HTTP/1.1 302 Found
Date: Thu, 21 Jan 2010 00:02:14 GMT
Server: Apache
Vary: accept-datetime
Location:
  http://arxiv.example.net/web/20010911203610/http://a.example.org/
Link: <http://a.example.org/>; rel="original",
  <http://arxiv.example.net/timemap/http://a.example.org/>
  ; rel="timemap"; type="application/link-format"
  ; from="Tue, 15 Sep 2000 11:28:26 GMT"
  ; until="Wed, 20 Jan 2010 09:34:33 GMT"
Content-Length: 0
Content-Type: text/plain; charset=UTF-8
Connection: close
```

Figure 11: Response from URI-G<>URI-R for Pattern 2.1

In a subsequent HTTP GET request, shown in Figure 12, the user agent can obtain the selected Memento by issuing an HTTP GET request against the URI-M that was provided in the "Location" header. The inclusion of the "Accept-Datetime" header in this request is not needed but will typically occur as the user agent is in datetime negotiation mode.

```
GET /web/20010911203610/http://a.example.org/ HTTP/1.1
Host: arxiv.example.net/
Accept-Datetime: Tue, 11 Sep 2001 20:35:00 GMT
Connection: close
```

Figure 12: User Agent Requests Selected Memento

The response has a "200 OK" HTTP status code. As per [Section 2.3](#), the use of Memento response headers and links in the response from URI-M is as follows:

- o A "Vary" header that includes an "accept-datetime" value MUST NOT be provided.
- o The response MUST include a "Memento-Datetime" header. Its value expresses the archival datetime of the Memento.
- o The "Link" header MUST be provided and it MUST contain at least a link with the "original" Relation Type that has the URI-R of the Original Resource as Target IRI. The provision of other links is encouraged and is subject to the considerations described in [Section 2.2](#).

The server's response to the request of Figure 12 is shown in

Figure 13. Note the provision of the recommended "timegate" and "timemap" links.

```
HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:02:15 GMT
Server: Apache-Coyote/1.1
Memento-Datetime: Tue, 11 Sep 2001 20:36:10 GMT
Link: <http://a.example.org/>; rel="original",
      <http://arxiv.example.net/timemap/http://a.example.org/>
      ; rel="timemap"; type="application/link-format",
      <http://arxiv.example.net/timegate/http://a.example.org/>
      ; rel="timegate"
Content-Length: 23364
Content-Type: text/html;charset=utf-8
Connection: close
```

Figure 13: Response from URI-M for Pattern 2.1

[4.2.2](#). Pattern 2.2 - URI-R<>URI-G ; 200-style negotiation ; distinct URI-M for Mementos

In case the TimeGate uses a 200 negotiation style, and each Memento has a distinct URI-M, the response to the user agent's request of Figure 10 has a "200 OK" HTTP status code, and the "Content-Location" header conveys the URI-M of the selected Memento. As per [Section 2.3](#), the use of Memento response headers and links in the response from URI-G is as follows:

- o The "Vary" header MUST be provided and it MUST include the "accept-datetime" value.
- o The response MUST include a "Memento-Datetime" header. Its value expresses the archival datetime of the Memento.
- o The "Link" header MUST be provided and it MUST contain at least a link with the "original" Relation Type that has the URI-R of the Original Resource as Target IRI. The provision of other links is encouraged and is subject to the considerations described in [Section 2.2](#).

The server's response to the request of Figure 10 is shown in Figure 14.


```
HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:09:40 GMT
Server: Apache-Coyote/1.1
Vary: accept-datetime
Content-Location:
  http://arxiv.example.net/web/20010911203610/http://a.example.org/
Memento-Datetime: Tue, 11 Sep 2001 20:36:10 GMT
Link: <http://a.example.org/>; rel="original",
      <http://arxiv.example.net/timemap/http://a.example.org/>
      ; rel="timemap"; type="application/link-format",
      <http://arxiv.example.net/timegate/http://a.example.org/>
      ; rel="timegate"
Content-Length: 23364
Content-Type: text/html; charset=UTF-8
Connection: close
```

Figure 14: Response from URI-G<>URI-R for Pattern 2.2

In a subsequent request, which is the same as Figure 10 but with HTTP GET instead of HEAD, the user agent can obtain the representation of the selected Memento. It will be provided as the entity-body of a response that has the same Memento headers as Figure 14.

4.2.3. Pattern 2.3 - URI-R<>URI-G ; 200-style negotiation ; no distinct URI-M for Mementos

In case the TimeGate uses a 200 negotiation style, but Mementos have no distinct URIs, the response to the user agent's request of Figure 10 has a "200 OK" HTTP status code, and it does not contain a "Content-Location" nor "Location" header as there is no URI-M of the selected Memento to convey. As per [Section 2.3](#), the use of Memento response headers and links in the response from URI-G is as follows:

- o The "Vary" header MUST be provided and it MUST include the "accept-datetime" value.
- o The response MUST include a "Memento-Datetime" header. Its value expresses the archival datetime of the Memento.
- o The "Link" header MUST be provided and it MUST contain at least a link with the "original" Relation Type that has the URI-R of the Original Resource as Target IRI. The provision of other links is encouraged and is subject to the considerations described in [Section 2.2](#).

The server's response to the request of Figure 10 is shown in Figure 15.


```

HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:09:40 GMT
Server: Apache-Coyote/1.1
Vary: accept-datetime
Memento-Datetime: Tue, 11 Sep 2001 20:36:10 GMT
Link: <http://a.example.org/>; rel="original",
      <http://arxiv.example.net/timemap/http://a.example.org/>
      ; rel="timemap"; type="application/link-format",
      <http://arxiv.example.net/timegate/http://a.example.org/>
      ; rel="timegate"
Content-Length: 23364
Content-Type: text/html; charset=UTF-8
Connection: close

```

Figure 15: Response from URI-G<>URI-R for Pattern 2.3

In a subsequent request, which is the same as Figure 10 but with HTTP GET instead of HEAD, the user agent can obtain the representation of the selected Memento. It will be provided as the entity-body of a response that has the same Memento headers as Figure 15.

4.3. Pattern 3 - The Original Resource is a Fixed Resource

This pattern does not involve datetime negotiation with a TimeGate but it can be implemented for Original Resources that never change state or do not change anymore past a certain point in their existence, meaning that URI-R and URI-M coincide either from the outset or starting at some point in time. This pattern is summarized in the below table. Examples are tweets or stable media resources on news sites.

Pattern	Original Resource	TimeGate	Memento	Negotiation Style
3	URI-R	-	URI-R	-

Table 4: Pattern 3

Servers that host such resources can support the Memento framework by treating the stable resource (FixedResource as per [W3C.gen-ont-20090420]) as a Memento. As per Section 2.3, the use of Memento response headers and links in responses from such a stable resource is as follows:

- o A "Vary" header that includes an "accept-datetime" value MUST NOT be provided.
- o The response MUST include a "Memento-Datetime" header. Its value expresses the datetime at which the resource became stable. Providing this value includes a promise that the resource has not changed since this datetime and will not change anymore beyond it.
- o The "Link" header MUST be provided and MUST have a link with the "original" Relation Type that has the URI-R of the stable resource itself as Target IRI.

Figure 16 shows a response to an HTTP HEAD request for the resource with URI-R `http://a.example.org/` that has been stable since March 20th 2009.

```
HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:09:40 GMT
Server: Apache-Coyote/1.1
Memento-Datetime: Fri, 20 Mar 2009 11:00:00 GMT
Link: <http://a.example.org/>; rel="original"
Content-Length: 0
Content-Type: text/plain; charset=UTF-8
Connection: close
```

Figure 16: Response from URI-R=URI-M for Pattern 3

4.4. Pattern 4 - Mementos without a TimeGate

Cases may occur in which a server hosts Mementos but does not expose a TimeGate for them. This can, for example, be the case if the server's Mementos result from taking a snapshot of the state of a set of Original Resources from another server as it is being retired. As a result, only a single Memento per Original Resource is hosted, making the introduction of a TimeGate unnecessary. But it may also be the case for servers that host multiple Mementos for an Original Resource but consider exposing TimeGates too expensive. In this case, URI-R and URI-M are distinct, but a TimeGate is absent. This case is summarized in the below table.

Pattern	Original Resource	TimeGate	Memento	Negotiation Style
4	URI-R	-	URI-M	-

Table 5: Pattern 4

Servers that host such Mementos without TimeGates can still support the Memento framework by providing the appropriate Memento headers and links. As per [Section 2.3](#), these are as follows for a response from URI-M:

- o A "Vary" header that includes an "accept-datetime" value MUST NOT be provided.
- o The response MUST include a "Memento-Datetime" header. Its value expresses the archival datetime of the Memento.
- o The "Link" header MUST be provided and it MUST have a link with the "original" Relation Type that has the URI-R of the associated Original Resource as Target IRI. The provision of other links is encouraged and is subject to the considerations described in [Section 2.2](#).

Figure 17 shows a response to an HTTP HEAD request for the Memento with URI-M

`http://arxiv.example.net/web/20010911203610/http://a.example.org/.`

Note the use of links: three links have the URI-M of the Memento as Target IRI and have respective Relation Types "memento", "first", and "last". This combination indicates that this is the only Memento for the Original Resource with Target IRI provided by the "original" link (`http://a.example.org/`) that the server is aware of. Note also that such a response does not imply that there is no server whatsoever that exposes a TimeGate; it merely means that the responding server neither provides nor is aware of the location of a TimeGate.

`HTTP/1.1 200 OK`

`Date: Thu, 21 Jan 2010 00:09:40 GMT`

`Server: Apache-Coyote/1.1`

`Memento-Datetime: Tue, 11 Sep 2001 20:36:10 GMT`

`Link: <http://a.example.org/>; rel="original",
 <http://arxiv.example.net/web/20010911203610/http://a.example.org/>
 ; rel="first last memento"
 ; datetime="Tue, 11 Sep 2001 20:36:10 GMT"`

`Content-Length: 0`

`Content-Type: text/plain; charset=UTF-8`

`Connection: close`

Figure 17: Response from URI-M<>URI-R for Pattern 4

4.5. Special Cases

4.5.1. Original Resource provides no "timegate" link

Cases exist in which the response from the Original Resource does not contain a "timegate" link, including:

- o The Original Resource's server does not support the Memento framework;
- o The Original Resource no longer exists and the responding server is not aware of its prior existence;
- o The server that hosted the Original Resource no longer exists.

In all these cases, the user agent SHOULD attempt to determine an appropriate TimeGate for the Original Resource, either automatically or interactively supported by the user.

4.5.2. Server exists but Original Resource no longer does

If possible, a server SHOULD also provide a "timegate" link in responses to requests for an Original Resource that the server knows used to exist, but no longer provides a current representation. This may allow access to Mementos for the Original Resource even if it no longer exists. A server's response to a request for the discontinued resource `http://a.example.org/pic` is illustrated in Figure 18.

```
HTTP/1.1 404 Not Found
Date: Thu, 21 Jan 2010 00:02:12 GMT
Server: Apache
Link:
  <http://arxiv.example.net/timegate/http://a.example.org/pic>
  ; rel="timegate"
Content-Length: 255
Connection: close
Content-Type: text/html; charset=iso-8909-1
```

Figure 18: Response from an Original Resource that not longer exists

4.5.3. Issues with Accept-Datetime

The following special cases may occur regarding the "Accept-Datetime" header when a user agent issues a request against a TimeGate:

- o If the value of the "Accept-Datetime" is either earlier than the datetime of the first Memento or later than the datetime of the most recent Memento known to the TimeGate, the first or most recent Memento SHOULD be selected, respectively.
- o If the value of the "Accept-Datetime" does not conform to the [rfc1123](#)-date construction rule of the BNF in Figure 1, the response MUST have a "400 Bad Request" HTTP status code. The use of Memento headers is as per [Section 2.3](#).
- o If a user agent issues a request against a TimeGate and fails to include an "Accept-Datetime" request header, the most recent Memento SHOULD be selected.

[4.5.4](#). Memento of a 3XX response

Cases exist in which HTTP responses with 3XX status codes are archived. For example, crawl-based web archives commonly archive responses with HTTP status codes "301 Moved Permanently" and "302 Found" whereas Linked Data archives hold on to "303 See Other" responses.

If the Memento requested by the user agent is an archived version of an HTTP response with a 3XX status code, the server's response MUST have the same 3XX HTTP status code. The use of other Memento headers is as described in the Memento column of [Section 2.3](#).

The user agent's handling of an HTTP response with a 3XX status code is not affected by the presence of a "Memento-Datetime" header. The user agent SHOULD behave in the same manner as it does with HTTP responses with a 3XX status code that do not have a "Memento-Datetime" header.

However, the user agent MUST be aware that the URI that was selected from the "Location" header of an HTTP response with a 3XX status code might not be that of a Memento but rather of an Original Resource. In the latter case it SHOULD proceed by looking for a Memento of the selected Original Resource.

For example, Figure 19 shows the response to an HTTP GET request for `http://a.example.org` issued on April 11 2008. This response is archived as a Memento of `http://a.example.org` that has as URI-M `http://arxiv.example.net/web/20080411000650/http://a.example.org`. The response to an HTTP GET on this URI-M is shown in Figure 20. It is a replay of the original response with "Memento-Datetime" and "Link" headers added, to allow a user agent to understand the response is a Memento. In Figure 20, the value of the "Location" header is the same as in the original response; it identifies an

Original Resource. The user agent proceeds with finding a Memento for this Original Resource. Web archives sometimes overwrite the value that was originally provided in the "Location" header in order to point at a Memento they hold of the resource to which the redirect originally led. This is shown in Figure 21. In this case, the user agent may decide it found an appropriate Memento.

```
HTTP/1.1 301 Moved Permanently
Date: Fri, 11 Apr 2008 00:06:50 GMT
Server: Apache
Location: http://b.example.org
Content-Length: 0
Content-Type: text/plain; charset=UTF-8
Connection: close
```

Figure 19: Response is a redirect

```
HTTP/1.1 301 Moved Permanently
Date: Thu, 21 Jan 2010 00:09:40 GMT
Server: Apache-Coyote/1.1
Memento-Datetime: Fri, 11 Apr 2008 00:06:50 GMT
Location: http://b.example.org
Link: <http://a.example.org>; rel="original",
      <http://arxiv.example.net/timemap/http://a.example.org>
      ; rel="timemap"; type="application/link-format",
      <http://arxiv.example.net/timegate/http://a.example.org>
      ; rel="timegate"
Content-Length: 0
Content-Type: text/plain; charset=UTF-8
Connection: close
```

Figure 20: Response is a Memento of a redirect; leads to an Original Resource


```
HTTP/1.1 301 Moved Permanently
Date: Thu, 21 Jan 2010 00:09:40 GMT
Server: Apache-Coyote/1.1
Memento-Datetime: Fri, 11 Apr 2008 00:06:50 GMT
Location:
  http://arxiv.example.net/web/20080411000655/http://b.example.org
Link: <http://a.example.org>; rel="original",
  <http://arxiv.example.net/timemap/http://a.example.org>
  ; rel="timemap"; type="application/link-format",
  <http://arxiv.example.net/timegate/http://a.example.org>
  ; rel="timegate"
Content-Length: 0
Content-Type: text/plain; charset=UTF-8
Connection: close
```

Figure 21: Response is a Memento of a redirect; leads to a Memento

[4.5.5](#). Memento of responses with 4XX or 5XX HTTP status codes

Cases exist in which responses with 4XX and 5XX HTTP status codes are archived. If the Memento requested by the user agent is an archived version of such an HTTP response, the server's response MUST have the same 4XX or 5XX HTTP status code. The use of other Memento headers is as described in the Memento column of [Section 2.3](#).

For example, Figure 22 shows the 404 response to an HTTP GET request for `http://a.example.org` issued on April 11 2008. This response is archived as a Memento of `http://a.example.org`, that has as URI-M `http://arxiv.example.net/web/20080411000650/http://a.example.org`. The response to an HTTP HEAD on this URI-M is shown in Figure 23. It is a replay of the original response with "Memento-Datetime" and "Link" headers added, to allow a user agent to understand the response is a Memento.

```
HTTP/1.1 404 Not Found
Date: Fri, 11 Apr 2008 00:06:50 GMT
Server: Apache
Content-Length: 0
Content-Type: text/plain; charset=UTF-8
Connection: close
```

Figure 22: Response is a 404


```
HTTP/1.1 404 Not Found
Date: Thu, 21 Jan 2010 00:09:40 GMT
Server: Apache-Coyote/1.1
Memento-Datetime: Fri, 11 Apr 2008 00:06:50 GMT
Link: <http://a.example.org>; rel="original",
      <http://arxiv.example.net/timemap/http://a.example.org>
      ; rel="timemap"; type="application/link-format",
      <http://arxiv.example.net/timegate/http://a.example.org>
      ; rel="timegate"
Content-Length: 0
Content-Type: text/plain; charset=UTF-8
Connection: close
```

Figure 23: Response is a Memento of a 404

4.5.6. Sticky "Memento-Datetime" value for Mementos

The provision of a "Memento-Datetime" in a response entails a promise that the response is frozen in time. As a consequence, the "Memento-Datetime" header associated with a Memento MUST be "sticky" in the following ways:

- o The server that originally assigns the "Memento-Datetime" header and value to a specific response MUST retain that header in all future responses to HTTP requests (with or without "Accept-Datetime" header) that occur against the Memento after the time of the original assignment of the header, and it MUST NOT change its associated value.
- o Applications that mirror Mementos at a different URI MUST retain the "Memento-Datetime" header and MUST NOT change its value unless mirroring involves a meaningful state change. This allows, for example, duplicating a web archive at a new location while preserving the value of the "Memento-Datetime" header of the archived resources. In this example, the "Last-Modified" header will be updated to reflect the time of mirroring at the new URI, whereas the value for "Memento-Datetime" will be maintained.

4.5.7. Intermediate Resources

An intermediate resource is a resource that issues a redirect to a TimeGate, to a Memento, or to another intermediate resource, and thus plays an active role in the Memento infrastructure. Intermediate resources commonly exist in web archives on the path from a TimeGate to an appropriate Memento.

A response of an intermediate resource has an HTTP status code indicative of HTTP redirection (e.g. 302) and uses Memento headers

and links that allow to recognize that the resource plays a role in the Memento framework:

- o A "Vary" header that includes an "accept-datetime" value MUST NOT be provided.
- o The response MUST NOT include a "Memento-Datetime" header.
- o The "Link" header MUST be provided and it MUST have a link with the "original" Relation Type that has the URI-R of the associated Original Resource as Target IRI. Links with "timegate", "timemap", and "memento" Relation Types are OPTIONAL and, if provided, MUST pertain to the Original Resource for which the user agent is trying to obtain a Memento.

A user agent SHOULD follow a redirection provided by an intermediate resource; multiple such redirections can be chained.

Consider the case where a user agent follows the "timegate" link provided in Figure 9 and engages in datetime negotiation with the assumed TimeGate in the manner shown in Figure 10. But instead of receiving a response as shown in Figure 11, it receives the one shown below in Figure 24. Such a response is unambiguously recognizable as coming from an intermediate resource.

```
HTTP/1.1 302 Found
Date: Thu, 21 Jan 2010 00:06:50 GMT
Server: Apache
Location:
  http://arxiv.example.net/new-timegate/http://a.example.org/
Link: <http://a.example.org>; rel="original"
Content-Length: 0
Content-Type: text/plain; charset=UTF-8
Connection: close
```

Figure 24: Redirecting Resource redirects to a TimeGate

5. TimeMaps

A TimeMap is introduced to support retrieving a comprehensive list of all Mementos for a specific Original Resource known to a server. The entity-body of a response to an HTTP GET request issued against a TimeMap's URI-T:

- o MUST list the URI-R of the Original Resource that the TimeMap is about;

- o MUST list the URI-M and archival datetime of each Memento for the Original Resource known to the server, preferably in a single document, or, alternatively in multiple documents that can be gathered by following contained links with a "timemap" Relation Type;
- o SHOULD list the URI-G of one or more TimeGates for the Original Resource known to the responding server;
- o SHOULD, for self-containment, list the URI-T of the TimeMap itself;
- o MUST unambiguously type listed resources as being Original Resource, TimeGate, Memento, or TimeMap.

The entity-body of a response from a TimeMap MAY be serialized in various ways, but the link-value format serialization described here MUST be supported. In this serialization, the entity-body MUST be formatted in the same way as the value of an HTTP "Link" header, and hence MUST comply to the "link-value" construction rule of "[Section 5](#). The Link Header Field" of [[RFC5988](#)], and the media type of the entity-body MUST be "application/link-format" as introduced in [[RFC6690](#)]. Links contained in the entity-body MUST be interpreted as follows:

- o The Context IRI is set to the anchor parameter, when specified;
- o The Context IRI of links with the "self" Relation Types is the URI-T of the TimeMap, i.e. the URI of the resource from which the TimeMap was requested;
- o The Context IRI of all other links is the URI-R of the Original Resource, which is provided as the Target IRI of the link with an "original" Relation Type.

In order to retrieve the link-value serialization of a TimeMap, a user agent SHOULD use an "Accept" request header with a value set to "application/link-format". This is shown in Figure 25.

```
GET /timemap/http://a.example.org/ HTTP/1.1
Host: arxiv.example.net
Accept: application/link-format;q=1.0
Connection: close
```

Figure 25: Request for a TimeMap

If the TimeMap requested by the user agent exists, the server's response has a "200 OK" HTTP status code and the list of Mementos is

provided in the entity-body of the response. Such a response is shown in Figure 26

```
HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:06:50 GMT
Server: Apache
Content-Length: 4883
Content-Type: application/link-format
Connection: close

<http://a.example.org>;rel="original",
<http://arxiv.example.net/timemap/http://a.example.org>
  ; rel="self";type="application/link-format"
  ; from="Tue, 20 Jun 2000 18:02:59 GMT"
  ; until="Wed, 09 Apr 2008 20:30:51 GMT",
<http://arxiv.example.net/timegate/http://a.example.org>
  ; rel="timegate",
<http://arxiv.example.net/web/20000620180259/http://a.example.org>
  ; rel="first memento";datetime="Tue, 20 Jun 2000 18:02:59 GMT"
  ; license="http://creativecommons.org/publicdomain/zero/1.0/",
<http://arxiv.example.net/web/20091027204954/http://a.example.org>
  ; rel="last memento";datetime="Tue, 27 Oct 2009 20:49:54 GMT"
  ; license="http://creativecommons.org/publicdomain/zero/1.0/",
<http://arxiv.example.net/web/20000621011731/http://a.example.org>
  ; rel="memento";datetime="Wed, 21 Jun 2000 01:17:31 GMT"
  ; license="http://creativecommons.org/publicdomain/zero/1.0/",
<http://arxiv.example.net/web/20000621044156/http://a.example.org>
  ; rel="memento";datetime="Wed, 21 Jun 2000 04:41:56 GMT"
  ; license="http://creativecommons.org/publicdomain/zero/1.0/",
...
```

Figure 26: Response from a TimeMap

5.1. Index and Paging TimeMaps

Cases exist in which a TimeMap points at one or more other TimeMaps:

- o Index Timemap - A TimeMap can merely point at other TimeMaps and not list any Mementos itself. This can happen when Mementos are spread across several archives that share a front-end. An example is shown in Figure 27.
- o Paging Timemap - The number of available Mementos can require introducing multiple TimeMaps that can be paged. An example is shown in Figure 28. Note that a Paging TimeMap contains links to other TimeMaps but actually also lists Mementos.

In both cases, including the "from" and "until" attributes for

"timemap" links is RECOMMENDED as a means to express the temporal span of Mementos listed in each TimeMap. Note that TimeMaps obtained by following a "timemap" link can contain links to further TimeMaps.

```
<http://a.example.org>;rel="original",
  <http://arxiv.example.net/timegate/http://a.example.org>
    ; rel="timegate",
  <http://arxiv.example.net/timemap/http://a.example.org>
    ; rel="self";type="application/link-format",
  <http://arxiv1.example.net/timemap/http://a.example.org>
    ; rel="timemap";type="application/link-format"
    ; from="Wed, 21 Jun 2000 04:41:56 GMT"
    ; until="Wed, 09 Apr 2008 20:30:51 GMT",
  <http://arxiv2.example.net/timemap/http://a.example.org>
    ; rel="timemap";type="application/link-format"
    ; from="Thu, 10 Apr 2008 20:30:51 GMT"
    ; until="Tue, 27 Oct 2009 20:49:54 GMT",
  <http://arxiv3.example.net/timemap/http://a.example.org>
    ; rel="timemap";type="application/link-format"
    ; from="Thu, 29 Oct 2009 20:30:51 GMT"
```

Figure 27: Index TimeMap

```
<http://a.example.org>;rel="original",
  <http://arxiv.example.net/timegate/http://a.example.org>
    ; rel="timegate",
  <http://arxiv.example.net/timemap/1/http://a.example.org>
    ; rel="self";type="application/link-format"
    ; from="Tue, 20 Jun 2000 18:02:59 GMT"
    ; until="Wed, 09 Apr 2008 20:30:51 GMT",
  <http://arxiv.example.net/timemap/2/http://a.example.org>
    ; rel="timemap";type="application/link-format"
    ; from="Thu, 10 Apr 2008 20:30:51 GMT"
    ; until="Tue, 27 Oct 2009 20:49:54 GMT",
  <http://arxiv.example.net/timemap/3/http://a.example.org>
    ; rel="timemap";type="application/link-format"
    ; from="Thu, 29 Oct 2009 20:30:51 GMT"
    ; until="Fri, 31 Aug 2012 12:22:34 GMT"
  <http://arxiv.example.net/web/20000620180259/http://a.example.org>
    ; rel="memento";datetime="Tue, 20 Jun 2000 18:02:59 GMT",
  <http://arxiv.example.net/web/20000621011731/http://a.example.org>
    ; rel="memento";datetime="Wed, 21 Jun 2000 01:17:31 GMT",
  <http://arxiv.example.net/web/20000621044156/http://a.example.org>
    ; rel="memento";datetime="Wed, 21 Jun 2000 04:41:56 GMT",
  ...
```

Figure 28: Paging TimeMap

5.2. Mementos for TimeMaps

A TimeMap can itself act as an Original Resource for which a TimeGate and Mementos may exist. Hence, as per [Section 2.3](#), the response from a TimeMap could include a "timegate" link to a TimeGate via which prior TimeMap versions are available. And, in cases where URI-T=URI-R=URI-G (a TimeMap is an Original Resource that acts as its own TimeGate), an "original" link pointing at the TimeMap URI-T would be included.

Therefore, caution is required in cases where a TimeMap for an Original Resource wants to explicitly express in a Link header for which Original Resource it is a TimeMap. It can do so by including a "timemap" link that has the URI-R of the Original Resource as Context IRI and the URI-T of the TimeMap as Target IRI.

Figure 29 shows the response to an HTTP HEAD request against a TimeMap that has `http://arxiv.example.net/timemap/http://a.example.org` as URI-T. This TimeMap provides information about Mementos for the Original Resource that has `http://a.example.org` as URI-R. The response includes an "original" link pointing to the Original Resource that this TimeMap is about. Note the use of the "anchor" attribute in this link to convey the URI-R of that Original Resource.

```
HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:06:50 GMT
Server: Apache
Link: <http://arxiv.example.net/timemap/http://a.example.org>
      ; anchor="http://a.example.org"; rel="timemap"
      ; type="application/link-format"
Content-Length: 0
Content-Type: application/link-format; charset=UTF-8
Connection: close
```

Figure 29: TimeMap links to the Original Resource it is about

6. IANA Considerations

This memo requires IANA to register the Accept-Datetime and Memento-Datetime HTTP headers defined in [Section 2.1.1](#) in the appropriate IANA registry.

This memo requires IANA to register the Relation Types "original", "timegate", "timemap", and "memento" defined in [Section 2.2](#) in the appropriate IANA registry.

This memo requires IANA to register the "datetime" and "license" attributes for the "memento" Relation Type, as defined in [Section 2.2.4](#), in the appropriate IANA registry.

This memo requires IANA to register the "from" and "until" attributes for the "timemap" Relation Type, as defined in [Section 2.2.4](#), in the appropriate IANA registry.

7. Security Considerations

Provision of a "timegate" HTTP "Link" header in responses to requests for an Original Resource that is protected (e.g., 401 or 403 HTTP response codes) is OPTIONAL. The inclusion of this Link when requesting authentication is at the server's discretion; cases may exist in which a server protects the current state of a resource, but supports open access to prior states and thus chooses to supply a "timegate" HTTP "Link" header. Conversely, the server may choose to not advertise the TimeGate URIs (e.g., they exist in an intranet archive) for unauthenticated requests.

The veracity of archives and the relationships between Original Resources and Mementos is beyond the scope of this document. Even in the absence of malice, it is possible for separate archives to have different Mementos for the same Original Resource at the same datetime if the state of the Original Resource was dependent on the requesting archive's user agent IP address, specific HTTP request headers, and possibly other factors.

Further authentication, encryption and other security related issues are otherwise orthogonal to Memento.

8. Changelog

v06 2013-02-14 HVDS MLN RS [draft-vandesompel-memento-06](#)

- o Major overhaul of the presentation of the specification.
- o Specification of patterns whereby URI-R=URI-G and with both 200 and 302 negotiation style.
- o Removal of Discovery section to increase focus on datetime negotiation aspects.

v05 2012-09-01 HVDS MLN RS [draft-vandesompel-memento-05](#)

- o Clarified the section on Memento Relation Types.
- o Re-introduced "license" attribute for "memento" Relation Type as it will become essential for IIPC.
- o Introduced from and until attributes for "timemap" links to accomodate paged TimeMap cases.
- o Introduced the notion of Redirecting Resource and inserted related information in various sections.
- o Added discovery of Mementos via host-meta.
- o Corrected ambiguous uses of the term "representation".

v04 2012-05-18 HVDS MLN RS [draft-vandesompel-memento-04](#)

- o Removed the possibility to use an interval indicator in an Accept-Datetime header as no one is implementing it.
- o Corrected typo in Other Relation Types table.
- o Added TimeMap examples to illustrate index of TimeMaps and TimeMap paging.
- o Changed Discovery component from using robots.txt with Memento-specific add-ons to well-known URI and host-meta.
- o Removed "embargo" and "license" attributes for links with a "memento" Relation Type because no one is using them.

v04 2011-12-20 HVDS MLN RS [draft-vandesompel-memento-03](#)

- o Added description of Mementos of HTTP responses with 3XX, 4XX and 5XX status code.
- o Clarified that a TimeGate must not use the "Memento-Datetime" header.
- o Added wording to warn for possible cache problems with Memento implementations that choose to have an Original Resource and and its TimeGate coincide.

v03 2011-05-11 HVDS MLN RS [draft-vandesompel-memento-02](#)

- o Added scenario in which a TimeGate redirects to another TimeGate.

- o Reorganized TimeGate section to better reflect the difference between requests with and without interval indicator.
- o Added recommendation to provide "memento" links to Mementos in the vicinity of the preferred interval provided by the user agent, in case of a 406 response.
- o Removed TimeMap Feed material from the Discovery section as a result of discussions regarding (lack of) scalability of the approach with representatives of the International Internet Preservation Consortium. An alternative approach to support batch discovery of Mementos will be specified.

v02 2011-04-28 HVDS MLN RS [draft-vandesompel-memento-01](#)

- o Introduced wording and reference to indicate a Memento is a FixedResource.
- o Introduced "Sticky Memento-Datetime" notion and clarified wording about retaining "Memento-Datetime" headers and values when a Memento is mirrored at different URI.
- o Introduced section about handling both datetime and regular negotiation.
- o Introduced section about Mementos Without TimeGate.
- o Made various changes in the section Relation Type "memento", including addition of "license" and "embargo" attributes, and clarification of rules regarding the use of "memento" links.
- o Moved section about TimeMaps inside the Datetime Negotiation section, and updated it.
- o Restarted the Discovery section from scratch.

v01 2010-11-11 HVDS MLN RS First public version
[draft-vandesompel-memento-00](#)

v00 2010-10-19 HVDS MLN RS Limited circulation version

2010-07-22 HVDS MLN First internal version

9. Acknowledgements

The Memento effort is funded by the Library of Congress. Many thanks to Kris Carpenter Negulescu, Michael Hausenblas, Erik Hetzner, Larry

Masinter, Gordon Mohr, Mark Nottingham, David Rosenthal, Ed Summers, James Anderson, Tim Starling, Martin Klein for feedback. Many thanks to Samuel Adams, Scott Ainsworth, Lyudmilla Balakireva, Frank McCown, Harihar Shankar, Brad Tofel for early implementations.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.
- [RFC4151] Kindberg, T. and S. Hawke, "The 'tag' URI Scheme", [RFC 4151](#), October 2005.
- [RFC4287] Nottingham, M., Ed. and R. Sayre, Ed., "The Atom Syndication Format", [RFC 4287](#), December 2005.
- [RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", [RFC 5785](#), April 2010.
- [RFC5829] Brown, A., Clemm, G., and J. Reschke, "Link Relation Types for Simple Version Navigation between Web Resources", [RFC 5829](#), April 2010.
- [RFC5988] Nottingham, M., "Web Linking", [RFC 5988](#), October 2010.
- [RFC6415] Hammer-Lahav, E. and B. Cook, "Web Host Metadata", [RFC 6415](#), October 2011.
- [RFC6690] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", [RFC 6690](#), August 2012.

10.2. Informative References

- [Fitch] Fitch, "Web site archiving - an approach to recording every materially different response produced by a website", July 2003,
<<http://ausweb.scu.edu.au/aw03/papers/fitch/paper.html>>.
- [I-D.masinter-dated-uri] Masinter, L., "The 'tdb' and 'duri' URI schemes, based on

dated URIs", [draft-masinter-dated-uri-10](#) (work in progress), January 2012.

[RFC1123] Braden, R., "Requirements for Internet Hosts - Application and Support", STD 3, [RFC 1123](#), October 1989.

[W3C.REC-aww-20041215]
Jacobs and Walsh, "Architecture of the World Wide Web",
December 2004, <<http://www.w3.org/TR/webarch/>>.

[W3C.gen-ont-20090420]
Berners-Lee, "Architecture of the World Wide Web",
April 2009, <<http://www.w3.org/2006/gen/ont>>.

Authors' Addresses

Herbert VandeSompel
Los Alamos National Laboratory
PO Box 1663
Los Alamos, New Mexico 87545
USA

Phone: +1 505 667 1267
Email: hvdsomp@gmail.com
URI: <http://public.lanl.gov/herbertv/>

Michael Nelson
Old Dominion University
Norfolk, Virginia 23529
USA

Phone: +1 757 683 6393
Email: mln@cs.odu.edu
URI: <http://www.cs.odu.edu/~mln/>

Robert Sanderson
Los Alamos National Laboratory
PO Box 1663
Los Alamos, New Mexico 87545
USA

Phone: +1 505 665 5804
Email: azaroth42@gmail.com
URI: <http://public.lanl.gov/rsanderson/>

