

Internet Engineering Task Force	H. VandeSompel	
Internet-Draft	Los Alamos National Laboratory	
Intended status: Informational	M. Nelson	
Expires: May 16, 2011	Old Dominion University	
	R. Sanderson	
	Los Alamos National Laboratory	
	November 12, 2010	

[TOC](#)

## **HTTP framework for time-based access to resource states -- Memento draft-vandesompel-memento-00**

### **Abstract**

The HTTP-based Memento framework bridges the present and past Web by interlinking current resources with resources that encapsulate their past. It facilitates obtaining representations of prior states of a resource, available from archival resources in Web archives or version resources in content management systems, by leveraging the resource's URI and a preferred datetime. To this end, the framework introduces datetime negotiation (a variation on content negotiation), and new Relation Types for the HTTP Link header aimed at interlinking resources with their archival/version resources. It also introduces an approach to discover and serialize a list of resources known to a server, each of which provides access to a representation of a prior state of a same resource.

### **Status of this Memo**

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of 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 May 16, 2011.

---

## Table of Contents

- [1. Introduction](#)
  - [1.1. Terminology](#)
  - [1.2. Purpose](#)
  - [1.3. Notational Conventions](#)
- [2. The Memento Framework, Datetime Negotiation component: HTTP headers, HTTP Link Relation Types](#)
  - [2.1. HTTP Headers](#)
    - [2.1.1. Accept-Datetime, Memento-Datetime](#)
      - [2.1.1.1. Values for Accept-Datetime](#)
      - [2.1.1.2. Values for Memento-Datetime](#)
    - [2.1.2. Vary](#)
    - [2.1.3. Location](#)
    - [2.1.4. Link](#)
  - [2.2. Link Header Relation Types](#)
    - [2.2.1. Memento Framework Relation Types](#)
      - [2.2.1.1. Relation Type "original"](#)
      - [2.2.1.2. Relation Type "timegate"](#)
      - [2.2.1.3. Relation Type "timemap"](#)
      - [2.2.1.4. Relation Type "memento"](#)
    - [2.2.2. Other Relation Types](#)
- [3. The Memento Framework, Datetime Negotiation component: HTTP Interactions](#)
  - [3.1. Interactions with an Original Resource](#)
    - [3.1.1. Step 1: User Agent Requests an Original Resource](#)
    - [3.1.2. Step 2: Server Responds to a Request for an Original Resource](#)
      - [3.1.2.1. Original Resource is an Appropriate Memento](#)
      - [3.1.2.2. Server Exists and Original Resource Used to Exist](#)
      - [3.1.2.3. Missing or Inadequate "timegate" Link in Original Server's Response](#)
  - [3.2. Interactions with a TimeGate](#)
    - [3.2.1. Step 3: User Agent Negotiates with a TimeGate](#)
    - [3.2.2. Step 4: Server Responds to Negotiation with TimeGate](#)
      - [3.2.2.1. Successful Scenario](#)
      - [3.2.2.2. Datetime Out of the Server's Range](#)
      - [3.2.2.3. Accept-Datetime Not Provided](#)
      - [3.2.2.4. Multiple Matching Mementos](#)
      - [3.2.2.5. Datetime Out of the User Agent's Range](#)
      - [3.2.2.6. Accept-Datetime Unparseable](#)
      - [3.2.2.7. TimeGate Does Not Exist](#)
      - [3.2.2.8. HTTP Methods other than HEAD/GET](#)

<a href="#">3.2.3.</a>	Recognizing a TimeGate
<a href="#">3.3.</a>	Interactions with a Memento
<a href="#">3.3.1.</a>	Step 5: User Agent Requests a Memento
<a href="#">3.3.2.</a>	Step 6: Server Responds to a Request for a Memento
<a href="#">3.3.2.1.</a>	Memento Does not Exist
<a href="#">3.3.3.</a>	Recognizing a Memento
<a href="#">4.</a>	The Memento Framework, Discovery Component
<a href="#">4.1.</a>	TimeMaps
<a href="#">4.2.</a>	Discovery of TimeMaps, TimeGates
<a href="#">5.</a>	IANA Considerations
<a href="#">6.</a>	Security Considerations
<a href="#">7.</a>	Changelog
<a href="#">8.</a>	Acknowledgements
<a href="#">9.</a>	References
<a href="#">9.1.</a>	Normative References
<a href="#">9.2.</a>	Informative References
<a href="#">Appendix A.</a>	Appendix B: A Sample, Successful Memento Request/Response cycle
<a href="#">§</a>	Authors' Addresses
<a href="#">§</a>	Intellectual Property and Copyright Statements

---

## 1. Introduction

[TOC](#)

---

### 1.1. Terminology

[TOC](#)

This specification uses the terms "resource", "request", "response", "entity", "entity-body", "entity-header", "content negotiation", "client", "user agent", "server" as described in [RFC 2616 \(Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," June 1999.\)](#) [RFC2616], and it uses the terms "representation" and "resource state" as described in [W3C.REC-aww-20041215 \(Jacobs and Walsh, "Architecture of the World Wide Web," December 2004.\)](#) [W3C.REC-aww-20041215]. In addition, the following terms specific to the Memento framework are introduced:

\*Original Resource: An Original Resource is a resource that exists or used to exist, and for which access to one of its prior states is desired.

\*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  $T_j$  is a resource that encapsulates the state that the Original Resource had at time  $T_j$ .

\*TimeGate: A TimeGate for an Original Resource is a resource that supports negotiation to allow selective, datetime-based, access to prior states of the Original Resource.

\*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. Purpose

[TOC](#)

The state of an Original Resource may change over time. Dereferencing its URI at any specific moment in time during its existence yields a representation of its then current state. Dereferencing its URI at any time past its existence no longer yields a meaningful representation, if any. Still, in both cases, resources may exist that encapsulate prior states of the Original Resource. Each such resource, named a Memento, has its own URI that, when dereferenced, returns a representation of a prior state of the Original Resource. Mementos may, for example, exist in Web archives, Content Management Systems, or Revision Control Systems.

Examples are:

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

\*<http://web.archive.org/web/19970107171109/http://www.ietf.org/>

\*<http://webarchive.nationalarchives.gov.uk/20080906200044/http://www.ietf.org/>

Mementos for Original Resource [http://en.wikipedia.org/wiki/Hypertext\\_Transfer\\_Protocol](http://en.wikipedia.org/wiki/Hypertext_Transfer_Protocol) :

\*[http://en.wikipedia.org/w/index.php?title=Hypertext\\_Transfer\\_Protocol&oldid=366806574](http://en.wikipedia.org/w/index.php?title=Hypertext_Transfer_Protocol&oldid=366806574)

\*[http://en.wikipedia.org/w/index.php?title=Hypertext\\_Transfer\\_Protocol&oldid=33912](http://en.wikipedia.org/w/index.php?title=Hypertext_Transfer_Protocol&oldid=33912)

\*[http://web.archive.org/web/20071011153017/http://en.wikipedia.org/wiki/Hypertext\\_Transfer\\_Protocol](http://web.archive.org/web/20071011153017/http://en.wikipedia.org/wiki/Hypertext_Transfer_Protocol)

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

\*<http://www.w3.org/TR/2004/PR-webarch-20041105/>

\*<http://www.w3.org/TR/2002/WD-webarch-20020830/>

\*<http://webarchive.nationalarchives.gov.uk/20100304163140/http://www.w3.org/TR/webarch/>

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

- \*Is fully distributed in the sense that resource versions may reside on multiple hosts, and that any such host is likely only aware of the versions it holds;
- \*Uses the global notion of datetime as a resource version indicator and access key;
- \*Leverages the following primitives of [W3C.REC-aww-20041215 \(Jacobs and Walsh, "Architecture of the World Wide Web," December 2004.\)](#) [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 a resource identified by URI-R as it existed at some time Tj. Note the relationship with the ability to identify the state of a resource at some datetime Tj by means of a URI as intended by the proposed Dated URI scheme [I-D.masinter-dated-uri \(Masinter, L., "The 'tdb' and 'duri' URI schemes, based on dated URIs," October 2010.\)](#) [I-D.masinter-dated-uri].
2. A bridge from the present to the past, consisting of:

- \*An appropriately typed link from a resource identified by URI-R to an associated TimeGate identified by URI-G, which is aware of (at least part of the) version history of the resource identified by URI-R;
- \*The ability to content negotiate in the datetime dimension with the TimeGate identified by URI-G, as a means to obtain a representation of the state that the resource identified by URI-R had at some datetime Tj.

3. A bridge from the past to the present, consisting of appropriately typed link from a resource identified by URI-M, which encapsulates the state a resource identified by URI-R had at some datetime Tj, to the resource identified by URI-R.

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

---

### 1.3. Notational Conventions

[TOC](#)

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 [RFC 2119 \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#) [RFC2119].

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

- \*URI-R is used to denote the URI of an Original Resource.

- \*URI-G is used to denote the URI of a TimeGate.

- \*URI-M is used to denote the URI of a Memento.

- \*URI-T is used to denote the URI of a TimeMap.

- \*When scenarios are described that involve multiple Mementos, URI-M<sub>0</sub> denotes the URI of the first Memento known to the responding server, URI-M<sub>n</sub> denotes the URI of the most recent known Memento, URI-M<sub>j</sub> denotes the URI of the selected Memento, URI-M<sub>i</sub> denotes the URI of the Memento that is temporally previous to the selected Memento, and URI-M<sub>k</sub> denotes the URI of the Memento that is temporally after the selected Memento. The respective datetimes for these Mementos is T<sub>0</sub>, T<sub>n</sub>, T<sub>j</sub>, T<sub>i</sub>, and T<sub>k</sub>; it holds that T<sub>0</sub> ≤ T<sub>i</sub> ≤ T<sub>j</sub> ≤ T<sub>k</sub> ≤ T<sub>n</sub>.

---

## 2. The Memento Framework, Datetime Negotiation component: HTTP headers, HTTP Link Relation Types

[TOC](#)

The Memento framework is concerned 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 HTTPS resources.

---

### 2.1. HTTP Headers

[TOC](#)

The Memento framework operates at the level of HTTP request and response headers. It introduces two new headers ("Accept-Datetime", "Memento-Datetime"), introduces new values for two existing headers ("Vary", "Link"), and uses an existing header ("Location") without modification. All these headers are described below. Other HTTP headers

are present or absent in Memento response/request cycles as specified by [RFC 2616 \(Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," June 1999.\)](#) [RFC2616].

---

#### **2.1.1. Accept-Datetime, Memento-Datetime**

[TOC](#)

The "Accept-Datetime" request header is used by a user agent to indicate it wants to retrieve a representation of a Memento that encapsulates a past state of an Original Resource. To that end, the "Accept-Datetime" header is conveyed in an HTTP GET/HEAD request issued against a TimeGate for an Original Resource, and its value indicates the datetime of the desired past state of the Original Resource. The "Accept-Datetime" request header has no defined meaning for HTTP methods other than HEAD and GET.

The "Memento-Datetime" response header is used by a server to indicate that the response contains a representation of a Memento, and its value expresses the datetime of the state of an Original Resource that is encapsulated in that Memento. The URI of that Original Resource is provided in the response, as the Target IRI (see [RFC5988 \(Nottingham, M., "Web Linking," October 2010.\)](#) [RFC5988]) of a link provided in the HTTP "Link" header that has a Relation Type of "original" (see [Section 2.2 \(Link Header Relation Types\)](#)).

The presence of a Memento-Datetime header and associated value for a given resource constitutes a promise that the resource is stable and that its state will no longer change. Therefore, the server that originally assigns the header and value, MUST retain the Memento-Datetime header in all responses to HTTP HEAD/GET requests (with or without "Accept-Datetime" header) that occur against the resource after the time of the original assignment of the header, and it MUST NOT change its associated value. Similarly, if an application is mirroring the resource at a different URI, it SHOULD retain the resource's Memento-Datetime header and value if mirroring the resource does not include a meaningful change to the resource's state. For example, this behavior allows duplicating a Web archive at a new location while preserving the Memento-Datetime values of the archived resources.

---

##### **2.1.1.1. Values for Accept-Datetime**

[TOC](#)

Values for the "Accept-Datetime" header consist of a MANDATORY datetime expressed according to the [RFC 1123 \(Braden, R., "Requirements for Internet Hosts - Application and Support," October 1989.\)](#) [RFC1123] format, which is formalized by the rfc1123-date construction rule of the BNF in [Figure 1 \(BNF for the datetime format\)](#), and an OPTIONAL

interval indicator expressed according to the iso8601-interval rule of the BNF in [Figure 1 \(BNF for the datetime format\)](#). The datetime MUST be represented in Greenwich Mean Time (GMT).  
Examples of "Accept-Datetime" request headers with and without an interval indicator:

```
Accept-Datetime: Thu, 31 May 2007 20:35:00 GMT
Accept-Datetime: Thu, 31 May 2007 20:35:00 GMT; -P3DT5H;+P2DT6H
```

The user agent uses the MANDATORY datetime value to convey its preferred datetime for a Memento; it uses the OPTIONAL interval indicator to convey it is interested in retrieving Mementos that reside within this interval around the preferred datetime, and not interested in Mementos that reside outside of it. Not using an interval indicator is equivalent with expressing an infinite interval around the preferred datetime.

The interval mechanism can be regarded as an implementation of the functionality intended by the q-value approach that is used in regular content negotiation. The q-value approach is not supported for Memento's datetime negotiation because it is well-suited for negotiation over a discrete space of mostly predictable values, not for negotiation over a continuum of unpredictable datetime values.

---

```
accept-dt-value = rfc1123-date *SP [ iso8601-interval ]
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"
iso8601-interval = ";" *SP "-" duration *SP ";" *SP "+" duration
duration = "P" ( dur-date | dur-week )
dur-date = ( dur-day | dur-month | dur-year ) [ dur-time ]
dur-year = 1*DIGIT "Y" [ dur-month ] [ dur-day ]
dur-month = 1*DIGIT "M" [ dur-day ]
dur-day = 1*DIGIT "D"
dur-time = "T" ( dur-hour | dur-minute | dur-second )
dur-hour = 1*DIGIT "H" [ dur-minute ] [ dur-second ]
dur-minute = 1*DIGIT "M" [ dur-second ]
dur-second = 1*DIGIT "S"
dur-week = 1*DIGIT "W"
```

**Figure 1: BNF for the datetime format**

---

#### 2.1.1.2. Values for Memento-Datetime

[TOC](#)

Values for the "Memento-Datetime" headers MUST be datetimes expressed according to the rfc1123-date construction rule of the BNF in [Figure 1 \(BNF for the datetime format\)](#); they MUST be represented in Greenwich Mean Time (GMT).

An example "Memento-Datetime" response header:

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

---

#### 2.1.2. Vary

[TOC](#)

The "Vary" response header is used in responses to indicate the dimensions in which content negotiation was successfully applied. This header is used in the Memento framework to indicate both whether datetime negotiation was applied or is supported by the responding server.

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

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

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

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

---

#### 2.1.3. Location

[TOC](#)

The "Location" header is used as defined in [RFC 2616 \(Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," June 1999.\)](#) [RFC2616]. Examples are given in [Section 3 \(The Memento Framework, Datetime Negotiation component: HTTP Interactions\)](#) below.

---

#### 2.1.4. Link

[TOC](#)

The "Link" response header is specified in [RFC5988 \(Nottingham, M., "Web Linking," October 2010.\)](#) [RFC5988]. The Memento framework introduces new Relation Types to convey typed links among Original Resources, TimeGates, Mementos, and TimeMaps. Already existing Relation Types, among others, aimed at supporting navigation among a series of ordered resources may also be used in the Memento framework. This is detailed in [Link Header Relation Types \(Link Header Relation Types\)](#), below.

---

### 2.2. Link Header Relation Types

[TOC](#)

The "Link" header specified in [RFC5988 \(Nottingham, M., "Web Linking," October 2010.\)](#) [RFC5988] is semantically equivalent to the "<LINK>" element in HTML, as well as the "atom:link" feed-level element in Atom [RFC 4287 \(Nottingham, M., Ed. and R. Sayre, Ed., "The Atom Syndication Format," December 2005.\)](#) [RFC4287]. By default, the origin of a link expressed by an entry in a "Link" header (named Context IRI in [RFC5988 \(Nottingham, M., "Web Linking," October 2010.\)](#) [RFC5988]) is the IRI of the requested resource.

---

#### 2.2.1. Memento Framework Relation Types

[TOC](#)

The Relation Types used in the Memento framework are listed in the remainder of this section, and their use is summarized in the below table. [Appendix A \(Appendix B: A Sample, Successful Memento Request/Response cycle\)](#) shows a Memento request/response cycle that uses all the Relation Types that are introduced here.

---

Relation Type	Original Resource	TimeGate	Memento
original	NA, except see <a href="#">Section 3.1.2.1 (Original Resource is an Appropriate Memento)</a>	REQUIRED, 1	REQUIRED, 1
timegate	RECOMMENDED, 0 or more	NA	RECOMMENDED, 0 or more
timemap	NA	RECOMMENDED, 0 or 1	RECOMMENDED, 0 or more

memento	NA, except see <a href="#">Section 3.1.2.1 (Original Resource is an Appropriate Memento)</a>	REQUIRED, 1 or more	REQUIRED, 1 or more
---------	--	---------------------	---------------------

**Table 1: The use of Relation Types**

For several of the Relation Types introduced in the Memento framework, the use of a "datetime" attribute is REQUIRED. The value for this attribute MUST be a datetime expressed according to the [RFC 1123 \(Braden, R., "Requirements for Internet Hosts - Application and Support," October 1989.\)](#) [RFC1123] format, which is formalized by the rfc1123-date construction rule of the BNF in [Figure 1 \(BNF for the datetime format\)](#); it MUST be represented in Greenwich Mean Time (GMT).

#### 2.2.1.1. Relation Type "original"

[TOC](#)

"original" -- A "Link" header entry with a Relation Type of "original" is used to point from a TimeGate or a Memento to their associated Original Resource. In all cases, an entry with the "original" Relation Type MUST occur exactly once in a Link header. Details for the entry are as follows:

\*Context IRI: URI-G, URI-Mj

\*Target IRI: URI-R

\*Relation Type: "original"

\*Use: REQUIRED

\*Cardinality: 1

#### 2.2.1.2. Relation Type "timegate"

[TOC](#)

"timegate" -- A "Link" header entry with a Relation Type of "timegate" is used to point both from an Original Resource or a Memento to a TimeGate for the Original Resource. In both cases, the use of an entry with the "timegate" Relation Type is RECOMMENDED. Since more than one TimeGate can exist for any Original Resource, multiple entries with a "timegate" Relation Type MAY occur, each with a distinct Target IRI. Details for the entry are as follows:

\*Context IRI: URI-R or URI-Mj

\*Target IRI: URI-G

\*Relation Type: "timegate"

\*Use: RECOMMENDED

\*Cardinality: 0 or more

---

#### 2.2.1.3. Relation Type "timemap"

[TOC](#)

"timemap" -- A "Link" header entry with a Relation Type of "timemap" is used to point from both a TimeGate or a Memento to a TimeMap resource from which a list of Mementos known to the responding server is available. Use of an entry with the "timemap" Relation Type is RECOMMENDED, and if used it MUST occur exactly once. This link MUST include a "type" attribute and its value MUST be "application/link-format", referring to the MIME type introduced in [I-D.ietf-core-link-format \(Shelby, Z., "CoRE Link Format," October 2010.\)](#)

[I-D.ietf-core-link-format]. Details for the entry are as follows:

\*Context IRI: URI-G or URI-Mi

\*Target IRI: URI-T

\*Relation Type: "timemap"

\*Target Attribute: "type"

\*Use: RECOMMENDED

\*Cardinality: 0 or more

---

#### 2.2.1.4. Relation Type "memento"

[TOC](#)

"memento" -- A "Link" header entry with a Relation Type of "memento" is used to point from both a TimeGate and a Memento to various Mementos for an Original Resource. This link 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. Use of entries with the "memento" Relation Type is REQUIRED and it MUST be as follows:  
For all responses to HTTP HEAD/GET requests issued against an existing TimeGate or Memento:

- \*One "memento" link MUST be included that has as Target IRI the URI of the temporally first Memento known to the responding server;
- \*One "memento" link MUST be included that has as Target IRI the URI of the temporally most recent Memento known to the responding server.

For all responses to HTTP HEAD/GET requests issued against a TimeGate or a Memento, in which a Memento is selected or served by the responding server:

- \*One "memento" link MUST be included that has as Target IRI the URI of the Memento that was selected or served;
- \*One "memento" link SHOULD be included that has as Target IRI the URI of the Memento that is previous to the selected Memento in the temporal series of all Mementos (sorted by ascending Memento-Datetime values) known to the server;
- \*One "memento" link SHOULD be included that has as Target IRI the URI the Memento that is next to the selected Memento in the temporal series of all Mementos (sorted by ascending Memento-Datetime values) known to the server.
- \*Other "memento" links MAY only be included if both the previous and next links are provided. Each of these OPTIONAL "memento" links MUST have as Target IRI the URI of a Memento other than the ones listed above.

Note that the Target IRI of some of these links may coincide. For example, if the selected Memento actually is the first Memento known to the server, only three distinct "memento" links may result. The value for the "datetime" attribute of these links would be the datetimes of the first (equal to selected), next, and most recent Memento known to the responding server.

The summary is as follows:

- \*Context IRI: URI-G, URI-Mj
- \*Target IRI: URI-M
- \*Relation Type: "memento"
- \*Target Attribute: "datetime"

\*Use: REQUIRED

\*Cardinality: 1 or more

---

### 2.2.2. Other Relation Types

[TOC](#)

Web Linking [RFC5988 \(Nottingham, M., "Web Linking," October 2010.\)](#) [RFC5988] allows for the inclusion of links with different Relation Types but the same Target IRI, and hence the Relation Types introduced by the Memento framework MAY be combined with others as deemed necessary. As the "memento" Relation Type focuses on conveying the datetime of a linked Memento, Relation Types that allow navigating among the temporally ordered series of Mementos known to a server are of particular importance. With this regard, the Relation Types listed in the below table SHOULD be considered for combination with the "memento" Relation Type. A distinction is made between responding servers that can be categorized as systems that are the focus of [RFC5829 \(Brown, A., Clemm, G., and J. Reschke, "Link Relation Types for Simple Version Navigation between Web Resources," April 2010.\)](#) [RFC5829] (such as version control systems) and others that can not (such as Web archives). Note that, in terms of [RFC5829 \(Brown, A., Clemm, G., and J. Reschke, "Link Relation Types for Simple Version Navigation between Web Resources," April 2010.\)](#) [RFC5829], the last Memento (URI-Mn) is the version prior to the latest (i.e. current) version.

---

Memento Type	RFC5988 system	non RFC5988 system
First Memento (URI-M0)	first	first
Last Memento (URI-Mn)	last	last
Selected Memento (URI-Mj)	NA	NA
Memento prior to selected Memento (URI-Mi)	predecessor-version	prev
Memento next to selected Memento (URI-Mk)	successor-version	next

**Table 2: The use of Relation Types**

---

---

### 3. The Memento Framework, Datetime Negotiation component: HTTP Interactions

[TOC](#)

This section describes the HTTP interactions of the Memento framework for a variety of scenarios. First, [Figure 2 \(Typical Memento request/response chain\)](#) provides a schematic overview of a successful request/response chain that involves datetime negotiation. Dashed lines depict HTTP transactions between user agent and server. [Appendix A \(Appendix B: A Sample, Successful Memento Request/Response cycle\)](#) shows these HTTP interactions in detail for the case where the Original Resource resides on one server, whereas both the TimeGate and the Mementos reside on another. Scenarios also exist in which all these resources are on the same server (for example, Content Management Systems) or on different servers (for example, an aggregator of TimeGates). Note that, in Step 2 and Step 6, the HTTP status code of the response is shown as "200 OK", but a series of "206 Partial Content" could be substituted without loss of generality.

---

```
1: UA --- HTTP GET/HEAD; Accept-Datetime: Tj -----> URI-R
2: UA <-- HTTP 200; Link: URI-G ----- URI-R
3: UA --- HTTP GET/HEAD; Accept-Datetime: Tj -----> URI-G
4: UA <-- HTTP 302; Location: URI-Mj; Vary; Link:
    URI-R,URI-T,URI-M0,URI-Mn,URI-Mi,URI-Mj,URI-Mk ----- URI-G
5: UA --- HTTP GET URI-Mj; Accept-Datetime: Tj -----> URI-Mj
6: UA <-- HTTP 200; Memento-Datetime: Tj; Link:
    URI-R,URI-T,URI-G,URI-M0,URI-Mn,URI-Mi,URI-Mj,URI-Mk -- URI-Mj
```

**Figure 2: Typical Memento request/response chain**

---

\*Step 1: In order to determine what the URI is of a TimeGate for an Original Resource, the user agent issues an HTTP HEAD/GET request against the URI of the Original Resource (URI-R).

\*Step 2: The entity-header of 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.

\*Step 3: The user agent starts the datetime negotiation process with the TimeGate by issuing an HTTP GET request against its URI-G thereby including an "Accept-Datetime" HTTP header with a value of the datetime of the desired prior state of the Original Resource.

\*Step 4: The entity-header of the response from URI-G includes a "Location" header pointing at the URI of a Memento (URI-Mj) for the Original Resource. In addition, the entity-header contains an HTTP "Link" header with a Relation Type of "original" pointing at the Original Resource, and an HTTP "Link" header with a Relation Type of "timemap" pointing at a TimeMap (URI-T). Also HTTP Links pointing at various Mementos are provided using the "memento" Relation Type, as specified in [Section 2.2.1.4 \(Relation Type "memento"\)](#).

\*Step 5: The user agent issues an HTTP GET request against the URI-Mj of a Memento, obtained in Step 4.

\*Step 6: The entity-header of the response from URI-Mj includes a "Memento-Datetime" HTTP header with a value of the datetime of the Memento. It also contains an HTTP "Link" header with a Relation Type of "original" pointing at the Original Resource, with a Relation Type of "timegate" pointing at a TimeGate associated with the Original Resource, and with a Relation Type of "timemap" pointing at a TimeMap. The state that is expressed by the representation provided in the response is the state the Original Resource had at the datetime expressed in the "Memento-Datetime" header. This response also includes HTTP Links with a "memento" Relation Type pointing at various Mementos, as specified in [Section 2.2.1.4 \(Relation Type "memento"\)](#).

The following sections detail the specifics of HTTP interactions with Original Resources, TimeGates and Mementos under various conditions.

---

### 3.1. Interactions with an Original Resource

[TOC](#)

This section details HTTP GET/HEAD requests targeted at an Original Resource (URI-R).

---

#### 3.1.1. Step 1: User Agent Requests an Original Resource

[TOC](#)

In order to try and discover a TimeGate for the Original Resource, the user agent MAY issue an HTTP HEAD or GET request against the Original Resource's URI. Use of the "Accept-Datetime" header in the HTTP HEAD/GET request is OPTIONAL.

[Figure 3 \(User Agent Requests Original Resource\)](#) shows the use of HTTP HEAD indicating the user agent is not interested in retrieving a representation of the Original Resource, but only in determining a TimeGate for it. It also shows the use of the "Accept-Datetime" header

anticipating that the user agent will set it for the entire duration of a Memento request/response cycle.

---

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

**Figure 3: User Agent Requests Original Resource**

---

### 3.1.2. Step 2: Server Responds to a Request for an Original Resource

[TOC](#)

The response of the Original Resource's server to the user agent's HTTP HEAD/GET request of Step 1, for the case where the Original Resource exists, is as it would be in a regular HTTP request/response cycle, but in addition MAY include a HTTP "Link" header with a Relation Type of "timegate" that conveys the URI of the Original Resource's TimeGate as the Target IRI of the Link. Multiple HTTP Links with a relation type of "timegate" MAY be provided to accomodate situations in which the server is aware of multiple TimeGates for an Original Resource. The actual Target IRI provided in the "timegate" Link may depend on several factors including the datetime provided in the "Accept-Datetime" header, and the IP address of the user agent. A response for this case is illustrated in [Figure 4 \(Server of Original Resource Responds\)](#).

---

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

**Figure 4: Server of Original Resource Responds**

---

Servers that actively maintain archives of their resources SHOULD include the "timegate" HTTP "Link" header because this link is an important way for a user agent to discover TimeGates for those resources. This includes servers such as Content Management Systems, Control Version Systems, and Web servers with associated transactional archives [Fitch \(Fitch, "Web site archiving - an approach to recording every materially different response produced by a website," July 2003.\)](#) [Fitch]. Servers that do not actively maintain archives of their resources MAY include the "timegate" HTTP "Link" header as a way to convey a preference for TimeGates for their resources exposed by a third party archive. This includes servers that rely on Web archives such as the Internet Archive to archive their resources. The server of the Original Resource MUST treat requests with and without an "Accept-Datetime" header in the same way:

- \*The response MUST either always or never include a HTTP "link" header with an entry that has a "timegate" Relation Type and the URI of a TimeGate as the Target IRI.

- \*The entity-body of the response MUST be the same, for user agent requests with or without a "Accept-Datetime" header.

---

#### 3.1.2.1. Original Resource is an Appropriate Memento

[TOC](#)

The "Memento-Datetime" header MAY be applied to an Original Resource directly as both an indication that the state of the Original Resource has not changed since the datetime conveyed in the "Memento-Datetime" header, and as a promise that it will not change anymore beyond it. This may occur, for example, for certain stable media resources on news sites. In case the user agent's preferred datetime is equal to or more recent than the datetime conveyed as the value of Memento-Datetime in the server's response in Step 2, the user agent SHOULD conclude it has located an appropriate Memento, and it SHOULD NOT continue to Step 3. [Figure 5 \(Response to a request for an Original Resource that was created stable\)](#) illustrates such a response to a request for the resource with URI `http://a.example.org/pic` that has been stable since it was created. Note the use of both the "memento" and "original" Relation Types for links that have as Target IRI the URI of the Original Resource.

---

```
HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:02:12 GMT
Server: Apache
Link:
  <http://a.example.org/pic>
    ; rel="original memento"
    ; datetime="Fri, 20 Mar 2009 11:00:00 GMT"
Memento-Datetime: Fri, 20 Mar 2009 11:00:00 GMT
Content-Length: 255
Connection: close
Content-Type: text/html; charset=iso-8909-1
```

**Figure 5: Response to a request for an Original Resource that was created stable**

Cases may also exist in which a resource becomes stable at a certain point in its existence, but changed previously. In such cases, the Original Resource may know about a TimeGate that is aware of its prior history and hence MAY also include a link with a "timegate" Relation Type. This is illustrated in [Figure 6 \(Response to a request for an Original Resource that became stable\)](#), where the "memento" and "original" Relation Types are used as in [Figure 5 \(Response to a request for an Original Resource that was created stable\)](#), and the existence of a TimeGate to negotiate for Mementos with datetimes prior to Fri, 20 Mar 2009 11:00:00 GMT is indicated.

```
HTTP/1.1 200 OK
Date: Thu, 21 Jan 2010 00:02:12 GMT
Server: Apache
Link:
  <http://a.example.org/pic>
    ; rel="original memento"
    ; datetime="Fri, 20 Mar 2009 11:00:00 GMT",
  <http://arxiv.example.net/web/timegate/http://a.example.org/pic>
    ; rel="timegate"
Memento-Datetime: Fri, 20 Mar 2009 11:00:00 GMT
Content-Length: 255
Connection: close
Content-Type: text/html; charset=iso-8909-1
```

**Figure 6: Response to a request for an Original Resource that became stable**

---

#### 3.1.2.2. Server Exists and Original Resource Used to Exist

[TOC](#)

Servers SHOULD also provide a "timegate" HTTP "Link" header in responses to requests for an Original Resource that the server knows used to exist, but no longer does. This allows the use of an Original Resource's URI as an entry point to representations of its prior states even if the resource itself no longer exists. A server's response for this case is illustrated in [Figure 7 \(Response to a request for an Original Resource that not longer exists\)](#).

---

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

**Figure 7: Response to a request for an Original Resource that not longer exists**

---

In case the server is not aware of the prior existence of the Original Resource, its response SHOULD NOT include a "timegate" HTTP Link. [Section 3.1.2.3 \(Missing or Inadequate "timegate" Link in Original Server's Response\)](#) details what the user agent's behavior should be in such cases.

---

#### 3.1.2.3. Missing or Inadequate "timegate" Link in Original Server's Response

[TOC](#)

A user agent MAY ignore the TimeGate returned in Step 2. However, when engaging in a Memento request/response cycle, a user agent SHOULD NOT proceed immediately to Step 3 by using a TimeGate of its own preference but rather SHOULD always start the cycle by issuing an HTTP GET/HEAD against the Original Resource (Step 1, [Figure 3 \(User Agent Requests Original Resource\)](#)) as it is an important way to learn about dedicated or preferred TimeGates for the Original Resource. Also, cases exist in

which the response in Step 2 will not provide a "timegate" link, including:

- \*The Original Resource's server does not support the Memento framework;
- \*The Original Resource does no longer exist and the responding server is not aware of its prior existence;
- \*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.

---

### 3.2. Interactions with a TimeGate

[TOC](#)

This section details HTTP GET/HEAD requests targeted at a TimeGate (URI-G).

---

#### 3.2.1. Step 3: User Agent Negotiates with a TimeGate

[TOC](#)

In order to negotiate with a TimeGate, the user agent MUST issue a HTTP HEAD or GET against its URI, its request MUST include the "Accept-Datetime" header to express its datetime preference, and the use of that header MUST be as described in [Section 2.1.1.1 \(Values for Accept-Datetime\)](#). The URI of the TimeGate may have been provided as the Target IRI of a "timegate" HTTP "Link" header in the response from the Original Resource (Step 2, [Figure 4 \(Server of Original Resource Responds\)](#)), or may have resulted from another discovery mechanism, for example, based on the aggregation of TimeMaps ([Section 4.1 \(TimeMaps\)](#)) or user interaction. Such a request is illustrated in [Figure 8 \(User agent negotiates with TimeGate\)](#).

---

```
GET /web/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 8: User agent negotiates with TimeGate**

---

---

### 3.2.2. Step 4: Server Responds to Negotiation with TimeGate

[TOC](#)

In order to respond to a datetime negotiation request (Step 3, [Section 3.2.1 \(Step 3: User Agent Negotiates with a TimeGate\)](#)), the server uses an internal algorithm to select the Memento that best meets the user agent's datetime preference, and redirects to it. The exact nature of the selection algorithm is at the server's discretion but SHOULD be consistent. A variety of approaches can be used including selecting the Memento that is nearest in time (either past or future) or nearest in the past relative to the requested datetime. Special cases for datetime negotiation with a TimeGate exist, and they are addressed in [Section 3.2.2.3 \(Accept-Datetime Not Provided\)](#) through [Section 3.2.2.7 \(TimeGate Does Not Exist\)](#).

---

#### 3.2.2.1. Successful Scenario

[TOC](#)

In cases where the TimeGate exists, and the datetime provided in the user agent's "Accept-Datetime" header can be parsed and is not out of range, the server selects a Memento based on the user agent's datetime preference. The response MUST have a "302 Found" HTTP status code, and the "Location" header MUST be used to convey the URI of the selected Memento. The "Vary" header MUST be provided and it MUST include the "negotiate" and "accept-datetime" values to indicate that datetime negotiation has taken place. The "Link" header MUST be provided and contain links with Relation Types subject to the considerations described in [Section 2.2 \(Link Header Relation Types\)](#). Such a response is illustrated in [Figure 9 \(Server of TimeGate responds\)](#).

---

```
HTTP/1.1 302 Found
Date: Thu, 21 Jan 2010 00:06:50 GMT
Server: Apache
Vary: negotiate, accept-datetime
Location:
  http://arxiv.example.net/web/20010911203610/http://a.example.org
Link: <http://a.example.org>; rel="original",
  <http://arxiv.example.net/web/timemap/http://a.example.org>
  ; rel="timemap"; type="application/link-format",
  <http://arxiv.example.net/web/20000915112826/http://a.example.org>
  ; rel="first memento"; datetime="Tue, 15 Sep 2000 11:28:26 GMT",
  <http://arxiv.example.net/web/20080708093433/http://a.example.org>
  ; rel="last memento"; datetime="Tue, 08 Jul 2008 09:34:33 GMT",
  <http://arxiv.example.net/web/20010911203610/http://a.example.org>
  ; rel="memento"; datetime="Tue, 11 Sep 2001 20:36:10 GMT",
  <http://arxiv.example.net/web/20010911203610/http://a.example.org>
  ; rel="prev memento"; datetime="Tue, 11 Sep 2001 20:30:51 GMT",
  <http://arxiv.example.net/web/20010911203610/http://a.example.org>
  ; rel="next memento"; datetime="Tue, 11 Sep 2001 20:47:33 GMT"
Content-Length: 0
Content-Type: text/plain; charset=UTF-8
Connection: close
```

**Figure 9: Server of TimeGate responds**

---

Note that 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.

---

#### 3.2.2.2. Datetime Out of the Server's Range

[TOC](#)

In case, in Step 3, a user agent's "Accept-Datetime" header does not convey an interval indicator, and conveys a datetime that is either earlier than the datetime of the first Memento or later than the datetime of the most recent Memento known to the server, the server's response MUST be as described in [Section 3.2.2.1 \(Successful Scenario\)](#), with a selection of the first or most recent Memento, respectively. This is illustrated in [Figure 10 \(A TimeGate's response to a request for a Memento with a datetime earlier than that of the first Memento\)](#) that shows the response from a TimeGate exposed by a MediaWiki server to a request by a user agent that has an "Accept-Datetime: Mon, 31 May

1999 00:00:00 GMT" header. Note that a link is provided with a "successor-version" Relation Type but not with a "predecessor-version" Relation Type.

---

```
HTTP/1.1 302 Found
Server: Apache
Content-Length: 709
Content-Type: text/html; charset=utf-8
Date: Thu, 21 Jan 2010 00:09:40 GMT
Location:
  http://a.example.org/w/index.php?title=Clock&oldid=1493688
Vary: negotiate, accept-datetime
Link: <http://a.example.org/w/Clock>; rel="original",
  <http://a.example.org/Special:TimeMap/http://a.example.org/w/Clock>
  ; rel="timemap",
  <http://a.example.org/w/index.php?title=Clock&oldid=1493688>
  ; rel="first memento"; datetime="Sun, 28 Sep 2003 01:42:00 GMT",
  <http://a.example.org/w/index.php?title=Clock&oldid=1493854>
  ; rel="successor-version memento"
  ; datetime="Tue, 30 Sep 2003 14:28:00 GMT",
  <http://a.example.org/w/index.php?title=Clock&oldid=337446696>
  ; rel="last memento"; datetime="Tue, 12 Jan 2010 19:55:00 GMT"
Connection: close
```

**Figure 10: A TimeGate's response to a request for a Memento with a datetime earlier than that of the first Memento**

---

#### 3.2.2.3. Accept-Datetime Not Provided

[TOC](#)

In case, in Step 3, a user agent issues a request to a TimeGate and fails to include an "Accept-Datetime" request header, the response MUST be handled as in [Section 3.2.2.1 \(Successful Scenario\)](#), with a selection of the most recent Memento known to the responding server.

---

#### 3.2.2.4. Multiple Matching Mementos

[TOC](#)

Because the finest datetime granularity expressible using the [RFC 1123 \(Braden, R., "Requirements for Internet Hosts - Application and Support," October 1989.\)](#) [RFC1123] format used in HTTP is seconds

level, cases may occur in which a TimeGate server is aware of multiple Mementos that meet the user agent's datetime preference. This may occur in CMS with very high update rates. The response in this case MUST be handled as in [Section 3.2.2.1 \(Successful Scenario\)](#), with the selection of one of the matching Mementos.

As an example, [Figure 11 \(A TimeGate's response to a request that has multiple Mementos with a matching datetime\)](#) shows a hypothetical response from a TimeGate on a MediaWiki server to a request for a Memento for the Original Resource `http://a.example.org/w/Clock` for which two Mementos exist for the user agent's preferred datetime.

---

```
HTTP/1.1 302 Found
Server: Apache
Content-Length: 705
Content-Type: text/html; charset=utf-8
Date: Thu, 21 Jan 2010 00:09:40 GMT
Vary: negotiate, accept-datetime
Location:
  http://a.example.org/w/index.php?title=Clock&oldid=322586071
Link: <http://a.example.org/w/Clock>; rel="original",
  <http://a.example.org/Special:TimeMap/http://a.example.org/w/Clock>
    ; rel="timemap"; type="application/link-format",
  <http://a.example.org/w/index.php?title=Clock&oldid=1493688>
    ; rel="first memento"; datetime="Sun, 28 Sep 2003 01:42:00 GMT",
  <http://a.example.org/w/index.php?title=Clock&oldid=337446696>
    ; rel="last memento"; datetime="Tue, 12 Jan 2010 19:55:00 GMT",
  <http://a.example.org/w/index.php?title=Clock&oldid=322586071>
    ; rel="memento"; datetime="Sun, 31 May 2009 15:43:00 GMT",
  <http://a.example.org/w/index.php?title=Clock&oldid=326164283>
    ; rel="memento successor-version"
    ; datetime="Sun, 31 May 2009 15:43:00 GMT"
  <http://a.example.org/w/index.php?title=Clock&oldid=326164283>
    ; rel="memento predecessor-version"
    ; datetime="Sun, 31 May 2009 15:41:24 GMT"
Connection: close
```

**Figure 11: A TimeGate's response to a request that has multiple Mementos with a matching datetime**

### 3.2.2.5. Datetime Out of the User Agent's Range

In case, in Step 3, a user agent conveys an interval indicator, and the responding server is not aware of any Mementos with datetimes within the expressed interval, the server's response MUST have a "406 Not Acceptable" HTTP status code. The use of the "Vary" header MUST be as described in [Section 3.2.2.1 \(Successful Scenario\)](#). The use of the "Link" header MUST be as described in [Section 2.2 \(Link Header Relation Types\)](#). Specifically, the use of links with a "memento" Relation Type MUST follow the rules for the case where no Memento is selected by the responding server, i.e. only "memento" links to the first and most recent Mementos MUST be provided ([Section 2.2.1.4 \(Relation Type "memento"\)](#)).

[Figure 12 \(User agent expresses interval of interest in Accept-Datetime header\)](#) shows a user agent using an "Accept-Datetime" header conveying an interval of interest starting 5 hours before and ending 6 hours after Tue, 11 Sep 2001 20:35:00 GMT. [Figure 13 \(A TimeGate's response indicating it has no Mementos within the interval of interest\)](#) shows the response from the TimeGate.

---

```
GET /web/timegate/http://a.example.org HTTP/1.1
Host: arxiv.example.net
Accept-Datetime: Tue, 11 Sep 2001 20:35:00 GMT; -P5H;+P6H
Connection: close
```

**Figure 12: User agent expresses interval of interest in Accept-Datetime header**

---

```
HTTP/1.1 406 Not Acceptable
Date: Thu, 21 Jan 2010 00:06:50 GMT
Server: Apache
Vary: negotiate, accept-datetime
Link: <http://an.example.org>; rel="original",
      <http://arxiv.example.net/web/timemap/http://a.example.org>
      ; rel="timemap"; type="application/link-format",
      <http://arxiv.example.net/web/20000915112826/http://a.example.org>
      ; rel="memento first"; datetime="Tue, 15 Sep 2000 11:28:26 GMT",
      <http://arxiv.example.net/web/20080708093433/http://a.example.org>
      ; rel="memento last"; datetime="Tue, 08 Jul 2008 09:34:33 GMT",
Content-Length: 1732
Connection: close
Content-Type: text/plain; charset=UTF-8
```

**Figure 13: A TimeGate's response indicating it has no Mementos within the interval of interest**

---

#### 3.2.2.6. Accept-Datetime Unparseable

[TOC](#)

In case, in Step 3, a user agent conveys a value for the "Accept-Datetime" request header that does not conform to the accept-dt-value construction rule of the BNF in [Figure 1 \(BNF for the datetime format\)](#), the TimeGate server's response MUST have a "400 Bad Request" HTTP status code. With all other respects, responses in this case MUST be handled as described in [Section 3.2.2.5 \(Datetime Out of the User Agent's Range\)](#).

---

#### 3.2.2.7. TimeGate Does Not Exist

[TOC](#)

Cases may occur in which a user agent issues a request against a TimeGate that does not exist. This may, for example, occur when a user agent uses internal knowledge to construct the URI of an assumed, yet non-existent TimeGate. In these cases, the response from the target server MUST have a "404 Not Found" HTTP status code, and SHOULD include a "Vary" header that includes the "negotiate" and "accept-datetime" values as an indication that, generally, the server is capable of datetime negotiation. The response MUST NOT include a "Link" header with any of the Relation Types introduced in [Section 2.2.1 \(Memento Framework Relation Types\)](#).

---

#### 3.2.2.8. HTTP Methods other than HEAD/GET

[TOC](#)

In the above, the safe HTTP methods GET and HEAD are described for TimeGates. TimeGates MAY support the safe HTTP methods OPTIONS and TRACE in the way described in [RFC 2616 \(Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," June 1999.\)](#) [RFC2616]. Unsafe HTTP methods (i.e. PUT, POST, DELETE) MUST NOT be supported by a TimeGate. Such requests MUST yield a response with a "405 Method Not Allowed" HTTP status code, and MUST include an "Allow" header to convey that only the HEAD and GET (and OPTIONALLY the OPTIONS and TRACE) methods are supported. In addition, the response MUST have a "Vary" header that includes the "negotiate" and "accept-datetime" values to indicate the TimeGate supports datetime negotiation. [Figure 14 \(Response from a TimeGate accessed with HTTP method other than HEAD/GET\)](#) shows such a response.

---

```
HTTP/1.1 405 Method Not Allowed
Date: Thu, 21 Jan 2010 00:02:12 GMT
Server: Apache
Vary: negotiate, accept-datetime
Allow: HEAD, GET
Content-Length: 255
Connection: close
Content-Type: text/html; charset=iso-8909-1
```

**Figure 14: Response from a TimeGate accessed with HTTP method other than HEAD/GET**

---

#### 3.2.3. Recognizing a TimeGate

[TOC](#)

When a user agent issues a HTTP HEAD/GET request against a resource of which it found the URI as the Target IRI of an entry in the "Link" header with a "timegate" Relation Type, it SHOULD NOT assume that the targeted resource effectively is a TimeGate and hence will behave as described in [Section 3.2.2 \(Step 4: Server Responds to Negotiation with TimeGate\)](#).

A user agent MUST decide it has reached a TimeGate if the response to a HTTP HEAD/GET request against the resource's URI contains a "Vary" header that includes the "negotiate" and "accept-datetime" values. If the response does not, the user agent MUST decide it has not reached a TimeGate and proceed as follows:

- \*If the response contains a redirection, the user agent SHOULD follow it. Note that even a chain of redirections is possible, e.g. URI-R -> URI-1 -> URI-2 -> ... -> URI-G

- \*If the response does not contain a redirection, or if the redirection (chain) does not lead to a TimeGate, the user agent SHOULD attempt to determine an appropriate TimeGate for the Original Resource, either automatically or interactively supported by the user.

Resources that are not TimeGates (i.e. do not behave as described in [Section 3.2.2 \(Step 4: Server Responds to Negotiation with TimeGate\)](#)) MUST NOT use a "Vary" header that includes the "accept-datetime" value.

---

### 3.3. Interactions with a Memento

[TOC](#)

This section details HTTP GET/HEAD requests targeted at a Memento (URI-M).

---

#### 3.3.1. Step 5: User Agent Requests a Memento

[TOC](#)

In Step 5, the user agent issues a HTTP GET request against the URI of a Memento. The user agent MAY include an "Accept-Datetime" header in this request, but the existence or absence of this header MUST NOT affect the server's response. The URI of the Memento may have resulted from a response in Step 4, or the user agent may simply have happened upon it. Such a request is illustrated in [Figure 15 \(User agent requests Memento\)](#).

---

```
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 15: User agent requests Memento**

---

### 3.3.2. Step 6: Server Responds to a Request for a Memento

[TOC](#)

If the Memento requested by the user agent in Step 5 exists, the server's response MUST have a "200 OK" HTTP status code (or "206 Partial Content", where appropriate), and it MUST include a "Memento-Datetime" header with a value equal to the archival datetime of the Memento, that is, the datetime of the state of the Original Resource that is encapsulated in the Memento. The "Link" header MUST be provided and contain links subject to the considerations described in [Section 2.2 \(Link Header Relation Types\)](#). The Target IRI and, when applicable, the datetime values in the "Link" header associated with the "memento" Relation Type SHOULD be the same as conveyed in Step 4, in case the TimeGate and the selected Memento reside on the same server. However, they MAY be different in case the TimeGate and the selected Memento reside on different servers.

[Figure 16 \(Server of Memento responds\)](#) illustrates the server's response to the request issued against a Memento in Step 5 ([Figure 15 \(User agent requests Memento\)](#)).

---

```
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/timemap/http://a.example.org>
      ; rel="timemap"; type="application/link-format",
      <http://arxiv.example.net/web/timegate/http://a.example.org>
      ; rel="timegate",
      <http://arxiv.example.net/web/20000915112826/http://a.example.org>
      ; rel="first memento"; datetime="Tue, 15 Sep 2000 11:28:26 GMT",
      <http://arxiv.example.net/web/20080708093433/http://a.example.org>
      ; rel="last memento"; datetime="Tue, 08 Jul 2008 09:34:33 GMT",
      <http://arxiv.example.net/web/20010911203610/http://a.example.org>
      ; rel="memento"; datetime="Tue, 11 Sep 2001 20:36:10 GMT",
      <http://arxiv.example.net/web/20010911203610/http://a.example.org>
      ; rel="prev memento"; datetime="Tue, 11 Sep 2001 20:30:51 GMT",
      <http://arxiv.example.net/web/20010911203610/http://a.example.org>
      ; rel="next memento"; datetime="Tue, 11 Sep 2001 20:47:33 GMT"
Content-Length: 23364
Content-Type: text/html;charset=utf-8
Connection: close
```

**Figure 16: Server of Memento responds**

---

The server's response MUST include the "Memento-Datetime" header regardless whether the user agent's request contained an "Accept-Datetime" header or not. This is the way by which resources make explicit that they are Mementos. Due to the sparseness of Mementos in most archives, 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.

---

#### **3.3.2.1. Memento Does not Exist**

[TOC](#)

Cases may occur in which a TimeGate's response (Step 4) points at a Memento that actually does not exist, resulting in a user agent's request (Step 5) for a non-existent Memento. In this case, the server's response MUST have the expected "404 Not Found" HTTP Status Code and it MUST NOT contain a "Memento-Datetime" header.

---

#### **3.3.3. Recognizing a Memento**

[TOC](#)

When following the redirection provided by a confirmed TimeGate (see [Section 3.2.3 \(Recognizing a TimeGate\)](#)), a user agent SHOULD NOT assume that the targeted resource effectively is a Memento and hence will behave as described in [Section 3.3.2 \(Step 6: Server Responds to a Request for a Memento\)](#).

A user agent MUST decide it has reached a Memento if the response to a HTTP HEAD/GET request against the resource's URI contains a "Memento-Datetime" header with a legitimate value. If the response does not, the following applies:

\*If the response contains a redirection, the user agent SHOULD follow it. Even a chain of redirections is possible, e.g. URI-G -> URI-X -> URI-Y -> ... -> URI-M.

\*If the response by a confirmed TimeGate does not contain a redirection, or if the redirection (chain) that started at a confirmed TimeGate does not lead to a resource that provides a "Memento-Datetime" header, the user agent MAY still conclude that it has likely arrived at a Memento. That is because cases exist in which archives and CMS are made compliant with the Memento framework "by proxy". In these cases TimeGates will redirect to Mementos in such systems, but the responses from these Mementos will not (yet) include a "Memento-Datetime" header.

---

## 4. The Memento Framework, Discovery Component

[TOC](#)

---

### 4.1. TimeMaps

[TOC](#)

A TimeMap resource is introduced to support retrieving a comprehensive list of all Mementos known to a responding server. The entity-body of a response to an HTTP GET request issued against a TimeMap's URI:

- \*MUST list the URI of the Original Resource that the response lists Mementos for;
- \*MUST list the URI of one or more TimeGates for the Original Resource;
- \*MUST list the URI and datetime of each Memento known to the responding server;
- \*SHOULD, for self-containment, list the URI of the TimeMap itself;
- \*MUST unambiguously type listed resources as being Original Resource, TimeGate, Memento, or TimeMap.

TimeMaps MAY be serialized in various ways, but the link-value format serialization MUST be supported. In this serialization, the entity-body MUST be formatted in the same way as the value of a HTTP "Link" header, and hence MUST comply to the "link-value" construction rule of "Section 5. The Link Header Field" of [RFC5988 \(Nottingham, M., "Web Linking," October 2010.\)](#) [RFC5988]. The media type of the entity-body MUST be "application/link-format", and the use of the Relation Types is subject to the considerations in [Section 2.2 \(Link Header Relation Types\)](#) with the following exceptions:

\*Instead of a Memento selected by the responding server, all Mementos known to the server MUST be listed;

\*Since no Memento was selected by the server, the entity-body MUST NOT contain links with "prev", "next", "predecessor-version", "successor-version" Relation Types.

In order to retrieve the link-value serialization of a TimeMap, a user agent SHOULD use an "Accept: application/link-format" header. This is shown in [Figure 17 \(Request for a TimeMap\)](#). The response from the TimeMap is shown in [Figure 18 \(Response from a TimeMap\)](#); for practical reasons the entity-body in the example has been abbreviated.

---

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

**Figure 17: Request for a TimeMap**

---

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

<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",
<http://arxiv.example.net/web/20000620180259/http://a.example.org>
  ; rel="first memento";datetime="Tue, 20 Jun 2000 18:02:59 GMT",
<http://arxiv.example.net/web/20091027204954/http://a.example.org>
  ; rel="last memento";datetime="Tue, 27 Oct 2009 20:49:54 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 18: Response from a TimeMap**

---

## 4.2. Discovery of TimeMaps, TimeGates

[TOC](#)

As described in [Section 3 \(The Memento Framework, Datetime Negotiation component: HTTP Interactions\)](#), TimeMaps and TimeGates can be discovered via HTTP Links with the "timemap" and "timegate" Relation Type, respectively. Additional discovery mechanisms are RECOMMENDED, including:

- \*The inclusion of HTML LINK elements with "timegate" and "timemap" rel types in Original Resources that provide an HTML response, e.g. `<link href="http://arxiv.example.net/timegate/http://a.example.org" rel="timegate">` ;
- \*The implementation of batch discovery mechanisms for TimeMaps using SiteMaps or feed technology.

---

[TOC](#)

## 5. IANA Considerations

This memo requires IANA to register the "Link" header Relation Types defined in [Section 2.2.1 \(Memento Framework Relation Types\)](#) in the appropriate IANA registry.

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

---

## 6. Security Considerations

[TOC](#)

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.

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

---

## 7. Changelog

[TOC](#)

\*v01 2010-11-11 HVDS MLN RS First public version

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

\*2010-07-22 HVDS MLN First internal version

---

## 8. Acknowledgements

[TOC](#)

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 for early feedback. Many thanks to Samuel Adams, Scott Ainsworth, Lyudmilla Balakireva, Frank McCown, Harihar Shankar, Brad Tofel for early implementations.

---

## 9. References

[TOC](#)

---

### 9.1. Normative References

[TOC](#)

[I-D.ietf-core-link-format]	Shelby, Z., " <a href="#">CoRE Link Format</a> ," draft-ietf-core-link-format-01 (work in progress), October 2010 ( <a href="#">TXT</a> ).
[RFC2119]	Bradner, S., " <a href="#">Key words for use in RFCs to Indicate Requirement Levels</a> ," BCP 14, RFC 2119, March 1997 ( <a href="#">TXT</a> , <a href="#">HTML</a> , <a href="#">XML</a> ).
[RFC2616]	Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, " <a href="#">Hypertext Transfer Protocol -- HTTP/1.1</a> ," RFC 2616, June 1999 ( <a href="#">TXT</a> , <a href="#">PS</a> , <a href="#">PDF</a> , <a href="#">HTML</a> , <a href="#">XML</a> ).
[RFC5829]	Brown, A., Clemm, G., and J. Reschke, " <a href="#">Link Relation Types for Simple Version Navigation between Web Resources</a> ," RFC 5829, April 2010 ( <a href="#">TXT</a> ).
[RFC5988]	Nottingham, M., " <a href="#">Web Linking</a> ," RFC 5988, October 2010 ( <a href="#">TXT</a> ).

---

### 9.2. Informative References

[TOC](#)

[Fitch]	Fitch, " <a href="#">Web site archiving - an approach to recording every materially different response produced by a website</a> ," July 2003.
[I-D.masinter-dated-uri]	Masinter, L., " <a href="#">The 'tdb' and 'duri' URI schemes, based on dated URIs</a> ," draft-masinter-dated-uri-07 (work in progress), October 2010 ( <a href="#">TXT</a> ).
[RFC1123]	Braden, R., " <a href="#">Requirements for Internet Hosts - Application and Support</a> ," STD 3, RFC 1123, October 1989 ( <a href="#">TXT</a> ).
[RFC4287]	Nottingham, M., Ed. and R. Sayre, Ed., " <a href="#">The Atom Syndication Format</a> ," RFC 4287, December 2005 ( <a href="#">TXT</a> , <a href="#">HTML</a> , <a href="#">XML</a> ).
[W3C.REC-aww-20041215]	Jacobs and Walsh, " <a href="#">Architecture of the World Wide Web</a> ," December 2004.

---

[TOC](#)

**Appendix A. Appendix B: A Sample, Successful Memento Request/Response cycle**

---

Step 1 : UA --- HTTP GET/HEAD; Accept-Datetime: Tj -----> URI-R

HEAD / HTTP/1.1

Host: a.example.org

Accept-Datetime: Tue, 11 Sep 2001 20:35:00 GMT

Connection: close

Step 2 : UA <-- HTTP 200; Link: URI-G ----- URI-R

HTTP/1.1 200 OK

Date: Thu, 21 Jan 2010 00:02:12 GMT

Server: Apache

Link: <http://arxiv.example.net/web/timegate/http://a.example.org>  
; rel="timegate"

Content-Length: 255

Connection: close

Content-Type: text/html; charset=iso-8859-1

Step 3 : UA --- HTTP GET/HEAD; Accept-Datetime: Tj -----> URI-G

GET /web/timegate/http://a.example.org

HTTP/1.1

Host: arxiv.example.net

Accept-Datetime: Tue, 11 Sep 2001 20:35:00 GMT

Connection: close

Step 4 : UA <-- HTTP 302; Location: URI-Mj; Vary; Link:

URI-R, URI-T, URI-M0, URI-Mn, URI-Mi, URI-Mj, URI-Mk ---- URI-G

HTTP/1.1 302 Found

Date: Thu, 21 Jan 2010 00:06:50 GMT

Server: Apache

Vary: negotiate, accept-datetime

Location:

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

Link: <http://a.example.org>; rel="original",  
<http://arxiv.example.net/web/20000915112826/http://a.example.org>  
; rel="first memento"; datetime="Tue, 15 Sep 2000 11:28:26 GMT",  
<http://arxiv.example.net/web/20080708093433/http://a.example.org>  
; rel="last memento"; datetime="Tue, 08 Jul 2008 09:34:33 GMT",  
<http://arxiv.example.net/web/timemap/http://a.example.org>  
; rel="timemap"; type="application/link-format",  
<http://arxiv.example.net/web/20010911203610/http://a.example.org>  
; rel="memento"; datetime="Tue, 11 Sep 2001 20:36:10 GMT",  
<http://arxiv.example.net/web/20010911203610/http://a.example.org>  
; rel="prev memento"; datetime="Tue, 11 Sep 2001 20:30:51 GMT",  
<http://arxiv.example.net/web/20010911203610/http://a.example.org>  
; rel="next memento"; datetime="Tue, 11 Sep 2001 20:47:33 GMT"

Content-Length: 0

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

Step 5 : UA --- HTTP GET URI-Mj; Accept-Datetime: Tj -----> URI-Mj

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

Step 6 : UA <-- HTTP 200; Memento-Datetime: Tj; Link: URI-R,  
URI-T, URI-G, URI-M0, URI-Mn, URI-Mi, URI-Mj, URI-Mk ---- URI-Mj

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/20000915112826/http://a.example.org>  
; rel="first memento"; datetime="Tue, 15 Sep 2000 11:28:26 GMT",  
<http://arxiv.example.net/web/20080708093433/http://a.example.org>  
; rel="last memento"; datetime="Tue, 08 Jul 2008 09:34:33 GMT",  
<http://arxiv.example.net/web/timemap/http://a.example.org>  
; rel="timemap"; type="application/link-format",  
<http://arxiv.example.net/web/timegate/http://a.example.org>  
; rel="timegate",  
<http://arxiv.example.net/web/20010911203610/http://a.example.org>  
; rel="memento"; datetime="Tue, 11 Sep 2001 20:36:10 GMT",  
<http://arxiv.example.net/web/20010911203610/http://a.example.org>  
; rel="prev memento"; datetime="Tue, 11 Sep 2001 20:30:51 GMT",  
<http://arxiv.example.net/web/20010911203610/http://a.example.org>  
; rel="next memento"; datetime="Tue, 11 Sep 2001 20:47:33 GMT"  
Content-Length: 23364  
Content-Type: text/html; charset=utf-8  
Connection: close

**A successful flow with TimeGate and Mementos on the same server**

---

## Authors' Addresses

[TOC](#)

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

	USA
Phone:	+1 505 667 1267
Email:	<a href="mailto:hvdsomp@gmail.com">hvdsomp@gmail.com</a>
URI:	<a href="http://public.lanl.gov/herbertv/">http://public.lanl.gov/herbertv/</a>
	Michael Nelson
	Old Dominion University
	Norfolk, Virginia 23529
	USA
Phone:	+1 757 683 6393
Email:	<a href="mailto:mln@cs.odu.edu">mln@cs.odu.edu</a>
URI:	<a href="http://www.cs.odu.edu/~mln/">http://www.cs.odu.edu/~mln/</a>
	Robert Sanderson
	Los Alamos National Laboratory
	PO Box 1663
	Los Alamos, New Mexico 87545
	USA
Phone:	+1 505 665 5804
Email:	<a href="mailto:azaroth42@gmail.com">azaroth42@gmail.com</a>

---

## Full Copyright Statement

[TOC](#)

Copyright © The IETF Trust (2010).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).