

HL7 hData RESTful Transport v0.19

Editor: Gerald Beuchelt

Contributors: Robert Dingwell, Andrew Gregorowicz, Marc Hadley, Mark Kramer, Samuel Sayer, Harry Sleeper

The MITRE Corporation
202 Burlington Rd.
Bedford, MA 01730
U.S.A.

© 2009-2011 The MITRE Corporation. All rights reserved.

1 Introduction

The hData RESTful API specification defines a network transport API for accessing components of a Health Record and sending messages to an EHR system.

A related specification, the hData Record Format (HRF) **Invalid source specified.**, describes the logical organization of the information in an electronic health record (EHR). Please refer to the HRF specification for more details on the HL7 hData Record Format and how it fits into the HL7 version 3 standards.

1.1 Namespaces

This document uses the following namespaces. This specification uses a number of namespace prefixes throughout; they are listed in Table 1. Note that the choice of any namespace prefix is arbitrary and not semantically significant.

Namespace Prefix	Namespace URI	Description
hrf	http://projecthdata.org/hdata/schemas/2009/06/core	Namespace for elements in this document
hrf-md	http://projecthdata.org/hdata/schemas/2009/11/meta	SectionDocument metadata

1.2 Glossary (Non-Normative)

HL7 hData Record Format (HRF) - this specification specifies an abstract hierarchical organization, packaging, and metadata for individual documents (referred to as “Section Documents” within the HRF specification). Section Documents can be of any type, either XML documents (such as CDA documents,

26 H7v3 messages, or simplified XML wire formats, etc.) or of other media types (such as e.g. MS Word
 27 documents or DICOM files). Also contained in this specification is the format for specifying the content
 28 that goes into an hData record, which is called the hData Content Profile (HCP) format.

29 **hData Record (HDR)** - an single instantiation of the HRF.

30 **HL7 hData Restful Transport** - this specification defines how the abstract hierarchical organization
 31 defined within the HRF specification is access and modified through a RESTful approach, using HTTP as
 32 the access protocol. It creates a unique mapping to an URL structure, and defines how HTTP verbs such
 33 as GET, PUT, DELETE, etc. affect the underlying information.

34 **hData Content Profile (HCP)** - a profile of the content of an HDR. The HRF specification contains the
 35 definition of the HCP format. RLUS – fill in!!

36 1.3 Notational Conventions

37 The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT",
 38 "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC](#)
 39 [2119](#).

40 When describing concrete XML schemas, this specification uses the following notation: each member of
 41 an element's [children] or [attributes] property is described using an XPath notation (e.g.,
 42 /x:MyHeader/x:SomeProperty/@value1). The use of {any} indicates the presence of an element
 43 wildcard. The use of @{any} indicates the presence of an attribute wildcard.

44 2 hData Record RESTful Transport

45 2.1 Overview

46 Any HDR can be represented as a set of Hypertext Transfer Protocol (HTTP 1.1, see [8]) resources in a
 47 canonical way. The entire HDR is referenced by a base URL that depends on the implementation. See
 48 IETF RFC 3986, section 5 for more details. This base URL will be denoted as *baseURL* throughout this
 49 document.

50 2.1.1 Out of Scope

51 While this specification does not dictate the format of the base URL, the base URL SHOULD NOT contain
 52 a query component. All content within an HDR MUST be expressible as a HTTP resource. In the
 53 following, the minimum version for HTTP is 1.1. This specification does not define any access controls to
 54 the web resources. It is RECOMMENDED that a comprehensive access control management system is
 55 always deployed with any hData installation (see Section 4)

56 2.1.2 General Conventions

57 Any HTTP GET, PUT, POST, DELETE, or OPTIONS operation (see [8], section 9) on a given resource that
 58 are either (i) unspecified or (ii) not implemented MUST return an HTTP response with a status code of

59 405 that includes an Allow header that specifies the allowed methods. All operations may return HTTP
60 status codes in the 5xx range if there is a server problem.

61 It is RECOMMENDED that all section document responses include a "Last-Modified" header. It is
62 RECOMMENDED that all document resources support the "If-ModifiedSince" and "If-Unmodified-Since"
63 headers to support conditional GET and optimistic concurrency.

64 **2.2 Operations on the Base URL**

65 **2.2.1 GET**

66 If there is no HRF at the base URL, the server SHOULD return a 404 - Not found status code.

67 The server MUST offer an Atom 1.0 compliant feed of all child sections, as identified in the
68 corresponding sections node in the root document. Each entry MUST contain a link to the resource for
69 each child section.

70 It is RECOMMENDED that the server also offers a web user interface that allows users to access and
71 manipulate the content of the HDR, as permitted by the policies of the system. Selecting between the
72 two can be achieved using standard content negotiation (HTTP Accept header). This is not necessary for
73 systems that are used by non-person entities only.

74 Status Code: 200, 404

75 **2.2.2 POST - Parameters: extensionId, path, name**

76 The request body is of type "application/x-www-form-urlencoded" and MUST contain the extensionId,
77 path, and name parameters. The extensionId parameter MUST be a string that is equal to the
78 extensionId attribute of one of the registered <extension> nodes of the root document of the HDR
79 identified by *baseURL*. The path MUST be a string that can be used as a URL path segment. If any
80 parameters are incorrect or not existent, the server MUST return a status code of 400.

81 The system MUST confirm that there is no other section registered as a child node that uses the same
82 path name. If there is a collision, the server MUST return a status code of 409.

83 If the extensionId is not registered as a valid extension, the server MUST verify that it can support this
84 extension. If it cannot support the extension it MUST return a status code of 406. It MAY provide
85 additional entity information. If it can support that extension, it MUST register it with the root.xml of
86 this record.

87 When creating the section resource, the server MUST update the root document: in the node of the
88 parent section a new child node must be inserted. If successful, the server MUST return a 201 status
89 code and SHOULD include the location of the new section. The name parameter MUST be used as the
90 user-friendly name for the new section.

91 Status Code: 201, 400, 406, 409

92 **2.2.3 PUT**

93 This operation is undefined by this specification.

94 Status Code: 405, unless an implementer defines this operation.

95 **2.2.4 DELETE**

96 This operation is undefined by this specification.

97 Status Code: 405, unless an implementer defines this operation.

98 **2.2.5 OPTIONS**

99 The OPTIONS operation on the *baseURL* is per [8], section 9.2, intended to return communications
 100 options to the clients. Within the context of this specification, OPTIONS is used to indicate which
 101 security mechanisms are available for a given *baseURL* and a list of hData content profiles supported by
 102 this implementation. All implementations MUST support OPTIONS on the *baseURL* of each HDR and
 103 return a status code of 200, along with:

- 104 • The X-hdata-security HTTP header defined in section 4 of this specification The security
 105 mechanisms defined at the baseURL are applicable to all child resources, i.e. to the entire HDR.
- 106 • An X-hdata-hcp HTTP header that contains a comma separated list of the identifiers of the
 107 hData Content Profiles supported by this implementation

108 The server MAY include additional HTTP headers. The response SHOULD not include an HTTP body. The
 109 client MUST NOT use the Max-Forward header when requesting the security mechanisms for a given
 110 HDR.

111 Status Code: 200

112 **2.3 *baseURL/root.xml***

113 **2.3.1 GET**

114 This operation returns an XML representation of the current root document, as defined by the HRF
 115 specification.

116 Status Code: 200

117 **2.3.2 POST, PUT, DELETE**

118 These operations MUST NOT be implemented.

119 Status Code: 405

120 **2.4 *baseURL/sectionpath***

121 **2.4.1 GET**

122 This operation MUST return an Atom 1.0 **Invalid source specified.** compliant feed of all section
 123 documents and child sections contained in this section. Each entry MUST contain a link to a resource

124 that uniquely identifies the section document or child section. If the section document type defines a
 125 creation time, is RECOMMENDED to set the Created node to that datetime.

126 For section documents, the Atom Content element MUST contain the XML representation of its
 127 metadata (see **Invalid source specified.**, Section 2.4.1).

128 Status Code: 200

129 **2.4.2 POST**

130 For creating a new sub section, three additional parameters are used, and the POST will create a new
 131 child section within this section. For new documents a document MUST be sent that conforms to the
 132 business rules expressed by the extension that the section has registered.

133 **2.4.2.1 Add new section – Parameters: extensionId, path, name**

134 The content type MUST equal “application/x-www-form-urlencoded” for the POST method to create a
 135 new sub section. The extensionId parameter is the URI in the root.xml document that identifies the
 136 Extension element. If the extensionId is not registered as a valid extension, the server MUST verify that
 137 it can support this extension. If it cannot support the extension it MUST return a status code of 406 and
 138 MAY provide additional information in the entity body. If it can support that extension, it MUST register
 139 it with the root.xml of this record. The path MUST be a string that can be used as a URL path segment.
 140 The name parameter MUST be used as the user-friendly name for the new section. If any parameters
 141 are incorrect, the server MUST return a status code of 400.

142 The system MUST confirm that there is no other section registered as a child node that uses the same
 143 path name and that it can create a new subsection identified by the path parameter. If there is a
 144 collision, the server MUST return a status code of 409.

145 When creating the section resource, the server MUST update the root document: in the node of the
 146 parent section a new child node must be inserted. The server MUST return a 201 status code. The
 147 extensionId and path parameters are REQUIRED, the name parameter is OPTIONAL.

148 Status Code: 201, 400, 406, 409

149 **2.4.2.2 Add new document**

150 When adding a new section document, the request Content Type MUST be “multipart/form-data” if
 151 including metadata. In this case, the content part MUST contain the section document. The content part
 152 MUST include a Content-Disposition header with a disposition of “form-data” and a name of “content”.
 153 The metadata part MUST contain the metadata for this section document. The metadata part MUST
 154 include a Content-Disposition header with a disposition of “form-data” and a name of “metadata”. It is
 155 to be treated as informational, since the service MUST compute the valid new metadata based on the
 156 requirements found in the HRF specification. The content media type MUST conform to the media type
 157 of either the section or the media type identified by metadata of the section document. For XML media
 158 types, the document MUST also conform to the XML schema identified by the extensionId for the

159 section or the document metadata. If the content cannot be validated against the media type and the
160 XML schema identified by the content type of this section, the server MUST return a status code of 400.

161 If the request is successful, the new section document MUST show up in the document feed for the
162 section. The server returns a 201 with a Location header containing the URI of the new document.

163 Status Code: 201, 400.

164 2.4.3 PUT

165 This operation is not defined by this specification.

166 Status Code: 405, unless an implementer defines this operation.

167 2.4.4 DELETE

168 This operation SHOULD be implemented, but special precaution should be taken: if a DELETE is sent to
169 the section URL, the **entire** section, its documents, and subsections are completely deleted. Future
170 requests to the section URL MUST return a status code of 404, unless the record is restored. If successful
171 the server MUST return a status code of 204.

172 Status Code: 204, 404

173 2.5 *baseURL/sectionpath/documentname*

174 2.5.1 GET

175 This operation returns a representation of the document that is identified by *documentname* within the
176 section identified by *sectionpath*. The *documentname* is typically assigned by the underlying system and
177 is not guaranteed to be identical across two different systems. Implementations MAY use identifiers
178 contained within the infoset of the document as *documentnames*.

179 If no document of name *documentname* exists, the implementation MUST return a HTTP status code
180 404.

181 Status Codes: 200, 404

182 2.5.2 PUT

183 This operation is used to update a document. The content MUST conform to the media type identified
184 by the document metadata or the section content type. For media type application/xml, the document
185 MUST also conform to the XML schema that corresponds to the content type identified by the
186 document metadata or the section. If the parameter is incorrect or the content cannot be validated
187 against the correct media type or the XML schema identified by the content type of this section, the
188 server MUST return a status code of 400.

189 If the request is successful, the new section document MUST show up in the document feed for the
190 section. The server returns a 200.

191 Status Code: 200, 400.

192 **2.5.3 POST**

193 This operation is used to replace metadata on a section document. This operation SHOULD NOT be used
194 unless necessary for replicating information within an organization. If a section document is copied from
195 one system to another, a new document metadata instance MUST be constructed from the original
196 metadata according to the rules in the HRF specification.

197 The request Media Type MUST be application/xml. The body MUST contain the document metadata. It
198 MUST conform to the XML schema for the document metadata, defined in **Invalid source specified..** If
199 the metadata is not of media type application/xml or it cannot be validated against the document
200 metadata XML schema, the server MUST return a status code of 400.

201 If the request is successful, the document metadata for the section document MUST be updated. The
202 server returns a 201.

203 Status Code: 201, 400.

204 **2.5.4 DELETE**

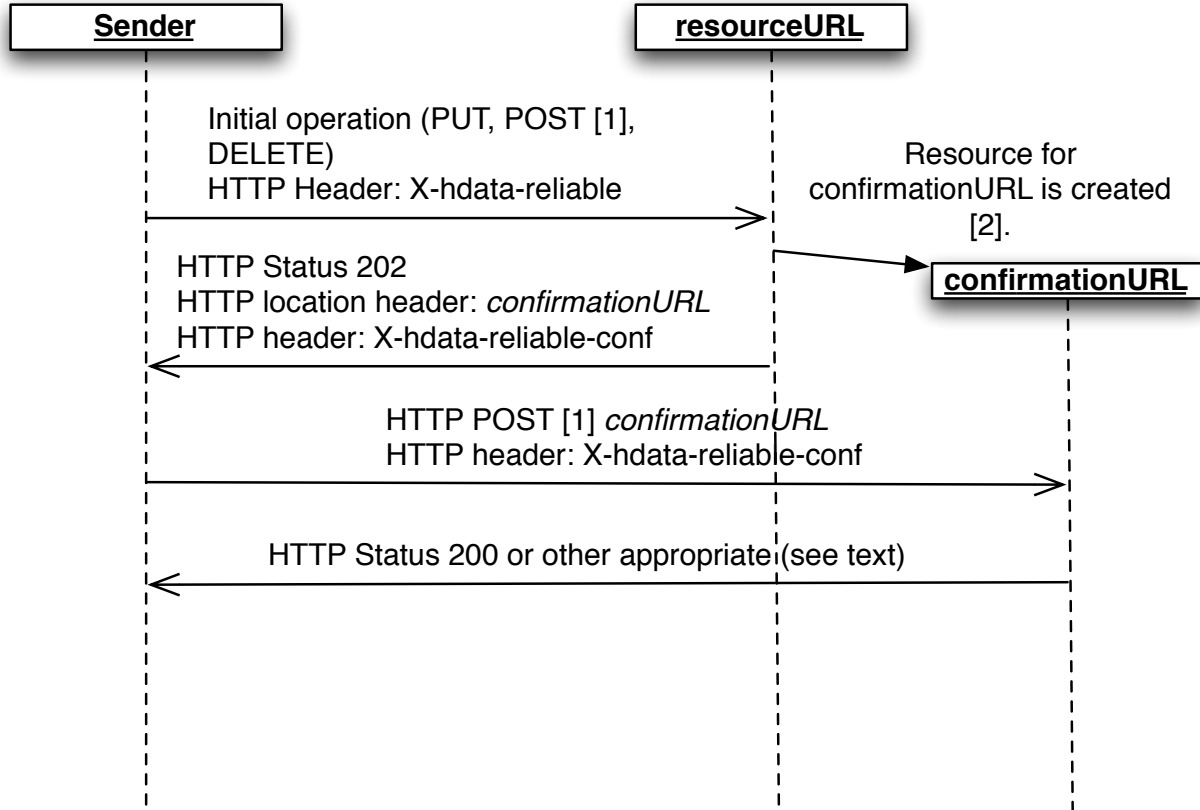
205 This operation SHOULD be implemented, but special precaution should be taken: if a DELETE is sent to
206 the document URL, the document is completely deleted. Future requests to the section URL MAY return
207 a status code of 410, unless the record is restored.

208 Status Code: 204, 410

209 **3 Reliable Operations**

210 This pattern is a complex multi step exchange, which is applicable to situations where a multi-phase
211 commit is required. This pattern MAY be combined when interacting with an hData Record or with other
212 message patterns, as long as there is no overloading of HTTP methods.

213 The use of the reliable operations pattern will be governed by the business requirements of the
214 business domain.



[1] All POST methods must be implemented to support idempotency, e.g. through mechanisms like "Post Once Exactly" (POE).

[2] The confirmationURL may be identical to the resourceURL for document transactions.

215 Please see the text for more details on the interactions.

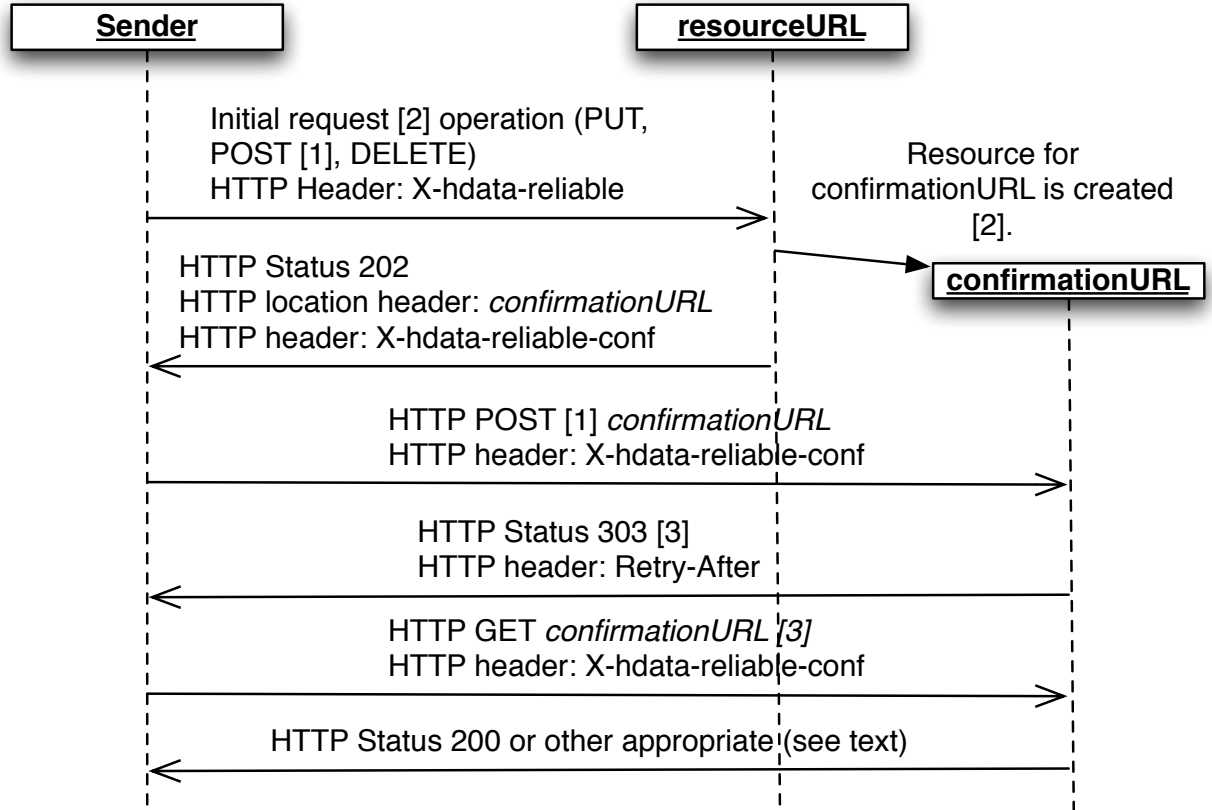
216 The flow of the patterns is as follows:

- 217 1. The sender accesses the *resourceURL* resource using PUT, POST, or DELETE. To indicate that it
218 wants to use the reliable operations pattern, it sets the HTTP message header "X-hdata-
219 reliable".
- 220 2. If the *resourceURL* is capable of performing the reliable operations pattern, it will create a new
221 resource for a message at *confirmationURL*, and return an HTTP status code of 202. The HTTP
222 result MUST contain the *confirmationURL* in the HTTP location header and a confirmation secret
223 in the "X-hdata-reliable-conf" header. This secret SHOULD be a simple string of sufficient length
224 to prevent guessing. The service MUST NOT process the message at this stage.
225 If the *resourceURL* does not implement the reliable operations pattern, it MUST return an HTTP
226 status code of 405 and discard the message.

- 227 3. The sender MUST then POST an empty request body to the resource at *confirmationURL* and set
228 the “X-hdata-reliable-conf” header to the value provided in step 2. Upon receipt, the service –
229 listening at the *confirmationURL* – MUST validate the confirmation secret. Once the GET secret
230 is validated, the service processor MUST process the message immediately.
- 231 4. If the validation is successful, the *confirmationURL* returns an HTTP result with the expected
232 status code for the initial operation. If the validation is not successful, the service MUST return
233 an HTTP status code of 409. The sender MUST retry the POST until it receives either a different
234 HTTP status code.

235 **Remarks:**

- 236 1. Since POST is not idempotent, the service MUST implement a safe guard against duplicity of
237 requests for all posts in this flow. It is RECOMMENDED that the service implements “POST Once
238 Exactly” (POE), as described in <http://www.mnot.net/drafts/draft-nottingham-http-poe-00.txt>.
- 239 2. The *confirmationURL* resource MAY be destroyed after the reliable message pattern flow is
240 complete. The service MAY maintain the *confirmationURL* after the pattern flow completes.
- 241 3. If the initial operation in step 1 above is an application-level request message or document, the
242 *confirmationURL* MAY provide an application-level response in step 4. The response MAY be
243 provided by returning the response body in the final step; the HTTP status code MUST NOT be
244 409. For asynchronous responses, the *confirmationURL* MAY return an HTTP status 303 with a
245 “Retry-After” header, indicating when the response will be available through a GET operation at
246 the *confirmationURL*.



[1] All POST methods must be implemented to support idempotency, e.g. through mechanisms like "Post Once Exactly" (POE).

[2] The request/response protocol is defined at the application level and not through this specification. The Sender and the service at the resourceURL will determine if the operation is a request.

[3] The 303/Retry-After step is optional. It MAY be used for asynchronous responses.

Please see the text for more details on the interactions.

247
248
249
250
251
252
253
254
255
256
257

This specification does not provide guidance to what constitutes an application-level request/response protocol. Implementers of this specification can decide if this mechanism is appropriate for their application.

4. There is no default for how long the *confirmationURL* resource is *available for* confirmation (step 3). The service MAY destroy the *confirmationURL* resource and discard the message if the sender does not complete step 3 of the pattern flow. It is strongly RECOMMENDED to advertise the maximum time for confirming the message to the developer of the sender in the documentation for the service. If the service discards the message after timing out *the confirmation* step, it MUST return a status code of 404 at the *confirmationURL* permanently. If

258 the service issued a “Retry-After” header in response (as indicated in Remark 3.), it MUST
259 provide the *confirmationURL* until after the expiration of the time indicated by this header.
260 5. For operations on hData Records (as described in section 2) special provision MUST be taken to
261 prevent alteration of the resource once the reliable message pattern is initiated. The service
262 MUST provide the old status of the resource until step 3 completes. It is RECOMMENDED to use
263 the resource URL (which is different from the URL for the metadata for the resource URL) also as
264 the *confirmationURL*.

265 4 Security Considerations

266 This transport and API specification can be used to transfer data in many different situations, for
267 example, inside organizations, between organizations, or from medical devices. As such, the
268 specification cannot provide a comprehensive security solution that addresses the needs of all possible
269 applications. However, this section describes a number of basic security mechanisms that hData
270 implementations MUST support. In addition, this section describes general web security considerations
271 and how additional security mechanisms and systems can be added to implementations of this
272 standard. Implementers of hData are advised to review their domain specific security requirements and
273 select or create appropriate security mechanisms. The section concludes with a discussion of risk
274 analysis, which is highly recommended prior to implementing and deploying any infrastructure for
275 clinical systems.

276 4.1 Security Mechanism Specification

277 To allow the support of multiple security mechanisms at a single HRF resource, clients MUST be able to
278 always access the *baseURL* through an HTTP OPTIONS request (see [8], section 9.2). If the resource
279 employs any security mechanism with the exception of transport security (see 4.2.1), it MUST include
280 the HTTP header X-hdata-security which MUST contain a comma separated list of URL-encoded URIs
281 that identify the supported security mechanism. Section 4.2 includes the URIs for the baseline security
282 mechanisms.

283 It is RECOMMENDED that hData Content Profiles include a detailed specification of any required custom
284 security mechanisms. If the custom security mechanism The URIs for identifying these additional
285 security mechanisms SHOULD be made unique by using the DNS domain name in the first part of the
286 URI.

287 Any new security mechanism specification that is compliant with this standard needs to provide the
288 following items. This SHOULD be done through a commonly readable text document, such as HTML. This
289 package provides implementers with the necessary security protocol information to create the security
290 mechanism for their system.

291 1. Common Name (REQUIRED) – free text, recommended to be less than 32 characters

- 292 2. Identifier (REQUIRED) – URI, recommended to include the originating organizations DNS domain
293 name for uniqueness. NOT REQUIRED for transport security (see 4.2.1). It is RECOMMENDED to
294 use a URL that resolves into the HTML representation of the security mechanism specification.
- 295 3. Exclusiveness (REQUIRED) – free text, describes if the mechanism can be combined with other
296 mechanism
- 297 4. Description (REQUIRED) – free text, includes a comprehensive description of all allowed
298 interaction patterns, parameters, and dependencies
- 299 5. State diagram (RECOMMENDED) – UML state diagram, identifies all actors and illustrates all
300 allowed interaction patterns
- 301 6. Business rules (RECOMMENDED) – free text, describes the business/domain justification and
302 rules for this security mechanism
- 303 7. Example (RECOMMENDED) – free text, recommended to include examples including packet
304 content for all interaction patterns
- 305 8. Other Content (OPTIONAL)

306 4.2 Baseline Security

307 The mechanisms described in this section MUST be supported by all implementation of this
308 specification. While transport security is always RECOMMENDED, there can be situations where
309 transport security is not required.

310 The versions of IETF standards selected within this specification are the minimal REQUIRED versions. It is
311 RECOMMENDED to use more modern versions, as long as these newer versions are backward
312 compatible.

313 4.2.1 HTTP Transport Security

314 Transport security is implemented within the network stack below the HTTP transport layer.

- 315 1. Common Name: HTTP Transport Security
- 316 2. Identifier: none – Not required because the identifier is encoded in the *baseURL* URL through
317 the https scheme.
- 318 3. Exclusiveness: This mechanism can be combined with all other security mechanism.
- 319 4. Description: Implementations MUST support TLS 1.1 or higher. This protocol is described in
320 detail in IETF RFC 4345 [2]. TLS supports both anonymous clients, as well as client
321 authentication. Implementations of this specification MUST support anonymous client, and
322 MUST support client authentication through TLS. If TLS client authentication is supported,
323 implementation MAY use the principal obtained from the exchange in their authentication and
324 authorization process.

325 4.2.2 Message Security

- 326 1. Common Name: S/MIME Message security
- 327 2. Identifier: <http://www.hl7.org/hdata/2011/03/security/smime-messages>
- 328 3. Exclusiveness: This mechanism can be combined with all other security mechanisms.

- 329 4. Description: Implementations MUST support S/MIME 3.2 or higher which is an IETF internet
330 standard described in IETF RFC 5751 [4]. S/MIME requires PKI certificates for sender and
331 receiver, and also a way for the sender to discover the public key certificate for the receiver. The
332 sender should include its own certificate in the S/MIME message. Implementations MUST use
333 SHA-256 and RSA for signature and encryption, respectively. To achieve confidentiality,
334 implementations MUST use the EnvelopedData content type [10], section 2.4.3. The hData
335 SectionDocument that becomes the MIME payload of the S/MIME message MUST be prepared
336 by the implementation according to the requirements of the S/MIME specifications, with special
337 consideration for the MIME content type.
- 338 While out of scope for this specification, there are a number of ways to discover the certificates:
- 339 • If the receiver offers any web resources through https, it is RECOMMENDED to use the
340 server certificate.
 - 341 • If any discovery services are available, it is RECOMMENDED that the metadata for the
342 endpoint includes the public key certificate.
 - 343 • If DNS CERT resource records (IETF 4398 [5]) are available, the sender MAY use the
344 certificate published.

345 4.2.3 Authentication

346 Authentication can be achieved through all of the mechanisms described in this section.

347 Implementations of this specification MUST support all described authentication mechanisms, but these
348 mechanisms MAY be disabled at deploy or runtime.

349 4.2.3.1 HTTP Basic Authentication

- 350 1. Common Name: HTTP Basic Authentication
- 351 2. Identifier: <http://www.hl7.org/hdata/2011/03/security/http-basic-auth>
- 352 3. Exclusiveness: This mechanism can be combined with all other security mechanisms. When
353 combining with other authentication mechanisms, it SHOULD use the other mechanism's
354 security principal for authentication and authorization.
- 355 4. Description: Implementations MUST implement HTTP Basic Authentication as specified in IETF
356 RFC 2617 [6], section 2.

357 4.2.3.2 HTTP TLS Authentication

- 358 1. Common Name: HTTP over TLS
- 359 2. Identifier: <http://www.hl7.org/hdata/2011/03/security/http-tls-auth>
- 360 3. Exclusiveness: This mechanism SHOULD NOT be combined with other authentication security
361 mechanisms. If combined with other security mechanisms, the principal of the client certificate,
362 as identified by the Common Name (CN) attribute of the certificate, SHOULD be used as the
363 security principal in all subsequent authentication and authorization decisions.
- 364 4. Description: Implementations MUST implement HTTP TLS Client Certificates as specified in IETF
365 RFC 2246 [7], section 7.4.6.

4.3 Specifying A Custom Security Mechanism

Additional security mechanisms that can be published through the X-hdata-security header can be created as needed by the behavioral model and the application domain. It is RECOMMENDED to include or reference security mechanisms necessary for a given hData Content Profile (HCP) within the HCP package. The security mechanism description MUST comply with the template specified in section 4.2.

4.4 General Web Security Considerations

Because hData is implemented using common web technology, it is subject to the same security considerations as other security-sensitive web applications and services. Because Internet threats and vulnerabilities are constantly evolving, hData implementations should apply current best practices to assure appropriate levels of security.

These security best practices should be considered not only at the software application layer, but also at lower layers such as the network layer and physical layer. For example, hData implementations MAY also support lower-level protection mechanisms, such as IPSEC or other bulk traffic encryption. Typically, such technologies have no direct impact on the application layer, and their use and implementation is determined by the networking infrastructure. Protection of critical infrastructure services such as DNS or DHCP MAY be necessary. Information security must be integrated with non-IT security as well:

- Any information processing systems must be protected from intentional and unintentional physical harm, both man-made as well as natural.
- Business processes and non-IT workflow must integrate with information security, and prevent circumvention of information security measures.
- System operators and end users must be cleared for access at the appropriate level.

The reader is advised to consult appropriate resources in this area for more information, such as NIST 800-12, NIST 800-14, ISA-99, and ISO 27002.

4.5 Risk Assessment Approach and Best Practices

It is highly RECOMMENDED to perform a comprehensive risk analysis prior to deploying any clinical application. Risk analysis is a systematic consideration of the threats, vulnerabilities, and consequences of gaps in security, as well as mitigation strategies for risks. Often, the threats and vulnerabilities are captured in terms of specific scenarios that can be re-used during security audits throughout the system's lifecycle. The reader is advised to consult appropriate resources for more information on cyber risk assessment, such as NIST 800-30, the IHE security cookbook [11], and ISO/TS 25238,

5 Realization of RLUS Profiles

5.1 Introduction

The Retrieve, Locate, Update Service (RLUS) Specification defines an HL7 framework for healthcare services. The hData RESTful Transport is a realization of RLUS Functional Profiles. The hData Content

400 Profile (HCP) [1], section 3, acts as such as a Semantic Profile in the sense of [5], section 6.1. Taken
 401 together, the two portions of the hData specification forms an RLUS Conformance profile. This section
 402 provides a mapping between the hData RESTful implementation and the RLUS framework.

403 5.2 Implementation of RLUS Interfaces

404 The RLUS specification defines a number of interfaces in [5], Section 5 “Detailed Functional Model”.
 405 These are mostly implemented by the hData specification, as detailed within the table below.

RLUS Interface	hData RESTful Implementation
Create RLUS Entry (5.1.1)	Section 2.4.2.2 describes how a new SectionDocument can be created.
Update RLUS Entry (5.1.2)	Section 2.5.2 describe how an existing SectionDocument can be updated
Delete RLUS Entry (5.1.3)	Section 2.5.4 describes how a SectionDocument can be deleted.
List Conformance Profiles (5.1.4)	Section 2.2.5 describes the X-hdata-hcp header which returns a list of hData content profiles
List Semantic Signifiers (5.2.1)	The root.xml at the baseUrl contains the list of supported elements within the Extensions node. The list of Extension elements represents the list of semantic signifiers, as required by [5] 5.2.1. (The HRF specification [1] recommends URLs as identifiers for each Extension, which should resolve into a RDDDL document describing the given Extension. This is consistent with the recommendation of [5] section 5.2.1 to provide an explanation for each semantic signifier.)
List Semantic Signifiers for Resource (5.2.2)	A SectionDocument resource within an HDR MUST conform to a single semantic signifier only. By default, this is the semantic signifier (<Extension>) that is defined in the root.xml for the Section in which the SectionDocument is located. This MAY be overridden in the metadata, as explained in [1], section 2.5.1.
Retrieve Semantic Signifier (5.2.3)	For any <Extension> that is a URL and resolves into a RDDDL document, the necessary description can be retrieved. Thus, if an hData implementation strives to be compliant to this interface, recommendation in [1] section 2.3 to use URLs and resolve into RDDDLs becomes a requirement.
Locate Resources by Resource Parameter (5.3.1)	Parameter-specific query may be implemented either over a single HDR or a collection of HDR by another specification. This is out-of-scope for the HRF and this specification.
Retrieve Resource (5.3.2)	This is implemented using a HTTP GET operation on the resource identified by its URL.
Retrieve Resource by Resource Parameter (5.3.3)	Parameter-specific query may be implemented either over a single HDR or a collection of HDR by another specification. This is out-of-scope for the HRF and this specification.

Create Resource (5.3.4)	Section 2.4.2.2 describes how a new SectionDocument can be created.
Update Resource (5.3.5)	Section 2.5.2 describe how an existing SectionDocument can be updated
Delete Resource (5.3.6)	Section 2.5.4 describes how a SectiondDocument can be deleted.

406

407 5.3 Conformance to RLU Functional Profiles

408 Section 6 of the RLU specification [5] defines five functional profiles:

Functional Profile	hData Realization	Comments
Location		

409

410

411 6 Bibliography

412 [1] G. Beuchelt et al., "hData Record Format", The MITRE Corporation, 2011.

413 [2] IETF RFC 4345 "Transport Layer Security (TLS) 1.1", online at <http://tools.ietf.org/html/rfc4346>

414 [3] IETF Network Working Group. (2005, Dec.) IETF. [Online]. <http://www.ietf.org/rfc/rfc4287.txt>

415 [4] IETF Network Working Group "S/MIME 3.2 Message Specification", online at

416 <http://tools.ietf.org/html/rfc5751>

417 [5] IETF Network Working Group, "Storing Certificates in the Domain Name System (DNS)", online at

418 <http://tools.ietf.org/html/rfc4398>

419 [6] IETF Network Working Group, "HTTP Authentication: Basic and Digest Access Authentication", online

420 at <http://tools.ietf.org/html/rfc2617>

421 [7] IETF Network Working Group "The TLS Protocol", online at <http://tools.ietf.org/html/rfc2246>

422 [8] IETF Network Working Group "Hypertext Transfer Protocol – HTTP 1.1", online at

423 <http://tools.ietf.org/html/rfc2616>

424 [9] HL7 Resource Location and Updating Service (RLUS), DSTU Release 1, Health Level Seven, Inc.,

425 December 2006

426 [10] "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Message Specification ", RFC

427 3851, The Internet Society, July 2004, online at <http://www.rfc-editor.org/rfc/rfc3851.txt>

428 [11] "Cookbook:Preparing the IHE Profile Security Section", IHE International, October 2008, online at
429 http://www.ihe.net/Technical_Framework/upload/IHE_ITI_Whitepaper_Security_Cookbook_2008-11-
430 [10.pdf](http://www.ihe.net/Technical_Framework/upload/IHE_ITI_Whitepaper_Security_Cookbook_2008-11-10.pdf)