

Integrating the Healthcare Enterprise



5

**IHE IT Infrastructure
Technical Framework Supplement**

10

IHE Appendix on HL7[®] FHIR[®]

HL7[®] FHIR[®] STU 3

15

Rev. 1.2 – Trial Implementation

20 Date: July 21, 2017
Author: IHE ITI Technical Committee
Email: iti@ihe.net

25

Please verify you have the most recent version of this document. See [here](#) for Trial Implementation and Final Text versions and [here](#) for Public Comment versions.

Foreword

30 This is a supplement to the IHE IT Infrastructure Technical Framework V14.0. Each supplement undergoes a process of public comment and trial implementation before being incorporated into the volumes of the Technical Frameworks.

This supplement is published on July 21, 2017 for trial implementation and may be available for testing at subsequent IHE Connectathons. The supplement may be amended based on the results of testing. Following successful testing it will be incorporated into the IT Infrastructure
35 Technical Framework. Comments are invited and can be submitted at http://www.ihe.net/ITI_Public_Comments.

This supplement describes changes to the existing technical framework documents.

“Boxed” instructions like the sample below indicate to the Volume Editor how to integrate the relevant section(s) into the relevant Technical Framework volume.

40 *Amend Section X.X by the following:*

Where the amendment adds text, make the added text **bold underline**. Where the amendment removes text, make the removed text **~~bold strikethrough~~**. When entire new sections are added, introduce with editor’s instructions to “add new text” or similar, which for readability are not bolded or underlined.

45

General information about IHE can be found at <http://ihe.net>.

Information about the IHE IT Infrastructure domain can be found at http://ihe.net/IHE_Domains.

Information about the organization of IHE Technical Frameworks and Supplements and the process used to create them can be found at http://ihe.net/IHE_Process and <http://ihe.net/Profiles>.

50 The current version of the IHE IT Infrastructure Technical Framework can be found at http://ihe.net/Technical_Frameworks.

CONTENTS

55

Introduction to this Supplement..... 4

 Open Issues and Questions 4

 Closed Issues 4

Appendices..... 5

60 Appendix Z – FHIR Implementation Material 5

 Z.1 Resource Bundles 5

 Z.2 Query Parameters..... 5

 Z.2.1 Query Parameter Modifiers..... 6

 Z.2.2 Token Parameters 6

65 Z.2.3 String Parameters 6

 Z.3 CapabilityStatement Resource 6

 Z.4 StructureDefinition Resource 6

 Z.5 Resource Reference URIs in FHIR..... 6

 Z.6 Populating the Expected Response Format 7

70 Z.7 Guidance on Access Denied Results 7

 Z.8 Mobile Security Considerations 8

 Z.9 FHIR Data Types 9

 Z.9.1 Identifier Type 9

 Z.9.1.1 Identifier and HL7 version 3 "root plus extension" 10

75 Z.9.1.2 XDS CXi mapped to FHIR Identifier Type 10

 Z.10 Profiling conventions for constraints on FHIR..... 11

Appendix E Usage of the CX Data Type in PID-3 Patient Identifier List..... 12

 E.3 FHIR Identifier Type 12

80

Introduction to this Supplement

Whenever possible, IHE profiles are based on established and stable underlying standards. However, if an IHE committee determines that an emerging standard offers significant benefits for the use cases it is attempting to address and has a high likelihood of industry adoption, it may develop IHE profiles and related specifications based on such a standard.

The IHE committee will take care to update and republish the IHE profile in question as the underlying standard evolves. Updates to the profile or its underlying standards may necessitate changes to product implementations and site deployments in order for them to remain interoperable and conformant with the profile in question.

This Technical Framework Supplement uses the emerging HL7^{®1} FHIR^{®2} specification. The FHIR release profiled in this supplement is STU 3. HL7 describes the STU (Standard for Trial Use) standardization state at <https://www.hl7.org/fhir/versions.html>.

85

Open Issues and Questions

None

Closed Issues

None

¹ HL7 is the registered trademark of Health Level Seven International.

² FHIR is the registered trademark of Health Level Seven International.

90

Appendices

Add the following Appendix to the Volume 2x Appendices

Appendix Z – FHIR Implementation Material

95

The HL7 FHIR standard has several overarching concepts, which should be profiled consistently throughout any mobile/lightweight IHE transactions using FHIR. IHE profiles FHIR, like any other standard, in ways that narrow the standard for specific use-cases. IHE profiles are intended to be proper subsets of the standard and are not intended to be incompatible.

We discuss here how IHE profiles the FHIR standard (Standard for Trial Use) such as Resources, Datatypes, Valuesets, Extensions, Transactions, Query Parameters, CapabilityStatement, etc.

100

References

HL7 FHIR	HL7 FHIR standard STU3 http://hl7.org/fhir/STU3/index.html
RFC2616	Hypertext Transfer Protocol – HTTP/1.1
RFC7540	Hypertext Transfer Protocol – HTTP/2
RFC3986	Uniform Resource Identifier (URI): Generic Syntax
RFC4627	The application/json Media Type for JavaScript Object Notation (JSON)
RFC6585	Additional HTTP Status Codes

Z.1 Resource Bundles

105

Any operation that results in, or requires submission of, a collection of resources is done via a Resource Bundle mechanism. A FHIR Bundle Resource is a collection of resources that are related, for example, the result of a search operation, or a collection of historical versions of a resource.

Bundles are described at FHIR STU3 <http://hl7.org/fhir/STU3/bundle.html>.

This section has no specific constraints.

Z.2 Query Parameters

110

FHIR STU3: <http://hl7.org/fhir/STU3/search.html> specifies a series of query parameter types which may be used when querying for a resource on a server. The representation of these query parameters within the HTTP request URL are intended to support a broad set of use cases and in some cases the behavior is unclear.

115

In this section, we discuss query parameters in the context of RESTful HTTP queries represented in the request URL within IHE profiles.

Query parameters not mentioned here are not constrained by this appendix.

Z.2.1 Query Parameter Modifiers

This section has no specific constraints.

Z.2.2 Token Parameters

120 A token type is a parameter that provides an exact match search, either on a string of characters, potentially scoped by a URI. It is mostly used against a code or identifier data type where the value may have a URI that scopes its meaning, where the search is performed against the pair from a Coding or an Identifier. Tokens are also used against other fields where exact matches are required. (See FHIR STU3 <http://hl7.org/fhir/STU3/search.html#token>.)

125 This section has no specific constraints.

Z.2.3 String Parameters

For a simple string search, a string parameter serves as the input for a case- and accent-insensitive search against sequences of characters. By default, a field matches a string query if the value of the field equals or starts with the supplied parameter value, after both have been
130 normalized by case and accent. (See FHIR STU3 <http://hl7.org/fhir/STU3/search.html#string>.)

IHE Profiles may choose to require support for the string search modifiers.

This section has no specific constraints.

Z.3 CapabilityStatement Resource

135 HL7 FHIR allows service implementers to publish a `CapabilityStatement` Resource describing the resources, transport, formats, and operations that can be performed on a series of resources for the service instance. The `CapabilityStatement` Resource is described in FHIR STU3 <http://hl7.org/fhir/STU3/CapabilityStatement.html>.

This section has no specific constraints.

Z.4 StructureDefinition Resource

140 HL7 FHIR allows service implementers to publish a `StructureDefinition` Resource describing the constraints, terminology bindings, extensions and search parameters supported for a resource type. This `StructureDefinition` Resource allows consumers to determine the capabilities and data requirements of a particular FHIR-based service. The
145 `StructureDefinition` Resource is described in FHIR STU3 <http://hl7.org/fhir/STU3/structuredefinition.html>.

This section has no specific constraints.

Z.5 Resource Reference URIs in FHIR

Many of the defined elements in a resource are references to other resources. Using these
150 references, the resources combine to build a web of information about healthcare. Resource References are described in FHIR STU3 <http://hl7.org/fhir/STU3/references.html>.

This section has no specific constraints.

Z.6 Populating the Expected Response Format

155 The FHIR standard provides for response message content encoded as either XML or JSON. The server actor shall support both message encodings, whilst the client actors shall support one and may optionally support both.

There are two methods for the client to indicate preference for encoding:

- the use of HTTP content negotiation and
- the `_format` query parameter

160 This is described in FHIR STU3 <http://hl7.org/fhir/STU3/HTTP.html#mime-type>. The server actor shall support both methods. Note that the value of the `_format` parameter must be a subset of the HTTP content negotiation.

A client actor shall indicate preference for response format, using at least one method, with at least one of the following values. A server actor may support other encodings. To enable simpler query encoding, the value of `_format` may be the short “json” or “xml”.

165

Table Z.6-1: Desired response encoding

Desired Encoding	mime-type Value
json	application/fhir+json
xml	application/fhir+xml

Z.7 Guidance on Access Denied Results

170 The server must choose the response carefully when an Access Denied condition exists. Returning too much information may expose details that should not be communicated. The Access Denied condition might be because of missing required Authentication, because the user is not authorized to access the endpoint, because the user is not authorized to access specific data, or because of other policy reasons.

To balance usability of the response against appropriate disclosure, the actual result method used needs to be controlled by policy and context.

175 Typical methods used are:

- **Return a Success with Bundle containing zero results** – This result is indistinguishable from the case where no data is known. When consistently returned on Access Denied, this approach will not expose which patients exist, or what data might be blinded. This method is also consistent with cases where some results are authorized while other results are excluded from the results. This can only be used when returning a Bundle is a valid result.
- 180

- **Return a 404 “Not Found”** – This approach also protects from data leakage, as it is indistinguishable from a query against a resource that does not exist. It does however leak that the user is authenticated.
- 185 • **Return a 403 “Forbidden”** – This approach communicates that the reason for the failure is an Authorization failure. It should only be used when the client and/or user is trusted to be given this information. Thus, this method is used mostly when the user is allowed to know that access is forbidden. It does not explain how the user might change things to become authorized. This approach may leak that content exists.
- 190 • **Return a 401 “Unauthorized”** – This communicates that user authentication was attempted and failed to be authenticated. This approach may leak that content exists.

When the server needs to report an error, it shall use HTTP error response codes and should include a FHIR OperationOutcome with more details on the failure. See FHIR STU3 <http://hl7.org/fhir/STU3/http.html> and <http://hl7.org/fhir/STU3/operationoutcome.html>

195 **Z.8 Mobile Security Considerations**

There are many security and privacy concerns with mobile devices, including lack of physical control. Many common information technologies use of HTTP, including REST, access far less sensitive information than health information. These factors present an especially difficult challenge for the security model. Application developers should perform a Risk Assessment during design of their applications, and organizations responsible for the operational environment should perform Risk Assessments on the design and deployment of the operational environment. See FHIR STU3 Security <http://hl7.org/fhir/STU3/security.html>.

Actors should not communicate any patient information unless proper authentication, authorization, and communications security have been performed.

205 There are many reasonable methods of securing interoperability transactions. These security models can be layered in without modifying the characteristics of the transaction. The use of TLS is encouraged, specifically the use of the ATNA Profile. User authentication on mobile devices is encouraged using Internet User Authorization (IUA) Profile. The IUA Profile is a profile of the OAuth protocol. IUA enables external Authorization providers, which can leverage pluggable authentication providers, such as OpenID Connect. The network communication security and user authentication are layered in at the HTTP transport layer and do not modify the interoperability characteristics defined in the transaction.

215 Security audit logging (e.g., ATNA) is recommended. Support for ATNA-based audit logging on the mobile health device may be beyond the ability of the client-constrained environment. For example, the client actor may only support HTTP interactions using JSON encoding, while the Record Audit Event [ITI-20] transaction requires the SYSLOG protocol and XML encoding. For this reason, the use of ATNA Audit Logging is not mandated. This means that the organization responsible for the operational environment must choose how to mitigate the risk of relying only on the service side audit logging.

220 Many transactions using HTTP REST will include query parameters that would be identifiers,
quasi-identifiers, or sensitive health topics. For example, it is common for patient identifier to be
a query parameter. With this URL pattern, the query parameters are typically visible in the server
audit log or browser history. The risk from this visibility should be mitigated in system or
operational design, by protecting the logs as sensitive data, or by designing other measures into
225 the system to prevent inappropriate exposure.

Z.9 FHIR Data Types

This section includes specific guidance and constraints that are common to use of FHIR Data types.

Z.9.1 Identifier Type

230 The HL7 FHIR standard uses the data type Identifier to express a business identifier that
uniquely identifies a thing or object (see FHIR STU3
<http://hl7.org/fhir/STU3/datatypes.html#identifier>) including document uniqueIds, medical
record numbers or patient identifiers. This concept is different than the resource identifier,
known as “logical id” or “id” in FHIR, which identifies a particular resource. (A resource
235 identifier may also be represented as an Identifier instance however.)

IHE adds constraints to the Identifier data type; requirements for populating its elements vary slightly depending on what actor is originating a transaction.

240 The FHIR Identifier type introduces a new mechanism for conveying the originating system of a
particular identifier. Whereas HL7 Version 2 and Version 3 messages identify an assigning
organization as an HD (Hierarchical Descriptor) or an OID in the “root” attribute respectively,
HL7 FHIR requires the use of a URI. This may necessitate some configuration on the part of
actors in IHE profiles to correctly map between a URI and an OID or HD to maintain
consistency with other actors which are not implementing the FHIR specification.

Both the `value` and `system` shall be populated.

245 When the `value` is a globally unique value, the `system` value shall be “urn:ietf:rhc:3986”.

A `value` that is an OID shall be represented as a URI with scheme “urn:oid:”, for example:

```
250 {  
  "system": "urn:ietf:rhc:3986",  
  "value": "urn:oid:1.2.826.0.1.3680043.2.1611.1.2.32884.10619.27943.27629.41504"  
}
```

A `value` that is a UUID shall be represented as a URI with a scheme “urn:uuid:”, for example:

```
255 {  
  "system": "urn:ietf:rhc:3986",  
  "value": "urn:uuid:13cc6fc6-55ef-4dbc-a426-e0e82dffbe42"  
}
```

Z.9.1.1 Identifier and HL7 version 3 "root plus extension"

In HL7 version 3, uniqueId can be expressed as a *root*, or as a *root plus extension*.

When converting an HL7 version 3 uniqueId to FHIR, if no *extension* is provided, the *root* shall be placed into the `Identifier.value`, and the `Identifier.system` shall be set to

260 "urn:ietf:rhc:3986". For example, the HL7 version 3 value

```
<identifier root="1.2.826.0.1.3680043.2.1611.1.2.32884.10619.27943.27629.41504" />
```

would be expressed in FHIR as

```
265 {
  "system": "urn:ietf:rhc:3986",
  "value": "urn:oid:1.2.826.0.1.3680043.2.1611.1.2.32884.10619.27943.27629.41504"
}
```

When an *extension* is provided, the *extension* shall be placed into the `Identifier.value`, and the `Identifier.system` shall be set to the *root*. For example, the HL7 version 3 value

270

```
<identifier root="1.2.826.0.1.3680043.2.1611.1.2.32884.10619.27943.27629.41504"
extension="84566" />
```

would be expressed in FHIR as

```
275 {
  "system": "urn:oid:1.2.826.0.1.3680043.2.1611.1.2.32884.10619.27943.27629.41504"
  "value": "84566"
}
```

Z.9.1.2 XDS CXi mapped to FHIR Identifier Type

In XDS, a subset of CX is defined as CXi.

The following mapping shall be used unless otherwise specified:

280 CXi.1 (id) = Identifier.value
 CXi.4 (assigning authority) = Identifier.system
 CXi.5 (identifier type code) = Identifier.type
 CXi.6 (homeCommunityId) = <not mapped>

Thus, a CXi value such as

```
285 2013001^^^&1.2.3.4.5.6&ISO^urn:ihe:iti:xds:2013:accession
```

would be expressed in FHIR as:

```
290 <identifier>
  <type>
    <coding>
      <system value="urn:ietf:rhc:3986"/>
      <code value="urn:ihe:iti:xds:2013:accession"/>
    </coding>
```

```
295      </type>
      <system value="urn:oid:1.2.3.4.5.6"/>
      <value value="2013001"/>
</identifier>
```

Z.10 Profiling conventions for constraints on FHIR

The following terms refer to the values used in the OPT column of tables in ITI Technical Framework Volumes 2, 3, and 4 that define constraints being profiled:

- 300 **R** Required. This element is required by FHIR. A sending application shall populate the element with a non-empty value. A receiving application may ignore the information conveyed by the element. A receiving application shall not raise an error solely due to the presence of the element, but may raise an error due to the absence of the element.
- 305 **R+** Required. This element is required by IHE profiling, but is not a required element by FHIR. This element shall be treated as "R", above.
- 310 **R2** Required if known. If the sending application has data for the element, it is required to populate the element with a non-empty value. If the value is not known, the element may be omitted. A receiving application may ignore the information conveyed by the element. A receiving application shall not raise an error solely due to the presence or absence of the element.
- 315 **O** Optional. At its discretion, a sending application may populate the element with a non-empty value. A receiving application may ignore the information conveyed by the element. A receiving application shall not raise an error solely due to the presence or absence of the element.
- 320 **C** Conditional. There is a stated condition associated with the element. If the condition is true, the element shall be treated as "R", above. If the condition is false, the element shall be treated as "O", above.
- 320 **X** Not supported. A sending application shall not populate the element. A receiving application may, if the element is received, ignore the information conveyed by the element, or may raise an error due to the presence of the element.

Appendix E Usage of the CX Data Type in PID-3 Patient Identifier List

Add the following new section to the end of ITI TF-2x: Appendix E Usage of the CX Data Type in PID-3 Patient Identifier List

E.3 FHIR Identifier Type

325 The HL7 FHIR standard uses the data type Identifier to express a business identifier that uniquely identifies a thing or object (see FHIR STU3 <http://hl7.org/fhir/STU3/datatypes.html#identifier>) including medical record numbers or patient identifiers. See Appendix Z.9.1 for general guidance on FHIR Identifier datatype. This concept is different than the resource identifier, known as “logical id” or “id” in FHIR, which identifies a resource. (A resource identifier may also be represented as an Identifier instance however.)

330

This section specifies how IHE profiles use the Identifier data type in FHIR resources.

IHE adds constraints to the Identifier data type; requirements for populating its elements vary slightly depending on what actor is originating a transaction.

335 The FHIR Identifier type introduces a new mechanism for conveying the originating system of a particular identifier. Whereas HL7 Version 2 and Version 3 messages identify an assigning organization as an HD (Hierarchical Descriptor) or an OID in the “root” attribute, respectively, HL7 FHIR requires the use of a URI. This may necessitate some configuration on the part of actors in IHE profiles to correctly map between a URI and an OID or HD to maintain consistency with other actors which are not implementing the FHIR specification.

340 IHE imposes the following restrictions on the FHIR Identifier datatype for a Patient:

- Both the `value` and `system` shall be populated. See Appendix Z.9.1 Identifier Type
- The `assigner` attribute may be populated (the name of the organization which assigned the identifier). When the assigning authority name is provided, the actor shall also populate the `display` attribute.

345