Briefing Paper
NEHTA FHIR Foundations: Canonical URIs and Structure Definitions

Date prepared:	26 November 2015
Prepared by:	Reuben Daniels
Prepared for:	NEHTA Architecture Review Board

Briefing point

It is recommended that the NEHTA Architecture Review Board (ARB):

- Endorse the following URI path for use in assigning canonical URIs for NEHTA owned and managed FHIR technical artefacts:

  `http://ns.electronichealth.net.au/fhir`

- Endorse the following pattern for assigning canonical URIs for NEHTA owned and managed FHIR Structure Definition technical artefacts:

  `http://ns.electronichealth.net.au/fhir/StructureDefinition/<path>/<id>/<version>`

- Endorse the approach for resolving NEHTA Structure Definition canonical URIs in traditional web browsers and FHIR client software applications
- Endorse the following foundational conventions for NEHTA owned and managed FHIR Structure Definition technical artefacts:
  - a convention for key metadata values for NEHTA Structure Definition technical artefacts
  - a convention for referencing Structure Definition resource instances in specifications and technical artefacts

Relevant background information

FHIR® Structure Definitions

Structure Definition resources are part of the FHIR specification and provide a technical means of defining FHIR resources, extensions, and data types as well as the associated rules of these structures (e.g. assigning cardinalities, binding to specific terminologies, etc.). Instances of the FHIR *StructureDefinition*[1] resource type are used to define Structure Definitions in FHIR. These resources allow the Structure Definitions to be shared or published in repositories, thus providing a basis for technical specifications, validation, software source code, reporting, or user interface generation. Part of the definition of these resources include assignment of Canonical URIs which uniquely identify the Structure Definition regardless of the environment it may be found in e.g. identification of the structure definition in multiple FHIR Servers.

The need for FHIR Structure Definitions in NEHTA

As part of the National Clinical Terminology Service (NCTS) program of work, NEHTA will be providing infrastructure and application services which implement parts of the FHIR Terminology Services[2] subsystem of the FHIR specification and adopting a number of core terminology-specific FHIR resource types (e.g. Value Sets, Concept Maps, and Conformance resources). FHIR profiles and extensions are required to ensure the core FHIR resources suit the NCTS's specific use cases and these need to be defined by NEHTA through definition of a number of Structure Definitions. In support or these definition, a good foundations of managing Structure Definition within NEHTA is required. These foundations may then

[1] http://hl7.org/fhir/structuredefinition.html
[2] http://hl7.org/fhir/terminology-service.html

be leveraged and reused in other areas of the business e.g. Clinical Informatics logical modelling, future Template Service Packages for payloads based on FHIR, etc.

Issues to be aware of

1. **A URI path for providing canonical URIs for NEHTA FHIR-specific technical artefacts**
   In FHIR, Universal Resource Identifiers (URIs) are used to identify namespaces and to provide canonical identifiers for technical artefacts e.g. code systems, identifiers, and FHIR resources. In CDA these are typically represented using Object Identifiers (OIDs). URIs are preferred in FHIR as these provide:
   - an indication of the owner/maintainer of the artefact (i.e. usually through resolution or evaluation of the hostname in Uniform Resource Locator (URL) form of URIs, or querying well maintained registries e.g. Internet Assigned Numbers Authority's (IANA) URN namespaces[3].
   - given the human-readable nature of URIs, when used thoughtfully, they can provide a good indication of the use/purpose of the URI
   - the URL form of URIs may be resolvable in a commodity web browser on the Internet which may provide further information (or metadata) about the URI

   NEHTA currently uses the `ns.electronichealth.net.au` Internet domain to assign URIs for technical artifacts e.g. **`http://ns.electronichealth.net.au/id/hi/ihi/1.0`** is the canonical URI for the Healthcare Identifier (HI) Service's Individual Healthcare Identifier (IHI).

   A decision is required regarding the creation of a new URI path under the **`http://ns.electronichealth.net.au`** URL for assignment of canonical URIs for NEHTA owned and managed FHIR-specific technical artefacts. FHIR-specific technical artefacts generally correspond (but are not limited) to infrastructural FHIR resource types such as Structure Definition, Naming System, Data Element, Conformance, Value Set, and Concept Map.

   Proposal:
   Assign the **`http://ns.electronichealth.net.au/fhir`** URI path as the start of a new URI hierarchy for FHIR-specific technical artefacts.

   Notes:
   - This URI will not be a canonical URI itself. However, all NEHTA canonical URIs for FHIR-specific technical artefacts will use this URI as a base.
   - While NEHTA FHIR-specific technical artefacts will be based on this URI, NEHTA specifications (and software implementations) may use or incorporate other canonical FHIR URIs e.g. those from the HL7 FHIR specification itself; or URIs that capture implementation-specific behavior.

2. **Canonical URIs for NEHTA Structure Definition resources**
   With a URI path for FHIR-specific technical artefacts established, a path (one level down in the hierarchy) is required for Structure Definition canonical URIs. Further, a pattern for assigning specific canonical URIs must be agreed and be flexible enough to support the current and future needs NEHTA may have for Structure Definition canonical URIs.

   Proposal:
   - Assign **`http://ns.electronichealth.net.au/fhir/StructureDefinition`** as the URI base path for FHIR structure definitions
   - Adopt the following pattern for all NEHTA Structure Definition canonical URIs
     `http://ns.electronichealth.net.au/fhir/StructureDefinition/<path>/<id>/<version>`
     where,
     `<path>` is a multi-component, human readable path to further contextualize or categorize the URI e.g. ncts, pcehr/template, ci/model, etc.
     `<id>` is a string of characters ensuring uniqueness when combined with the URI components preceding it

[3] http://www.iana.org/assignments/urn-namespaces/urn-namespaces.xhtml

<version> is string of characters representing a version identifier based on an orderable versioning scheme e.g. 1.0.0, 20151231, etc.

Examples:

http://ns.electronichealth.net.au/fhir/StructureDefinition/ncts/profile/ValueSet/1.0.0
http://ns.electronichealth.net.au/fhir/StructureDefinition/ci/scs/DischargeSummary/3.0.1
http://ns.electronichealth.net.au/fhir/StructureDefinition/ci/extension/multipleBirthIndicator/1.0.2
http://ns.electronichealth.net.au/fhir/StructureDefinition/nash/SignatureProfile/1.0.0

Notes:

- The URI pattern should adhere to the syntax rules for URIs as specified in RFC 3986[4]
- The **http://ns.electronichealth.net.au/fhir/StructureDefinition** URI is not a canonical URI itself. It is a path common to all canonical URIs for NEHTA Structure Definitions.
- The **<path>** URI component shall include at lease two path components and should be descriptive to users. The first component shall indicate an business area or team within NEHTA who has management of all URIs starting with this component e.g. `ci`, `ncts`, `nash`. The second (and further path components) provide a human-readable description of the technical artefact e.g. `profile/ValueSet`
- The **<version>** part does not prescribe any specific semantics relating to versions. It is recommended that versions be human readable and ordered. The following approaches to versioning are recommended:
    a. x.y.z according to Semantic Versioning 2.0.0[5]
    b. YYYYMMDD-X where YYYYMMDD is the date associated with the version in which YYYY indicates year, MM indicates numerical month, DD indicates day (all with leading zeroes as required to ensure an YYYYMMDD is always eight digits) and X is a positive integer starting with 1 indicating the number of the release on that date (i.e. 1 is the first release, 2 the second, etc.)
- Although instantiations of the proposed canonical URI pattern closely resemble FHIR RESTful API endpoints, they are not. These canonical URIs serve as FHIR Business Identifiers[6] and may be referenced as values in fields intended for business identifiers in FHIR resource instances. Accordingly, the actual Structure Definition FHIR resources (`StructureDefinition` resources) may exist in FHIR Servers and available through FHIR endpoints with different Logical and Version Id resource metadata[7].
- Groups in NEHTA who create or maintain Structure Definitions and their canonical URIs may decide appropriate policies related to the use of this the **<path>** and **<version>** components of their canonical URIs. Further, these groups should ensure that the broader NEHTA policies and guidelines regarding NEHTA URIs and Namespaces are adhered to e.g. NEHTA Namespace Management Policy[8] and the NEHTA Technical Identifiers URI Scheme[9]

3. Approach for resolving NEHTA FHIR-specific canonical URIs

Although not strictly required, it is often useful (to implementers and other interested parties) for URL-form http URIs to be resolvable outside of their primary technical use e.g. when entered into a traditional web browser the canonical URI for W3C XML Schema Instance (`http://www.w3.org/2001/XMLSchema-instance`) resolves to an HTML web page providing a link to the associated technical specification. This is currently supported for most of NEHTA's (`ns.electronichealth.net.au`) canonical URIs.

For FHIR technical artefacts (defined using machine readable FHIR resources) it is also useful to have these URIs resolve to the associated FHIR resource definition when requested by FHIR client software applications.

In practice, a FHIR client application would request the canonical URI as an endpoint via HTTP with an `Accept` HTTP request header provided as either the `application/xml+fhir` or `application/json+fhir` media types. The responding HTTP Server then returns a permanent or

[4] https://tools.ietf.org/html/rfc3986
[5] http://semver.org/spec/v2.0.0.html
[6] http://hl7.org/fhir/resource.html#identifiers
[7] http://hl7.org/fhir/http.html#2.1.0.2
[8] http://confluence.nehta.net.au/download/attachments/47156618/Namespace%20Management%20Policy%20v03.doc?api=v2
[9] http://confluence.nehta.net.au/download/attachments/47156618/Technical%20Identifiers%20URI%20Scheme%20v0.3.1.docx?version=1&modificationDate=1338360791000&api=v2

temporary HTTP re-direct to an authentic FHIR Server RESTful endpoint which returns the FHIR resource as either XML or JSON based on the HTTP request Supporting these interactions facilitates retrieval of the associated FHIR technical artefacts and promotes enhanced tooling solutions and sharing of artefacts.

Proposal:

- When entered into a traditional web browser (creating an HTTP request with `Accept` header provided as `text/html`) canonical URIs for NEHTA FHIR-specific technical artefacts shall resolve to a web page. The body of the web page shall contain, at a minimum, hyperlinks (`<a>` HTTP inks) to the following:
    - links download the associated FHIR technical artefact (as a FHIR resource) in both XML and JSON formats.
    - links to all related specifications which describe the semantics of the artefact

- When used as endpoints by a FHIR client software applications, canonical URIs for NEHTA FHIR-specific technical artefacts must be processed by an HTTP Server which redirects the client to an authentic FHIR Server (as described above) which returns the associated FHIR resource definition in either XML and JSON format (dependent on the request).

4. Conventions for metadata and referencing NEHTA Structure Definitions

As part of adopting FHIR, NEHTA needs to establish a number of foundations to facilitate use of the specification in eHealth specifications. Accordingly, a number of foundational conventions are required to ensure creation of Structure Definitions by NEHTA staff is carried out in a considered, manner to establish a good balance of best practice, flexibility, and uniformity. This will ensure increased adoption of these Structure Definition resources internally and by our stakeholders.

Specific conventions for Structure Definitions have been identified:

- Values for metadata fields[10] in NEHTA Structure Definition resource definitions
- Referencing NEHTA Structure Definition resources in technical specifications and from other FHIR resources

Proposal:

- The following Structure Definition metadata fields should always be assigned when creating NEHTA Structure definitions

Metadata Field[11]	Convention
`url`	The canonical URI for the Structure Definition using the proposed pattern
`identifier`	Additional `identifier` elements may be provided where they exist or are necessary e.g. a NEHTA Object Identifier corresponding to the canonical URI.
`version`	The `version` should be excluded as the version of the Structure Definition is part of the `url`

- The following conventions should be adopted when referencing a NEHTA Structure Definition resources

Reference Type	Convention
Technical Specification or Design	The full canonical URI should be provided

[10] http://hl7.org/fhir/structuredefinition.html#6.17.3.1
[11] http://hl7.org/fhir/structuredefinition-definitions.html#6.17.8

A profile reference in FHIR Resource metadata	The full canonical URI should be provided e.g.
	`<meta>`
	`  <profile="`**`<canonical URI>`**`"/>`
	`</meta>`

Reference to an extension in a FHIR resource	The full canonical URI should be provided e.g.
	`<extension url="`**`<canonical URI>`**`">`
	`  <!-- extension content -->`
	`</extension>`