



# Designing URI Sets for the UK Public Sector.

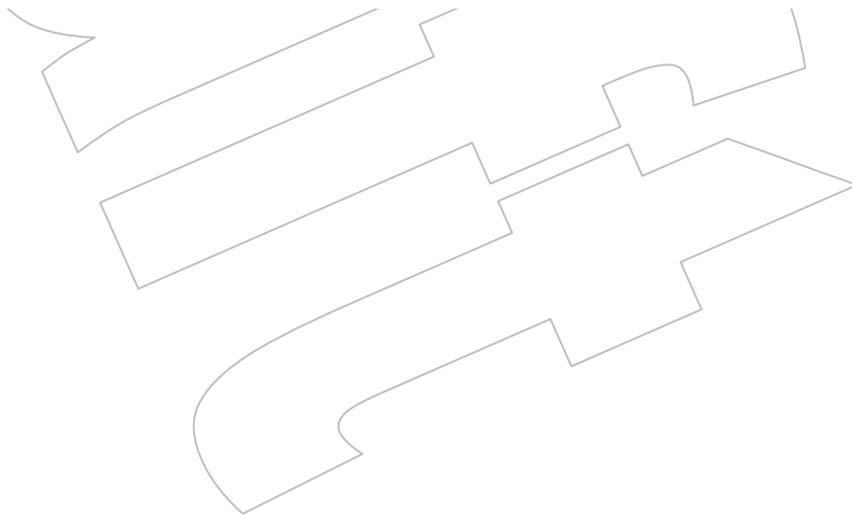


A report from the Public Sector Information Domain of the CTO Council's cross Government Enterprise Architecture.

Draft – for consideration by the Information Domain

Version 2.0b

May 2010





## Change Control

Version	Date	Editor	Note
1.0	09/10/2009	Paul Davidson CIO Sedgemoor District Council	This document is the result of a series of workshops organised by the Chief Technology Officer (CTO) Council's Information Domain during July and August 2009, and wider feedback.
2.0a	22/04/2010	Paul Davidson CIO Sedgemoor District Council	Updates the guidance basen on the Chief Technology Officer (CTO) Council's Information Domain workshops during March and April 2010, and feedback to the first published version.
2.0b	04/05/2010	Paul Davidson CIO Sedgemoor District Council	Incorporates feedback to 2.0a

### Changes from Version 1.0

- Consistent use of the term 'set'.
- 'Universal Resource Identifier' changed to be 'Uniform Resource Identifier'
- Change use of the term 'concept' to 'class'
- New section for 'Providing Information about a URI'  
New diagram showing how to combine information from many sources
- Extra definition of Representation URI
- New section introducing Temporary or 'Proxy' URI Sets
- Set URI  
Remove the suggestion to use the word the **set** in the URI. **Id/class** is now used to name the set which can be dereferenced to the document providing set characteristics and optionally a list of identifier URIs in the set.
- List URI  
Removed by combining it into the Set URI
- Advice of selecting a Reference to use for a set.
- Subscribing to a URI Set  
Extra advice on why this is recommended.
- Governance  
New table showing the responsibilities of each level of governance
- Contributors  
Added Oxford University, Talis, Birmingham City Council, Sedgemoor District Council, Scottish Government, Welsh Assembly, Home Office, Central Office of Information, Epimorphics

## Contents

Introduction	1
Definitions, frameworks and principles.	2
Designing a URI set	4
Choosing the right domain for URIs	6
The path structure for URIs	6
Providing Information about the 'Things' that URIs name	8
Looking up a URI	8
Coping with change and the passage of time	9
Examples of URIs within a set	9
Machine and human-readable formats	10
Governance arrangements	11
Appendix 1: Contributors	11



## Introduction

1. This paper is a part of the family of policy and guidance documents associated with the cross-Government Enterprise Architecture (xGEA) and the Government ICT Strategy. Further guidance will be added to cover related topics such as
  - Defining classes to share meaning
  - Publishing Linked Open Data
  - Querying Linked Open Data
2. This document defines the design considerations and guidance by which UK Public Sector Uniform Resource Identifier (URI) sets should be developed and maintained. They are designed both to encourage those that definitively own reference data to make it available for re-use, and to give those that have data that could be linked, the confidence to re-use a URI set that is not under their direct control.
3. This document will be of direct interest to
  - Owners of reference data in the UK public sector
  - Data owners who wish to improve the re-use of their data by incorporating URIs that they do not control.
  - Solution providers to the UK Public Sector
4. Some definitions and frameworks are laid out to define the types of resources that URIs can name, and the relationships between those types. A number of principles are then proposed against which, a series of design considerations are made.

These include:

- Choosing the right domain for URI sets
- The path structure for URIs
- Coping with change and the passage of time
- How to 'Look up' a URI
- The quality characteristics that apply to all URIs within a set
- Machine-readable and human-readable formats
- The governance arrangements necessary to bring the confidence to use and re-use UK Public Sector URIs

### ***How URIs can be used to publish public sector reference data***

5. Typically, government departments and agencies keep a list for each type of 'Thing' that they are responsible for, or handle in some way, and associate an identifying reference to each entry on the list. They then make use of that identifier as they make statements about the 'Thing' in their data. The lists therefore contain the 'Reference Data' that provide a common meaning and common identifier to refer to the same 'Thing' within that department or agency.
6. URI sets provide an opportunity to share common meaning and common identifiers across the public sector, and with the public, to join-up otherwise disparate data from many sources. Those that have confidence that the set is fit for their purpose are then likely to re-use it, rather than create their own.
7. Uniform Resource Identifiers (URIs), a component of the World Wide Web, provide one means of uniquely naming a 'Thing' (or 'Resource'). The principles of 'Linked Open Data' rely on the RDF data model, where statements are made about resources, identified by URI(s).
8. URI sets will be an integral component of a UK Public Sector Information Architecture that supports many goals including the release of government data, reduced duplication, and increased information sharing towards transforming government services.
9. URI sets can be published by the UK public sector to provide comprehensive and reliable identifiers for 'Things' such as schools, roads, legislation, locations, projects, events and so on. Where the quality of these sets can be described consistently, other data owners will have the confidence to re-use them in their own data, leading to a web of data that can be linked, queried, and aggregated.
10. The existing UK Public Sector standards for metadata and 'findability' work well when applied to documents, but are not sufficient to support a 'Web of Data', where each individual statement can be queried and linked.

### ***The need for design rules and guidance***



- When this guidance was first issued in October 2009, there were only a handful of early adopters of URI sets in the UK Public Sector, such as
  - The BBC
  - Ordnance Survey
  - The Office of Public Sector Information

It was noticeable that while each faced similar design issues and choices, the implementations were quite different. This guidance has served to provide a common pattern for solving common design issues; consequently, more UK Public Sector URI Sets are being created with a consistent look and feel, which will lead to confident take up and deployment.

- Much of the design is based on established and emerging good practice, whereas some implementation decisions are made to meet the specific needs of the UK Public Sector. In brief, these include
  - Use of **data.gov.uk** as the domain to root those URI sets that are promoted for re-use.
  - Organisation of URI sets into 'sectors' ( e.g. education, transport, health ) with a lead department or agency.
  - Consistent use of metadata to describe the quality characteristics of each URI set.

### Evolution of the design considerations

- This document was originally produced via a series of workshops organised by the Chief Technology Officer (CTO) Council's Information Domain during July and August 2009, and wider feedback to early drafts. This version has been produced via a Chief Technology Officer (CTO) Council's Information Domain workshops in March and April 2010, taking on feedback from early adopters, and comments received by the Cabinet Office's Write-to-Reply web facility.
- Further work is underway to interpret this guidance for specific contexts. For example, the Information Domain is working with the UK Location Council to produce a document titled "Designing URI Sets for Location".
- A web based community of practice will be used to provide supporting ...
  - Technical implementation guidance and bindings
  - Worked examples
  - Glossary of terms and definitions
  - Links to the published material that was used to support this guidance
  - Links to further material and good practice

## Definitions, Frameworks and Principles

### Types of URI

- URIs can be used to name, ...

Type of Resource	Type of URI to name the resource	Definition / Scope
Real-World 'Things'	Identifier URI	<p>These are the physical and abstract 'Things' that may be referred to in statements.</p> <p><u>Example of Physical Real-World 'Things'</u> A school, a person, a road</p> <p><u>Example of Abstract 'Things'</u> A government sector, an ethnic group, an event</p> <p>Documents or 'works' are also examples of Real-world things that can be named in this way as distinct from the content that they contain.</p> <p>Real-World Things can be referred to as 'Things' ( with a capital 'T' )</p> <p>A Real-World 'Thing' cannot be found on the web, whereas information about it can. It is important therefore to be able to distinguish between a Real-World 'Thing', as distinct from information about it, when making</p>



		statements that refer to it.
<b>Information on the web about Real-World 'Things'</b>	Document URI	These name the documents that are located on the web which are explicitly linked by the publisher of each 'Identifier URI' to provide information about Real World things.
	Representation URI	Where the publisher of a 'Document URI' provides more than one format, each format may be separately named by a Representation URI.  Depending on the formats, some Representation URIs may name documents which are machine-readable, and can therefore provide further links about the named resource.  Representation URI refer to variants of a more generic resource identified by a Document URI. Variant resources provide access to more restrictive set of web representations than are accessible via the generic Document URI. Typically a Representation URI will designate a variant resource that has a single available web representation.
<b>Definitions of classes</b>	Ontology URI	Whereas a 'Real-World Thing' identifies an individual instance of that thing, there is also a need to provide a definition of the class.  The 'Ontology URI' could be looked up to give that definition.
<b>Relationships between things</b>	Ontology URI	Each part of an RDF statement can be named using a URI. This includes the relationship between 'Real-World Things'.  The 'Ontology URI' will then give a link into an ontology that can provide further reasoning about the relationship and the classes that can be related using it.

### What is a URI Set?

- For the purposes of this document, the term 'URI Set' is to mean a collection of reference data published using URIs, about a single class, governed from a single source.
- An additional type of URI is required to name the URI Set itself.

<b>A URI Set</b>	Set URI	A type of Identifier URI that names the URI set and can be resolved to provide <ul style="list-style-type: none"> <li>the quality characteristics of the set</li> <li>optionally, a list of the Identifier URIs that are contained within a set</li> </ul>
------------------	---------	--

### Design principles for Public Sector URI Sets.

- The following principles have been derived from existing good practice and revised to meet the challenges for UK Public Sector URI sets.

Principle	
Use HTTP so that URIs can be resolved.	MUST
Use a consistent path structure to explicitly indicate the type of URI.	RECOMMEND
The publisher will make it clear whether the set is promoted for re-use by other parts of government and/or the public.	MUST
Public Sector URI sets should publish their expected longevity, and potential for re-use.	MUST
Those Public Sector URI sets that are promoted for re-use should be designed to last for at least 10 years.	RECOMMEND
Where more than one Representation URI is available, provide a Document URI where Content-Negotiation can be used to provide the most appropriate representation	RECOMMEND



Avoid exposing the technical implementation of a URI in its structure.	RECOMMEND
As a minimum, provide a machine-readable Representation URI	MUST
If appropriate, provide a human-readable Representation URI in HTML	RECOMMEND
Provide a means of discovering each of the available Representation URIs for a single Document URI	RECOMMEND
A URI Set will publish its Authorisation, Authentication, and Data Quality characteristics using a common vocabulary.	MUST
A URI structure will not contain anything that could change, such as session IDs.	MUST
A URI path structure will be readable so that a human has a reasonable understanding of its contents.	RECOMMEND

5. When considering which part of the UK public sector should set up URI sets ...

Principle	Considerations
The department or agency responsible for a 'Real World Thing' should also be responsible for defining it, and naming instances of it, on behalf of the appropriate sector.	URIs should be organised into sectors with a lead department or agency.  Lead departments/agencies should engage with stakeholders to ensure that the set is of sufficient quality to meet a wide range of purposes.
URIs from a Set that is promoted for re-use should not contain the name of the department or agency currently responsible for it.	This copes with machinery of government changes where a department or agency may cease or have its scope changed.

## Designing a URI Set

### What will a URI Set contain?

- A URI Set will contain
  - A URI to name the set and describe its quality characteristics
  - Each of the Identifier URIs for the Real-World 'Things' in a single class.
  - Optionally, Ontology URI to define the set's class, and relationships.
  - Optionally, the Set URI to list the Identifier URIs contained in the set.

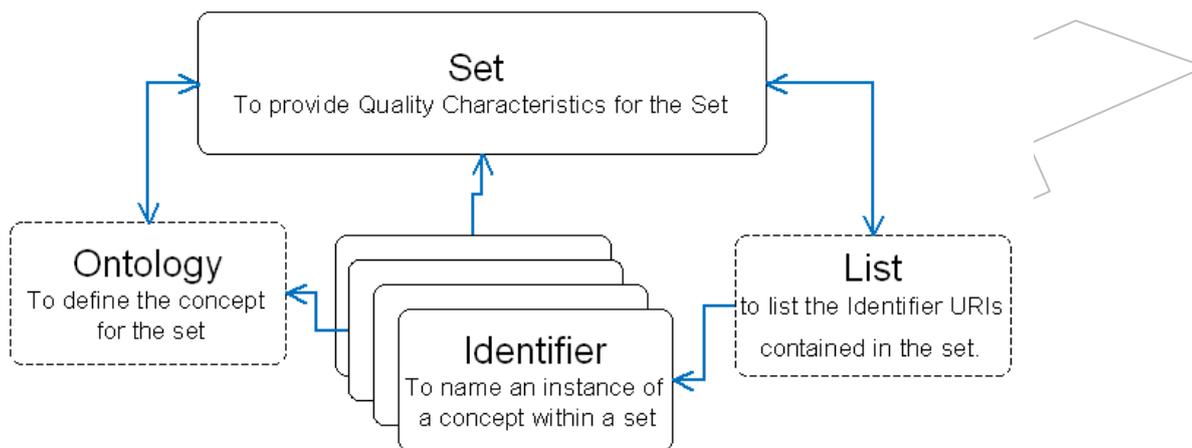


Figure 1: URIs that make up a Set

- Each type of URI will provide a means of looking up data on the web about the resource that the URI names, using
  - A Document URI
  - At least one machine readable Representation URI for each Document URI
- The publisher of a URI Set will implement facilities to provide mechanisms to
  - Lookup an Identifier URI and be redirected to its Document URI



- Discover each of the Representation URIs that are available for a Document URI
  - Get the most appropriate Representation URI for a specified application.
4. Further guidance about how data can be looked up via Document and Representation URIs is provided later in this document.

### **Quality characteristics to publish for a URI Set**

5. The following information should be provided as a part of the metadata to describe the URI Set

<b>Class definition</b>	Either by <ul style="list-style-type: none"> <li>• an Ontology URI that resolves to an machine readable definition</li> <li>• as human readable metadata</li> </ul>
<b>Relationships to other URI sets</b>	To highlight other associated URI sets
<b>Provenance</b>	To describe the source and purpose of the reference data
<b>Official Status</b>	To describe the range of statuses of the identifiers that that are contained in the set.
<b>Accuracy</b>	To describe the closeness to the truth that the set attempts to achieve.
<b>Completeness</b>	To describe the degree to which the Identifier URIs are a complete set against the definition of the class.
<b>Timeliness</b>	To describe the time-lag between a change to a 'Real-World Thing' being applied to the URIs in the set.
<b>License Terms</b>	To describe the terms of use for the URI set.
<b>Intended Longevity</b>	To provide a guarantee of persistence of the set.
<b>Intended Audience</b>	To describe who may confidently use the set. This provides a means of marking the set as being promoted for re-use.
<b>Representations available</b>	To describe the range of file formats of Representation URIs in the set.

6. Guidance as to the bindings and vocabularies used to express this metadata will be provided via the associated community of practice. This is likely to take the form of an RDF binding to an application profile of the e-Government Metadata Scheme ( eGMS ) and/or Dublin Core properties and classes, or to sector specific metadata profiles for example, INSPIRE as applied to location.

### **Temporary or 'Proxy' URI Sets**

7. There may be occasions where a URI Set is created other than by the sector that has authority over assigning identifiers for that class. This may be because:
- The set does not exist yet.
  - Individual URIs are not present in the definitive URI Set.
  - There is disagreement about the definition of the set.
  - The quality characteristics of a definitive set are not sufficient for a purpose.
8. Sets such as these should not be promoted for re-use, and should be marked as such using the set's quality characteristics. Apart from that, they should follow the same pattern as other URI Sets.
9. The sector that does have authority should be encouraged to provide a URI Set to give authoritative standing on web.

### **How should URI sets be discovered?**

10. A consistent set of metadata will be required to support the expectation that a single query can discover the UK Public Sector URI Sets that are promoted for re-use, for a given subject.
11. Publishers may wish to offer a facility to 'subscribe' to their set. This has many benefits including:
- Knowledge of dependencies help to justify resourcing
  - Alerting consumers to changes, and improvements
  - Automating processing required when alterations are applied to a set
12. Guidance as to discovery metadata and subscription methods will be provided via the associated community of practice.

## Choosing the right domain for URIs

### General

1. When considering the domain to root a URI Set in ...
  - the publisher will require content control of the sub-domain that it ultimately resolves to
  - the domain will have appropriate service-levels and scalability for resilience and performance

### Requirements for URI sets that are promoted for re-use.

2. In addition, where a URI set is promoted for re-use, the following considerations apply to find a balance for central and federated components.
  - Flexibility & readability
  - Administrative burden
  - Infrastructure costs
3. In particular, the domain will ...
  - Expect to be maintained in perpetuity
  - Not contain the name of the department or agency currently defining and naming a class, as that may be re-assigned.
  - Support a direct response, or redirect to department/agency servers.
  - Ensure that classes do not collide
  - Require the minimum of central administration and infrastructure costs
  - Be scalable for throughput, performance, resilience.
4. The choice of domain should provide the confidence to the consumer, that the URI set has met minimum quality criteria, including implementing these design considerations. In other words, the domain itself should convey an assurance of quality and longevity.
5. Due to the drive to rationalise websites and also to separate presentation of data from its location, UK Public Sector URIs will be based around the **data.gov.uk** domain, split by 'sectors' as sub-domains. When looking up a URI, the data.gov.uk servers either provide the response themselves, or DNS is used to redirect enquiries to the appropriate Department or Agency server.
6. A sector is NOT a department name. Sectors should be understandable by the public, rather than reflecting how government is currently organised. New departments taking over all or part of a sector are required to maintain the URI sets. The community of practice will provide further information about the use of sectors which are likely to be aligned to other initiatives such as Directgov.
7. Using 'education' as an example sector for a URI Set promoted for re-use, gives:

<http://education.data.gov.uk>

This:

- Shows that the set is a part of the education sector
- Puts it in the data.gov.uk collection of UK Public Sector URIs promoted for re-use
- Can be redirected using DNS to a departmental server for the content.
- Is from the 'data.gov.uk' domain and therefore not confused with a presentation web site

## The path structure for URIs

### General

1. The path structure of a URI may contain elements to:

Identify the set class	Class
	A word or string to capture the essence of the real-world 'Thing' that the set names.  e.g. school



	<ul style="list-style-type: none"> <li>• Lower case</li> <li>• Words separated by hyphens</li> <li>• Singular ( e.g. 'school', not 'schools' )</li> </ul>
<b>Identify an individual instance of a real-world 'Thing'</b>	<p><b>Reference</b></p> <p>A string that is used by the set publisher to identify an individual instance of the class.</p> <p>The reference should match the way that it is used in normal use.</p> <p>In some circumstances, a name may be appropriate as the Reference, e.g. 'England'. Where the name may change, or becomes overly verbose, a code may be more appropriate.</p> <p>When selecting the pattern of Reference to use for a set, the following considerations should be taken into account</p> <ul style="list-style-type: none"> <li>• Uniqueness      References may not clash within a Set</li> <li>• Persistence      A Reference should not change where the resource that is being named does not change even if its title does.</li> <li>• Reuse /      Using existing coding schemes</li> <li>• Naturalness      A unique attribute of the resource being named</li> <li>• Readability      There is value in a human being able to 'read' a URI, but URIs are for machines to interpret so readability should be considered in that light. Be aware of spelling and language differences.</li> <li>• Ability to Generate      A proven process to create a reference</li> </ul>

2. Examples of class/reference pairs.

- road/M5
- school/123

3. The class/reference construct may be repeated as necessary, for example:

- road/M5/junction/24
- school/123/class/5

4. Other components of a URI path may

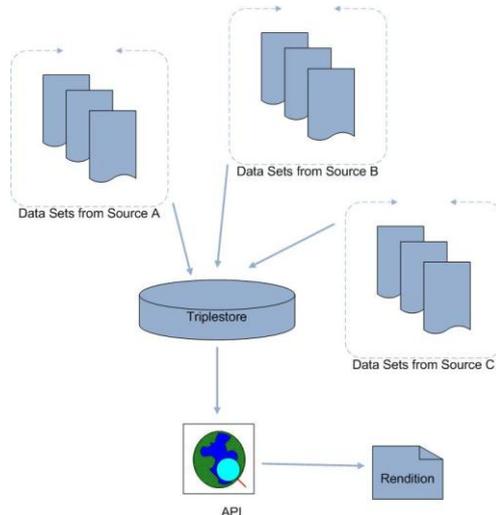
<b>Identify the type of URI</b>	<p>URI type, for example one of</p> <ul style="list-style-type: none"> <li>• id – Identifier URI</li> <li>• doc – Document URI, Representation URI</li> <li>• def – Ontology URI</li> </ul>
<b>Identify a file format</b>	<p>File-extension, to indicate the format of a document that will be returned.</p> <p>doc.{ext}</p> <p>e.g. doc.rdf, doc.html</p>

### ***The path structure for each type of URI***

5. Machines should not deconstruct the path of a URI to discover other related types. However, for consistency and readability, the following can be used as a guide.

<b>Identifier URI</b>	<p>Contains the string 'id' to show that it is an Identifier URI.</p> <p>Contains pairs of Class/Reference</p>
<b>Document URI</b>	<p>The same as its associated Identifier URI, but replacing the 'id' string with 'doc' or removing it completely.</p>
<b>Representation URI</b>	<p>The same as its associated Document URI but with a file extension</p>
<b>Ontology URI</b>	<p>{domain}/def/{class}</p>
<b>Set URI</b>	<p>Contains the name of the class. i.e. {domain}/id/{class} and {domain}/doc/{class}</p>

## Providing Information about the 'Things' that URIs name



### General

1. A URI may be used to link data without ever having to look it up. Looking up a URI is the process of 'resolving' it (also known as dereferencing).
2. The purpose of resolving a URI is to ...
  - Gain greater definition of the resource that it names
  - Discover other related resources
  - Gain some basic information which a typical enquirer is likely to find useful
3. Figure 2 proposes how the publisher of the set could arrange for information from many sources to be combined to provide a consistent document URI for identifiers across a set.

Figure 2: Combining information from many sources for a Document URI

4. This approach enables the publisher of a set to:
  - cope with the use of different identifiers and terms at each contributing dataset.
  - construct a comprehensive URI set by the union of entities described across many datasets

## Looking up a URI

### General

5. As this guidance prescribes HTTP as the protocol for UK Public Sector URIs, resolving a URI is via an HTTP GET. The client requesting the HTTP GET, may be a machine, or a human.
6. Referring to Figure 1, the consumer should not rely on syntactic manipulation of URIs to locate the other associated set components. The publisher should provide the ability to look-up between the components, for example, given just the URI of an instance of a class (say a particular school), to find the corresponding ontology document; thence locate information about the set which may include a catalogue of all the known class instances. The blue arrows in the figure represent the dereferencing of a web reference taking into account any further redirection to return with a representation of the data.
7. Some of the issues to be considered when providing mechanisms to resolve a URI are given below. The community of practice will provide greater definition of the protocols/headers/status codes etc used when resolving each type of URI.

### Resolving Identifier URIs

8. In a linked data architecture, it must be possible to resolve the URI for a real-world 'Thing' to a Document that contains information about that thing. In other words, it must be possible to resolve an Identifier URI to a Document URI.
9. There are two patterns for resolving Identifier URIs to provide a Document URI, as described in *Cool URIs for the Semantic Web*<sup>1</sup>:
  - Hash URIs
  - 303 responses.

<sup>1</sup> <http://www.w3.org/TR/2007/WD-cooluris-20071217/>



10. Either pattern meets the requirements of separating Identifier URIs from Document URIs. While 303 responses provide the most flexibility, hash URIs may prove necessary when there is limited access to the server configuration. The pattern selected by the publisher has no impact on the capability required of consumers.
11. The community of practice will provide greater definition of the protocols/headers/status codes etc used for both patterns, and implementation guidance to support a decision as to which to deploy.

## Resolving Document URIs

12. A Document URI will resolve to the most appropriate Representation URI to provide information in a format as requested by the client. Where more than one Representation URI is available containing the same information in different formats, this may be achieved using Content Negotiation.
13. It will often be useful to version the information that is available about a particular Thing, to indicate when the information contained in a Document is valid. It is also often helpful to journey back in time to view information that was valid in the past. To do this, there need to be 'dated' Document URIs that contain information valid on or from a particular date.
14. The community of practice will provide greater definition of the protocols/headers/status codes etc used when resolving Document URIs, plus advice on how to support 'dated' Document URIs, and Content Negotiation.

## Coping with change and the passage of time

### General

1. Once created, a URI should persist unaltered.
2. The essence of a real-world 'Thing' is unlikely to change, whereas a description of it at a point in time may change. For example, the essence of the M5 motorway remains unchanged from its original conception, through to its completion. If it were to be extended, it remains the M5 motorway. Consequently, a single Identifier URI for the M5 motorway may have a number of versions of its associated Document URI that might be versioned by a status, and/or date and so on.
3. Where the essence of a real-world 'Thing' naturally passes through various stages, those stages could be designed into the structure of the URI path, thus creating a separate URI for each stage. For example  
`act/1985/67/enacted`
4. The URI set publisher may provide a URI alias to the current version.
5. A URI set will explain how it copes with change and the passage of time and give advice to consumers about which 'alias URIs' are appropriate to use in their statements.
6. The community of practice will provide further scenarios of how URI sets may have to cope with 'change' and guidance as to how they can be designed to cope.

## Examples of URIs within a Set

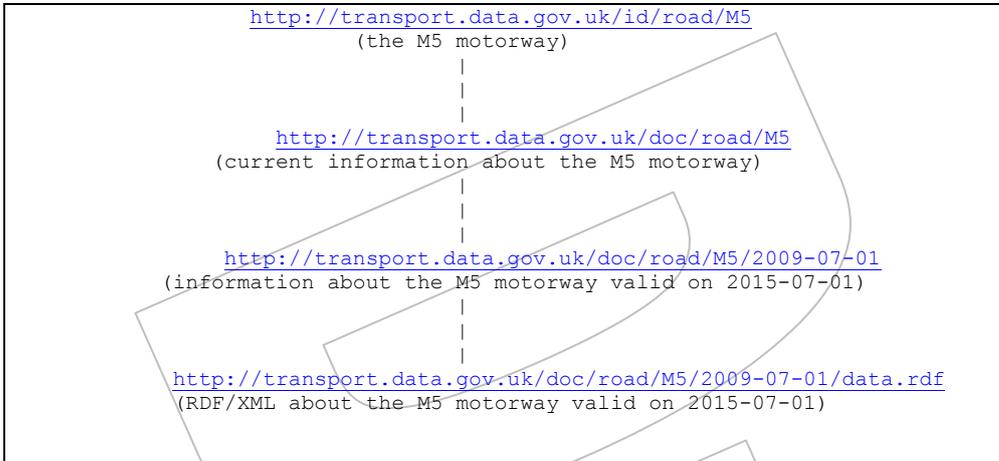
### Example of URIs within a set

URI Type	URI structure	Examples
<b>Identifier</b>	<a href="http://domain/id/class/reference">http://domain/id/class/reference</a> or <a href="http://domain/class/reference#id">http://domain/class/reference#id</a>	<a href="http://education.data.gov.uk/id/school/78">http://education.data.gov.uk/id/school/78</a> <a href="http://education.data.gov.uk/school/78#id">http://education.data.gov.uk/school/78#id</a> <a href="http://transport.data.gov.uk/id/road/M5/junction/24">http://transport.data.gov.uk/id/road/M5/junction/24</a>
<b>Document</b>	<a href="http://domain/doc/class/reference">http://domain/doc/class/reference</a>	<a href="http://education.data.gov.uk/doc/school/78">http://education.data.gov.uk/doc/school/78</a>
<b>Representation</b>	<a href="http://domain/doc/class/reference/doc.file-extension">http://domain/doc/class/reference/doc.file-extension</a>	<a href="http://education.data.gov.uk/doc/school/78/doc.rdf">http://education.data.gov.uk/doc/school/78/doc.rdf</a>
<b>Definition of the Set Class</b>	<a href="http://domain/def/class">http://domain/def/class</a>	<a href="http://education.data.gov.uk/def/school">http://education.data.gov.uk/def/school</a>



<b>Set</b>	<a href="http://domain/id/class">http://domain/id/class</a> dereferencing to <a href="http://domain/doc/class">http://domain/doc/class</a>	<a href="http://education.data.gov.uk/id/school">http://education.data.gov.uk/id/school</a> <a href="http://education.data.gov.uk/doc/school">http://education.data.gov.uk/doc/school</a>
------------	--	--

## Example of how URIs within a Set could be resolved



## Machine and human-readable formats

### General

1. When a URI is resolved and a document returned, a variety of formats may be used. The following is guidance as to what format to use in what circumstance.

### Formats for Representation URIs

2. Documents may be available in multiple formats. Each possible representation of the Document should have a distinct representation URI. For example:
  - RDF/XML at <http://transport.data.gov.uk/doc/road/M5/junction/24/doc.rdf>
  - Turtle at <http://transport.data.gov.uk/doc/road/M5/junction/24/doc.ttl>
  - HTML at <http://transport.data.gov.uk/doc/road/M5/junction/24/doc.html>
  - JSON at <http://transport.data.gov.uk/doc/road/M5/junction/24/doc.json>
3. Having distinct representation URIs for each possible representation makes it possible to access a particular representation without changing the headers that are sent with a request, which is useful for certain clients. Each document should also include pointers to the other available formats, and lists of resources should also include lists of possible representations. For example:
  - in HTML, use the <link> element with a suitable type attribute
  - in RDF, use a dct:hasFormat property
  - in Atom, use the <atom:link> element with a suitable type attribute
4. At least one representation must be machine readable in a way that enables the construction of an RDF graph. It is recommended that this representation is one of:
  - RDF/XML<sup>2</sup> (preferred, as this is supported by the largest number of tools)
  - XHTML with a GRDDL<sup>3</sup> transformation
  - XHTML with embedded RDFa<sup>4</sup>
5. It is also acceptable to provide the following formats for the construction of an RDF graph:
  - Turtle<sup>5</sup>

<sup>2</sup> <http://www.w3.org/TR/rdf-syntax-grammar/>

<sup>3</sup> <http://www.w3.org/TR/grddl/>

<sup>4</sup> <http://www.w3.org/TR/rdfa-syntax/>

- N3<sup>6</sup>
6. Other formats are also useful. In particular, it is useful to provide other machine-readable formats such as:
    - JSON
    - CSV
  7. It is also useful to provide human-readable versions of the documents in HTML, if appropriate.
  8. The community of practice will provide further examples to illustrate how information can be returned from looking up a Representation URI.

## Governance arrangements

### General

1. For these design considerations to be effective and credible, governance will need to be established at various points:

<b>Set</b>	<ul style="list-style-type: none"> <li>• Designing the set</li> <li>• Disambiguation of references so that same identifier does not point to two different resources.</li> <li>• Joining and transforming disparate datasets to provide a consistent representation of a URI.</li> <li>• Setting and monitoring quality characteristics</li> <li>• Providing subscription and alert services</li> </ul>
<b>Sector</b>	<ul style="list-style-type: none"> <li>• Engaging with stakeholders for a sector</li> <li>• Determining and prioritising URI Sets within each sector</li> <li>• Avoiding class naming collisions</li> <li>• Commitment to quality and longevity of URI Sets</li> <li>• Monitoring quality</li> </ul>
<b>UK Public Sector</b>	<ul style="list-style-type: none"> <li>• Maintaining this guidance</li> <li>• Definition of sectors</li> <li>• Allocation of lead departments/agencies for sectors</li> <li>• Accreditation of URI Sets to be rooted at data.gov.uk</li> <li>• Infrastructure</li> </ul>

## Appendix 1: Contributors

1. The CTO Council is grateful for the participation from those representing organisations including, The BBC, Ordnance Survey, The Office of Public Sector Information, The London Gazette, The Stationery Office, University of Southampton, Oxford University, Scottish Government, Welsh Assembly Government, Department for Transport, Department for Communities and Local Government, Department for Children, Schools and Families, Department for Environment, Food and Rural Affairs, the UK Location Programme, National Policing Improvement Agency, Home Office, Central Office of Information Sedgemoor District Council, Birmingham City Council, the esd-Toolkit (Local Government), the Local e-Government Standard Body (LeGSB), The Cabinet-Office, Fujitsu, Hewlett Packard, Talis, Epimorphics.

---

<sup>5</sup> <http://www.w3.org/TeamSubmission/turtle/>

<sup>6</sup> <http://www.w3.org/DesignIssues/Notation3>