Interoperability Solutions for European Public Administrations

CORE VOCABULARIES SPECIFICATION

CORE BUSINESS VOCABULARY

CORE PERSON VOCABULARY

CORE LOCATION VOCABULARY

JOINING UP GOVERNMENTS

DOCUMENT METADATA

Property	Value
Release date	*16/02/2012*
Status:	*For public review*
Version:	*0.2*

Statement of copyright:

This specification was prepared for the ISA programme by:
PwC EU Services

Disclaimer:

The views expressed in this draft specification are purely those of the authors and may not, in any circumstances, be interpreted as stating an official position of the European Commission.

The European Commission does not guarantee the accuracy of the information included in this study, nor does it accept any responsibility for any use thereof.

Reference herein to any specific products, specifications, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favouring by the European Commission.

All care has been taken by the author to ensure that s/he has obtained, where necessary, permission to use any parts of manuscripts including illustrations, maps, and graphs, on which intellectual property rights already exist from the titular holder(s) of such rights or from her/his or their legal representative.

This draft specification can be downloaded from the ISA website:
https://joinup.ec.europa.eu/asset/core_person/release/02
https://joinup.ec.europa.eu/asset/core_location/release/02
https://joinup.ec.europa.eu/asset/core_business/release/02

TABLE OF CONTENTS

1. INTRODUCTION

1.1 ABOUT THE ISA PROGRAMME

This specification has been created as part of Action 1.1 [A1.1] of the Interoperability solutions for European public administrations (ISA) programme of the European Commission (EC). This programme funds initiatives to foster the efficient and effective cross-border electronic interactions between European public administrations. Action 1.1 of this programme is targeted towards improving the semantic interoperability of European e-Government systems. It addresses these by encouraging the sharing and reuse of semantic assets. As part of Action 1.1, the ISA Programme intends to build consensus on a number of e-Government Core Vocabularies for public sector information exchange.

1.2 MOTIVATION

The **metadata** and **reference data** used in electronic public services across Europe most often has a very specific context. Attaining consensus on common metadata and reference data for these electronic services is a *critical* step towards semantic interoperability. Unfortunately, consensus building is hindered by the diverse cultural, multi-lingual, legal, and organisational contexts of these e-Government services. To alleviate this problem, consensus building should start at a higher level of abstraction that surpasses the contexts of individual electronic public services, and thus the cultural, lingual, legal, and organisational differences of individual countries. In particular, consensus can be more easily attained on the semantics of a *small* set of fundamental concepts, for which less divergent opinions exist [EGOV-CV]. These concepts are what we describe as Core Vocabularies.

1.3 TERMINOLOGY

This document uses the following terminology:
Semantic interoperability is defined as the ability of information and communication technology (ICT) systems and the business processes they support to exchange data and to enable the sharing of information and knowledge: *Semantic Interoperability enables systems to combine received information with other information resources and to process it in a meaningful manner* (European Interoperability Framework 2.0[1]). It aims at the mental representations that human beings have of the meaning of any given data.

A **Core Vocabulary** is a simplified, reusable, and extensible data model that captures the fundamental characteristics of an entity in a context-neutral fashion [EGOV-CV]. Well known examples of existing Core Vocabularies include the Dublin Core Metadata Set [DC]. Such Core Vocabularies are the starting point for agreeing on new semantic interoperability assets and defining mappings between existing assets. Semantic interoperability assets that map to or

[1] http://ec.europa.eu/isa/documents/isa_annex_ii_eif_en.pdf

extend such Core Vocabularies are the *minimum required* to guarantee a level of cross-domain and cross-border interoperability that can be attained by public administrations.

1.4 THIS SPECIFICATION

The Core Vocabularies defined here present concepts and terms that can be used to describe three things that are closely related:

- People
- Locations
- Businesses

This is obviously a very short list, one that excludes many other domains that are no less interconnected such as, accounts and payments, projects and processes, buildings, education and healthcare organisations. Any of these and many more could have been included and were considered. However, it was decided to focus, initially at least, on these three areas as there was evidence from some public sector data repositories and broader discussions that these were the highest priority areas. Furthermore, the links between them are particularly strong so that the three create a single coherent set.

In developing these three core vocabularies the working group consciously limited its work further. The vocabulary for describing a person relates to the **natural person**, i.e. the individual as opposed to any role they may play in society or the relationships they have to other people, organisations and property; all of which contribute significantly to the broader concept of *identity*. In describing businesses, the working group focussed solely on **legal entities**, that is, trading bodies that are formally registered with the relevant authority and that appear in business registers. This excludes sole traders, virtual organisations and other types of 'agent' that are able to do business. The broadest vocabulary is that for location which covers **places**, **addresses**, and geographical **geometries**. Specific vocabulary terms that were considered but not included are described in Section 5.

1.5 RELATED WORK

Care has been taken to ensure that the defined terms are most likely to match those used by others, thus maximising the interoperability and re-usability of data. A minimum number of new terms has been defined with the majority of those called for in the conceptual model being re-used from existing, widely used, vocabularies. The existing vocabularies referred to include:

- the UN Centre for Trade Facilitation and e-Business [UN/CEFACT]
- Friend of a Friend [FOAF]
- INSPIRE Addresses Technical Guidelines [INADR]
- W3C Point of Interest Core [POI]
- xEBR Core reference Taxonomy created by XBRL Europe [xEBR]
- schema.org

Further information on the approach taken is provided in Section 8 Background and Objectives.

2. CONFORMANCE STATEMENT

A publisher using the vocabularies can choose whether to publish using either RDF or XML as their technology and may use any of the terms defined in this document.

A consumer of data published using the vocabularies must understand all the terms defined below in one of three ways which should be declared when claiming conformance:

> RDF conformance, meaning that RDF data published using any term in the vocabularies will be consumed and processed accordingly;
>
> XML conformance, meaning that XML data published using any term in the vocabularies will be consumed and processed accordingly;
>
> XML and RDF conformance, meaning that both RDF and XML data published using any term in the vocabularies will be consumed and processed accordingly.

3. NAMESPACES

At the time of publication, an active discussion is ongoing concerning the namespaces to be used for the core vocabularies. Factors such as long term stability and change control are paramount. For the time being, we are simply using 'example.org' as a place holder, to be replaced in the near future.

With that in mind, we define the namespaces and suggested prefixes for the three core vocabularies thus:

Prefix	Namespace
person	http://example.org/ns/person#
legal	http://example.org/ns/legal#
locn	http://example.org/ns/locn#

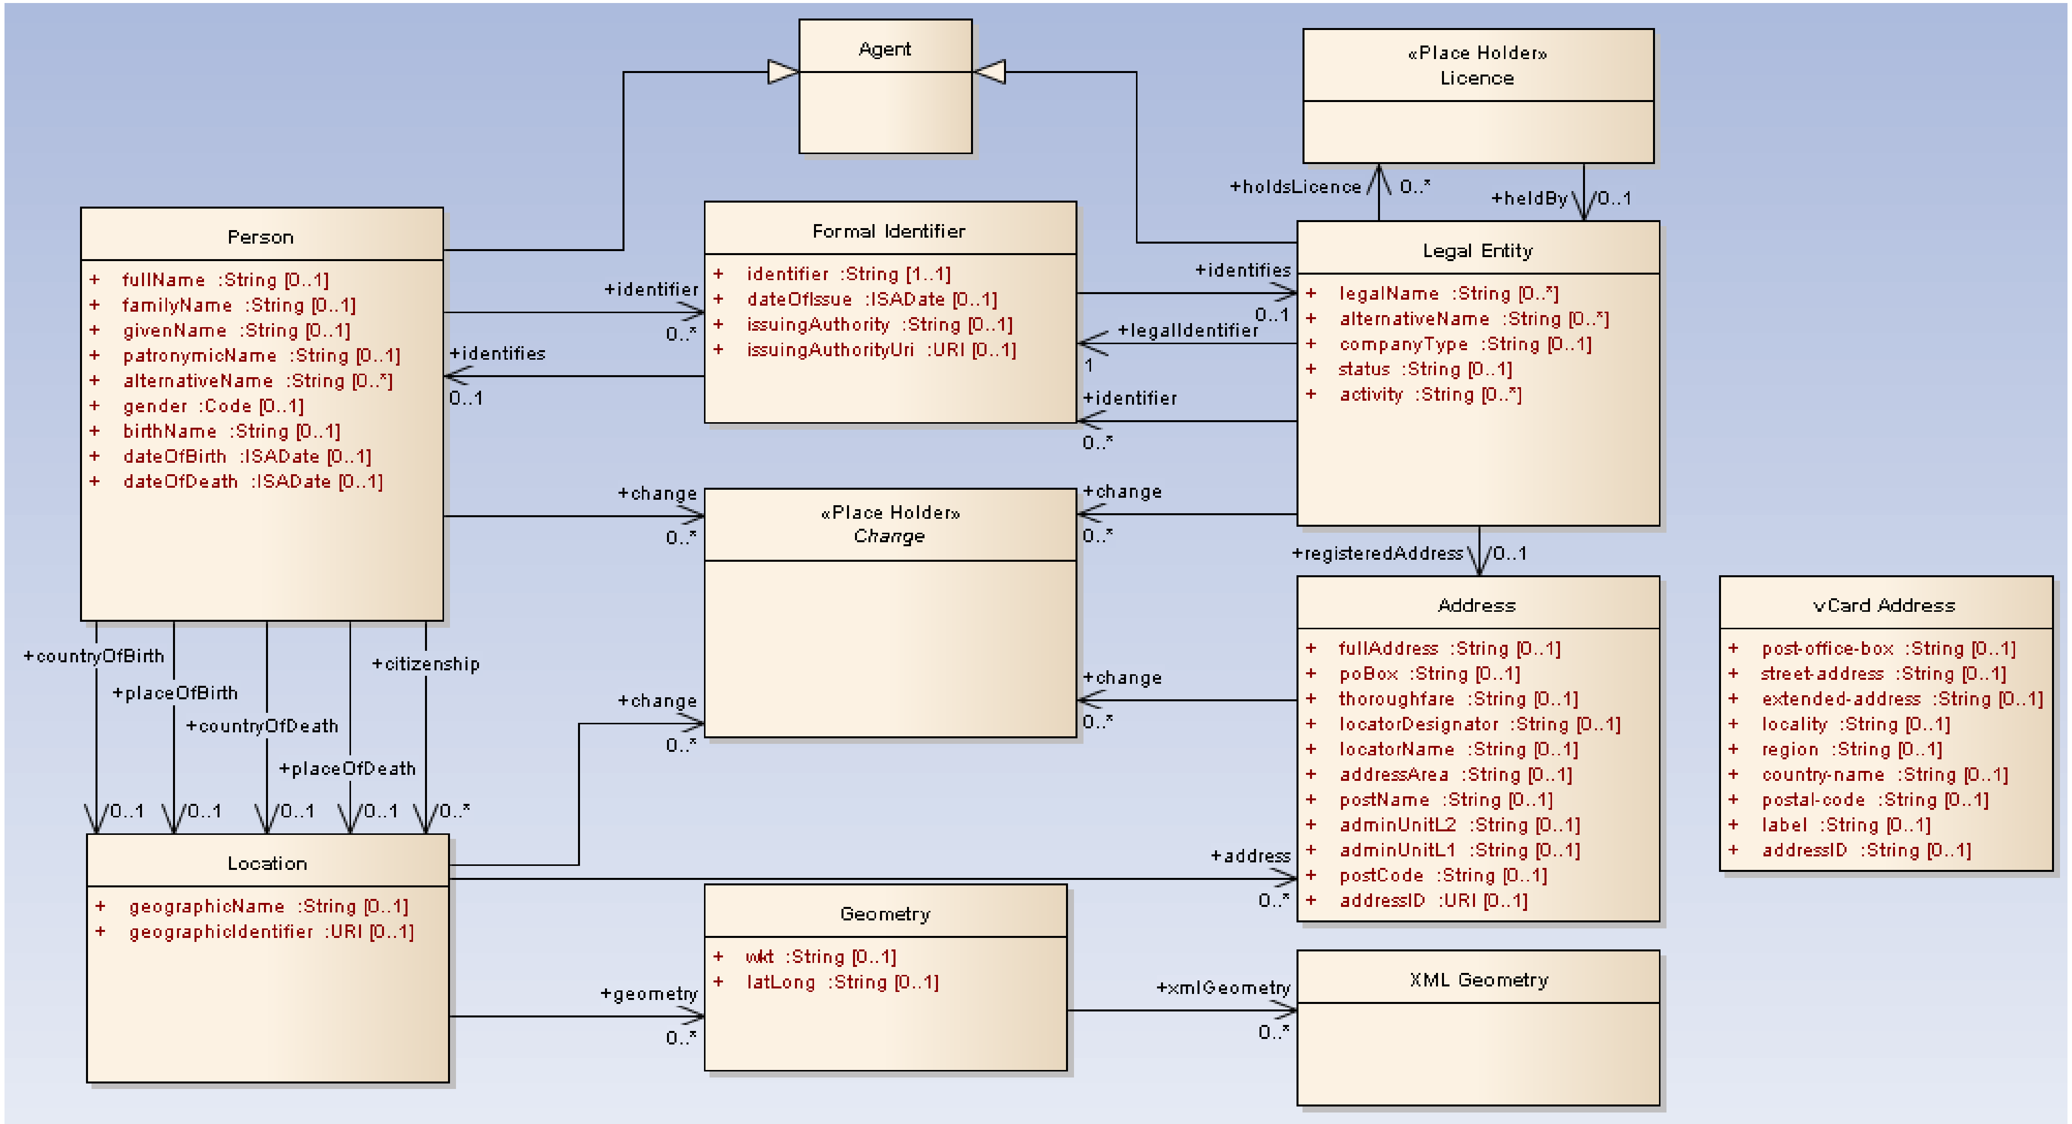

Figure 1 UML diagram for the ISA Core Vocabularies

4. CONCEPTUAL MODEL

The conceptual model presented in Figure 1 is independent of any technology that may be used to represent it. It describes the minimal set of classes, relationships and properties necessary to describe a natural person, a business that has legal entity status, and a location. Very few cardinality restrictions have been defined to ensure maximum flexibility.

In the following sections, each class is defined, followed by definitions of its properties and relationships.

4.1 THE PERSON CLASS

The Person Class is a sub class of the more general 'Agent' class that encompasses organisations, legal entities, groups etc. - any entity that is able to carry out actions.

For avoidance of doubt, the word 'relationships' is used throughout this document in its technical sense. Section 5.1 makes clear that this vocabulary does not cover the societal relationships that an individual person has with, for example, other members of their family or their employer. The concern here is with the key aspects that define the natural person.

4.1.1 Full Name

Attribute	Abstract Data Type	Cardinality
fullName	String	[0..1]

The attribute “Full Name” contains the complete name of a person as one string. In addition to the content of Given Name, Family Name and, in some systems, Patronymic Name, this can carry additional parts of a person’s name such as titles, middle names or suffixes like “the third” or names which are neither a given nor a family name.

Reasoning

A name sticks with a person for a long time period. In most European countries[1] a name may only be changed according to certain laws and life events, e.g. marriage. The name denominates a natural person even if he/she changes their address. Documents like birth certificate or diploma usually don’t carry an address but always the name. Thus the name is one of the core attributes. However it is not sufficient to identify a person since there are combinations of very common names like Smith in the UK, Meier in Germany, or Li in China.

4.1.2 Given Name

Attribute	Abstract Data Type	Cardinality
givenName	String	[0..1]

A given name, or multiple given names, are the denominator(s) that identify an individual within a family. These are given to a person by his or her parents at birth or may be legally recognised as 'given names' through a formal process. All given names are ordered in one field so that, for example, the Given Name for Johan Sebastian Bach is 'Johan Sebastian.'

4.1.3 Family Name

Attribute	Abstract Data Type	Cardinality
familyName	String	[0..1]

A family name is usually shared by members of a family. This attribute also carries prefixes or suffixes which are part of the Family Name, e.g. “de Boer”, “van de Putte”, “von und zu Orlow”. Multiple family names, such as are commonly found in Hispanic countries, are recorded in the single Family Name field so that, for example, Miguel de Cervantes Saavedra's Family Name would be recorded as "Cervantes Saavedra."

4.1.4 Patronymic Name

Attribute	Abstract Data Type	Cardinality
patronymicName	String	[0..1]

Patronymic names are important in some countries. Iceland does not have a concept of 'family name' in the way that many other European countries do, for example. In Bulgaria and Russia, patronymic names are in every day usage, for example, the Sergeyevich in 'Mikhail Sergeyevich Gorbachev.'

4.1.5 Alternative Name

Attribute	Abstract Data Type	Cardinality
alternativeName	String	[0..*]

Any name by which an individual is known. Many individuals use a short form of their name, a 'middle' name as a 'first' name or a professional name. For example, the British politician and former UN High Representative for Bosnia and Herzegovina, Jeremy John Durham Ashdown, Baron Ashdown of Norton-sub-Hamdon, is usually referred to simply as 'Paddy Ashdown' or 'Lord Ashdown.'

Note

It is not the role of the alternative name field to record nick names, pet names or other 'familiar names' that will be of no consequence in public sector data exchange.

4.1.6 Gender

Attribute	Abstract Data Type	Cardinality
gender	Code	[0..1]

A code specifying the current gender of a person. ISO 5218 specifies 4 categories in widespread use:

0 = not known;
1 = male;
2 = female;
9 = not applicable.

However, some governments have defined additional gender designations such as Scotland's "8 - Other Specific" to cover intersex, transgender and 'third gender' categories. At the time of this publication the working group is seeking input on this topic but, taking the Scottish example on board, offers the following proposed list of controlled values:

- not known (0)
- male (1)
- female (2)
- other specific (8)
- not applicable (9)

Textual values are preferred over numerical values as this is much easier for application developers to work with.

Note

Quoting from ISO 5218: No significance is to be placed upon the fact that “Male” is coded “1” and “Female” is coded “2”. This standard was developed based upon predominant practices of the countries involved and does not convey any meaning of importance, ranking or any other basis that could imply discrimination.

4.1.7 Birth Name

Attribute	Abstract Data Type	Cardinality
birthName	String	[0..1]

All data associated with an individual are subject to change. Names can change for a variety of reasons, either formally or informally, and new information may come to light that means that a correction or clarification can be made to an existing record. Birth names tend to be persistent however and for this reason they are recorded by some public sector information systems. There is no granularity for birth name - the full name should be recorded in a single field.

4.1.8 Date of Birth, Date of Death

Attribute	Abstract Data Type	Cardinality
dateOfBirth dateOfDeath	ISA Date	[0..1]

A date that specifies the birth date (death date) of a person. See section 4.8 for details of The ISA Date Data Type.

4.1.9 Place of Birth, Place of Death

Relationship	Target Class	Cardinality
placeOfBirth placeOfDeath	Location	[0..1]

The Place of Birth and Place of Death are given using the Location class which is associated via the appropriate relationship. As detailed in Section 4.4, the Location Class has two properties:

- The geographic name of the place, which is given as a string such as "Amsterdam" or "Valetta";
- an identifier, such as a geonames URI http://sws.geonames.org/2759794 (which identifies Amsterdam) or http://sws.geonames.org/2562305 (which identifies Valetta).

The use of identifiers is preferred as these are unambiguous, however, public sector data typically uses simple names to record places and this is fully supported.
Although highly unlikely in practice, it is possible to define a place of birth or death in terms of an address or, even less likely, a geometry given as coordinates.

4.1.10 Country of Birth, Country of Death

Relationship	Abstract Data Type	Cardinality
countryOfBirth countryOfDeath	Location	[0..1]

A person's Country of Birth and Death is defined using the Location class which is associated via the appropriate relationship.

The Location Class, defined in Section 4.4, has two properties: a Geographic Name and a Geographic Identifier. Plain codes like "DE" should be provided as value for Geographical Names whereas URIs should be provided as value of the Geographical Identifier. Ideally, provide both.

Providing a simple country name is problematic and should be avoided whereas using a standardised system that allows the use of a code list for country names has a lot of potential for increasing semantic interoperability.

Known diversity that one has to deal with when exchanging country names between different communication partners without relying on an agreed code list are:

a. Official name vs. common name (e.g. “Federal Republic of Germany” vs. Germany),
b. Different languages (Italy vs. Italia),
c. Historic name vs. Current name (Burma vs. Myanmar),
d. Ambiguity of similar sounding countries (“Republic of the Congo” vs. “Democratic Republic of the Congo”)

The Publications Office of the European Union recommends the use of ISO 3166-1 codes for countries in all cases except two:

- use 'UK' in preference to the ISO 3166 code GB for the United Kingdom;
- use 'EL' in preference to the ISO 3166 code GR for Greece.

4.1.11 Citizenship

Relationship	Abstract Data Type	Cardinality
citizenship	Location	[0..*]

The citizenship relationship links a Person to a Location which identifies a country that has conferred citizenship rights on the individual such as the right to vote, to receive certain protection from the community or the issuance of a passport. Multiple citizenships are recorded as multiple instances of the citizenship relationship.

Reasoning

A person has one, multiple or even no citizenship. Citizenship is information needed by many cross-border use cases and is a legal status as opposed to the more culturally-focussed and less well-defined term "nationality." As with Country of Birth and Death, using plain text names of countries is strongly discouraged.

4.1.12 Identifier

Relationship	Abstract Data Type	Cardinality
identifier	Formal Identifier	[0..*]

For many systems, the identifier is the key piece of information about an individual and therefore an important part of the core person vocabulary. However, all identifiers are context-specific and when exchanging data between systems it is important to provide additional information that makes this explicit. An individual may have any number of identifiers. For this reason the Formal Identifier class is used (defined in Section 4.3).

The Formal Identifier class allows an identifier to be given as well as information about which authority issued it and, if needed, when.

4.2 LEGAL ENTITY CLASS

This is the key class for the Business Core Vocabulary and represents a business that is legally registered. In many countries there is a single registry although in others, such as Spain and Germany, multiple registries exist. A Legal Entity is able to trade, is legally liable for its actions, accounts, tax affairs etc.

This makes legal entities distinct from the concept of organisations, groups or sole traders. Many organisations exist that are not legal entities yet to the outside world they have staff, hierarchies, locations etc. Other organisations exist that are an umbrella for several legal entities (universities are often good examples of this). This vocabulary is concerned solely with registered legal entities and does not attempt to cover all possible trading bodies.

Legal Entity is a sub class of the more general 'Agent' class that does encompass organisations, natural persons, groups etc. - i.e. an Agent is any entity that is able to carry out actions.

4.2.1 Legal Name

Attribute	Abstract Data Type	Cardinality
legalName	String	[0..*]

The legal name of the business. A business might have more than one legal name, particularly in countries with more than one official language. In such cases, and where the encoding technology allows, the language of the string should be identified (see Section 4.9).

4.2.2 Alternative Name

Attribute	Abstract Data Type	Cardinality
alternativeName	String	[0..*]

Some jurisdictions recognise concepts such as a trading name or alternative forms of a legal entity's name. The Alternative Name property can be used to record such names but should not be used to record translations of the primary legal name. Where more than one legal name exists and where they have equal standing but are expressed in different languages, identify the language used in each of the multiple legal names (see previous section).

It is notable that some jurisdictions regard the use of any name other than the primary Legal Name as suspicious.

4.2.3 Company Type

Attribute	Abstract Data Type	Cardinality
companyType	String	[0..*]

This property records the type of company. Familiar types are SA, PLC, LLC, GmbH etc. At the time of publication, there is no agreed set of company types that crosses borders. The term 'SA' is used in Poland and France for example although they mean slightly different things. The UK's LLP and Greece's EPE provide further example of close, but not exact, matches.

That said, each jurisdiction will have a limited set of recognised company types and these should be used in a consistent manner. It is to be hoped that further work can be carried out in this area in the near future.

4.2.4 Status

Attribute	Abstract Data Type	Cardinality
companyStatus	String	[0..1]

Recording the status of a company presents the same issues as its type. The terms 'insolvent', 'bankrupt' and 'in receivership,' for example, are likely to mean slightly different things with different legal implications in different jurisdictions. Again, at the time of publication, there is no agreed vocabulary for recording the status of a company although XBRL Europe [xEBR] is in the process of defining such a vocabulary.

Taking the existing XBRL work as a starting point, however, the term 'Normal Activity' does appear to have cross-border usefulness and this should be used in preference to terms like 'trading' or 'operating.'

Best Practice for recording various other status levels is to use the relevant jurisdiction's terms and to do so in a consistent manner.

4.2.5 Activity

Attribute	Abstract Data Type	Cardinality
companyActivity	String	[0..1]

The activity of a company should be recorded using a controlled vocabulary. Several such vocabularies exist, many of which map to the UN's ISIC codes [ISIC4]. Where a particular controlled vocabulary is in use within a given context, such as SIC codes in the UK, it is acceptable to use these, however, the preferred choice for European interoperability is NACE [NACE].

4.2.6 Legal Identifier

Relationship	Target Class	Cardinality
legalIdentifier	Formal Identifier	[1..1]

The legal status of a business is conferred on it by an authority within a given jurisdiction. The Legal Identifier is therefore a fundamental relationship between a legal entity and the authority with which it is registered. The details of the registration are provided as properties of the Formal Identifier class. There is no restriction on the type of legal identifier. In many countries, the business register's identifier is the relevant data point. The tax number often fulfils this function in Spain.

The cardinality of this relationship is 1..1, i.e. a legal entity **must** have a legal identifier.

In RDF terms, legalIdentifier is a sub property of identifier (see next section) with specific semantics.

Rationale

Core Vocabularies do not normally demand that any specific field (property/relationship) is defined. However, a legal entity can only have that status if it is given by a suitable authority. Therefore the working group resolved that a Legal Entity Class could not sensibly be defined without the legal identifier.

4.2.7 Identifier

Relationship	Target Class	Cardinality
identifier	Formal Identifier	[0..1]

The identity relation links a resource to any formally issued identifier for that resource other than the one that confers legal status upon it.

Legal Entities, people, organisations and other Agents may have any number of identifiers (but only one legal identifier). For example, in many jurisdictions, a business will have one or more tax numbers associated with them which do not, by themselves, confer legal entity status. An individual may be issued with identifiers for everything from social security to club membership. The identifier relationship must not be used to link to the identifier issued by the authority that conferred legal entity status on a business.

In RDF terms, the identifier property is not constrained by its domain so that any resource can be linked to a Formal Identifier.

4.2.8 Registered Address

Relationship	Target Class	Cardinality
registeredAddress	Address	[0..1]

In almost all jurisdictions, legal entities must register a postal address. This may or may not be the actual address at which the legal entity does its business, it is commonly the address of their lawyer or accountant, but it is the address to which formal communications can be sent.

In RDF terms, registeredAddress is a sub property of the more general address property that links any resource to an associated address without further semantics. See Section 4.4.3.

4.2.9 Legal Entity

Relationship	Target Class	Cardinality
legalEntity	Legal Entity	[0..*]

The legal entity relationship can be used to link any resource to a Legal Entity Class. This is useful, for example, where an organisation includes one or more legal entities. The Dublin Core term isPartOf is a suitable inverse of this relationship [DC].

4.3 FORMAL IDENTIFIER CLASS

The Formal Identifier class represents any identifier issued by any authority, whether a government agency or not. It captures the identifier itself and details of the issuing authority, the date on which the identifier was issued and so on.

4.3.1 Identifier

Attribute	Abstract Data Type	Cardinality
identifier	String	[1..1]

The role of the Formal Identifier Class is to provide and contextualise the identifier. A formal Identifier Class without the actual identifier cannot be considered as being defined and therefore this property is mandatory.

4.3.2 Issuing Authority URI

Attribute	Abstract Data Type	Cardinality
issuingAuthorityUri	URI	[0..1]

If the Issuing Authority has a recognised URI it should be used. There are three principal advantages to this over providing the name as text (see next section).:

- it is significantly less error-prone;
- if stored in RDF, it is only stored once, no matter how many records point to it;
- further information can be found about the authority by following linked data items.

4.3.3 Issuing Authority

Attribute	Abstract Data Type	Cardinality
issuingAuthority	String	[0..1]

An issuing authority can be identified by its name as well as its URI. As this field takes a string there is considerable room for error so publishers are urged to use a consistent form of the name.

4.3.4 Date of Issue

Attribute	Abstract Data Type	Cardinality
dateOfIssue	ISA Date	[0..1]

The date on which the identifier was issued. The ISA Date data type is defined in Section 4.8.

4.3.5 Identifies

Relationship	Target Class	Cardinality
identifies	Resource	[0..1]

The identifies relationship links the class to the resource it identifies. In the context of the Business Core Vocabulary this will be a Legal Entity but the class can be used to identify any resource, including any sub class of Agent.

4.4 LOCATION

ISO 19112 defines a location as "an identifiable geographic place." With this in mind, "Eiffel Tower", "Madrid" and "California" are all locations and this is a common way of representing locations in public sector data, i.e. simply by using a recognised name. Such identifiers are common although they can be highly ambiguous as many places share the same or similar names.

In addition to a simple (string) label or name for a Location, this vocabulary defines three further ways in which it can be identified:

- by URI (such as a GeoNames or DBpedia URI);
- by address;
- by geometry, that is, a point, line or polygon expressed using coordinates in some coordinate reference system.

The latter two are defined using the Address and Geometry Classes.

4.4.1 Geographic Name

Attribute	Abstract Data Type	Cardinality
geographicName	String	[0..*]

Again quoting from ISO 19112, a geographic name is a "spatial reference in the form of a label or code that identifies a location. “Spain” is an example of a country name; “SW1P 3AD” is an example of a postcode. Both are geographic names.

The country codes defined in ISO 3166 are further examples of geographic names. As noted in Section 4.1.10, the Publications Office of the European Union recommends the use of ISO 3166-1 codes for countries in all cases except two:

- use 'UK' in preference to the ISO 3166 code GB for the United Kingdom;
- use 'EL' in preference to the ISO 3166 code GR for Greece.

The 'place' may be an abstract location such as a jurisdiction as opposed to a physical location. The geographic name might be a recognised name, such as 'Paris,' which may or may not be ambiguous in context, or it might be descriptive (such as 'Rhine km 203'). Although geographic names are always useful, it is recommended that an additional method be used to identify a location.

Places may have multiple geographic names in a single language or multiple languages. Where a name is language-specific, that language should be identified along with the name (see Section 4.9). There is no cardinality restriction on the number of geographic names that can be used to describe a single location.

4.4.2 Geographic Identifier

Attribute	Abstract Data Type	Cardinality
geographicIdentifier	URI	[0..*]

A URI that identifies the location.

GeoNames.org provides stable, widely recognised identifiers for more than 10 million geographical names that can be used as links to further information. For example, http://sws.geonames.org/593116/ identifies the Lithuanian capital Vilnius. Unfortunately these URIs cannot easily be automatically deduced since the URI scheme uses simple numeric codes. Finding a GeoNames identifier for a Location is almost always a manual process. Where such identifiers are known or can be found, however, it is recommended that they be used.

Where the Location Class is used to identify a country, if the geonames URI is not known, the recommendation is to use DBpedia URIs of the form http://dbpedia.org/resource/ISO_3166-2:XX where XX is the ISO 3166 two character code for the country [DBpedia].

The use of a URIs has added advantages:

a. it can be used by automated systems to look up additional data (linked data);
b. a triple store will store only one copy of the URI, whereas if a string is used, a copy of that string is stored for each and every person in the database. Thus, in large data sets, the saving on memory capacity and the improvement in transmission efficiency can be substantial.

Note

DBpedia uses ISO 3166-1 to describe populated places and ISO 3166-2 to refer to countries. The actual country codes are identical so that the DBpedia URI that identifies the populated place 'Sweden' is http://dbpedia.org/resource/ISO_3166-1:SE and the country itself is identified by http://dbpedia.org/resource/ISO_3166-2:SE.

DBpedia does not currently provide identifiers for sub-divisions of countries that are part of ISO 3166-2, so, for example, the English county of Suffolk, for which the ISO 3166-2 code is GB-SFK, does not have an identifier.

Finally, the EU Publication Office's use of EL and UK for Greece and the United Kingdom respectively is not supported by DBpedia (use GR and GB).

4.4.3 Address

Relationship	Target Class	Cardinality
address	Address	[0..*]

The 'address' relationship associates any resource with the Address Class. There is no cardinality restriction on the relationship when used to link a Location with an Address. In other words, a single location may be associated with multiple Address classes.

4.4.4 Geometry

Relationship	Target Class	Cardinality
geometry	Geometry	[0..*]

The 'geometry' relationship associates any resource with the Geometry Class. There is no cardinality restriction on the relationship when used to link a Location with a Geometry. In other words, a single location may be associated with multiple Geometry classes.

See also 4.6.3 XML Geometry

4.5 ADDRESS

The representation of addresses varies widely from one country's postal system to another. Even within countries, there are almost always examples of addresses that do not conform to the stated national standard. At the time of publication, work is progressing on ISO 19160-1 that defines a method through which different addresses can be converted from one conceptual model to another.

For this specification we have drawn heavily on the INSPIRE Address Representation data type [INADR]. The relevant properties, all of which can be found in the INSPIRE Address Representation Class, are detailed in the following sections but in brief they are:

- PO Box (a specialisation of locator designator)
- Thoroughfare (a road, a waterway etc.)
- Locator designator (a building number, entrance number etc.)
- Locator name (a proper name for a building or room within a building)
- Address area (usually a city area or village)
- Locality (usually a town)
- Admin unit (usually a county or state)
- Post Code

To this very granular set of properties we add two further properties:

- full address (the complete address as a formatted string)
- addressID (a unique identifier for the address)

The first of these allows publishers to simply provide the complete address as one string, with or without formatting. This is analogous to vCard's label property.

The addressID is a hook that can be used to link the address to an alternative representation, such as vCard (see Section 6.1).

4.5.1 Full Address

Attribute	Abstract Data Type	Cardinality
fullAddress	String	[0..1]

The complete address written as a string, with or without formatting. Use of this property is recommended as it will not suffer any misunderstandings that might arise through the breaking up of an address into its component parts.

This property is analogous to vCard's label property but with two important differences:

1. formatting is not assumed so that, unlike vCard label, it may not be suitable to print this on an address label.
2. vCard's label property has a domain of vCard Address. The fullAddress property has no such restriction.

4.5.2 PO Box

Attribute	Abstract Data Type	Cardinality
poBox	String	[0..1]

The Post Office Box number. INSPIRE's name for this is "postalDeliveryIdentifier" for which it uses the locator designator property with a type attribute of that name. This vocabulary separates out the Post Office Box for greater independence of technology.

4.5.3 Thoroughfare

Attribute	Abstract Data Type	Cardinality
thoroughfare	String	[0..1]

Quoting from the INSPIRE guidelines, a thoroughfare is: "an address component that represents the name of a passage or way through from one location to another. A thoroughfare is not necessarily a road, it might be a waterway or some other feature.

4.5.4 Locator Designator

Attribute	Abstract Data Type	Cardinality
locatorDesignator	String	[0..1]

The locator designator is defined by the INSPIRE guidelines as "a number or a sequence of characters that uniquely identifies the locator within the relevant scope(s). The full identification of the locator could include one or more locator designators." In simpler terms, this is the building number, apartment number, etc. For an address such as "Flat 3, 17 Bridge Street", the locator is "flat 3, 17."

4.5.5 Locator Name

Attribute	Abstract Data Type	Cardinality
locatorName	String	[0..1]

Again quoting the INSPIRE guidelines, locator name is defined as: "Proper noun(s) applied to the real world entity identified by the locator. The locator name could be the name of the property or complex, of the building or part of the building, or it could be the name of a room inside a building.

The key difference between a locator and a locator name is that the latter is a proper name and is unlikely to include digits. For example, "Shumann, Berlaymont" is a meeting room within the European Commission headquarters for which locator name is more appropriate than locator.

4.5.6 Address Area

Attribute	Abstract Data Type	Cardinality
addressArea	String	[0..1]

Taking the definition from INSPIRE, the address area is: the name or names of a geographic area or locality that groups a number of addressable objects for addressing purposes, without being an administrative unit." This would typically be part of a city, a neighbourhood or village.

4.5.7 Post Name

Attribute	Abstract Data Type	Cardinality
postName	String	[0..1]

The key postal division of the address, usually the city. (INSPIRE's definition is "One or more names created and maintained for postal purposes to identify a subdivision of addresses and postal delivery points.)

4.5.8 Admin Unit Level 2

Attribute	Abstract Data Type	Cardinality
adminUnitL2	String	[0..1]

The region of the address, usually a county, state or other such area that typically encompasses several localities.

4.5.9 Admin Unit Level 1 (Country)

Attribute	Abstract Data Type	Cardinality
adminUnitL1	String	[0..1]

The uppermost administrative unit for the address, almost always a country.

Best practice is to use the ISO 3166-1 code but if this is inappropriate for the context, country names should be provided in a consistent manner to reduce ambiguity. For example, either write 'United Kingdom' or 'UK' consistently throughout the data set and avoid mixing the two.

4.5.10 Post Code

Attribute	Abstract Data Type	Cardinality
postCode	String	[0..1]

The post code (a.k.a postal code, zip code etc.). Post codes are common elements in many countries' postal address systems and so this is often a particularly important piece of data.

4.5.11 address ID

Attribute	Abstract Data Type	Cardinality
addressId	URI	[0..1]

The concept of adding a globally unique identifier for each instance of an address is a crucial part of the INSPIRE data spec. A number of EU countries have already implemented an ID (a UUID) in their Address Register/gazetteer, among them Denmark.

It is the address Identifier that allows an address to be represented in a format other than INSPIRE whilst remaining conformant to the core vocabulary. See Section 6.1.

4.6 GEOMETRY

The Geometry Class provides the means to identify a Location as a point, line or polygon expressed using coordinates in some coordinate reference system. There are several well implemented methods of doing so and the Location Core Vocabulary is designed to support as many of them as possible. It is likely that only one geometry representation would be used in a practical situation but it is permissible to use any combination.

4.6.1 wkt

Attribute	Abstract Data Type	Cardinality
wkt	String	[0..1]

A WKTLiteral representation of a geometry as defined in the evolving GeoSPARQL specification [GEOSPRARQL].

A WKTLiteral is a string consisting of an optional URI identifying the coordinate reference system followed by one or more spaces (Unicode U+0020 character) as a separator and a Well Known Text (WKT) string as described in ISO 19125-1 [ISO 1925].

For WKTLiterals, the initial URI identifies the spatial reference system for the geometry. The Open Geospatial Consortium maintains a set of CRS URIs under the `http://www.opengis.net/def/crs/` namespace. This leading spatial reference system URI is optional and in its absence, the following spatial reference system URI is assumed: `http://www.opengis.net/def/crs/OGC/1.3/CRS84`.

RDF Example:
<http://www.opengis.net/def/crs/EPSG/0/4326/>
Point(33.95 -83.38)"^^<http://www.opengis.net/def/dataType/OGC-SF/1.0/WKTLiteral>

(this gives both the coordinate reference system and the point itself as a string which is then data typed).

4.6.2 Lat/Long

Attribute	Abstract Data Type	Cardinality
latLong	String	[0..1]

A comma separated lat/long pair as defined in WGS84 (see W3C RDF vocabulary for this [W3CWGS]).

4.6.3 XML Geometry

Attribute	Abstract Data Type	Cardinality
xmlGeometry	XML Literal	[0..1]

Several XML-based methods exist for defining a geometry, that typically refers to a location on the Earth's surface. These XML dialects offer self-contained methods of describing a location and current key examples include GML [GML] and the W3C's Point of Interest specification [POI] for which specific sub relations are provided.

The xml geometry relationship can link either a Location Class or a Geometry Class to the XML Literal.

4.6.4 Location as GML

Attribute	Abstract Data Type	Cardinality
asGML (sub property of xmlGeometry)	GML Literal	[0..1]

A self-contained XML document describing a location using GML [GML]. The asGML property is not shown on the conceptual model as it is simply a sub property of xmlGeometry.

4.6.5 Location as GML_CE_Geometry

Attribute	Abstract Data Type	Cardinality
asGML_CE (sub property of xmlGeometry)	GML_CE_Geometry Literal	[0..1]

A self-contained XML document describing a location using the GML_CE_Geometry as defined by the W3C Point of Interest Core [POI]. The asGML_CE property is not shown on the conceptual model as it is simply a sub property of xmlGeometry.

4.7 UNDEFINED CLASSES

The model shows two classes that are not yet defined. The working group strongly recommends that a new working group is formed to define them

4.7.1 Change

Any record can only be seen as a snapshot of the data available to the publisher at a particular point in time. This is true for data related to all areas, not just people, places and businesses. The Change class is envisaged as one that captures any change in a data point. This might be a change in address, a change of name or any other update. It is likely that such a record would make it clear what data element was changed as well as when, why and by whom. This area - versioning and provenance - is a complex one that the working group hopes will be tackled by a subsequent activity as a matter of high priority.

The envisioned Change class entails a relationship:

Relationship	Abstract Data Type	Cardinality
change	Change	[0..*]

The change relationship links a Class to a record of change in a data field associated with the Class. For example, a change of name or a person's their death. It is assumed that an inverse relationship will also be defined.

4.7.2 Licence

In the context of a business register, the term licence refers to a permit to carry out a particular activity such as banking or mining. The Licence class is therefore envisaged as being similar to the Formal Identifier Class in which a specific body is registered with a recognised authority. As with the Change class, the working group hopes that near future work can be done to ensure that such a class is well defined.

Two inverse relationships are entailed by the Licence Class:

Relationship	Abstract Data Type	Cardinality
holdsLicence	Licence	[0..*]

Relationship	Abstract Data Type	Cardinality
heldBy	Legal Entity	[0..1]

4.8 THE ISA DATE DATA TYPE

Dates recorded within public sector data sets exist in many different formats. In order to make those dates interoperable, it is important that they should be represented in a regular manner wherever possible. The most widely used data types are those defined in the XML Datatypes [XSD] and dates should be recorded as formatted strings that are conformant with that standard with the relevant data type declaration. XML datatypes are fully compatible with ISO 8601 [ISO 8601].

Known dates are recorded as literals in the form `yyyy-mm-dd` and then typed as `xsd:date`.

Where the full date and time are both known, and this can be important in records of death for example, use the `xsd:dateTime` format of `yyyy-mm-ddThh:mm:ss` with or without timezone information.

Where just a year is known use `xsd:gYear` (yyyy) and where just the month is known use xsd:`gYearMonth` (`yyyy-mm`). The 'g' is for Gregorian. Non-Gregorian dates are not covered by XSD.

For emphasis, the datatype should be explicitly stated so that software can correctly interpret the data. If a date is provided in a string that is conformant with a standard other than XSD, and that is not covered by XSD, then it should be typed accordingly. It should also be noted that the UK Government provides a URI scheme for time intervals that may be useful [DGUT].

However, in some public sector data sets, full dates are not always known. For example, dates such as "some time in the 1920s" or "between 1925 and 1932" are not uncommon. There is no agreed formal way to record such uncertain dates and therefore an un-typed plain string should be used. In other words, a string like "between 1925 and 1932" cannot be improved upon.

Some data sets use other methods to express uncertainty such as using dates. Sometimes '00' or '??' is used to indicate uncertainty so that "August 1986" might be recorded as either 1986-00-00 or 1986-00-??. This practice is unnecessary and **strongly discouraged**, particularly if such dates are typed as XSD dates. If a string such as 1986-08-00 is fed to software expecting an xsd:date then it will either reject the data, in which case data is lost, or it might convert it to the nearest actual date 1986-08-01. In the latter case, the accurate yet imprecise date of "August 1986" has now been converted to a more precise date of 1st August 1986 that may be anything up to 30 days wide of the truth.

In summary:

- use the appropriate XSD data type wherever possible (dateTime, date, gYearMonth, gYear);
- where data is missing, don't invent it or try to fool the system - just give the date as an un-typed string.

4.9 LANGUAGES

Where data such as names exist in multiple languages, each version of the data should be included and each one associated with the relevant language identifier. RFC 3066 [RFC 3066] provides a commonly used set of identifiers for natural languages. This is the set recognised by UN/CEFACT and XML Schema.

Languages are represented by two character codes, optionally followed by a locale definition such as "de" meaning German and "de-at" meaning "German as spoken in Austria."

XML Example:

```
<Location>
  <geographicName xml:lang="en">London</geoGraphicName>
  <geographicName xml:lang="fr">Londres</geoGraphicName>
</Location>
```

RDF Example:

```
[] a locn:Location ;
  locn:geographicName "London"@en ;
  locn:geographicName "Londres"@fr .
```

5. PROPERTIES CONSIDERED AND EXCLUDED

5.1 RELATIONSHIPS

The Core Person conceptual model does not include the relationships that may exist between the described person and other entities such as family members, employers or property. Those relationships are clearly important in defining the role an individual plays in society and are often important in the definition of an individual's identity, however, they are not core to the identification of an individual in public sector data exchange.

5.2 PREFIXES, HONORARY TITLES AND SUFFIXES

Many family names include prefixes such as 'de' and 'von' and there is no simple way to decide whether or not such prefixes should be taken into account, for example, when sorting alphabetically. Herman Van Rompuy should be sorted under V whereas Ludwig van Beethoven should be sorted under B. Given the almost infinite variety of names and naming structures, there is no simple way to define the parts of names that will work across all possibilities. Such rules may be definable in a specific application or in a given country but that is out of scope for the Core Person vocabulary.

Therefore the working group resolved not to try and provide more granular properties for parts of names. It is perhaps noteworthy that the group very nearly resolved only to include the single property of Full Name with no further granularity. An excellent overview of different name conventions around the world is available from the W3C [PNAW].

5.3 ROLE

A specific relationship proposed as part of the Person Core Vocabulary was that of role. This could be used to link a record of any individual with the role they played, such as patient, company director or any other. It was narrowly decided not to include this relationship in this vocabulary. However, it is noted that in any given context, such a relationship can be expressed and the Person Core Vocabulary used. For example, a hospital patient record could point to a Person record in which case the 'patient' role is defined by the publisher of the patient record.

5.4 COUNTRY OF ORIGIN

The country or jurisdiction of origin is not included in the Business Core Vocabulary as this information will be obvious from the Formal identifier Class associated through the Legal Identifier.

5.5 HOLDS LICENCE & HELD BY

As noted in section 4.7.2, the business vocabulary sees a distinct need for a small vocabulary to describe licenses issued to companies to permit them to carry out certain activities. The immediate use case is banking but there are many more highly regulated industries and the working group did not feel sufficiently knowledgeable to be able to capture the essential details. For example, the holder of a mining licence may be restricted to specified geographical areas, or times of day when they can operate; a waste disposal company may be restricted in the types of waste it can handle. A vocabulary is likely to need to support such details and this was deemed out of scope for the current work.

5.6 SPATIAL RELATIONSHIPS

The Location Working Group briefly considered including spatial relationships such as 'contains', 'borders' or 'overlaps' within this vocabulary. It was decided not to do so as where such properties are needed, publishers are likely to be familiar with more specialised vocabularies. However, it is worth noting that the Dublin Core term dcterms:relation or a sub property thereof may be appropriate.

6. USAGE GUIDELINES

The Core Vocabularies have been designed with maximum flexibility in mind with almost no restrictions on how any specific term can be used. In an RDF context, no new Domain and Range restrictions have been defined.

[To Do - add some specific examples]

6.1 ALTERNATIVE ADDRESS REPRESENTATIONS (VCARD ETC.)

The INSPIRE method of representing addresses is very detailed, designed primarily for use in databases of addresses. Many datasets that include address data as one piece of information about something else are likely to have that data in simpler formats. These might be tailored to the specific need of the dataset, follow a national norm, or make use of a standard like vCard [VC].

To provide maximum flexibility in the core vocabulary, whilst remaining interoperable with INSPIRE Address Guidelines (which EU Member States are obliged to use), the Location Core Vocabulary provides two extra properties: full address (Section 4.5.1) and addressID (section 4.5.11). As vCard is a massively used representation of addresses, we provide an example of how it can be fitted into the model but the same general method can be applied to any alternative address representation.

For our example we'll use the following (entirely fictitious) address:

15 Acacia Avenue
New Town
Edinburgh
UK
EH1 1AA

to which we'll assign an identifier of http://address.example/id/uk/eh11aa.

Using the Address Class defined in this specification the various properties and values would be as show in Table 1.

Property	Value
fullAddress	15 Acacia Avenue New Town Edinburgh UK EH1 1AA
locatorDesignator	15
thoroughfare	Acacia Avenue
addressArea	New Town
postName	Edinburgh
adminUnitL1	UK
postCode	EH1 1AA
addressID	http://address.example/id/uk/eh11aa

Table 1 An Edinburgh Address Represented using the Location Core Vocabulary

The mapping in Table 2 can be used to derive a vCard Address from this data. Note the inclusion of the addressID property which is not part of the vCard specification but is necessary so that the two representations can be associated.

As is clear, in this mapping it is really only the separation of the thoroughfare name from the locator properties that marks a difference beyond the particular terms used. It should be noted, however, that other mappings are possible and may be more appropriate in different contexts. The INSPIRE and vCard models are different, were designed for different purposes, and there are likely to be edge cases where the mapping doesn't work effectively.

Mapping from granular data to the more coarse representation is relatively easy. However, the reverse is not possible in an automated way *unless* there are clear parsing rules which is unlikely. Therefore in situations where the available address data does not include, for example, separate locator designator and thoroughfare information, publishers should simply provide the fullAddress and addressID properties and then use the ID to associate the address with the vCard representation or whatever other representation is appropriate.

vCard	Location Core Vocabulary
street-address	concatenated string comprising: locatorDesignator locatorName thoroughfare
extended-address	addressArea
locality	postName
region	adminUnitL2
country-name	adminUnitL1
postal-code	postcode
label	fullAddress
addressID	http://address.example/id/uk/eh11aa

Table 2 An Edinburgh Address represented in vCard

7. RDF AND XML SCHEMAS FOR THE CORE VOCABULARIES

This section will describe the RDF and XML schemas expressing this Core Vocabularies specification. These will be added as soon as a decision has been taken on the namespaces to be used.

The RDF and XML representations will re-use existing vocabulary terms where possible. The actual schemas will be included in an annex and will be made available at an appropriate URL for public access. In parallel to the Public Comment period of the Core Vocabularies specification, expressions of the Core Vocabularies in RDF and XML will be further developed. Initial versions will be publicly available:

https://joinup.ec.europa.eu/asset/core_person/release/02
https://joinup.ec.europa.eu/asset/core_location/release/02
https://joinup.ec.europa.eu/asset/core_business/release/02

8. BACKGROUND AND OBJECTIVES

As noted in the introduction, these Core Vocabularies are one of a set produced under Action 1.1 of the ISA Programme [A1.1]. Further background is available in "e-Government Core Vocabularies [EGOV-CV] which offers an overview and context for the work.

The natural course of action for any practitioner or team given the task of recording information about a natural person is to write a list of the data elements they need (or already have) and to work within the specific context of their project. This often works in that it demonstrably meets the project's needs. The problems only arise when one team wants to exchange data with another. It's at that point that the choice of, say, 'first name' over 'given name' and 'surname' over 'family name' becomes an obstacle. Such terms are well defined in a variety of vocabularies and their use cannot be regarded as 'wrong', however, it's easy for simple choices to lead to unintended difficulties further down the road.

The aim of providing Core Vocabularies via the Joinup Platform is not to force teams to use a particular set of terms, or to require the re-engineering of data sets to use them (which can be prohibitively expensive). Rather the aim is to make it easy to see and use the terms that crop up across multiple domains; terms that, when used by public sector agencies, will make data more interoperable.

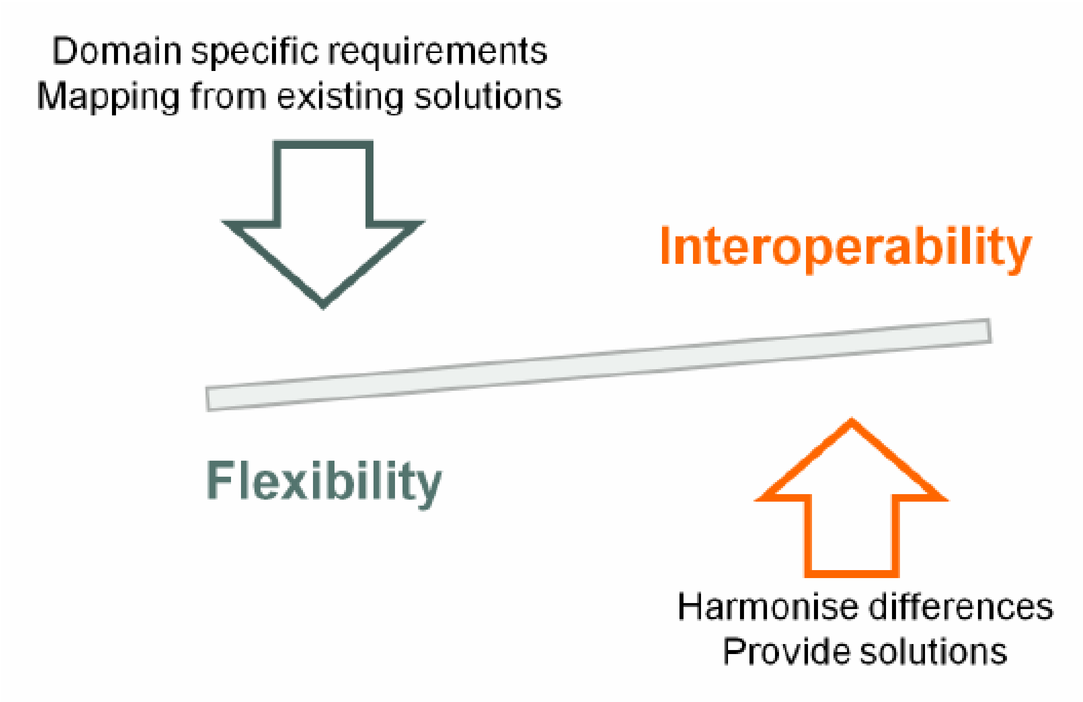

Figure 2 The struggle between enabling interoperability and giving flexibility

Identifiers are a case in point. In an international context, a person's passport number is likely to be critically important. This is not so within a university where the likelihood is that an in-house identifier will be assigned. Rather than one agency defining a term for 'passport number' and another for 'student number', both can use the core vocabulary term 'identifier', preferably with some additional contextual information. As Figure 2 illustrates, there is a balance to be struck between flexibility and interoperability.

9. APPROACH & COMMUNITY

Early work under the ISA Programme, known as SEMIC, produced an initial Core Person vocabulary specification [SCP]. This specification draws directly on that work and extends it slightly for the Person Class (all the terms in the SEMIC Core Person Vocabulary exist within this update). The Business and Location vocabularies are 'new' however, as is clear, they both draw heavily on existing work, particularly the Location vocabulary which is driven by work on the INSPIRE directive.

The process and methodology followed in the development is set out in detail in the Process and Methodology for Developing Core Vocabularies [PMDCV].

Specific acknowledgement is due to:

- Joseph Azzopardi, Malta Information Technology Agency
- Morten Borrebæk, Norwegian Mapping Authority
- Ignacio Boixo, Banco de España
- Stijn Goedertier, PwC EU Services
- Christophe Guéret, Vrije Univerisiteit Amsterdam
- Martin Herzog, Init
- Anja Hopfstock, Federal Agency for Cartography and Geodesy (BKG), Germany
- Anne Gro Hustoft, Agency for Public Management and eGovernment, Norway
- Michaela Elisa Jackson, CSI-Piemonte
- Morten Lind, Danish Enterprise and Construction Authority
- Michael Lutz, Joint Research Centre
- Piotr Madziar, DG MARKT
- Andras Micsik, MTA Sztaki
- Aine Mitchell, EBR
- Andrea Perego, Joint Research Centre
- Greg Potterton, Eurojust
- Raj Singh, Open Geospatial
- Sebastian Sklarß, Init.
- Paul Smits, Joint Research Centre
- Chris Taggart, Open Corporates
- Thomas Verdin, XBRL Europe
- Luk Vervenne, Synergetics, Belgium
- Philippe Vlerick, EC, DG Justice
- Neven Vrček, University of Zagreb
- Sebastian Sklarß, Init
- Rob Walker, Independent
- Stuart Williams, Epimorphics
- Peter Winstanley, Scottish Government

10. CHANGE CONTROL

The Core Vocabularies are published by the ISA Programme. Review comments and requests for changes can be made via the mailing list which is archived at http://joinup.ec.europa.eu/mailman/private/isa-cv-person-wg/.

11. REFERENCES

[A1.1] Action 1.1 Improving semantic interoperability in European eGovernment systems http://ec.europa.eu/isa/actions/01-trusted-information-exchange/1-1action_en.htm

[DBpedia] DBpedia is a community effort to extract structured information from Wikipedia and to make this information available on the Web. http://dbpedia.org/

[DC] DCMI Metadata Terms, Dublin Core Metadata Initiative. http://dublincore.org/documents/dcmi-terms/

[DGUT] data.gov.uk Time Intervals. Linked data for every time interval and instant into the past and future, from years down to seconds. This is an infinite set of linked data. It includes government years and properly handles the transition to the Gregorian calendar within the UK. http://thedatahub.org/dataset/data-gov-uk-time-intervals

[EGOV-CV] e-Government Core Vocabularies: The SEMIC.EU approach. Retrieved from European Commission - Directorate-General Informatics: http://www.semic.eu/semic/view/documents/egov-core-vocabularies.pdf

[FOAF] Friend of a Friend http://xmlns.com/foaf/spec/

[GML] Geography Markup Language, Open Geospatial Consortium Standard. http://www.opengeospatial.org/standards/gml

[INADR] D2.8.I.5 INSPIRE Data Specification on Addresses – Guidelines. Infrastructure for Spatial Information in Europe (INSPIRE). http://inspire.jrc.ec.europa.eu/documents/Data_Specifications/INSPIRE_DataSpecification_AD_v3.0.1.pdf

[ISA] Interoperability Solutions for European Public Administrations, http://ec.europa.eu/isa/

[ISIC4] International Standard Industrial Classification of All Economic Activities, Rev.4, United Nations Statistics Division. See http://unstats.un.org/unsd/cr/registry/regcst.asp?Cl=27

[ISO 3166-1] ISO 3166 code lists, ISO. http://www.iso.org/iso/iso_3166_code_lists

[ISO 5218] Codes for the representation of human sexes http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=36266

[ISO 8601] Data elements and interchange formats -- Information interchange -- Representation of dates and times, ISO 8601:2004. http://www.iso.org/iso/catalogue_detail?csnumber=40874

[ISO 19125] ISO 19125-1:2004 Geographic information -- Simple feature access -- Part 1: Common architecture See http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=40114)

[ISO 19112] Geographic information -- Spatial referencing by geographic identifiers (ISO 19112:2003) http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=26017

[JOINUP] The Joinup Platform is operated by the European Commission designed to enable the sharing and reuse open-source software, semantic assets and other interoperability solutions for public administrations. See http://joinup.ec.europa.eu/

[NACE] Statistical Classification of Economic Activities in the European Community. Eurostat. http://ec.europa.eu/eurostat/ramon/nomenclatures/index.cfm?TargetUrl=LST_NOM_DTL&StrNom=NACE_REV2&StrLanguageCode=EN

[PMDCV] Process and Methodology for Developing Core Vocabularies, 22 November 2011. https://joinup.ec.europa.eu/elibrary/document/isa-deliverable-process-and-methodology-developing-core-vocabularies

[PNAW] Personal names around the world, Richard Ishida, W3C. http://www.w3.org/International/questions/qa-personal-names

[POI] Points of Interest Core W3C Recommendations Track Working Draft http://www.w3.org/TR/poi-core/.

[RFC 3066] H. Alvestrand, ed. RFC 3066: Tags for the Identification of Languages 1995. Available at: http://www.ietf.org/rfc/rfc3066.txt

[SCP] SEMIC Core Person. Available via http://joinup.ec.europa.eu/asset/semic-eu_core_person_17/home

[TOGM] Towards Open Government Metadata, Vassilios Peristeras, DG DIGIT, ISA Unit, September 2011 https://joinup.ec.europa.eu/sites/default/files/towards_open_government_metadata_0.pdf

[UN/CEFACT] UN Centre for Trade Facilitation and e-Business http://www.unece.org/cefact/

The UN/CEFACT Core Components Library (CCL) can be found at http://www.unece.org/cefact/codesfortrade/unccl/CCL_index.html

[UPU] Universal Postal Union (UPU), "International Postal Address Components and Templates", UPS SB42-4, July 2004

[VC] vCard Specifications, Internet Mail Consortium. See http://www.imc.org/pdi/

[W3CDTF] Date and Time Formats. W3C Note 15 September 1997. http://www.w3.org/TR/NOTE-datetime

[W3CWGS] WGS84 Geo Positioning: an RDF vocabulary, W3C Note 2003 http://www.w3.org/2003/01/geo/wgs84_pos

[WKT] OGC GeoSPARQL - A Geographic Query Language for RDF Data Under development at the Open Geospatial Consortium (http://www.opengeospatial.org)

[xEBR] XBRL Europe http://www.xbrl.org/

[XSD] XML Schema Part 2: Datatypes Second Edition. W3C Recommendation 28 October 2004. http://www.w3.org/TR/xmlschema-2/#date

12. THE CORE VOCABULARIES WORKING GROUP

Given name	Family name	Organisation	Country	Task Force
Segun	Alayande	SmartInformatics Limited - british airports	UK	Core Person, Core Location, Core Business
Miguel	Alvarez Rodríguez	Ministerio de Política Territorial y Administración Pública	ES	Core Person
Adam	Arndt	The Danish Agency for Digitization	DK	Core Person, Core Location, Core Business
Joseph	Azzopardi	Malta Information Technology Agency	MT	Core Person, Core Location, Core Business
Peter	Bjerghave	Danish National Police	DK	Core Person
Lyubomir	Blagoev	USW Ltd	BG	Core Person, Core Location, Core Business
Ignacio	Boixo	Banco de España	ES	Core Business
Morten	Borrebæk	Norwegian Mapping Authority	NO	Core Location
Dan	Brickley	Vrije Universiteit Amsterdam	NL	Core Person, Core Location, Core Business
Heath	Bunting	Irational.org	UK	Core Person, Core Location, Core Business
Serena	Coetzee	University of Pretoria	ZA	Core Location
Paul	Farrell	Enterprise Registry Solutions / European Business Register	IE	Core Business
Corine	Frappart	Publications Office of the European Union	BE	Core Person, Core Location, Core Business
Marc	Fresko	Inforesight Ltd	UK	Core Person, Core Business
Andreas	Gehlert	Federal Ministry of Interior - Germany	DE	Core Person, Core Location, Core Business
Giorgios	Georgiannakis	DG SANCO	BE	Core Person, Core Location, Core Business
Christophe	Guéret	Vrije Universiteit Amsterdam	NL	Core Person, Core Location

Anja	Hopfstock	Federal Agency for Cartography and Geodesy (BKG), Germany	DE	Core Location
Anne Gro	Hustoft	Statistics Norway	NO	Core Person
Andreas	Illert	Federal Agency for Cartography and Geodesy (BKG), Germany	DE	Core Location
Michaela Elisa	Jackson	CSI-Piemonte - Governo e Gestione	IT	Core Person, Core Location
Brigitte	Joerg	German Research Center for Artificial Intelligence	DE	Core Person, Core Business
Anne	Kauhanen-Simanainen	Ministry of Finance	FI	Core Person, Core Location, Core Business
Thomas	Knape	Data Fusion/IEEE	IE	Core Person, Core Location, Core Business
Pim	Keizer	Ministry of Justice – The Netherlands	NL	Core Business
Madeleine	Kiss	Publications Office of the European Union	BE	Core Person, Core Location, Core Business
Morten	Lind	Danish Enterprise and Construction Authority	DK	Core Location
Nikos	Loutas	DERI (Digital Enterprise Research Institute)	IE	Core Person, Core Location, Core Business
Michael	Lutz	JOINT RESEARCH CENTRE INSPIRE	IT	Core Location
Matthias	Lüttgert	ENDA GmbH & Co. HG	DE	Core Business, Core Location
Martin	Määrits	Centre of Registers and Information Systems	EE	Core Business
Piotr	Madziar	DG MARKT	BE	Core Business
Charles-Henri	Menseau	Département Titres Identités	FR	Core Person
Andras	Micsik	MTA SZTAKI	HU	Core Person, Core Location
Aine	Mitchell	Enterprise Registry Solutions / European Business Register	IE	Core Business
Natalie	Nickel	e-Codex	DE	Core Person

Priit	Parmakson	Estonian Information Systems Authority	EE	Core Person
Andrea	Perego	Joint Research Centre (European Commission)	IT	Core Location
Vassilios	Peristeras	DG DIGIT	GR	Core Person, Core Location, Core Business
Greg	Potterton	Eurojust	BE	Core Person, Core Location, Core Business
Peter	Reichstädter	Egovlabs Austrian Federal Chancellery	AT	Core Person, Core Location, Core Business
Børge Krogh	Samuelsen	Local Government Denmark	DK	Core Person
Peter	Schmitz	Publications Office of the European Union	BE	Core Person, Core Business, Core Location
Raj	Singh	Open Geospatial Consortium	USA	Core Location
Sebastian	Sklarß	]init[AG	DE	Core Person, Core Business, Core Location
Ronald	Smit	e-Codex	NL	Core Person
Dan	Smith	LinkedGov	UK	Core Person, Core Business, Core Location
Paul	Smits	Joint Research Centre (European Commission	IT	Core Location
Ernst	Steigenga	e-Codex	NL	Core Person
Chris	Taggart	Open Corporates	UK	Core Business
Athina	Trakas	OGC (Open Geospacial Consortium)	GR	Core Location
Maris	Tudre	Centre of Registers and Information Systems	EE	Core Business
Thomas	Verdin	XBRL	BE	Core Person, Core Business
Luk	Vervenne	Synergetics	BE	Core Person
Philippe	Vlerick	DG JUST	BE	Core Person
Neven	Vrček	University of Zagreb	HR	Core Person, Core Business
Rob	Walker	Independent Consultant	UK	Core Location
Andreas	Weller	European Banking Authority (EBA)	LU	Core Business
Stuart	Williams	Epimorphics	UK	Core Location
Peter	Winstanley	Scottish Government	UK	Core Person, Core Location, Core Business

Glyn	Wintle	LinkedGov	UK	Core Person, Core Location, Core Business
Piotr	Piotrowski	Universal Postal Union	CH	Core Location
Uwe	Ewald	International Justice Analysis Forum	DE	Core Person
Peter	Burian	DG DIGIT	EU	Core Business, Core Location, Core Person
Muriel	Foulonneau	Tudor Research Centre	LU	Core Person

The SEMIC Team:

- Phil Archer, W3C (Core Vocabularies Editor)
- Makx Dekkers, AMI Consult
- Michiel De Keyzer, PwC
- Débora Di Giacomo, PwC
- João Rodrigues Frade, PwC
- Stijn Goedertier, PwC
- Thomas Rössler, W3C
- Niels Van Hee, PwC
- Rigo Wenning, W3C