

Identifiers in Electronic Public Procurement

SVERRE BAUCK, ROBIN WILTON



Sverre Bauck is Senior Advisor in the Agency for Public Management and eGovernment

The aim of this paper is to give a summary of the use of electronic identifiers in public e-procurement applications, and to suggest ways in which that use can be standardised to the benefit of all parties concerned. All procurement systems are using identifiers in business documents that are structured according to international standards, but they contain identifiers that have formats and qualities defined within a group, like a national market. The paper suggests an approach to achieve open cross border exchange of electronic procurement documents and its contained identifiers.

One key to borderless and open electronic procurement is to design standardised systems for the handling of identifiers and related trustworthy assertions. We examine the relative roles of identifiers and assertions, to identify some of the design criteria which such systems would need to satisfy.

Characteristics and Use of Electronic Identifiers

Interoperability across Organizational Boundaries

E-procurement in general is characterised by a need for electronic instructions (orders, invoices, agreements, etc.) to be exchanged both within and between organisations. For instance, within an organisation, a purchase request might originate in one department and be sent to a centralised function to generate a purchase order. That purchase order would then be sent out of the originating organisation to the supplier. A corresponding exchange might follow in the opposite direction when the goods are shipped (packing note, delivery slip, invoice, etc.).

At each stage in this sequence, the instruction passes through the hands of different individuals, or through different computer systems and applications, each of which – for that part of the process – considers itself to be the ‘owner’ of the instruction. Nevertheless, particularly in large and/or automated systems, the continuity of the instruction has to be maintained from start to finish – so typically the originator assigns an identifier to the instruction by the originator (for instance, an Order Number). When the instruction crosses organisational boundaries, the first identifier-related issue crops up; the receiving organisation may have different conventions for assigning identifiers (for instance, it may need an identifier which includes the date of receipt of the instruction). Thus, in exchanges of correspondence between organisations, it is not unusual to see a single letter with both ‘Your Reference:’ and ‘Our Reference:’ identifiers in it.

The second issue is that for each ‘owner’ during the process, the end-to-end identifier may not be the most important thing. For instance, suppose the instruction

is a purchase order for office supplies; the fulfilling department will want to generate a ‘picking list’ with item numbers and quantities for each of the items in the order. Those item numbers may or may not be known to the originator of the order, and may or may not have been included in the instruction. Nevertheless, the fulfilling department needs them if it is to be able to complete the order, manage its stock of items, and so on. Thus, the processing of the instruction through and between organisations may involve the use of multiple identifiers, which may apply to only a part of the complete transaction chain.

So much for the instruction itself and the items it concerns. There is also the matter of identifiers for the people/entities who deal with the instruction. For instance, it may be important to know that a purchase order originated from a person with the authority to issue it – and similarly, the resulting invoice and payment need to be clearly and correctly associated with the appropriate person/entity.

We introduce the distinction between people and entities at this stage because of the need to cater for legal as well as natural persons. For instance, if a limited

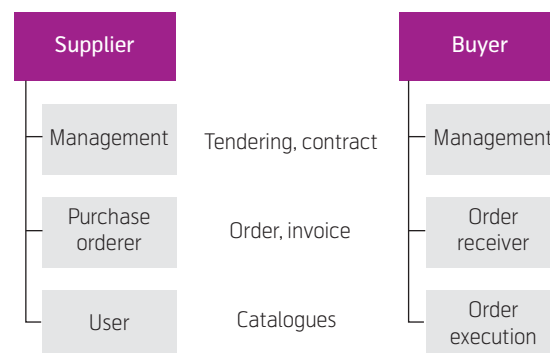


Figure 1 E-procurement processes across organizational boundaries

company orders goods, it is the company which is liable for payment, rather than a natural person employed by that company. In many cases, the use of an identifier is key to establishing the legal responsibility of the party concerned (for instance, using identifiers such as a company's registered address, registration number, VAT number and so on).

In general, identifiers for legal and natural persons are national: that is, they establish the uniqueness of that entity within a given (national) population. As a consequence, verifying the quality and possible attributes of such an identifier across borders calls for measures beyond those which suffice within the original national context, as well as additional dedicated efforts and resources.

Identifiers and Associated Information

It is important to bear in mind that often, the important characteristic of a given identifier is not necessarily reflected in the identifier itself – but may be to do with some associated pieces of information to which the identifier is merely an index. For example, someone's name at the bottom of a letter does not 'state' that they are a company director – but given the name, one could establish whether or not that person was a director by other means.

Similarly, the instructions (as described above) which pass between organisations are unlikely to include the full terms and conditions which apply to the transactions they represent. Rather, the instruction will serve as an implied acceptance of contractual terms which are specified elsewhere (for instance, there may be a default procurement agreement in place between the organisations in question, set up at the time of the first transaction).

There can also be a high degree of interpretation involved in mapping natural persons onto legal responsibilities. For instance, although there may be a legal requirement for the directors of a company to be named in the company's registration filing, it is unlikely that, say, the company's purchasing officer(s) will be similarly registered. Thus, there is an interpretive element in assuming that an order arriving on the company's headed paper (or via email) is valid, even though it is signed by an individual whose name does not appear in a public record of authorised purchase officers.

The degree to which such interpretation is necessary (or legal) may well vary from one jurisdiction to

another. Similarly, the extent to which a contract may be considered to have been formed based on implied rather than explicit information may also vary from one jurisdiction to another. For instance, the legal status of an instruction might vary depending on whether the sender has:

- digitally signed it;
- typed their name at the bottom;
- typed their initials at the bottom;
- set a default 'signature' text using their email client.

This topic has been examined in some depth in UK rulings by Judge Pelling QC – commentaries from Pinsent Mason Ltd law firm¹⁾, and Robin Wilton²⁾.

Persistence and Visibility of e-Procurement Information

In pre-digital bureaucratic processes, the path of a transaction (simply put) was usually reflected by the passage of a single piece of paper. Photocopying made it easier to take and keep copies of the original instruction, and electronic equivalents can now generally be copied, stored and distributed – to all intents – almost indefinitely.

In some senses, of course, this is a benefit: each actor in a transaction can, if required, generate an audit trail of their part in the process – which, in turn, makes it easier to deal with errors, complaints, repudiation and so on. On the other hand, the ability to generate, store and process detailed data in such quantities is neither risk- nor privacy-neutral.

For instance, if a third party is able to access the audit logs of such transactions, they may be able to retrieve inherently sensitive information (such as payment card details), or data which can be used to infer sensitive information (such as purchasing history, behavior patterns and so on). In terms of vulnerability analysis, it is now far easier to 'lose' (or inadvertently disclose) such records in vast quantities than was ever the case where physical paper records were concerned.

As one security specialist put it, at the time of the UK's HMRC data breach: "Losing 25 million records on CD was easy. Losing 25 million paper files requires a great deal of effort and ingenuity."³⁾

Even in a paper-based bureaucracy, it was not enough just to have effective processes for capturing paper records and being able to retrieve them on demand

¹⁾ Court rules that an email address is not a signature. <http://www.out-law.com/page-6839>

²⁾ Signed emails and the law. http://blogs.sun.com/racingsnake/entry/signed_emails_and_the_law

³⁾ A Collings, *E-Commerce Associates Ltd*.

(functions which, in a digital age, we take entirely for granted). Explicit means had to be put in place to dispose of unwanted or unnecessary information: any retention process had to be complemented by a 'weeding' process, through which data was discarded which otherwise served no purpose and merely made the efficient use of the desired data more difficult.⁴⁾

In electronic systems, the issue of appropriate 'weeding' is complicated by a number of factors:

1. The sheer volume of data it is now possible to accumulate without significant cost or effort (for instance, the number of emails in the average user's inbox is now typically on a scale which would be physically and administratively impossible with paper mail);
2. The cost of selectively deleting data, relative to indiscriminately retaining it;
3. The lack of appropriate meta-data which would allow 'weeding' to take place automatically rather than manually.

Concerning this last point, consider the gap between most digital file systems and the kind of meta-data which data controllers ought, by default, to be able to associate with any given piece of personal information they retain. Fair information principles state that the processing of personal data should be governed by a number of factors, such as the purpose for which the data was collected, the period for which the data ought to be retained, the people who can legitimately access the data, and so on. And yet the majority of information systems either ignore this kind of meta-data, or provide no means for it to be associated with information in the first place.

Summary of Current Situation

Thus, to summarize this introductory section: any given online transaction is likely to involve the use of multiple identifiers (even if it is something as simple as the sender's and recipient's email address), and these may include identifiers whose scope extends across the whole transaction, as well as identifiers whose scope is limited to a particular phase or a particular actor in the over-all sequence.

Identifiers may explicitly convey the information required (for instance, an address) or they may simply be a link to other information which defines specific terms and conditions relevant to this transaction (such as a company number). Furthermore, identifiers may not explicitly convey relevant information, but require interpretation (for example, "does the name on this email correspond to an authorized purchaser, and does the email therefore amount to a legitimate purchase order?").

All of these factors may have implications which differ from organisation to organisation, or jurisdiction to jurisdiction, in part because the means of expressing a contract continue to evolve through technological advances.

Those technological advances bring advantages for the management of transactions, but introduce risks and vulnerabilities which need to be mitigated. This, in turn, gives rise to requirements for new management disciplines, based on new kinds of data about identifiers and the data associated with them.

Next Steps

From Closed to Open Systems

Electronic business documents are normally structured by use of international standards like CEFACT, EDIFACT, UBL, and business partners need to share implementation of the structures and identifiers. The latter and their qualities are shared between partners of user groups. Some are national, like those used for both natural and legal persons; further, most of the most used ones were derived and established long before open on-line verification in issuers' registries was feasible or foreseen.

The principle of verifiable identifiers for use in open systems is now widely accepted and accepted as expressed in the standards ISO 6523 (ISO 6523 defines a *Structure for the Identification of Organisations (SIO)*). This is a syntax for uniquely identifying organizations in computer data interchange⁵⁾.

The principle is in use for e-mail systems on Internet, for bank accounts, for Geographic Location Numbers (GLN) and an increasing number of other systems. LDAP, DNS and OCSP are frequently used acronyms for systems used to retrieve up-to-date information

⁴⁾ For example, Filing Practice in the Civil Service – *Jefferies, KS (1946)* refers to the process of 'weeding' files to select papers which may be thrown away. *Jefferies* also observes that the process of any given government department reduces, essentially, to a description of its registry processes for filing, indexing and managing papers.

⁵⁾ http://en.wikipedia.org/wiki/ISO_6523) and ISO 15459 *Unique Identification of Objects*.

from the issuer of an identifier or attribute⁶⁾. Identifiers and needed attributes for legal and natural person are maintained to meet the needs of the issuing registry; OCSP systems are run by financial institutions to make public key certificate status information available to remote applications, DNS helps internet requests to be routed correctly, translating between IP (Internet Protocol) addresses and URLs (Uniform Resource Locators), and LDAP to find their subscribing addressees or other hierarchically-stored attributes.

Closed systems show their strength by having identifiers which are internal and maintained according to the local specifications, but identifiers conceived for closed systems can hinder the application of international document standards to local usage. The technical standards (mentioned above) for registries and services give a platform for open system and maintenance based on subsidiarity. However, this implies that the policies and procedures of the registrar will affect the quality, reliability and/or trustworthiness of an identifier; relaxed revocation policies might be acceptable in some cases, but catastrophic in others. Building open services on existing standardized identifier requires the conscious specification of identifier-related requirements along a set of dimensions:

1. Main aspect of assignment: legal, logistical, technical or organisational.
2. Should it be possible to verify its existence or other qualities? (A positive answer could give rise to the requirement for an ISO 6523 or ISO 15459 compliant solution.)
3. Should it display information understandable by humans, like date of birth, sex or age?
4. Should it be shared for reuse by other systems:
 - a. Within the user's organisation
 - b. Within a closed group, like a trading relationship
 - i. Domestically
 - ii. Across European borders
 - iii. Globally
5. Time: Should it be valid for a single process, for a period of time or a defined longer term?

The answers to these questions will help to conclude whether an existing identifier should be selected, whether a new one needs to be defined and agreed

within its user community, and what factors should govern its subsequent processing.

That processing should take into account concerns of integrity (of identity data); consistency; correctness and non-repudiation of identity-based assertions, but still have appropriate regard to the privacy of the user.

Characteristics of Identifiers, and Selection Criteria

In the introductory section of this document we looked at how identifiers have come to be used in typical procurement systems. However, in order to help suggest how to answer the questions above, we should take a short detour at this point to look at some of the fundamental principles of identifiers – how they relate to (and differ from) identities, what purposes they are best suited to, and what this implies for the design of systems which consume identifiers and related attributes.

It may be tempting to think that there is no distinction between an identity and an identifier. After all, if you are asked to prove your identity, or to show what the French refer to as 'une pièce d'identité', what is it that you show? Is it your identity, an identifier, proof of identity, or something else?

We would like to propose a simple model for distinguishing between these various possibilities.

One way to consider the word 'identity' is to consider the word 'identical' ... if we say that two things are identical, another way to express that is to say that there is a relation of identity between the two things. (This reduces the concept of identity to theoretical terms which would be recognisable to a philosopher like Leibniz⁷⁾ or a mathematician like Russell). In human terms, about the only equivalent use of the term 'identical' would be when referring to 'identical twins' ... though even there, we would have to admit that we are not referring to the twins as being 'the same thing' – only that they are the same in genetic make-up, birth and usually appearance.

However, there is a way in which the idea of 'identical' is useful when talking about human identities of the kind we want to characterise here. When I present my passport at a national border, I am asserting a relationship of identity between me (here, now) and the person to whom the passport was issued (elsewhere, earlier). There are various factors which make

⁶⁾ LDAP: Lightweight Directory Access Protocol; DNS: Domain Name Server; OCSP: Online Certificate Status Protocol. All are forms of directory or directory access protocol.

⁷⁾ Leibniz' Law of the Indiscernibility of Identicals. http://en.wikipedia.org/wiki/Identity_of_indiscernibles#Identity_and_indiscernibility

that assertion more reliable. For instance, does the photograph in the passport look like me? Can I produce, on demand, a signature which matches the one in the passport? Do I match other characteristics recorded in the passport (such as gender, height, eye colour, etc)?

There are also factors which can make the assertion of identity less reliable: for instance, could the passport officer be sure that I am not one of identical twins, and on that basis, which one of us is presenting the passport? Or, imagine that I have average or unremarkable appearance and a very common name. Is it possible that, among the possible population of passport-holders, there is someone else sufficiently like me that they could pass for the valid holder of my passport? In practice, the safeguard against this happening is that, in the process of issuing the passport, the issuer tries to assemble robust proof of the personal characteristics I am claiming (for instance, what evidence can I present which substantiates my claimed date of birth, name, nationality and so on?).

This suggests three relevant and useful things:

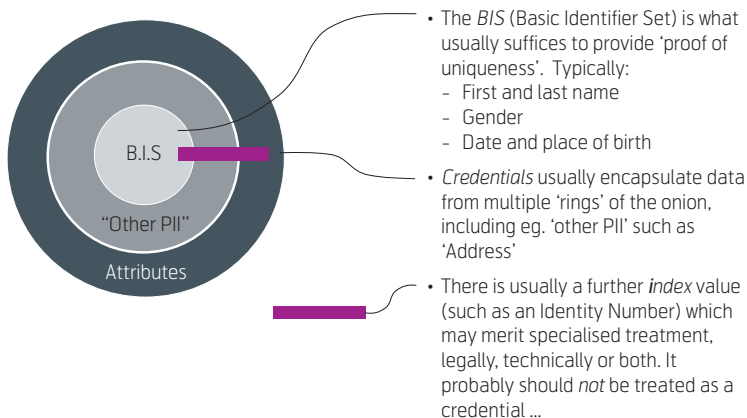
- First, that there is a distinction between ‘identities’ and ‘credentials’. In this case, the goal of the credential (the passport) is to encapsulate enough information about me to support a trustworthy assertion of the fact that the person presenting the passport is identical with the person to whom it was issued. There is a lot about my identity, however, which is not recorded in my passport (for example, what language I speak, who I am married to, even what I smell like ...). One quick way to sum it up might be “identity is something I have, inherently; a credential is something someone else can give to me which reflects some aspects of my identity”.
- Second, that a core purpose of a credential is to provide evidence in support of a claim of uniqueness within a given population. I use the word ‘population’ broadly here: in the case of a passport, I mean ‘population’ in its normal sense of the inhabitants of a nation state. However, in the case of a driving licence, say, ‘population’ might just mean the set of people authorised to drive (in other words, a subset of the national population). All the same, if my driving licence does not uniquely identify me among those who are entitled to drive, it is likely to fail to meet some of its core functional requirements.
- Third, it may be tempting to conclude, at this point, that a credential is the same as an identifier. I would warn against that temptation. My suggestion

is that most credentials include identifiers, but that in many cases, those identifiers are only a subset of the information the credential may hold.

To see why, let us look again at that requirement to ‘uniquely identify the holder, in a given population’. I know a man called David Walker. His is not an uncommon name, and, as it happens, another boy was born in the same hospital as him, on the same day, and was christened David Walker. The ‘normal’ pieces of information (name, gender, date of birth, place of birth), therefore, are not enough to distinguish between these two David Walkers – and if just those pieces of data were encapsulated in a credential, it would not be enough to distinguish between them either.

However, most credentials like driving licences, passports and so on include a further piece of data which is not a personal characteristic of the holder, but serves to ‘index’ the credential. Thus passports tend to have a serial number which is allocated in such a way as to ensure that two credentials cannot be issued with the same one (US Social Security numbers are issued using a similar principle). Thus, even if all the other data encapsulated in the credential is non-unique, the ‘index’ should be, and should thus make two otherwise indistinguishable credentials different.

Indeed, many credential-issuing systems contain, if one looks for it, some kind of index value which is not necessarily reflected in the credential itself, but which has the effect of making it possible to distinguish between records which otherwise contain identical values.



One ‘good practice’ approach to digital credentials suggests that they should ‘gravitate’ towards the centre of the onion - ie. not be overloaded with attribute data, but provide the means to link to it. The architecture described is compatible with this approach

Figure 2 The ‘Onion’ Model of Identity Data

That said, credentials such as driving licences and passports are, one might argue, just more complex than most others ... and one could argue that credentials such as email addresses are indeed simple identifiers. They serve to identify a user uniquely within the population of users of that mail service, and may well contain no data other than what is needed to perform that function.

Nevertheless, over the last few years, this model has proved a useful and productive way to draw out some of the defining characteristics of 'identity', 'identifiers' and 'credentials', and therefore how best to design and manage the systems which consume them. The diagram in Figure 2 illustrates the model, showing the various kinds of Personally Identifiable Information (PII) and other attributes which may be encapsulated in a credential such as a driving licence or passport.

Open Cross-Border Procurement and the Requirement for Standardized Systems for Identifiers

One key implication of 'open' procurement (particularly in a cross-border context) is that there is no coupling (tight or otherwise) between the issuing of a credential and the subsequent use of that credential in a given procurement transaction. In the simplest 'closed' case, Organisation A accredits John Smith as a procurement officer, and exchanges that information with Organisation B. When John Smith presents his credentials, Organisation B recognises the issuer and the significance of the credentials – namely, that John Smith is entitled to execute procurement transactions.

In an open case (within or across borders) one cannot assume that Organisations A and B have exchanged information about John's status, or that they have a mutually recognised credential. We may also want to increase the openness of the system by de-coupling the following two things:

1. Credentials which establish John's identity (for instance, as a citizen);
2. Credentials which establish John's status as a procurement officer.

One reason for this is that it makes it possible for John to be issued with a single credential (which establishes his identity as a citizen, say), but which can then, under the right circumstances, be used in support of many other kinds of assertion (that John is a licensed driver, that he is a qualified engineer, or – in this case – that he is a procurement officer).

This has implications for the design of such a system – in particular, it implies a highly distributed architecture in which the following are loosely coupled, but still reliable to all parties concerned:

- The issuing of identity credentials;
- The issuing of other trustworthy assertions;
- Reliance on those trustworthy assertions.

To sketch out such an architecture, let us define the following roles:

- **Registry:** The Registry is an organisation responsible for issuing citizens with a robust credential, capable of being used online to make assertions of identity (that is, proof of uniqueness within the national population);
- **Qualification Authority:** The Qualification Authority is responsible for certifying that John is a procurement officer. In that case, it is most probably his employer (ie. the organisation on behalf of whom he executes procurement transactions). In other cases (such as certifying that John is a qualified engineer), it might be a professional body or industry association, for example.
- **Relying Party:** The Relying Party is the organisation which wants to do business with John's employer, but only through transactions which are executed by a genuine procurement officer.

So:

- John is issued with a generic identity credential by the Registry; this credential allows John to assert his identity uniquely in the national population.
- John goes to the Qualification Authority, which accredits him as a procurement officer. The Qualification Authority uses John's identity credential to 'bind' an assertion of his identity to an assertion that he is an accredited procurement officer.

There are then two design options:

1. The Qualification Authority gives John a copy of that 'joint' assertion, which John simply presents to the Relying Party in support of a given procurement transaction;
2. The Qualification Authority stores a record of the 'joint' assertion; when John needs to prove his accreditation to the Relying Party, he just asserts his identity. The Relying Party knows who the Qualification Authority is (either because they want to transact with John's employer, or because

they know what qualification they are seeking to establish), so they issue a request to check John's status. The Qualification Authority responds with the 'joint' assertion, and the Relying Party is able to match the identity element of that against John's generic identity credential.

As this simple example illustrates, it is possible to design for a distributed architecture in which

- There is no requirement for specific credentials to be issued to cater for every use-case (ie. every new combination of individual/role, Qualification Authority and Relying Party);
- A generic credential can be issued and then used to generate assertions which are specific to a given combination of individual/role, Qualification Authority and Relying Party);
- This reduces the cost and effort required to cater for multiple and new instances of such combinations;
- It places the quality/accountability burden for the issuing of identity credentials where it belongs (that is to say, at the Registry);
- It places the quality/accountability burden for the issuing of professional accreditations where it belongs (that is to say, at the Qualification Authority).

A generic architecture is illustrated in Figure 3. NB – this diagram only represents the topology and a simplified 'flow'; I have not attempted to show the process of 'binding' John's identity assertion with his qualification assertion. In a highly distributed ('Web 2.0-style') implementation, the relying party might also simply ask John for an identity-based authentication, then send successive requests to the Registry (for confirmation of identity) and the Qualification Authority (for confirmation of accreditation).

Candidate Requirements for Pilot Projects

We conclude that open electronic procurement requires standards for documents and identifiers for entities and persons involved in the business processes; the latter needs to have sufficient quality and be subject to a defined and accountable maintenance regime. Those standards can be specified in ways which cater for highly distributed architectures, without sacrificing the ability to place accountability with the most appropriate stakeholder. In pursuit of such standards, open electronic procurement – especially

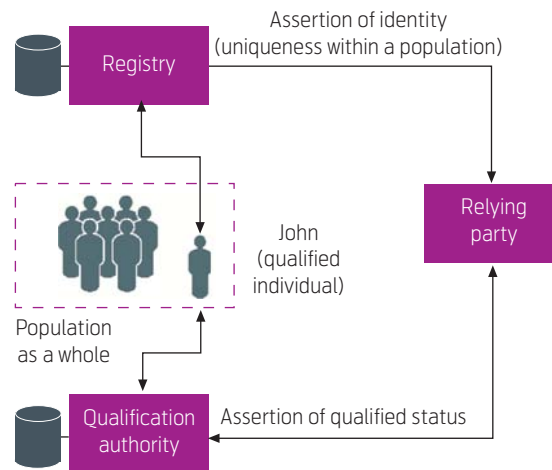


Figure 3 Distributed authentication/accreditation architecture

across borders – calls for extended investigation of the topic of identifiers.

Our thesis is that some pilot projects should be established to address these issues, and to generate suggestions for guidance and good practice in the use of electronic identifiers in e-procurement systems.

We therefore suggest a particular focus on the systematic definition of requirements for the design, selection and use of such identifiers, as set out earlier in this document. Specifically:

- What is to be identified, legal responsibility, site, account?
- Usage and legal, semantic, technical or organisational constraints.
- Should it be possible to verify its existence or other qualities? A positive answer should probably lead to a need for an ISO 6523 or ISO 15459 compliant solution with the use of maintained code lists.
- Should it display information understandable for humans, like date of birth, sex or age?
- Should it be shared for reuse by other systems?
 - a. Internally
 - b. Within a closed group, like a trading relationship
 - i. Domestically
 - ii. Across European borders
 - iii. Globally

The answers to these questions help to conclude whether an existing identifier should be selected, whether a new one needs to be defined and agreed

within its user community, and what factors should govern its subsequent processing.

We further believe that some of these requirements are best met through the definition of metadata, and of methods for associating metadata with identifiers so as to produce the kind of 'joint' assertion referred to above. For example, options might include:

- Identifiers which imbed metadata explicitly;
- Identifiers which are combined with metadata (for instance, by digital signing);
- Identifiers which are combined with a pointer to metadata which resides elsewhere.

Standards for the structure of electronic documents have existed for decades; open ones derived by UN/CEFACT, UN/EDIFACT, OASIS/UBL are in use, but differences in implementations and use of identifiers obstruct interoperability.

Identifiers have been considered technical objects used within systems; the present work shows that identifiers should be subject for detailed analyses and specifications in order to become essential tools for optimization of interoperability between systems across legal borders. This proposal results from our analysis of our observations, which pilot projects could verify, modify and further inform.

Sverre Bauck obtained his PhD in biophysics at the University of Oslo (1974) and was Postdoc at the National Institutes of Health in Washington DC, USA (1976-77). He has been working with digital systems and solutions for nearly forty years, including digital electronics, measurement technology, software development, solutions, security, teaching and standardization on a national and international level.

From 1986 to 1988 he worked at the Norwegian Directorate of Customs and Excises with the management of development and implementation of the electronic declaration system TVINN. The following four years he worked as a representative of the EFTA countries in the Western European EDIFACT Board Secretariat in the European Commission coordinating the development of EDIFACT standards.

Later on he has served as senior adviser in Statskonsult, ErgoGroup, the Brønnøysund Register Centre and currently with the Agency for Public Management and eGovernment working with analyses and solutions for public and private clients. Between 1995 and 1999 he led the EU project TAPPE – Telematics Applications: Public Procurement in Europe, and presently he is addressing the topic of identifiers in the EU project PEPPOL (www.peppol.eu). He is a member of the Kantara Initiative eGov group and has been writing articles and giving presentations on related subjects over the last two decades.

email: svb@difi.no

Robin Wilton is the founder and director of Future Identity Ltd., an independent company set up in January 2009 to provide structured consultancy on digital identity, privacy and public policy. Future Identity's clients have included: the Liberty Alliance; the UK VOME project (Visualisation and Other Methods of Expression); Internet Society; a UK Central Government department and the UK Information Commissioner's Office.

Robin is Director of Privacy and Public Policy at the Kantara Initiative – a world-wide consortium on interoperable digital identity – and has worked on privacy-enhancing applications of the Liberty Alliance's SAML implementations. He is an expert reviewer on two European FP7 projects relating to the Critical Financial Infrastructure, and is on the advisory boards of the European PrimeLife Project on privacy and identity management, and the UK's EnCoRe project on Consent and Revocation. Robin graduated from Oxford University in 1984 with an MA (Joint Honours) in Philosophy and Modern Languages; he worked for IBM for 12 years in systems engineering, technical support and consultancy; he left IBM to join an internet start-up, JCP Trustbase, which was acquired by Sun Microsystems in 2000. Robin spent nine years at Sun, the last four as Corporate Architect for Federated Identity in Sun's CTO team. He is a Fellow of the British Computer Society, with Chartered IT Professional status, and has also published papers on aspects of digital identity and privacy.

email: futureidentity@fastmail.fm