



DI1.3: Standardisation Strategy

Written by:

James Brusey, University of Cambridge

Mark Harrison, University of Cambridge

Duncan McFarlane, University of Cambridge

DELIVERABLE NO	DI1.3: Standardisation Strategy
DATE	14 May 2006
WORK PACKAGE NO	WP I1
VERSION NO.	1.4
ELECTRONIC FILE CODE	DI1.3_Standardisation_Strategy.doc
CONTRACT NO	507100 PROMISE A Project of the 6th Framework Programme Information Society Technologies (IST)
ABSTRACT:	This report revisits and redefines the PROMISE Standardisation Strategy.

STATUS OF DELIVERABLE		
ACTION	BY	DATE (dd.mm.yyyy)
SUBMITTED (author(s))	J. Brusey	15.05.2006
VU (WP Leader)	D. McFarlane	15.05.2006
APPROVED (QIM)	D. Kiritsis	15.05.2006

Revision History

Date (dd.mm.yyyy)	Version	Author	Comments
21.04.2006	1.0	James Brusey	IETF process and apply template
23.04.2006	1.1	Mark Harrison	EPCglobal process
04.05.2006	1.2	James Brusey	Finalising action plan
10.05.2006	1.3	Mark Harrison	Some corrections and clarifications
14.05.2006	1.4	James Brusey	Changes according to review comments

Author(s)' contact information

Name	Organisation	E-mail	Tel	Fax
James Brusey	University of Cambridge	Jpb54@cam.ac.uk	+441223765605	+441223338076
Mark Harrison	University of Cambridge	Mgh12@cam.ac.uk	+441223338178	+441223338076
Duncan McFarlane	University of Cambridge	Dm144@cam.ac.uk	+441223338069	+441223338076



Table of Contents

1	INTRODUCTION.....	3
1.1	PURPOSE OF THIS REPORT	3
1.2	OVERVIEW	3
1.3	REPORT STRUCTURE	3
2	BACKGROUND	4
2.1	STANDARDISATION WORK WITHIN THE PROMISE PROJECT TO-DATE	4
2.2	WHAT IS STANDARDISATION?	4
2.3	OPEN STANDARDS	5
2.3.1	<i>What are Open Standards?</i>	<i>5</i>
2.3.2	<i>Why Does PROMISE Want Open Standards?</i>	<i>6</i>
2.3.3	<i>How Can PROMISE Achieve Open Standards?</i>	<i>6</i>
2.4	EXISTING PROCESSES	7
2.4.1	<i>IETF Standards Process</i>	<i>7</i>
2.4.2	<i>EPCglobal Standards Development Process</i>	<i>8</i>
2.5	SUMMARY	10
3	PROMISE STANDARDISATION STRATEGY	11
4	ACTION PLAN.....	13
4.1	COLLECTION & ASSESSMENT OF EXISTING RELEVANT STANDARDS	13
4.1.1	<i>Develop a common architecture across demonstrators.....</i>	<i>Error! Bookmark not defined.</i>
4.1.2	<i>Examine specific standards activities.....</i>	<i>13</i>
4.1.3	<i>Examine specifications from the Research Cluster</i>	<i>13</i>
4.2	RECOMMENDATIONS FOR STANDARDS/SPECIFICATIONS.....	13
4.3	EVALUATION AND IDENTIFICATION.....	13
4.4	PROMOTION OF STANDARDS	14
4.4.1	<i>Establish relationships with key standards bodies</i>	<i>14</i>
5	CONCLUSIONS	15
6	APPENDIX.....	16
6.1	LESSONS LEARNT FROM THE EPCGLOBAL STANDARDS PROCESS	16
7	REFERENCES.....	21

Abbreviations:

Abbreviations used in this document:

IETF	Internet Engineering Task Force
ISO	International Standards Organisation
IEC	International Electrotechnical Commission
SGML	Standard Generalised Markup Language
RFC	Request for Comment
W3C	World Wide Web Consortium
EPC	Electronic Product Code
SAG	Software Action Group (part of EPCglobal)
HAG	Hardware Action Group (part of EPCglobal)
BAG	Business Action Group (part of EPCglobal)
UPnP	Universal Plug'n'Play
PLCS	Product Life Cycle Services
OASIS	Organisation for the Advancement of Structured Information Services
OGC	Open GIS Consortium



WS-I	Web Services Interoperability organisation
ATA	Air Transport Association
ITU	International Telecommunications Union
SOAP	Simple Object Access Protocol
WSDL	Web Services Description Language
RFID	Radio Frequency Identification
SNMP	Simple Network Management Protocol
FMCG	Fast Moving Consumer Goods
HLS	Health and Life Sciences
TSC	Technical Steering Committee
ARC	Architectural Review Committee
IPR	Intellectual Property Rights
LCWD	Last Call Working Draft

1 Introduction

1.1 Purpose of this Report

This report is largely a response to recommendation 14 of the first (Month 12) review (Anderl *et al.*, 2005), which is restated here:

Recommendation 14: Define a clear standardisation strategy

PROMISE should decide if it intends to be a follower or a driver in standardisation, clearly defining its strategy and action plan. PROMISE should state clearly if it intends to go for standardisation as a support to the industrial exploitation of its solution (as indicated in the DoW), or if they intend just to extend some standards (which ones?) or finally if they intend only to adopt existing standards to facilitate technology acceptance.

The aim of this report, therefore, is to define a clear standardisation strategy for the PROMISE project.

1.2 Overview

This report will argue that standardisation plays a central role in the PROMISE project. Despite this, the complete development and ratification of new PROMISE standards is unlikely to occur during the period of the project. Instead, the PROMISE project should aim to produce a candidate technical specification¹ for a new standard (or set of standards) and to submit it to an international standards body. Where appropriate, PROMISE should be a *driver* in standardisation.

Nevertheless, it is clear that there are a number of existing or emerging standards in this area. PROMISE should not reinvent existing technology, or try to compete with existing standards bodies. Therefore, as far as possible, PROMISE should adopt or extend existing standards. Where appropriate, PROMISE should *follow*.

In either case, the PROMISE consortium is not a standards body and will not produce standards. Instead, it will produce *candidate technical specifications* and submit them to an international standards body as the basis for a future standard.

Regardless of which mechanisms are used, a guiding principle for PROMISE is to ensure that the technology produced is based on *open standards*.

1.3 Report Structure

The report opens with a review of relevant background, including a review of previous work by the PROMISE consortium in this area. Section 3 proposes a standardisation strategy for PROMISE in light of this background. The fourth section recommends an action plan to ensure this strategy is followed. Finally, section 5 concludes the report.

¹ The term “candidate specification” in this document is intended to refer to an end product (or perhaps deliverable) as far as PROMISE is concerned, but perhaps will still be an initial working document as far as an international standards body is concerned. The main point to emphasise here is that any specification developed by PROMISE will have yet to undergo a process of achieving consensus amongst the wider community before it can become a standard.

*“The nice thing about standards is that there
are so many to choose from”*

– Andrew S. Tanenbaum

2 Background

2.1 Standardisation Work Within the PROMISE Project To-Date

In addition to the recommendation mentioned in the introduction, the first interim technical assessment report (Anderl and Bonetti, 2005) also made a similar recommendation regarding standardisation. The recommendation is restated here:

Recommendation 9: Focussing the PROMISE standardisation view

Standardisation in PROMISE needs to have a more clear and focussed view, defining the areas of standardisation where the project intends to contribute and have an influence. New standardisation activities called “PLM services” which have been standardised by OMG and others with respect to visualisation should be considered. Implications for PROMISE deriving from the standardisation scenario described in the WPI1 report should be understood. The overall policy and strategic objectives that PROMISE should aim to achieve in terms of support to standardisation, “limiting the variety” at the right level, should be made explicit in a supplementary specific report to be prepared and submitted at the next Month 12 review meeting. This standardisation policy report should be combined with the previous IMS related policy paper.

In response, a brief report (Stromberg and Barasic, 2005) was provided to the review committee at month 12. The main conclusions of the report were that PROMISE should work with standards in the following ways:

1. Use an existing standard where applicable
2. If (1) is not possible, extend an existing standard
3. If (1) and (2) are not possible, find a best practice definition

In addition, the report referred to the outcome of the associated deliverable DI1.2 (Stromberg *et al.*, 2005) and noted that four standardisation domains had been identified. Namely:

- Content domain,
- Device interoperation domain,
- PROMISE middleware domain,
- Backend domain.

2.2 What is Standardisation?

Wikipedia contributors (2006a) give a modern definition of standardisation:

“Standardisation ... is the process of establishing a technical standard among competing entities in a market, where this will bring benefits without hurting competition ...

... it is a universally agreed upon set of guidelines for interoperability”

Standardisation, according to this definition, sets a technical standard without introducing barriers to entry into the market. By being guidelines for interoperability, they may allow many competing

products to interoperate and thus compete on an even footing. Universal agreement is necessary to ensure that the standard is optimal from all points of view. Without this, a standard might suit some stakeholders but not others, and thus provide an unfair benefit to those that it suits.

The term “standard” has historically been the province of international bodies such as ISO (International Standards Organisation) or the IEC (International Electrotechnical Commission). The end of the 20th century has seen the growth of industrial consortiums devoted to standardisation. These consortiums have generally been much faster to market and produced less complex standards.

Compare, for example, the Standard Generalised Markup Language (SGML) and the HyperText Markup Language (HTML). SGML, ISO standard 8879:1986, was a standardisation of IBM’s GML, originally developed in the 1960s. Although it was widely used in printing and publishing, its complexity prevented it from being used more generally (Wikipedia contributors, 2006b). HTML, originally standardised by the W3C consortium (but now also an ISO standard) borrows several aspects of SGML but is much simpler and less rigorous. HTML was much better suited to the domain of serving web pages, and is now universal and popular. In the fast moving field of computer and Internet technologies, traditional bodies such as ISO tend to be too slow to allow the rapid development of standards that this type of industry requires. Industry consortiums can respond much more quickly and achieve many of the desirable effects of standardisation.

Although these “lesser” standardisation bodies have, in many ways, had the same effect as the more formal bodies (such as ISO), the tendency is to not refer to their output as a true “standard”. For example, the IETF publish “Requests for Comment” (RFCs), and the W3C publish “Recommendations”. Perhaps the point is that such recommendations or specifications are not intended to become legally binding. Instead, various parties adopt such “standards” because it is practical, convenient and helpful. They are adopted *de facto* – parties tend to behave as if there was a law regulating the standard even though no such law exists.

To ensure that standards are available to all, and to minimise inequity that might be produced by adopting a standard, a standard should be “open”. Open standards are discussed in more detail in the next section.

2.3 Open Standards

2.3.1 What are Open Standards?

Open standards are aimed at minimising the barriers to entry to all participants. West (2004) defines an open standard as one that provides equal rights to all economic actors, not just rights to the sponsor. Perens (no date) provides some more detailed requirements for open standards, from an open source perspective.

Availability – the standards documents should be readily accessible

Maximize end-user choice – the standard does not impose unnecessary or restrictive conditions on how the technology is used.

No royalty – there is no fee per use (although there may be fees associated with conformance testing).

No discrimination – anyone may access and use the standard regardless of, for example, ethnicity, or national allegiance.

Extension or subset – the standard should allow extension or subset.

Predatory practices – the standard may include license conditions to avoid “embrace and extend” tactics that would otherwise close the standard.

Open standards such as HTML (Hypertext Markup Language) can be contrasted against proprietary standards such as Microsoft Word. In the case of the open standard, many companies have been able to create HTML browsers and editors. In the case of the proprietary standard, only Microsoft has been successful at creating browsers and editors for Word. Users have plenty of choice when choosing an HTML browser and the software is generally available for no charge. Users have no choice when choosing a Word editor and the software is expensive. Microsoft may argue that the fee charged for Word encourages them to innovate. It is not clear that this is actually the effect.

Open standards have documentation that is available to all, whereas proprietary standards may have no documentation or documentation that is kept secret.

Note that Perens sets a requirement for there being no royalty associated with the standard. More often, standards consortiums require that royalties are “reasonable and non-discriminatory”. The definition of what is “reasonable” is ambiguous, but *any* royalty can be an insurmountable obstacle to the development of open source software.

Nonetheless, it is worth noting that some organisations that develop open standards (such as EPCglobal Inc.) have been unable to guarantee that their standards are free of royalties or possible infringement of patents, even though their intellectual property policy aims to encourage royalty-free standards where possible (Roberti, 2005).

2.3.2 Why Does PROMISE Want Open Standards?

Open standards are desirable for PROMISE since it is intended that the benefits of this project be available to the European Union as a whole.

If open standards were not an aim of PROMISE, a proliferation of proprietary solutions might result. These solutions would not necessarily work together. Without a clear standard, other companies or individuals would have difficulty in reproducing the work. This may be the case, even with a standard, if use of the standard is restricted.

In the short term, proprietary or closed standards, or even no standard at all may seem desirable to some companies (such as Microsoft) as it provides them with the maximum short-term benefit and minimises their competition. In the longer term, however, clearly defined open standards will enable wide adoption of the technology and allow the market to grow.

The PROMISE project is undertaken with the assumption that novel and useful technologies will be the result. This combination of novelty and usefulness will surely lead to new standards being formed (de facto or otherwise). Nevertheless, if no action is taken to ensure that the standards produced are open, they might easily be closed, proprietary ones.

2.3.3 How Can PROMISE Achieve Open Standards?

Krechmer (2006) proposes ten requirements for enabling open standards:

1. Open meeting – all may participate in the standards process
2. Consensus – all interests are discussed
3. Due Process – the use of ballots and provision for appeals
4. Open IPR – holders of copyright, patents, and trademarks relevant to the standard must make them available
5. One World – same standard for the same capability globally
6. Open Change – all changes are presented and agreed to in a forum following the above requirements

7. Open Documents – committee drafts and completed standards are easily available for use
8. Open Interface – interfaces (in particular, software interfaces and communication protocols) are not hidden and support extension.
9. Open Access – conformance to the standard is tested objectively
10. On-going Support – support for a standard continues even if the original sponsor is no longer involved.

The PROMISE consortium may not be able to completely develop standards. In fact, this report will argue that this should not be our goal. Nonetheless, some of the above requirements remain relevant. Given an intention to allow open standards to be created, it is necessary to ensure that the possibility of creating an open standard is not precluded. Specifically, action should be taken now to ensure that licensing restrictions on patents or trademarks do not become a barrier to exploitation; that specifications are produced in a manner that allows for global use, not just European use; and that software interfaces support extension. In addition, PROMISE should begin to plan the process for handing over candidate specifications to a standards body that supports open standards.

2.4 Existing Processes

There are a large number of standards bodies, and while their processes are similar, they are not the same. Two approaches, namely those of the IETF, and EPCglobal Inc., are examined below.

2.4.1 IETF Standards Process

The Internet Engineering Task Force (or IETF) uses a formal standards process based around Request for Comment (RFC) documents. Revision 3 of the IETF's standard process (Bradner, 1996) describes the process in some detail. The Internet Engineering Steering Group (or IESG) manages the standards process and is responsible for review and approval of standards. The main goals of the process are to ensure technical excellence, to have clear and concise documentation, to be open and fair, and to be timely. In addition, Bradner notes that prior implementation and testing is important.

As implied by the term “request for comment”, the process starts with “Internet drafts” or draft specifications that are widely circulated for review and revision. Drafts that are not updated, or recommended by the IESG for publication as RFCs, are dropped after six months. Once published, a standards-track RFC proceeds through three maturity levels: proposed standard, draft standard, and finally Internet standard. At the “proposed standard” phase it is not necessary for a working implementation to have been developed, although this is certainly desirable. To reach the draft stage it is necessary for there to be two, independently developed implementations. The final stage of “Internet standard” requires significant implementation and successful, real-world operational experience.

Note that RFCs are not always “standards-track”. Non-standards-track RFCs are designated experimental, informational, or historic. Experimental RFCs denote specifications for technologies that are being tested by one or more laboratory. Informational RFCs denote less formal documents that are nevertheless useful to the whole community. Historic RFCs are ones that are obsolete, usually due to the introduction of a newer RFC.

As with many other standards bodies, the IETF standards process has several requirements to do with ensuring that intellectual property associated with a standard is sufficiently open (see requirement 4 of Krechmer (2006)). For example, copyright owners must grant worldwide, royalty free, perpetual license for contributed material. Where relevant patents are known about, the IESG will attempt to obtain a written assurance from the patent holder that they will license

the patented technology to any implementer under reasonable and non-discriminatory terms. Although availability or validity of such assurance does not block the standard process directly, the requirement for two independent implementations to reach the draft standard stage is a useful test to ensure that the technology can be used in practice.

In summary, the IETF standards process is designed to ensure that the standards are implemented and well tested prior to being accepted. Gaynor and Bradner (2001) also argue that this type of process tends towards self-organisation, where different options must compete for survival in the marketplace and only those that thrive become accepted as full standards. Gaynor and Bradner offer three general rules to help limit the risk involved with standardisation:

1. Standards should have a modularized architecture to maximise the possibilities for experimentation
2. Standards introduction is best done in an evolutionary fashion, starting with simple elements and then growing in complexity.
3. Implementing the standard helps to show that it is at least feasible to do so. In addition, there should be at least two independently developed versions, in order to test the completeness and clarity of the standard.

2.4.2 EPCglobal Standards Development Process

EPCglobal Inc., is a not-for-profit subsidiary of GS1 (formerly the Uniform Code Council (UCC) and EAN International) and is responsible for the development of standards for the Electronic Product Code (EPC) Network architecture. It was formed to carry on the standardisation process begun by an industry-academic research project known as the Auto-ID Centre, which finished in 2003.

The Auto-ID Centre identified a number of technology components that would be required, including:

- A unique identifier or Electronic Product Code (EPC)
- Edge-level filtering to reduce volumes of data flowing
- Networked databases to store additional information about the objects
- A lookup mechanism to find the relevant networked databases for a given EPC identifier

The EPCglobal approach to standardisation is driven by the business use cases and requirements of the end-users who will use the technology in their businesses (EPCglobal Inc., 2006). This is to ensure that the needs of end users are fully considered in the technical standards that are ultimately developed. This is a different emphasis from some standards organizations (such as the IETF) where the technical feasibility is considered first.

EPCglobal make use of volunteers from member organisations to staff and run a number of action groups including: the Business Action Groups (BAGs); the Software Action Group (SAGs); and the Hardware Action Group (HAGs). BAGs are responsible for end-user needs, whereas SAGs and HAGs look at technical issues from a software and hardware perspective, respectively. Participation in an action group is subject to agreeing to the EPCglobal intellectual property policy and abidance by the code of conduct. In particular, the code of conduct includes an anti-trust caution, which guards against anti-competitive practices.

BAGs are formed for various industry sectors, the first two being for Fast Moving Consumer Goods (FMCG) and Healthcare and Life Sciences (HLS). Work groups are formed to address particular issues and collect requirement specifications (typically in the form of use cases). They examine these for commonalities and differences, and then extract the technical requirements.

The technical requirements are then passed to the Technical Steering Committee (TSC), which together with advice from the Architecture Review Committee (ARC) will either identify an existing HAG or SAG technical work group or form a new one.

All EPCglobal standards begin with a charter, which tightly specifies the scope of the standard and what it should aim to achieve. Figure 1 shows the EPCglobal submission track process for identifying the need for a new standard, resulting in the chartering of a new standard.

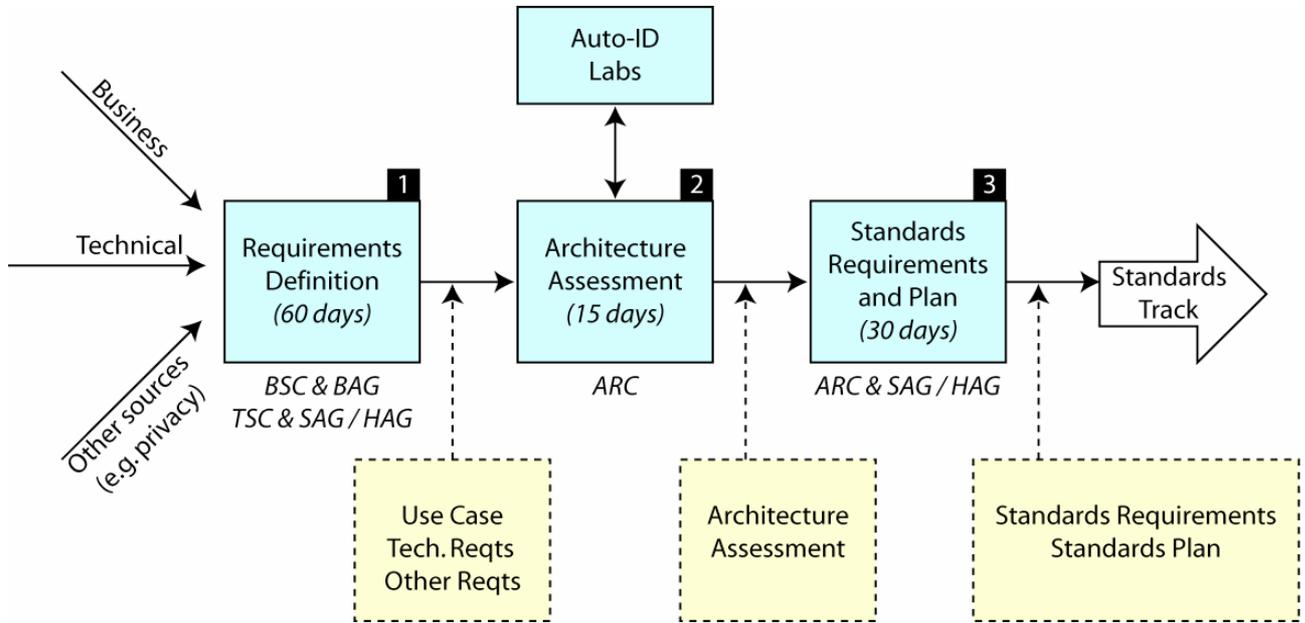


Figure 1 EPCglobal submission track process (requirements planning prior to standards development). Adapted from EPCglobal Inc., (2006)

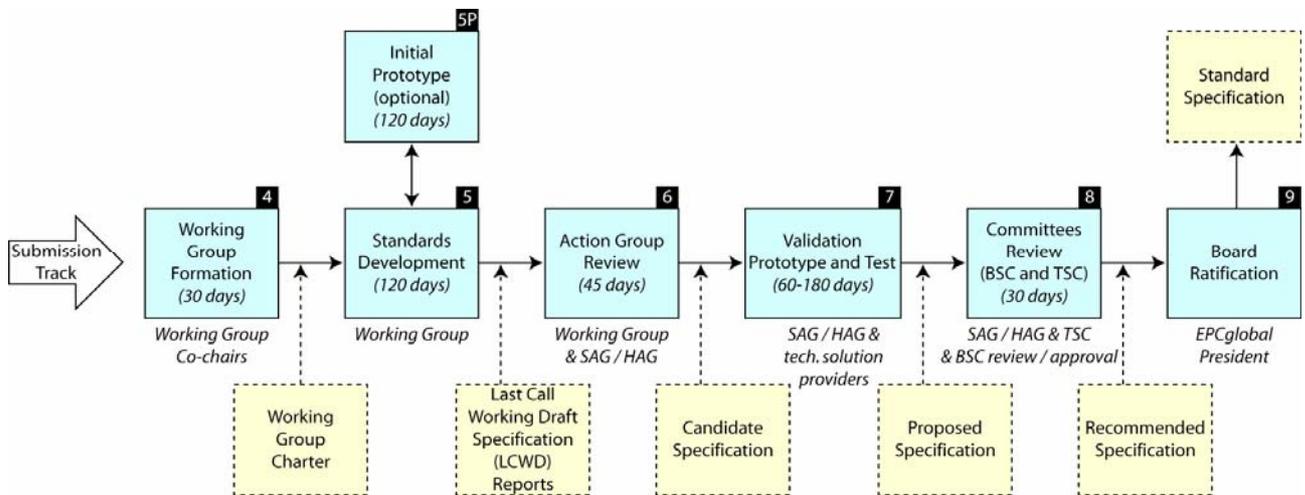


Figure 2 EPCglobal standards development track. Adapted from EPCglobal Inc., (2006)

Once a charter has been approved, a work group is formed to undertake the actual work of drafting a new standard. Figure 2 shows the series of stages that the standard must go through. The time periods shown are nominal values. For example, it may take between a few months to a couple of years to reach a stage at which consensus is reached within the work group. The draft standard is first reviewed within the work group, and then submitted as Last Call Working Draft (LCWD). All members have an opportunity to provide comments on the draft during a period of 45 days. Following this, the work group must review the comments received and make any necessary modifications to the draft. Further reviews are made of the technical content as well as possible patent infringements as the draft passes through Candidate Specification stage. A period

of validation of the standard by means of prototyping and testing then takes place, involving many of the technology solution providers. This is similar to the IETF standards process where two implementations are required prior to advancing beyond the proposed standard stage. The aims of the validation phase are to:

1. Extract the conformance requirements and test that working prototypes conform to these
2. Test the specification for ambiguous interpretation by different vendors

Next, the BAG that submitted the original requirements specification is given an opportunity to formally review whether it fully meets those requirements. The final stage is ratification by the EPCglobal Board of Governors.

The entire process from chartering to ratification takes a minimum of 9 months, although it is not unusual for the process to take 2-3 years for some standards.

EPCglobal standards are designed to be neutral with regard to operating system, hardware platform, programming language and terminology. The aim is to ensure that the standards defining the technical architecture are not biased towards any particular industry.

Further discussion of lessons learnt from the authors' experience with this relatively young standardisation body are given in the appendix.

2.5 Summary

A three-tiered approach has previously been recommended for PROMISE that tries to minimise the creation of new standards by looking to existing standards as much as possible. This approach is still relevant, however there is more to be said.

If PROMISE is novel, useful and successful, new standards are inevitable. Without action to ensure that they are open, they may be proprietary, de facto standards.

Examining existing standards processes, it becomes clear that, even for the relatively speedy industry consortiums, standardisation is a slow and risky process.

3 PROMISE Standardisation Strategy

The previous section established the fundamentals of standardisation and discussed why it is an important part of the PROMISE project. This section proposes a standardisation strategy for PROMISE.

The proposed standardisation strategy is based on a three-tiered perspective:

- **Adopt** existing standards where possible
- **Recommend** to extend / adapt existing standards
- **Initiate** development of new standards where a gap exists.

In addition, PROMISE will **interact** with standards bodies to

- Better understand their existing standards
- Influence the development of the standards to support PROMISE technology
- Leverage existing efforts to develop standards or specifications.

Finally, the overarching principle should be a commitment to ensuring that PROMISE technology is based on **open standards**.

The strategy seems at first glance to be the same as that proposed by Stromberg and Barasic (2005). As with their approach, the preference is to make use of existing standards where possible. In comparison with their approach, the above strategy emphasises interaction with standards bodies and clearly states that PROMISE will be involved in the initiating of new standards rather than producing best practice documents. These changes reflect two issues: First, that PROMISE is a research project and does not intend to become a standards organisation. Therefore, if new standards are needed, PROMISE must rely on existing bodies to perform the consensus process that leads to the full development of a standard. Second, creating new approaches and technologies with no view to standardisation may lead to proprietary standards being created *de facto*. To avoid this situation, PROMISE needs to actively interact with existing standards bodies.

The process of initiating new standards development is deserving of further discussion. PROMISE must interact with a standards body to perform this development. The scope of work for PROMISE might look something like figure 3. Requirements development outputs (such as architecture documents from WP A0, or working demonstrators from the application cluster) must achieve consensus to be approved as candidate specifications. Candidate specifications might immediately be submitted to the associated standards body as an input into their process. In some cases, PROMISE may find that it needs to revise the specifications. Another possibility is that the standards body may reject the specification and ask for it to be reworked. These two possibilities are shown as dashed arrows in the diagram.

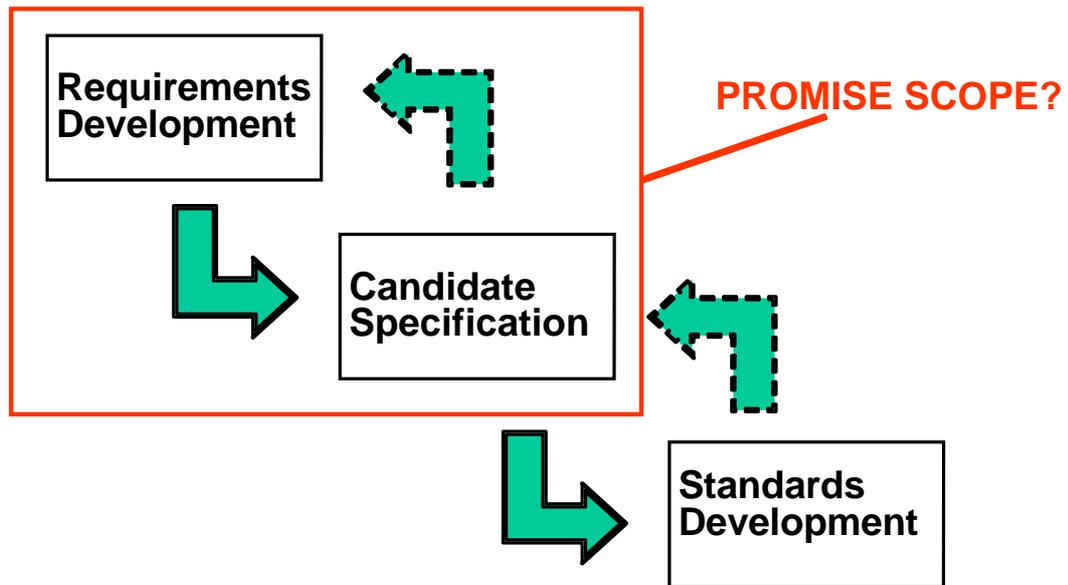


Figure 3 Scope of PROMISE work in creating new standards.

4 Action Plan

This report proposes the following plan to support the strategy proposed in the previous section.

4.1 Collection & assessment of existing relevant standards

The task of collecting and assessing existing standards is aimed at ensuring that PROMISE does not attempt to reinvent existing standards and maximises opportunities for making PROMISE interoperable with other technologies. This task can be broken down into the following sub-tasks:

- Examine specific standards activities
- Examine specifications from the Research Cluster

These sub-tasks are examined in further detail below.

4.1.1 Examine specific standards activities

There are a number of standards activities that may provide useful input to the project and the standardisation project. Some possible candidates include the EPC Network (because it involves a global standard for connecting RFID tags to back-end systems) and PLCS (because it has a focus on product life-cycle systems). There may be many others that are relevant to this task.

4.1.2 Examine specifications from the Research Cluster

Specifications developed as part of PROMISE Research Cluster activities need to be evaluated according to the three-tiered strategy (see section 3). This means that PROMISE must decide where existing standards should be adopted, which components should be recommended to a standards body as a possible extension, and which should be promoted as a new standard to a standards body.

4.2 Recommendations for standards/specifications

Where new standards are required, the initial stage will be to recommend that the associated specifications be drafted as standards to be submitted to a standards body. This task can be further divided into the following sub-tasks:

- Define the requirements of the interface or protocol that needs to be standardized. The requirements can be documented in a concise “charter” for a new standard.
- If resources allow, a full-length candidate specification can be developed. The candidate specification should include concrete examples of schema, interfaces, methods or functions, protocols, error messages and a description of each of these, together with precise clarifications about which features are mandatory and which are optional.
- Identify an appropriate international standards body and submit the charter and candidate specification to them. Greater credibility will be achieved if a number of the demonstrators implement the candidate specification and thereby achieve a high degree of interoperability between them, despite being developed independently.

4.3 Evaluation and identification

For each standards area

- Identify lead partner

- Identify appropriate standards body
- Produce draft specification
- Evaluate specification with regard to PROMISE objectives and demonstrators
- Submit draft to appropriate standards body
- Work with standards body to initiate process
- Produce final PROMISE standards blueprint

4.4 Promotion of Standards

Promotion of standards is an on-going task designed to promote standards within the consortium. It involves the following sub-tasks:

- Establish relationship with key standards bodies
- Presentation on PROMISE outputs and standards
- Introduce Standards material into Training

The sub-task of establishing relationships with key standards bodies is discussed in more detail below.

4.4.1 Establish relationships with key standards bodies

Various PROMISE consortium partners have already committed to interaction with a variety of standardisation bodies. These include:

- OASIS Product Life Cycle Services (PLCS) (EPFL)
- EPCglobal EPC Network (Cambridge)

In addition, there are several other standards organisations that might be considered. These include: ZigBee, IETF, UPnP, W3C, OGC, IEEE, ISO / IEC, WS-I, ATA, and ITU.

The nature of the interaction varies from PROMISE partners acting as members or chairing a working group to regular contact with an existing member of the standards organisation or one of its working groups.

5 Conclusions

This report has presented a standardisation strategy for PROMISE. The main outcomes of this work are

- Standardisation is central to PROMISE
- Closed (proprietary) standards should be avoided; PROMISE should commit to open standards
- PROMISE will need the help of existing standardisation bodies if PROMISE technology is to result in new standards
- Interaction with such bodies should also help to improve the PROMISE consortium's understanding of existing standards efforts
- Standardisation is risky and PROMISE should avoid trying to create too many new standards or creating standards for elements that do not need to be standardised.

6 Appendix

6.1 Lessons Learnt from the EPCglobal Standards Process

Figure 4 shows the EPC Network architecture as it stood in October 2003 at the time when the former Auto-ID Centre transitioned to EPCglobal and Auto-ID Labs. It is predominantly component-based. Almost all of the Auto-ID Centre specifications have since become obsolete by new EPCglobal standards (already ratified or in progress) or have otherwise been radically overhauled by EPCglobal to focus on standardizing interfaces.

The problem was that standards based on implementations were limited by the quality of the design of that implementation. Software often contains many internal interfaces as well as external ones. Finding the right interface level at which to standardise can tend to be a trial and error process. In this case, standardising based on the external interfaces of the implemented systems was much more successful and allowed implementers greater flexibility.

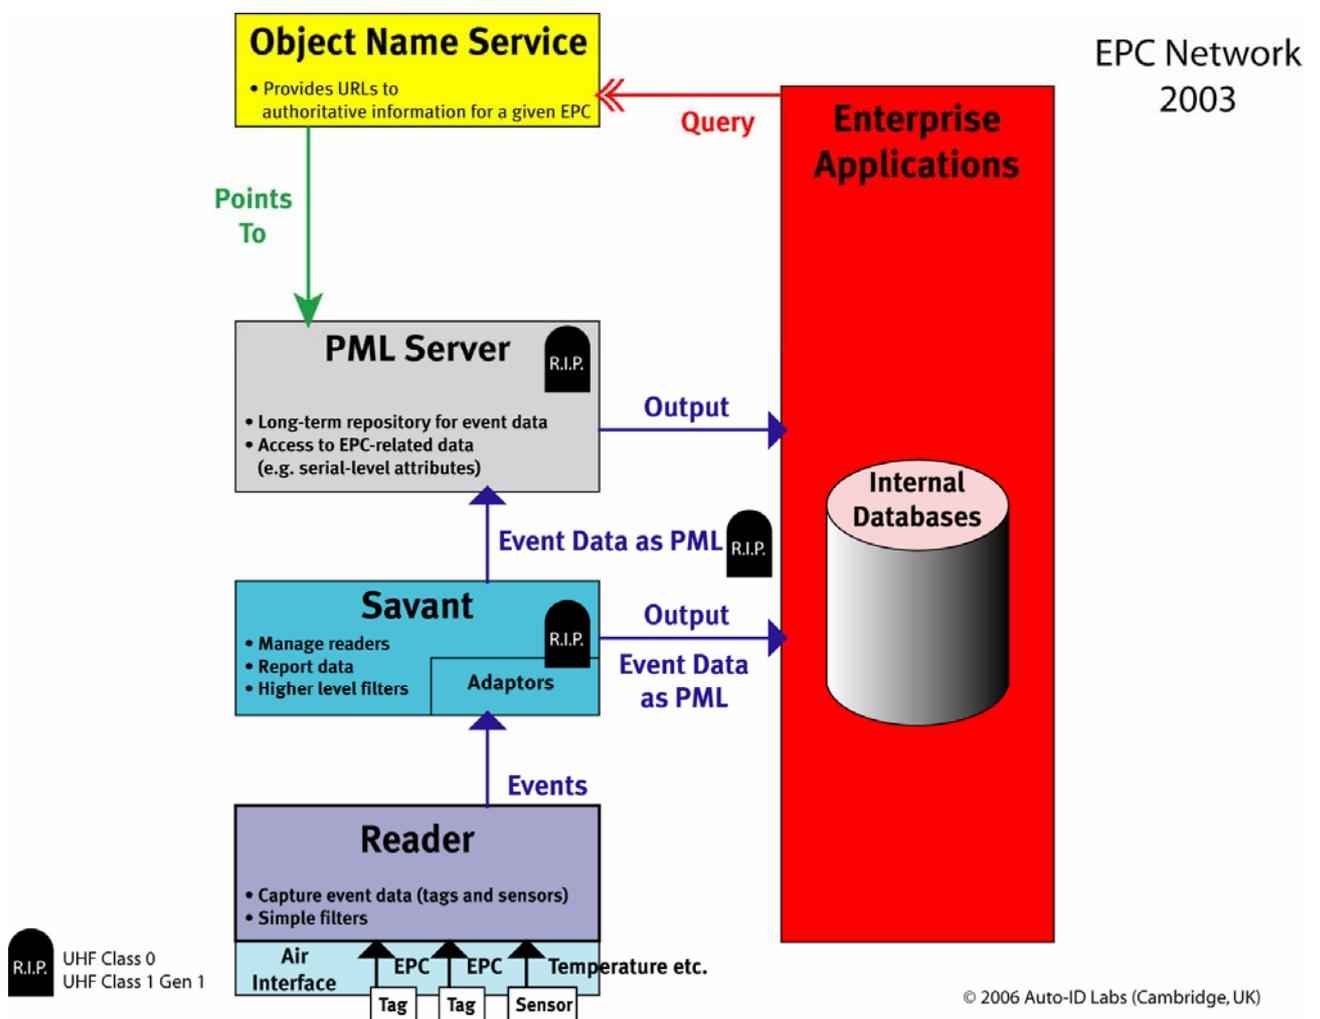


Figure 4 EPC Network architecture as of 2003. RIP indicates obsolescence.

Figure 5 shows a corresponding representation of the current EPC Network architecture (EPCglobal Inc., 2006).

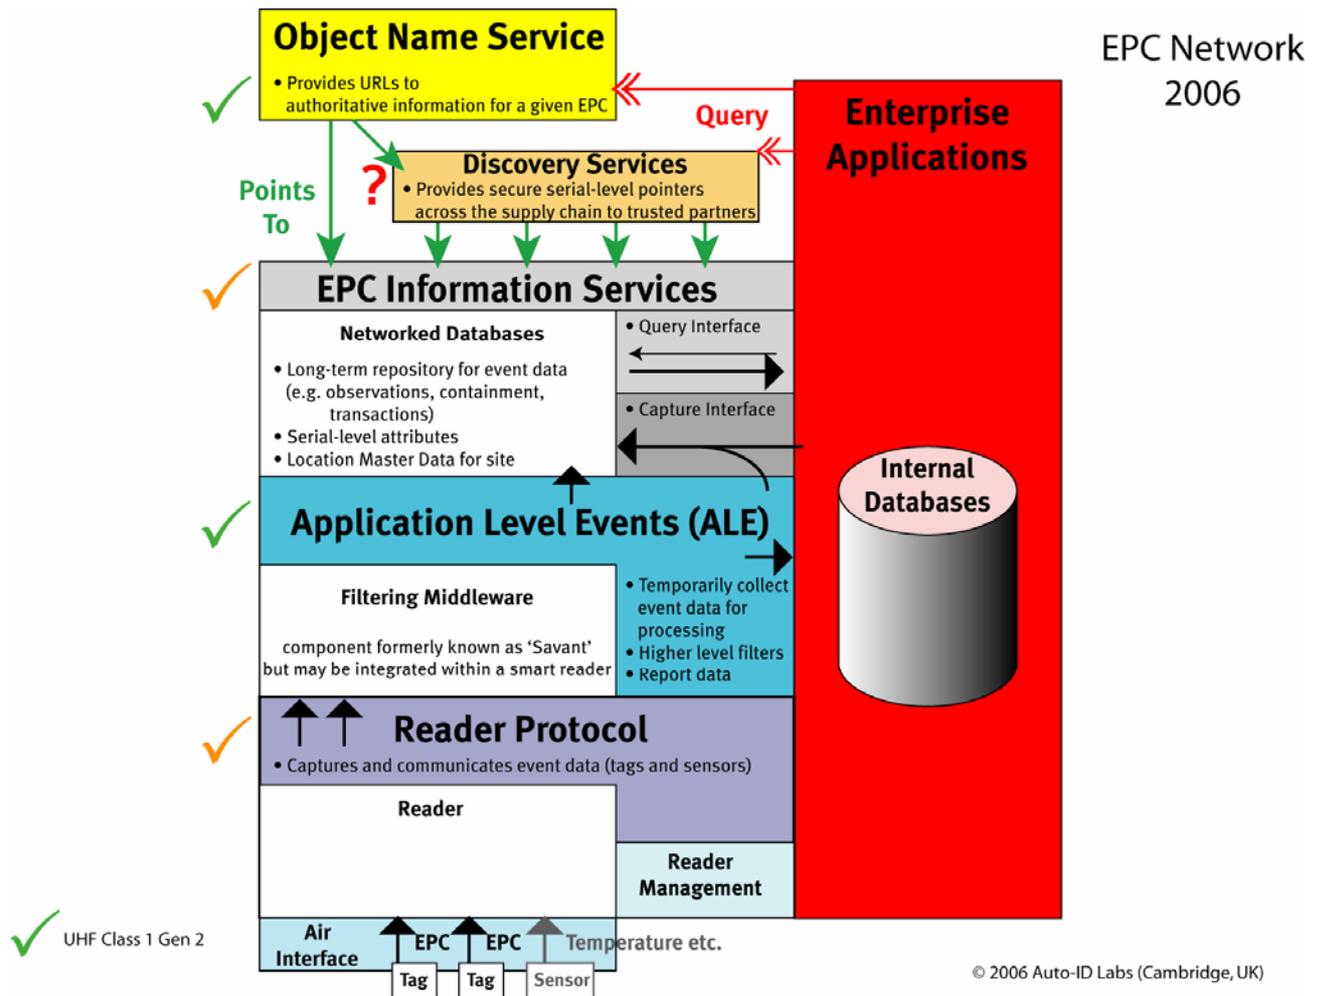


Figure 5 EPC Network architecture as of 2006 (EPCglobal Inc., 2006). Green ticks indicate ratified standards. Orange ticks indicate elements undergoing standardisation, for which ratified standards are expected in 2006. Red question marks indicate elements for which formal standards development work has not yet begun.

Some other issues include:

- 1) The technical architecture for Networked RFID can always be improved upon, as experience with prototypes and real world business use cases is fed back into the design – in this sense, the architectural design is initially an iterative process, where the grand visions of academics and technology providers is checked by end users from industry.
- 2) It is important to consider both forwards compatibility and backwards compatibility in the design of standards – (e.g. so that data from existing versions can be understood by implementations of future versions of the standard (new equipment can understand data from old equipment) – and also so that core features of data from implementations of future versions of the standard can nevertheless be understood by existing equipment (current equipment can understand the core data from new equipment, even if it does not understand the extensions which the newer version of the standard introduced).
- 3) It is important that the standards provide a mechanism for extensions – either for future versions of the standard – or for extensions by technology vendors or end-users.
- 4) There is greater value in standardizing interfaces between technology components than in prescribing the implementation of those components themselves. This allows the technology developers greater freedom to innovate and improve the efficiency of their

solutions, while maintaining a standardized interface for interoperability between technology from multiple vendors

- 5) Certification for compliance with standards is important for both technology vendors (who wish to have their products certified with global standards) and also for end users, who are seeking a multi-vendor competitively priced market, in which certification of compliance with the relevant standards is an assurance that products from different vendors should be compatible with each other.
- 6) Interfacing with existing systems (whether information systems such as ERP, WMS systems or identifier schemes (such as existing barcode formats)) must not be overlooked. End user companies have made large investments in their existing systems – and it may be very difficult to displace these and introduce a whole new way of identifying objects or a whole new way of exchanging information.
- 7) Standards should be designed to be neutral regarding operating system or platform (e.g. UNIX, Linux, Windows, Mac OS), programming languages (e.g. Java, C, C#, etc.) and should support one or more appropriate open transport bindings.
- 8) When standards are created as a result of publication of an open API by a vendor, this is preferable to a closed system of proprietary solutions and can be useful for driving initial adoption but may not lead to a standard that is sufficiently extensible or useful in the long term. For example, UHF Class 1 Generation 1 was an Auto-ID Center specification resulting from the publication of the RFID protocol developed by Alien Technology Inc. for user-programmable tags. Following this, their competitor, Matrics (now acquired by Symbol Technology Corp.) published their protocol for factory-programmed UHF tags as an Auto-ID Center specification, UHF Class 0. Both have since been deprecated by EPCglobal, replacing these with the UHF Class 1 Generation 2 standard as a result of a consensus standards development process, resulting in not only better performance and security features, but also serving as a single foundation for future UHF RFID protocols for more sophisticated EPC tags (e.g. with built-in sensors, additional user memory, encryption)
- 9) The existence of two alternative standards in a particular level of the architecture can be very problematic since it leads to confusion in the marketplace about which products to buy, especially in an open-loop environment. For example, the co-existence of UHF Class 0 and Class 1 Gen 1 tags and readers was problematic because it was unclear to manufacturers which retailers would be using Class 0 readers and which would be using Class 1 readers – so for a while, they were faced with the prospect of “double-tagging” to ensure that their products could be read, whatever the choice of reader. This added to the cost and complexity of tagging – and was one of the main drivers for a single UHF standard for EPC tags and readers.
- 10) There is benefit in developing a layered architecture or using “encapsulation” to allow end users to change low-level implementation configurations without needing to reconfigure higher-level applications and control systems that act upon the data. For example, the end-user may need to swap one or more physical RFID readers for another one (or a different combination of readers) to monitor a particular location, without having to reconfigure all of the applications which make use of RFID data from that location. The EPC Network uses indirection or encapsulation at multiple levels to provide the end users with flexibility to make changes to the hardware configuration without having to worry about the details being hard-coded into the applications. Some examples of encapsulation and indirection follow:

- a) The URN representations of an EPC insulates the information systems and filtering middleware from caring whether the unique identifier was read using RFID or barcode equipment or human input, as well as specifics of which barcode symbology or RFID frequency band or protocol.
 - b) The Reader Protocol standard insulates filtering middleware from specifics of communicating with different makes of RFID readers from different vendors, since it provides a standardized software interface for low-level communications with readers to enable reading and writing of tags.
 - c) The Application Level Events (ALE) protocol uses “logical readers” to insulate applications from hard-coding IDs of physical readers – applications instead request filtered data from logical readers which may consist of one or more physical readers. The association of particular physical reader IDs to particular logical readers is outside the scope of ALE but is implemented in some manner by filtering middleware
 - d) The EPC Information Services (EPCIS) protocol insulates applications from needing to know the details of the underlying database (whether XML, relational, object-oriented, its native query language (e.g. SQL, or XQuery) its schema (structure of tables, columns, fields or document schema), database drivers, database names, connection strings and addresses – or whether the data is even stored in-house or hosted at an external data hosting company. All of these implementation details are hidden from clients of EPC Information Services.
 - e) The Object Name Service (ONS) provides a level of indirection between the unique identifier (EPC) and the URL of the corresponding information services. It therefore allows manufacturers to change the address of their information services at any time, without impacting on IDs of tagged objects already in circulation (compare with DIALOG system or ID@URI in the case where the URI is a URL rather than a URN). It also provides applications with flexibility to choose the appropriate type of information service, since multiple service types may be listed in the ONS record for a given EPC
 - f) Discovery Services will in the future provide a further level of indirection about where lifecycle data is stored, by providing a massively scaleable one-to-many resolution service to link to data sources fragmented across an entire supply chain which may be unique to each individual object. This avoids the need to centralize all the information about an object in one location (politically unacceptable for most trading partners) and avoids the situation where either all of the lifecycle data is concentrated with either the manufacturer (since they can usually be identified from the ONS lookup of the EPC) or with the furthestmost downstream supply chain parties (as in the case of pedigree documents which accompany the object downstream)
- 11) One major feature of the Auto-ID Centre approach was the drive for rapid adoption of low-cost RFID technology, especially as many of the project’s sponsors were frustrated with the cost of RFID equipment and the length of time which formal standardisation approaches (e.g. ISO) were taking to develop their standards. Although the Auto-ID Centre identified a need for standards across the entire networked RFID architecture and did much to raise the profile of the vision of networked RFID, there may have been a tendency to give greater priority to expedience over rigour. For various reasons, almost all of the Auto-ID Centre specifications have since become obsolete (or substantially modified) either because:

- a) they were not developed in a true consensus development process but were simply the result of a very small number of vendors publishing their protocol or API openly (e.g. UHF Class 0, UHF Class 1 Gen 1) resulting in
 - b) a proliferation of non-compatible open APIs from competing vendors operating in the same domain caused confusion in the marketplace and threatened to increase costs and complexity for end users
 - c) they did not integrate with existing industry practices (the EPC identifier was originally a new number which did not attempt to embed the barcode identifiers such as GTIN-14 (UPC or EAN-13), SSCC-18 etc.)
 - d) they did not fully consider integration with existing information systems (Product Markup Language, PML Server)
 - e) they were too complex in design (Object Name Service, specifically the pre-processing of the binary EPC into the ONS lookup)
 - f) identifier formats were not sufficiently abstracted from the device used to carry the identifier (EPCglobal now recommends the communication of a URN representation of EPC between information systems, rather than a raw binary or hexadecimal string)
 - g) they lacked features which industry found useful (e.g. logistic filter bits in the EPC to indicate packaging level, for ease of discriminating between pallets, cases, items etc. – particularly important for more efficient reading of just the containers)
 - h) they prescribed implementation of components rather than standardizing interfaces (Savant filtering middleware – replaced by Reader Protocol and Application Level Events (ALE))
 - i) the specifications were too aligned with one particular vendor (a reference implementation of the Savant filtering middleware was originally commissioned by the Auto-ID Centre from OATSystems, a spin-off from MIT)
 - j) they did not sufficiently embrace existing or emerging internet standards or software design principles (Service Oriented Architectures) nor support existing industry standards (such as GTIN, AS2 EDI messaging)
- 12) Wherever possible, the standards should make use of existing state-of-the-art standards (e.g. for various high-level internet protocols such as web services), although it should be noted that there may be more than one competing standard in each domain of the architecture, developed by different standards organizations or different industry consortia.

It is important to have enough technology vendors who compete with each other in the same domain to be able to firstly identify the potential rival standards, monitor convergence of these competing standards, examine compatibility and overlap and the potential for a piece of technology to support two or more competing standards in situations where neither is overwhelmingly dominant over the other. Having two few competing vendors for each aspect of the architecture invariably results in less than optimal standards that are too closely aligned with one particular vendor's existing commercial solution and may not foster sufficient competition in the marketplace.

7 References

(Note that some of the on-line references are to the PROMISE e-Room, which requires a user name and password.)

Anderl, R., A. Bonetti (2005) *Interim technical assessment report*, IST-507100 PROMISE, 17 June, Brussels, Belgium, https://project.sintef.no/eRoom/indman/PROMISE/0_234aa

Anderl, R., A. Bonetti, F. Jammes, R. Kool (2005) *1st Review Report*, IST-507100 PROMISE, 13-14 December, Brussels, Belgium, https://project.sintef.no/eRoom/indman/PROMISE/0_2a7ba

Bradner, S. (1996). *The Internet standards process – revision 3*. IETF RFC 2026, Available online <http://rfc.sunsite.dk/rfc/rfc2026.html>

EPCglobal Inc. (2006). *The EPCglobal standards development process (version v1.2)*. Available online http://www.epcglobalinc.org/standards_technology/BoG%20Approved%20Mar%208th%202006.pdf

Gaynor, M., and S. Bradner (2001). *The Real Options Approach to Standardization*. Proc. Hawaii Intl. Conf. on Systems Sciences. IEEE.

Krechmer (2006). *Open standards requirements*. The International Journal of IT Standards and Standardisation Research, Vol. 4 No. 1, January - June.

PROMISE (2005). *Annex 1: Description of work – revision 4.0*. IST-507100 PROMISE, https://project.sintef.no/eRoom/indman/PROMISE/0_1b158

Roberti, M. (2005) *Intermec withdraws IP licensing plan*. RFID Journal. <http://www.rfidjournal.com/article/articleview/1387>

Stromberg G., D. Barasic (2005) *PROMISE standardisation strategy*, IST-507100 PROMISE, 14 November, https://project.sintef.no/eRoom/indman/PROMISE/0_1b178

Stromberg G., D. Barasic, G. Hackenbroich, M. Neugebauer, J. Anke, J. Brusey, D. Potter, A. Edler, A. Metin (2005) *DII.2: PROMISE standardisation domains*, IST-507100 PROMISE, 31 October, https://project.sintef.no/eRoom/indman/PROMISE/0_2951f

West, J. (2004). *What are Open Standards? Implications for Adoption, Competition and Policy*. Presented at the Standards and Public Policy conference. Chicago, Illinois.

Wikipedia contributors (2006a). *Standardization*. Wikipedia, The Free Encyclopedia. Retrieved 16:18, April 30, 2006 from

http://en.wikipedia.org/w/index.php?title=Standard_Generalized_Markup_Language&oldid=46686172

Wikipedia contributors (2006b). *Standard Generalized Markup Language*. Wikipedia, The Free Encyclopedia. Retrieved 18:27, April 30, 2006 from

<http://en.wikipedia.org/w/index.php?title=Standardization&oldid=50271117>