I-TEX: A Marketplace for Information Technology Standards in the 21st Century

BAA #99-07
Technical Topic Area: Information Technology Standards

Technical PoCs:

Dr. Kevin L. Mills
NIST
100 Bureau Drive Stop 8920
Gaithersburg, Maryland 20899-8920
Tel. 301.975.3618
Fax 301.590.0932
Email kmills@nist.gov

Dr. Peter Lucas
MAYA Design Group, Inc
2100 Wharton Street
Suite 702
Pittsburgh PA 51203
Tel. 412.488.2900
Fax 412.488.2940
Email lucas@maya.com

Dr. Steven S. Wildman
Northwestern University
Email s-wildman@nwu.edu

Dr. Sri Kumar
NIST
Email sri.kumar@nist.gov

Administrative PoC:

Mrs. Pam Davis
NIST
100 Bureau Drive Stop 8900
Gaithersburg, Maryland 20899-8900
Tel. 301.975.2735
Fax 301.590.0932
Email pam.davis@nist.gov

Mr. Ed Condrick
MAYA Design Group, Inc
2100 Wharton Street
Suite 702
Pittsburgh PA 15203
Tel. 412.488.2900
Fax 412.488.2940
Email condirck@maya.com

A. Innovative Claims

Successful technical standards for information technology provide a platform for economic growth and for an increasing array of products and services that improve the productivity of our workforce and the quality of life for our citizens. Evidence abounds: Internet standards, specifications from the W3C, de facto APIs from Microsoft, Java technology licensed from Sun, have all fostered related technical advances, improved productivity, fueled economic growth, and changed our lives. Yet, the marketplace that sets such technical standards suffers from extremely low efficiency in the current era of fast paced innovation, and likely costs our society many missed opportunities for even greater benefits from technical standards. For example, the machinery of standards-setting organizations grinds very slowly, and produces a diverse set of competing, often incompatible, standards that must be distilled further by market forces. Even standards coming from industry consortia suffer from bloat and poor focus, a result of the inevitable "compromise-by-superset" mentality needed to achieve consensus within consortia. De facto standards setting, ala Microsoft, allows numerous innovative small companies to be squeezed from the market over time as their ideas are absorbed into a growing set of software shipped from one source. With increasing frequency, licensed technologies, such as Java, appear headed for standard setting through adjudication in the courts.

A review of today's information technology marketplace suggests that the ponderous and inefficient processes currently in play could be counterbalanced by a complementary approach: standards setting by individual contributors. Some evidence exists that such an approach could succeed. For example, the open source movement (http://www.opensource.org) espouses the view that "when programmers on the Internet can read, redistribute, and modify the source for a piece of software, it evolves. People improve it, people adapt it, [and] people fix bugs. And this can happen at a speed that…seems astonishing." Perhaps you recognize some products that adhere to this vision: Linux, Apache, Perl, Tcl/Tk, Alice, and Curl. In each of these cases, champions and sponsors of certain standards, or potential standards, attempt to harness the collective potential inherent in a widely distributed, but organizationally unconnected, group of experts who willingly apply their expertise to develop and share source. Such standards setting through source code distribution appears akin to shareware on steroids. This leads us to our vision of a marketplace for information technology standards in the 21st century.

We **envision a highly efficient marketplace for technical standards**, where suppliers, whether individual contributors, consortia, or companies, electronically publish technical material, whether software, firmware, protocols, APIs, or VHDL specifications, through a distributed electronic exchange, and where customers, whether developers, value-added providers, system integrators, or end users, can locate, assess, exploit, and enrich material on the exchange. On the exchange, which we call I-TEX, standards setting will progress at Internet speed. The best technology will rise to the top, the weak and irrelevant will whither away. Technology from government-funded research transitions directly from the laboratory to the exchange, and technology from the exchange feeds further research. The exchange gives unrestricted market access to millions of individual contributors who would otherwise have little chance for influence.

While several diverse and uncoordinated efforts (e.g., open source, shareware, corporate web sites for product support) inch ahead, three serious impediments block the path toward our vision. First, suppliers provide diverse, often non-comparable, information describing the salient characteristics of their offerings. This diversity impedes potential customers from finding, comparing, and assessing available technology. Second, potential customers must judge the credibility of technical information and claims. Often, such judgments rely solely on the perceived reputation of the supplier. How can the credibility of information from millions of individual contributors be judged effectively? Third, with millions of contributors in the marketplace, a wide range of incentives must be accommodated. Profit motives must be protected, but other motives, such as fame, fun, and altruism, must also be enabled. Mechanisms must protect and compensate suppliers appropriately for their intellectual property.

In our expedition, we will create the technical underpinnings for a standards marketplace by constructing a web-accessible trading system for XML objects, for Java Beans or Jini Services, and for FPGA designs. At the same time, we will use the technology we develop as a laboratory to investigate institutional changes that might be required to support a trading system for standards. Building on technology produced for the DARPA Information Management program, we will: (1) determine the technical information needed by the marketplace, (2) provide mechanisms for finding, assessing, and enriching market information, and (3) incorporate innovative mechanisms for compensation. Importantly, we will construct the prototype so that others can easily extend it. We envision convening three teams of contributors, one for each technology (XML objects, Java Beans or Jini Services, and FPGA designs). Each

team will consist of contributors from a university, a large company, a small company, and a government research laboratory. Initially, we will use the team members as customers who will help us identify essential attributes for our trading system. At the end of the first year, we will deliver the trading system to DARPA, and we will make it available to the participating contributors and to the public.

We envision optional tasks to experiment with novel mechanisms for information enrichment, for compensation, and for assessment of technical specifications. We also envision optional tasks to study the effectiveness and efficiency of the trading system, as compared against other available market mechanisms. A final optional task addresses processes for extending the demonstrated framework to an arbitrary space of standards domains.

B. Technical Rationale

We believe that standards for information technology provide the key to future economic growth and to improved productivity. With complex products and technologies, the coordination among many and diverse parties required to effectively adapt these technologies to economic and social purposes is at best cumbersome and costly without well-established and reliable standards to facilitate the process. In the mainframe era of single-source monolithic systems, standards were dictated top-down by the systems developers, and there was little need for public involvement in the process. In the current era of distributed, networked PCs, it is increasingly clear that a lack of adequate standards retards technical progress. Looking forward to a future of vastly diverse, radically distributed information devices of all description, it seems clear that the emergence of high-quality, broadly accepted standards will be *the* critical issue in sustaining progress.

The legacy machinery for setting standards exhibits substantial inefficiencies. Because committee-based standards-creation efforts are often fraught with logistical and political difficulties, they tend to be monumental, time-consuming efforts. Moreover, precisely because they *are* so difficult and expensive, such efforts tend to be viewed as once-and-for-all "moon shot" feats--taking place in committee rooms far from the realities of day-to-day engineering practice. The resulting standards prove much too broad and inclusive to permit incremental, iterative refinement, driven by feedback from the field. As a substitute for the focussing effects of real engineering efforts, standards committees attempt to anticipate all possible needs, leading to overly-broad, massively complex standards documents that often fail to influence the market for information technology products. These inefficient standards setting processes lead to missed opportunities. These missed opportunities can be avoided when dominant players impose de facto standards; however, in such cases the effective result tends to inhibit beneficial innovation.

We believe that the best antidote to these dilemmas is to enable and encourage an entrepreneurial standards process. A vast and growing pool of individual contributors exists. If the efforts of those contributors could be harnessed, then increased efficiencies and increased opportunities might be achieved. One of the most significant but least-discussed aspects of the evolution of the ARPAnet, and later the Internet, was the early application of the new technology to its own development in the form of the RFC (Request for Comment) mechanism. The RFC mechanism allowed the initially tiny community of network designers to employ the new medium as an open, peer-reviewed forum for airing, critiquing and consensus building. This mechanism led in a startlingly efficient manner to the creation and successful deployment of complex and novel technologies. The challenge is to capture the spirit of this mechanism in ways that can support much larger and more diverse communities of practice dealing with a wider range of technical and policy issues.

Current efforts to harness the pool of individual talent appear to be diverse and uncoordinated, leading to inefficiencies in the market. Aspects of the situation that lead to these inefficiencies include the following:

- The results of formal standards efforts are often so ponderous and complex that, as a practical matter, they are inaccessible to practicing engineers who can often not even afford the time to read and decipher the standards documents, much less adopt them.
- The less formal "grass roots" standards efforts are chaotic, of uneven quality, often lacking the editorial mechanisms necessary to produce high-quality documentation, and often inaccessible due to the sheer diversity of activity. The effort involved in simply locating a relevant candidate for reuse often exceeds the effort needed to reproduce the functional equivalent in an ad hoc manner.
- The financial and legal mechanisms necessary to adequately reward the extra effort and cost to offer up even a completed piece of work for broader use are often lacking and not well

understood. As a result, much work is kept proprietary even when broader usage would be mutually beneficial to all parties.

We **envision a highly efficient marketplace for technical standards**, where suppliers, whether individual contributors, consortia, or companies, electronically publish technical material, whether software, firmware, protocols, APIs, or VHDL specifications, through a distributed electronic exchange, and where customers, whether developers, value-added providers, system integrators, or end users, can locate, assess, exploit, and enrich the material on the exchange. This marketplace will not merely support the reuse of particular bits of technology (although it will certainly support this), but it will be designed to encourage the rapid evolution of consensus standards of practice and of technical specification. This is not to say that the other mechanisms currently employed in setting standards will necessarily wither away. Economic transactions in other goods and services are typically coordinated by a variety of institutional arrangements, with the forms employed determined by a process of natural selection. These market exchanges depend on a complex of institutional arrangements for which analogs have yet to develop for the setting of information technology standards. Consequently, we are currently denied the potential advantages of markets in standards that could rapidly sort and evaluate alternatives, as well as speed their evolutionary development.

The halting steps toward developing markets in shareware and open standards, described briefly earlier, suggest that there is a vast reservoir of untapped talent, insight and energy that, with better institutional arrangements could dramatically accelerate the evolution of information technologies and the societal benefits to which they contribute. Tapping into this reservoir will require appropriate regulatory, legal, and technical infrastructures to support vigorous and more spontaneous markets in standards development and application. We think the real barriers here are at least as much institutional as technical. What we seek is a set of institutional arrangements that will improve our ability to exploit the reservoir of talent that now is barely being tapped because of the inefficiency of current arrangements. We propose to build a prototype of the technical underpinnings that might support the needed institutional arrangements, and then to investigate and evaluate the most promising directions for changes to institutional arrangements regarding legal requirements for information disclosure and for protection of intellectual property.

C. Technical Approach

Assumptions

Our vision will build upon several assumptions not true today. First, we assume a globally scalable system of distributed, uniquely identifiable, replicable objects. Second, we assume a public infrastructure for authentication and encryption. Third, we assume an accepted system for electronic payments.

Base Expedition

Our base expedition consists of five tasks, as described below. The first four tasks involve design of architectural elements of I-TEX, a trading exchange for information technology standards, while the fifth task consists of implementing an I-TEX prototype using Jini, new Java-based software for building distributed systems. Other elements of our proposed expedition are described as options that use this prototype exchange as a laboratory to investigate specific issues surrounding the institutional changes needed for our exchange model to succeed ultimately.

Task 1 – Information Architecture for Describing Technical Standards

We propose to define an information architecture, consisting of entities and relationships that define the meta-data required to support a marketplace in technical standards. Before we can design systems to support the creation of standards, we must first introduce an exact, operational definition of "standard". We shall adopt the following definition: "*A standard is a specification sufficient to create a collection of components that can interoperate.*" The terms "component" and "interoperate" used in this definition are necessarily very general at this level of analysis. In large part, the process of creating meta-standards that define a "standards space" will be one of providing domain-specific definitions of these terms. For example, in the space of Java beans, "component" would be equated with "bean", and "interoperate" might be something like "demonstrate the ability to exchange meaningful information with another bean conforming to this standard." Note that the Java Bean standard itself meets this requirement in one sense,

but that additional layers of standards (with more rigorous definitions of "meaningful") are necessary in order to support the emergence of useful ensembles of fully composable beans.

Once low-level interoperability is established, the next taxonomic step will involve a careful enumeration and documentation of the precise transactions that each component will support vis-à-vis its peers within the context of the standard. Thus, type systems, message formats, and specific protocols must be enumerated. The appropriate taxonomy for characterizing such properties of components will be a major topic of our research.

In this manner, we will build up an information architecture that ensures that each proposed standard is well formed and well defined. Requirements will include validation suites sufficient to operationally define interoperability; pointers making explicit relationships with superordinate and other related standards; and where possible, quantitative and qualitative metrics that can be used to order standards along dimensions of the "standards space."

Task 2 – Assessment Architecture for Evaluating Technical Specifications and Components

We propose to define mechanisms that enable potential customers to evaluate the characteristics of standards components available on the exchange. These mechanisms could include tests and test harnesses that can be executed, simulation models, as well as links to executable demonstrations.

The design of these mechanisms will emphasize contributor-driven evaluation criteria. Our design will also emphasize the fact that no new standard is "free." Even ignoring the cost and effort to create and document a new standard, each standard--like each new feature in a user interface--bears a cost of comprehension and complexity. That is, each new addition to a standards space reduces the effective accessibility and value of existing standards. With standards, it is not "the more the merrier" -- quite the opposite is true.

In light of this, it should be the burden of the contributor of a new standard to motivate the creation by objectively demonstrating its value. Thus, for example, a contributor might be required to propose some quantitative measures to empirically demonstrate the superiority of the standard over existing competitors, against at least one metric. Or else the contributor might define a new metric against which the proposed standard is superior. This process is one means by which the standards space will evolve.

Task 3 – Enrichment Architecture for Enhancing Technical Standards

We propose to define mechanisms for potential customers to annotate the available information in the exchange with their own assessments, to add elements to the assessment architecture, and to extend the information architecture itself. Recall that the primary goal of this proposal is to apply market forces to the evaluation of technical standards. Efficient communication and transparency are the linchpins of efficient markets, so our goal will be to provide a structured but open forum for comparing competing standards. This forum will allow consumers to share the results of their evaluations, and will enable producers to defend their contributions against misleading comparisons. Although the latter activity is self-motivated, consumer contributions are somewhat dependent on altruism. For this reason, we must keep the costs of such contributions as low as possible.

Task 4 – Protection and Compensation Architecture

We propose to define mechanisms for protection and compensation appropriate for the class of supplier. We will provide mechanisms for paid compensation per copy, for postcard-ware, and for exchange-wide publication of pseudo-value (akin to stock exchange valuation). We will explore a credit/debit scheme built into the exchange (like "money" in a MUD). We will consider schemes such as super distribution, which charges by usage rather than by copying.

Task 5 – Prototype Trading Exchange

We propose to implement the trading exchange for three classes of standard components: XML objects, Java Beans or Jini Services, and FPGA designs. We also plan to distribute the software publicly. Initiating the use of the trading system is covered under Option 1, Trading Expedition.

Option 1 – Trading Expedition

We propose to engage one or more communities implementing Java Beans or Jini Services, XML objects, or FPGA designs in pilot studies using our prototype exchange, and to evaluate the effectiveness and efficiency of the approach. We also propose to publicize the availability of the exchange system in an effort

to encourage others to adopt it for their technical standards or component developments. Using the three pilot studies we propose to investigate issues related to the design of efficient standard exchanges, understanding the similarities and differences, and how lessons learned in one context can be useful in another. The issues to be investigated will include specific information demands of participants and the credibility of the system.

There are at least three aspects to ensuring credibility. First, who is putting forth information and for what purpose? In this context, the variety of disclosures needed for specific purposes is an issue. Second, who owns the data in different parts of the system? For this pilot study, NIST can serve as a "neutral" party, and can administer the exchange from a central site with necessary security and authentication. However, when the exchange scales to millions of users and to a large product space, the system must be distributed. Through the pilot studies we propose to investigate and document the issues involved in credibly administering such a large-scale, distributed system. Third, can appropriate safeguards be implemented to avoid misuse? The generation of spurious and unnecessary information can increase processing costs, detract from process, create frustration, and reduce system credibility. Mechanisms for self-regulation may be context dependent; we expect our study to reveal useful insights.

Option 2 - Enrichment Expedition

We propose to design, implement, and evaluate some alternate mechanisms for enriching the information and assessment mechanisms available on the exchange. Given a candidate standard, consumers, in addition to product quality and attributes that can be verified, would generally like information that leads to a good assessment of value. Factors that influence this assessment include interoperability with other similar products, compatibility with complementary technologies and products that go with the candidate, and the candidate's ability to improve the benefits in a networked environment (network externality). Favorable experience of early adopters, along with information on future prospects of a better technology, can also reduce uncertainty and consumer fear of becoming stranded. Additionally, consumers may find it useful to know how the market is likely to evolve, as well as to comprehend incumbent products and market structure. Prospects of future competition will reduce fear of getting locked-in to a single vendor. Consumers may also wonder whether competition will be on features or price, and whether the trends are towards a bundled system or towards competition in components that permits mix and match.

Suppliers, on the other hand, would desire information about other potential suppliers, intellectual property restrictions, and patent portfolios, possible blocking standards and coalitions and partnering among participants. They would also like to know about consumer response and how to stimulate consumer interest.

These anticipated concerns of suppliers and consumers outlined above are broad, and what is feasible depends on the specific context. We propose to use these broad concerns as the initial prompting factors for the generation of shared information and for exchange enrichment. Later evolution of the enrichment will be driven by the concerns of the participants.

Option 3 – Protection and Compensation Expedition

The core of intellectual property law evolved in an era where print provided the primary form of information technology. While intellectual property law has evolved somewhat in response to new technologies (e.g., the application of the fair use doctrine to photocopies of copyrighted books and articles), more significant changes might be required to better encourage and enable innovation in standards and applications for modern information technologies. The exact nature of the required changes remains in dispute. We will examine whether the current intellectual property rights regime provides protections sufficient to encourage the vigorous trading in standards and applications envisioned for I-TEX. Specifically, we will study the trading that occurs in our I-TEX prototype, and we will conduct interviews with trading participants. We seek to determine the form of intellectual property protection sufficient to encourage participation in the exchange without undue risk that ideas shared will be misappropriated. We also seek a protective regime that is not so cumbersome as to significantly retard the flow of ideas and the adoption of innovations.

Option 4 – Assessment Expedition

We plan to assess benefits from I-TEX for individual participants and for society. Although assessment via controlled experimentation may not be feasible, we can still assess individual benefits by following, quantifying, and collecting data on user benefits and costs over time and by studying the determinants of

their returns through statistical analyses. If all users have positive benefits, then clearly there are enhanced social benefits. In case user benefits differ, we will study the determinants of difference, and study potential ways of aggregating individual benefits to arrive at overall social returns.

Option 5 – Extensibility Expedition

We propose to design and implement a mechanism of templates to extend the three initial standards spaces to encompass additional content areas and types of standards. We anticipate that this process will proceed both by the establishment of new content areas (e.g., communications protocols and device interconnection standards), and by the successive differentiation of existing categories (e.g., graphical Java Beans and communications Beans). We envision the emergence of a growing mosaic of "standards spaces" within which developers and users can navigate in order to locate the appropriate venue for the technology they wish to contribute or locate. If this space is to remain tractable, the system will eventually require the emergence of a set of "meta-standards" in order to impose some order onto what will eventually evolve into a large and diverse space of standards categories. In addition to designing mechanisms for contributors to contribute to the fleshing out of this growing mosaic of meta-standards, we will investigate techniques for visualizing and efficiently navigating the standards space.

References

1. Shapiro, Carl and Hal Varian. Information Rules. HBS Press, 1999.
2. "Collected Papers on Knowledge Management", *Harvard Business Review*, HBS Press, 1998.
3. Jim Waldo, Jini Architecture Overview, Sun Microsystems, 1998. http://java.sun.com/products/jini/whitepapers/architectureoverview.pdf
4. Andy Quinn, JavaBeans, Sun, 1998. http://java.sun.com/docs/books/tutorial/javabeans/
5. Tim Bray, Jean Paoli, and C. M. Sperberg-McQueen, The Annotated XML Specification, W3C, 1998. http://www.xml.com/axml/axml.html
6. Field Programmable Gate Arrays: An Enabling Technology, Virtual Computer Corporation, 1998. http://www.vcc.com/fpga.html
7. The Handle System, Corporation for National Research Initiatives, 1998. http://www.handle.net/
8. William Y. Alms, Christhope Blanchi, and Edward A. Overly, An Architecture for Information in Digital Libraries, D-Lib Magazine, Feb. 1997. http://www.dlib.org/dlib/february97/cnri/02arms1.html
9. The CNRI Registry, preliminary information, Corporation for National Research Initiatives, August 1997. http://www.cnri.reston.va.us/registry.html
10. "Rights Protection", Digital Object Infrastructure, Corporation for National Research Initiatives, 1997. http://www.cnri.reston.va.us/home/dobi.html#kss
11. "Classification and Indexing", Digital Object Infrastructure, Corporation for National Research Initiatives, 1997. http://www.cnri.reston.va.us/home/dobi.html#dobi
12. Jeremy Hylton, "Creating Collections with a Distributed Indexing Infrastructure", *W3C Distributed Searching and Indexing Workshop*, 1996. http://www.cnri.reston.va.us/home/dobi/collections.html