



Attorneys at Law

Consortium Standards Bulletin

A ConsortiumInfo.org publication

June 2005

Vol IV, No. 6

EDITORIAL

WEB OF DREAMS

Andrew Updegrove

Anyone who has ever had the pleasure of hearing Tim Berners-Lee speak knows that he is not the most patient of men. The words tumble out in a profusion of enthusiastic conviction, propelled by the type of urgency that comes from seeing the future, and wanting to make it exist *today*. It is thus ironic that the path that he is forced to walk to achieve that goal is one of the most tedious imaginable: standard setting. Not only is it a lengthy process, but one that is based upon consensus, requiring not just great labor, but a constant process of cajoling, inspiring and simply charging into the face of battle, challenging others to follow by sheer force of example.

If the rigor of the technical process was the only impediment to success, it would be challenge enough. But standards must not only be created to be worthwhile – they must be adopted and implemented as well. Thus, a second burdensome responsibility must be assumed by the champions of any new standard, which is to convince those for whom a standard has been created to not only accept the gift, but (figuratively speaking) to also prominently display it on the mantelpiece, rather than relegate it to the dark corners of the attic.

Never has this challenge been greater for Tim Berners-Lee than it is now, as he seeks to persuade the world that it is time to embrace his vision of the Semantic Web. To Berners-Lee, it is not so much a matter of convincing users to adopt a new type of Web, but of accepting a feature that was planned, but not included, in Web Version 1.0. But to much of the rest of the world, the Semantic Web is a difficult to understand abstraction, the utility of which is not universally grasped.

Perhaps worst of all, while those in industry generally accept the usefulness of the Semantic Web concept, they are not yet convinced that committing their resources to supporting it will provide greater economic gains than directing the same resources to other purposes.

For others, the question is not whether the Semantic Web is a good idea per se, but whether the technical framework conceived by Berners-Lee and his staff at the W3C is the right one. Among some, the debate centers over whether it is too comprehensive or too limited? Too inflexible or not rigid enough? Is it insensitive to cultural values? Is it an infeasible concept that is doomed to be ignored? Does it have an insoluble “chicken and egg” problem that can only be solved by browser developers buying in?

The views on the Semantic Web expressed today as the initial enabling standards work approaches completion seem almost as numerous as those that hold an opinion. And it must be confessed that in researching this issue with standards professionals in both the consortium as well as the corporate worlds, those that expressed pessimism for the eventual, pervasive implementation of the Semantic Web outnumbered the optimists.

The result is that Berners-Lee has found it necessary to spend an increasing amount of his time barnstorming the world to garner support for the Semantic Web. Given that the Semantic Web activity was launched in 2001 (a precursor project, the Metadata Activity, dates back to 1998), it would try the patience of a saint, much less someone with the energy level of Berners-Lee, to face the need to hit the stump after so many years of arduous conceptual and technical work. And yet, he is patient – reminding the world that this is how it happened before, with people scratching their heads and doubting, but ultimately “getting it” and climbing on board.

The difficult watershed of opinion that Berners-Lee straddles today may be summarized by two catch phrases that captured popular attention more than twenty years apart, each of which sought to convey a sense of the unimaginable.

The first was coined in the 1960's, when protestors of the Vietnam war challenged society to imagine a world where the concept of armed conflict had been rejected, and where the existence of war had therefore been rendered impossible. That smirking but memorable phrase was, "What if you gave a war, and nobody came?"

The second slogan was popularized by the 1989 hit movie, "Field of Dreams," in which the central character (Kevin Costner) is haunted by a voice, heard only by him, that intones, "If you build it, they will come." Against the disbelief of family, friends and strangers, Costner's character invests all he has in a seemingly romantic and hopeless quest, but nonetheless prevails and is victorious.

So which of these metaphors will apply to the labors of Berners-Lee, and the hundreds that have worked with him to create the standards they hope will be used to create the Semantic Web? Will it be, "What if you gave a Semantic Web, and nobody came?" Or, now that the core standards have been built, will a world of adopters play upon the field that has been prepared for them?

We believe that the latter will (eventually) prove to be the right metaphor, but for a reason that is common to both phrases.

Today, there are uncounted millions of technically adept, imaginative, and motivated individuals that are as comfortable with code, architectural concepts and new ideas as their grandparents were with home improvement projects, hand tools, and tried-and-true methodologies. Today's generation is eager to experiment with virtually every new information technology capability that is been made available to it, resulting in myriad, and often surprising results.

One need look no farther than phenomena such as open source software and music download file sharing to appreciate the potential for the viral, peer to peer distributed uptake of new concepts. And the same examples demonstrate that the multinational corporations that had no interest in these concepts at the outset can become enthusiastic adopters when the potential for profiting from the same innovations becomes demonstrable.

Today, the attitude of most major corporations towards the Semantic Web is more supportive in principal than evidenced by actual product plans. True, the W3C Web site lists multiple statements of support <www.w3.org/2004/01/sws-testimonial#twenty>, but the absence of some major corporations is as conspicuous as is the presence of those that are included.

Consider also the results of the following Google searches on three new areas of business opportunity, each of which is enabled by standards that are at roughly the same state of development vis-à-vis their suitability as a basis for productization:

	"Web Services"	"RFID"	"Semantic Web"
Google/News	2,180	2,130	62
Google/Web	63,400,000	14,000,000	5,380,000

Search performed on June 17, 2005

In the case of Web services, the search results reflect the fact that many of the biggest IT vendors believe that great returns on investment may result from the development and deployment of Web services-based products for the supply chain – more of a "Killer Application Platform" than a single "Killer App", but equally effective. In consequence, they are placing huge bets on Web services, resulting in significantly more attention being paid by both the technical as well as the business and financial press to the technology – over two thousand press releases and other stories in the preceding month.

The example of RFID is even more telling. Prior to the announcements by Wal-Mart and the United States Department of Defense that each would require RFID tagging from large numbers of vendors, the deployment of RFID technology was more in the "if" than the "when" category. Now, RFID is generating

news stories at a rate comparable to Web services, riding on the credibility that these major adoption notices bestowed. In the case of RFID, therefore, it was not a “Killer App” that turned the tide, but a “Killer Customer.” The announcements by Wal-Mart, and then the DoD, persuaded additional industry players that a critical mass of buyers and vendors would indeed enter the field. As a result, the prices of RFID tags would fall, and profits would eventually (if not immediately) reward those that decided to invest in providing RFID-based products and services.

Does this mean that the Semantic Web is doomed to die before it is born, or that announcements comparable to those that launched RFID as a credible technology investment must be obtained before progress can be made?

While many think that the answer to that question is “yes”, we disagree. We believe that the Semantic Web will come into being (albeit gradually), simply because it can.

One reason for this prediction is that while a “Killer App” is hugely useful to legitimate and drive adoption of a new technology, better mousetraps have been a big business from time immemorial. 28.8k modems did not enable more applications than their 14.4k predecessors, nor did nominally 56k modems create dramatic new product opportunities over 28.8k modems. And yet they were developed, and the quest for broader bandwidth still continues at a fevered pace. Even if the Googles and the Yahoos of today do not display current interest in Semantic searching, others will see the potential for agents that search the Web, for cataloguing existing data in databases for Semantic purposes in intranets, and so on.

Eventually, there will be competitors to Google and Yahoo, because the profit potential for search engine-based advertising is so great. When this happens, the attitudes of the then-dominant browsing service providers will doubtless change. And, given Google’s penchant for rolling out new, secretly developed Beta services, do we really know today that they are on the sidelines after all?

One of the wonders that the Internet has made possible is that everyone can participate in the process of conception, creation and enjoyment of new technology and the fruits of that technology (assuming that the technology is not encumbered, on which more below). As a result, it is no longer necessary for huge corporations to commit to a given technical approach in order for that approach to become widely implemented. RSS syndication is an excellent example of a useful technique that arose not from a corporation, or even from a standards organization, but which nonetheless has become pervasive. Nor is it now usually necessary for new types of infrastructure to be created at great cost, as historically often constrained innovation in the case of areas such as telecommunications.

Today, with over a billion connected individuals, any good idea, and any robust tool, can be taken up and utilized if it becomes known, and if there are no impediments to its use. Tim Berners-Lee and the other true believers at the W3C have taken care of that last element as well, with the hard-won adoption of a Patent Policy in May of 2003 that makes it as difficult as possible for any standard adopted by the W3C to run afoul of a blocking patent or the requirement to pay a licensing fee.

Already, the evidence is that grass-roots interest is building for the Semantic Web, just as occurred with the Web before it. Today, a wide variety of communities, both ad-hoc, academic, and in some cases sectoral (e.g., life sciences) have begun to explore the potential of Semantic techniques to address real problems and attain defined goals. As happened ten years ago, individuals, and individual communities, are realizing possibilities and opportunities that can be enabled using the standards, tutorials, white papers, and other materials that the W3C has generated. These activities will be organic, opportunistic, and rooted in achieving real goals.

And yes, there are examples of utilization of Semantic Web standards by some corporations as well, such as the Semantic Web tools included by IBM in its alphaWorks open source offerings, and the capabilities built by Adobe into its Creative Suite software, which will automatically include Semantic Web descriptions of all files that are created by Adobe applications.

But what of the critics that say that the Semantic Web structure is too rigid, or too flexible; too neutral and not cultural? Are they totally wrong?

Yes and no. Yes, in that the Semantic Web, just like the Web before it, is a best-guess charting of a voyage into the unknown. Those that opt to build out the next generation of the Web will go where they will, and the standards that enable the creation of that Web will evolve to match the places where these pioneers choose to travel, and be optimized to achieve the tasks that they undertake. As with all standards, there will be things that can be done, and things that (initially) can't be done. And then there will be refinements, and new standards, to enable doing those things, too.

But the important thing is that voyages into the Semantic Web have now been made possible, and that they have begun. Ultimately, as with the original Web before and open source more recently, the bigger players will come on board when it appears to be in their best interests to do so. There is nothing wrong, and perhaps even much to gain, by the Semantic Web becoming real in that order.

So perhaps the critics will prove to be right, in that the Semantic Web may look different in ten years than it was envisioned by Tim Berners-Lee ten years ago, or as it has been enabled by the work of the W3C today. We believe that the detractors will be proven wrong when they contend that the Semantic Web will not happen, simply because the design of the standards system is not to their liking, or because no Killer App has yet been announced. And perhaps they will be right in that the Semantic Web may not prove to be as explosively utilized as is HTML. But, then again, it does not have to be so pervasive in order to be incredibly useful.

Still, Berners-Lee has a last, long, arduous lap to go. He has already conceived and shared the vision for the Semantic Web, provided the leadership to see the enabling standards become real, and fought the battle to ensure that those standards may be utilized by anyone without cost or troubling licensing restrictions. Now he is putting his reputation on the line to provide the credibility that will make it easier for those who are working on Semantic Web tools, techniques and encoding to do their work, so that the rest of us may eventually come to believe in his vision as well.

We believe that we are living in a time of democratization of technology that will result in an explosion of innovation, and that may well see a reordering of the forces of technological evolution. Rather than living in a world regulated by "top down" vendor-driven decisions on products, architectures and services, a "bottom up," neural, global, real-time, self-calibrating and adjusting process of collaborative innovation will become pervasive, perhaps out-competing the capabilities of corporate research and development, and the power of corporations to mandate from the beginning what technical and architectural outcomes will succeed in the end.

Instead, corporations may become the opportunistic beneficiaries of this free research and development, with the most savvy and open-minded enjoying the most commercial success. After all – which bottom line would look better – one that is burdened with the costs of research and development, market research, and missionary selling, or one that is driven by free technology and the production of products that the market has already shown it wishes to buy?

So it is that we conclude that, now that Berners-Lee has led the process of building the enabling standards for the Semantic Web, "they will come". In fact, creative, open-minded and adventurous players are already taking their places on the field. We live in a world today where, when confronted with the potential for something as valuable as the Semantic Web to become possible, it makes as little sense to ask "could it be stopped?" as "will it happen?"

Comments? updegrove@consortiuminfo.org

Copyright 2005 Andrew Updegrove