Microformats

The Next (Small) Thing on the Semantic Web?

Rohit Khare • *CommerceNet*

"Designed for humans first and machines second, microformats are a set of simple, open data formats built upon existing and widely adopted standards."

— *Microformats.org*

When we speak of the "evolution of the Web," it might actually be more appropriate to speak of "intelligent design" — we can actually point to a living, breathing, and actively involved Creator of the Web. We can even consult Tim Berners-Lee's stated goals for the "promised land," dubbed the Semantic Web. Few presume we could reach those objectives by randomly hacking existing Web standards and hoping that "natural selection" by authors, software developers, and readers would ensure powerful enough abstractions for it.

Indeed, the elegant and painstakingly interlocked edifice of technologies, including RDF, XML, and query languages is now growing powerful enough to attack massive information challenges in disciplines such as bioinformatics. All the same, incremental, messy innovation continues to take hold in fits and starts within such narrower ecological niches of the Web as blogging.

A prime example of this phenomenon is *microformats*, a new approach to encoding semistructured information in ordinary XHTML. Clever application of several existing XHTML elements and its powerful class attribute system can make it easier to describe people, places, events, and other common types of semistructured information in human-readable form.

Microformats are better-adapted to the blogosphere for some seemingly minor reasons. In Fred Brooks' classic "No Silver Bullet" essay, he distinguishes between accidental and essential features of a problem. It might well be that the essential challenge in publishing a social network is precisely encoding the great variety of personal, professional, and genealogical relationships between people and organizations. By contrast, an accidental challenge is that any blogger with some knowledge of HTML can add microformat markup to a text-input form, but uploading an external file dedicated to machine-readable use remains forbiddingly complex with most blogging tools.

So, although any intelligent designer ought to be able to rely on the long-established facility of file transfer to publish the "right" model of a social network, the path of least resistance might favor adding one of a handful of fixed tags to an existing indirect form — the "blogroll" of hyperlinks to other people's sites.

Sure, the XHTML Friends Network (XFN) microformat might be a weaker abstraction than the RDF-based Friend-of-a-Friend (FOAF) format,[1] but choosing to merely "pave the cow path" of existing writing styles as simply as possible could actually lead to significant adoption, which is what really matters for standards of any stripe.

In this column, I'll take a more detailed look at some examples of microformats, the general principles by which they can be constructed, and how a community of users is forming around these seemingly ad hoc specifications to advance the cause of what some call an alternative to the Semantic Web — the "lowercase semantic web."

The Lowercase Semantic Web

Suppose you wanted to publicize an upcoming lecture. The existence of vCalendar[2] should be the end

of the debate on how to do so: it's a widely acclaimed international standard that thoroughly addresses such calendaring and scheduling concerns as time zones, recurrences, owners, and presenters. You'd simply link to a `myEvent.vcs` file that looked something like this example from Ryan King's primer on the topic[3]:

```
BEGIN:VCALENDAR
BEGIN:VEVENT
SUMMARY:Microformats: What the Hell Are They \
and Why Should I Care?
DTSTART:20050926T000000Z
LOCATION:Balder Room
DTEND:20050926T010000Z
DESCRIPTION:Ryan King will explain why microfor-
mats are important and how you can mark up specific
kinds of content in ways that make it easier for the right
people to find your stuff.
END:VEVENT
END:VCALENDAR
```

Given that this doesn't look very human-readable, you'd assume it was well-suited to machine-readability. Given that today's machines prefer angle brackets to colon-delimited header-value pairs, however, RDF Calendar[4] is a "better" alternative for use with the Semantic Web. For example, Masahide Kanzaki's RDFical-a-matic tool (www.kanzaki.com/docs/sw/rdfical-a-matic.html) would generate the following output for the same event listing:

```
<rdf:RDF
 xmlns:rdf='http://www.w3.org/
1999/02/22-rdf-syntax-ns#'
 xmlns='http://www.w3.org/2002/
12/cal/ical#'>
 <Vcalendar>
  <prodid>-//kanzaki.com//RDFCal 1.0//EN
  </prodid>
  <version>2.0</version>
  <method>PUBLISH</method>
  <component>
   <Vevent>
    <dtstart rdf:parseType='Resource'>
     <dateTime>2005-09-26T00:00:00Z
     </dateTime>
    </dtstart>
    <dtend rdf:parseType='Resource'>
     <dateTime>2005-09-26T01:00:00Z
     </dateTime>
    </dtend>
```

```
    <summary>Microformats: What the Hell Are They
and Why Should I Care?</summary>
    <description>Ryan King will explain why
microformats are important and how you can mark up
specific kinds of content in ways that make it easier for
the right people to find your stuff.</description>
    <location>Balder Room</location>
    <dtstamp>20051012T061505Z</dtstamp>
    <uid>1129097705622@kanzaki.com</uid>
   </Vevent>
  </component>
 </Vcalendar>
</rdf:RDF>
```

Both of these alternatives must be stored as external resources and hyperlinked into the original Web page from HTML anchor text such as:

```
<a href="/myEvent.vcs">
 <b>Microformats: What the Hell Are They and Why
Should I Care?</b>
 <p>Ryan King will explain why microformats are
important and how you can mark up specific kinds of
content in ways that make it easier for the right people
to find your stuff.</p>
 <small>September 25th, 2005, 5-6PM in the
     <i>Balder Room</i></small>
</a>
```

In modern usage, of course, this HTML is an abomination of inline formatting. Applying Cascading Stylesheets (CSS) is simpler, more flexible, and more accessible than relying on such relics of the "browser wars" as `<small>`. You could use an external stylesheet like the following to define each string's look and feel:

```
<div class="vcalendar vevent">
 <span class="summary">Microformats: What the
Hell Are They and Why Should I Care?</span>
 <p class="description">Ryan King will explain
why microformats are important and how you can mark
up specific kinds of content in ways that make it easier
for the right people to find your stuff.</p>
 <abbr class="dtstart" title=
"20050926T050000-0700">September 25th, 2005, 5
</abbr>– <abbr class="dtend" title=
"20050926T060000-0700">6PM</abbr> in the
<span class="location">Balder Room</span>
</div>
```

The payoff for choosing the class names I used in this example (shown in blue) is that they elimi-

nate the need for linking to an external representation in the first place. The inline style information is as sufficient as the other formats for encoding the same information — especially when combined with some of the lesser-used corners of the XHTML specification to "abbreviate" the machine-readable ISO-8601 timestamps[5] with human-readable phrases.

Unlocking XHTML's Power

"But a Web full of XML documents of arbitrary application — 'plain XML'? That future never happened."

—*David Janes (Blogmatrix.com)*

Despite this seemingly ad hoc process, there's actually a fairly principled transformation for encoding event metadata into XHTML. Let's look at how it works now, and return to the "why" in the next section.

When XML was new, CSS was scarce, and the browser wars raged, HTML was often cast as a hopeless muddle. The "Web of HTML" was poised to give way to a "Web of XML" in which each publisher used its own tags and presentation logic to empower a new generation of browsers. Today, users have access to fairly full-featured XML+XSL browsers on the desktop, but it's too late. Like Java, XML's niche turned out to be on the server side, out of sight.

In the meantime, HTML grew up and became a proper XML application, XHTML, offering all the rigor and modularization that an information architect could ask for. Similarly, CSS support in browsers, on PCs and handhelds, and in print matured to the point that authors and designers adopted it broadly. This was the key ecological change that triggered a resurgence of experimentation with "plain old" HTML.

If XML's essential strength — decentralized evolution of new tag sets — was also its essential weakness, then there's little to be gained by simply renaming the problem of Babel by encouraging random mutation of new CSS selector names (classes) instead. Technologically, XHTML class attributes do add a critical degree of freedom, insofar as they can accommodate multiple values in space-separated lists.

Socially, however, the key insight that the microformats community is taking advantage of is "appealing to authority" — stealing selector names outright from well-established standards rather than reinventing the wheel.

Rather than creating a new calendaring speci-

fication out of thin air, for example, the hCalendar microformat (www.microformats.org/wiki/hcalendar) reuses the names, objects, properties, values, types, hierarchies, and constraints from IETF's iCalendar[5] (itself a pared-down profile of vCalendar). It doesn't even insist on a clever prefix: it still uses the `vevent` field because that's the case-insensitive transliteration of the label in the original specification.

In the same way, spaces become dashes and plurals are reduced to singular instances (for example, `categories` in iCalendar becomes `category` in hCalendar). The latter rule works by expanding lists of values into multiple sibling elements in the Document Object Model (DOM). Similarly, hierarchical containment relationships in the originals are represented by nesting the corresponding microformatted XHTML elements.

Finally, we can abbreviate particularly ugly data using the `<abbr>` construct. This is useful for dates, for example, not merely because the ISO-8601 format is longer than the original text, but also because the constraints of machine readability can appear inconsistent: a conference ending on the 7th must actually be marked up with a `dtend` field on the 8th because it's an exclusive range delimiter.

Each of these rules elaborates one basic theme: use the most semantically appropriate XHTML element in the first place. The preceding examples included `<div>`s and `<span>`s, but those are actually the last resort. Better choices are to use existing list, dictionary, link, or quote constructs.

hCard How-To

Continuing in a practical vein, let's deconstruct the hCard microformat (www.microformats.org/wiki/hcard). Start with an IETF specification for using vCard in email.[7] Applying the rule of using the most appropriate XHTML element, the URL data field becomes a class on an anchor (`<a class="url" href="...">...</a>`); EMAIL becomes a `mailto:` link; and PHOTO becomes a class on images.

Some data fields can occur more than once, or have further internal structure. Singular keys such as a formatted name (`fn`) are resolved by using only the first matching descendant element; given that a person can have multiples of things like telephone numbers, however, every instance of a descendant element with class `tel` should be preserved, each with its own additional flags such as `home, work, fax,` or `pref` ("preferred").

Finally, we must evaluate the results of apply-

ing these transliteration rules for how well they balance human- and machine-readability. For example, the information about whether a phone line supports faxing or goes to the person's residence should be kept visible. Putting it in a class attribute would hide it from the reader, so hCard adds an extra layer of indirection using the class `type`, with any of the vCard list of telephone types inserted as visible text instead.

The astute reader with an eye toward localization might note that this overloads the use of a previously machine-readable flag. The four-letter string "home" might also be meaningful in English, but we're still waiting for real-world experience with the complications of, say, German: `<span class="type home">Haupttelefon</span>`.

Mapping the `tel` field to separate attributes forces the introduction of a `value` class to delimit where the actual digits begin:

```
<span class="tel">
  Our <span class="type">work</span> number:
  <span class="value">
  +1.415.555.WORK</span>
</span>
```

Breaking `tel` into a type and value keeps the qualifier "`work`" visible and separates it from the actual phone number.

Other elements once considered mandatory, such as `prodid`, `version`, and `source`, became less relevant and were dropped. Over time, hCard users also gained enough experience to suggest author-friendly optimizations, such as using the words in a formatted name to derive the implied `given-name` and `family-name` of the corresponding compound name property, `n`; preferring `organization-name` when an `organizational-unit` isn't mentioned; and assuming that an hCard represents corporate contact information when `fn` and `org` both appear on the same element.

Overall, transforming an existing schema into a microformat is a relatively straightforward process. Once we can express it in HTML, it becomes easier to adopt for authors comfortable with HTML and, even more so with the advent of template-driven content-management systems such as blogging tools that natively support microformats. Transforming into XHTML also offers positive feedback in the form of easy-to-apply CSS stylesheets that make contact information more attractive.

The same ability to extract information from XHTML using class selectors has enabled a wide range of tools and utilities. The hCard Creator Javascript form (www.tantek.com/microformats/hcard-creator.html) and a similar editor for hCalendar can even be inserted into existing editable text areas using a bookmarklet.) The X2V transformation service applies an Extensible Stylesheet Language Transformations (XSLT) template to export hCards and hCalendar entries found on a given Web page into an `.ics` file suitable for import into any standard address book application. Brian Suda currently operates this service on a voluntary basis (http://suda.co.uk/projects/X2V/), but Technorati will also soon be operating a large-scale public transformation service using this technology (http://feeds.technorati.com/events/, followed by the URL of the page to analyze). Unfortunately, because the vCard legacy format predated the wide availability of unicode, the X2V script requires users to be aware of which international character set and encoding to use.

Overall, transforming an existing schema into a microformat is a relatively straightforward process.

Even more exciting are applications that use Greasemonkey (www.diveintogreasemonkey.org) to find, edit, and share microformats found while surfing (see p. 3 for more on Greasemonkey). So-called "user scripts," such as Mark Pilgrim's MagicLine (www.mozdev.org/pipermail/greasemonkey/2005-August/004738.html) and Monkey Do (www.mozdev.org/pipermail/greasemonkey/2005-August/005030.html), are already detecting, parsing, storing, sharing, and searching snippets of structured data captured from Web pages.

The h* Effect

Skeptics might well note that the examples I've presented so far are just two facets of the same specification. A cynic might go further and ask whether microformats are simply a matter of slapping an "h" in front of existing specifications: call it the "h* effect."

Microformat advocates might celebrate such criticism because it underscores the philosophy of "reduce, reuse, and recycle" — shorthand for several design principles that contrast strongly with existing standards bodies and processes.
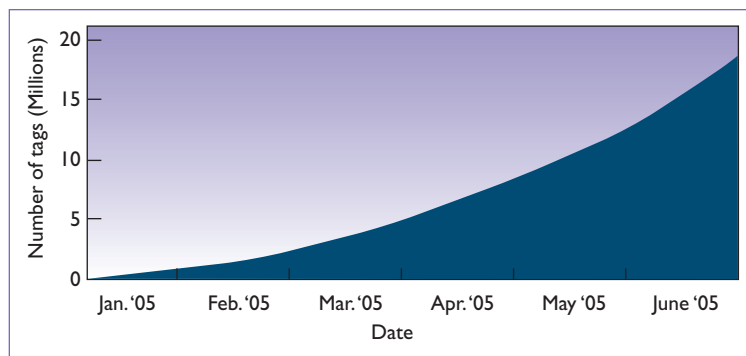
Figure 1. Microformat adoption rate. Data from Technorati and analysis by researchers at Carnegie Mellon University show that blog posters adopted the relTag microformat extremely rapidly.

Reduce

The microformats community and process serves to focus attention on a specific problem ("How can we point to licensing terms for blog posts?") and favors the simplest solutions. This is so central to its culture that it is enshrined in an official manifesto of sorts on the microformats.org Web site.

Reuse

Never proceed from a priori reasoning alone; work from experience and favor examples of current practice. Always keep in mind Picasso's dictum, "lesser artists borrow; great artists steal." Avoid the "not invented here" syndrome and embrace existing, widely adopted schemas.

In fairness, XML advocates would readily salute the same flag. Creating a new document-type definition (DTD) might be easy, but most would prefer to reuse existing standards. The key point of departure with the Semantic Web community is the (lowercase) semantic web community's rallying cry: "design for humans first, machines second."

Recycle

Make sure the results make it easy to decentralize innovation by encouraging modularity and the ability to embed. By ensuring that microformats are always valid XHTML, you can carry them in blog posts, Atom and RSS feeds, and anywhere else you can access the Web.

Presentable and Parsable

Sam Ruby's postulate states, "The accuracy of metadata is inversely proportional to the square of the distance between the data and the metadata" (www.intertwingly.net/slides/2004/devcon/68.html).

Combining this with the venerable maxim "out of sight, out of mind" explains why the microformats community insists on keeping semistructured information in-band and visible.

Thinking about Linking

"If one Web site links to another, the link doesn't carry any information about why the sites are linked. But what if it did?"

—*Tantek Çelik*[8]

The most successful microformat of all is *tagging*. The general idea of attaching short words to describe an item is ancient, but the recent "Web 2.0" enthusiasm for tags can be traced to the del.icio.us shared bookmarking service. Users started choosing tags that weren't just keywords but also labeled groups and roles ("to-read"). By *surfacing*, or displaying, items to which other users have applied the same tag, or by suggesting additional tags for a given item, tags power an exciting bottom-up process for collaboratively organizing information.

Another startup, Flickr.com, tagged photographs in the same way, to great effect. In retrospect, it seems inevitable that the approach would be applied to blogs, too. Technorati.com, which hosted the original developer's wiki for microformats, reported that it was tracking 20 million tagged posts within six months of the `relTag` microformat's introduction. Today, nearly a third of all blog entries include tags.[9]

Figure 1 is based on a frame from a short video generated by the Art and Computer Science research group at Carnegie Mellon University to visualize the dramatic adoption rate of relTag. (The 60-second video is well worth watching: www.our-media.org/node/37881.) How could tags be retrofitted to such a wide range of blogs (and blogging tools) so rapidly? RelTag (www.microformats.org/wiki/reltag) couldn't be much simpler: it's a value for a hyperlink's REL attribute. To indicate that a diary entry relates to ice cream, for example, you'd insert `<a href="http://technorati.com/tags/ice-cream" rel="tag">Ice Cream!</a>`. The convention for tagging in HTML is to use the last component of the URL path as the tag name for further indexing, thus letting users cite or create any tag vocabulary they'd like.

Typed link relations are a mainstay of hypertext theory, but they've generally been overlooked on the Web. Consider the social networking phenomenon of *blogrolls* — lists of authors' favorite

blogs, presented as lists of links in the margins of their homepages. In contrast to more ambitious efforts, such as the RDF FOAF, which more completely capture social-network relationships, the XHTML Friends Network (XFN) took the approach of adding link relationships to existing blogrolls.[10] Table 1 shows XFN's arbitrary vocabulary for human relationships. Although incomplete in any theoretical sense, the XFN vocabulary aims to solve "80 percent" of the problem.

The last of those (`rel=me`) is the most intriguing: it allows authors to link together all of the resources on the Web that represent themselves. It might seem superfluous, but it provides a pivot point for integrating fragmented digital identities, which are currently decentralized across multiple independent, isolated Web sites.

Microstandards?

Link-based microformats emerged first. These include vote links that express authors' opinions of the linked page, which can be tallied into instant polls by search engines; the `nofollow` link, designed to avoid influencing search engines' ranking algorithms when links occurred in comments; and license links that indicate the copyrights that apply to a given document. Today, these are known as *elemental* microformats, in contrast to the inline semantic markup of text, presented earlier, which are known as *compound* microformats.

Figure 2 illustrates the stack of standards on which microformats build. Like a sedimentary fossil record, it also happens to be a roughly accurate historical timeline of their emergence. Successful adoption of new standards at each layer has been essential for later diversification and complication in the layers above. The "standards process," such as it is, also evolved during these phases.

XFN was a product of the Global Multimedia Protocols Group (GMPG; www.gmpg.org), a self-proclaimed club of a few designers whose name they borrowed from Neal Stephenson's *Snow Crash* (1992, Bantam Spectra). The same team defined XHTML Metadata Profiles (XMDP),[11] a format for describing standards like XFN. Although it sounds like a mouthful, XMDP is just a clever way to tell readers a list of the class names and `rel\rev` link-attribute values that a particular microformat uses. XMDP declarations are linked in from the lesser-known `profile` attribute of the `<head>` element in HTML 4.01.

Although it's not as ambitious as other, more powerful schema description languages, XMDP

Table 1. XHTML Friends Network (XFN) vocabulary of human relationships for annotating links between blogs and home pages.

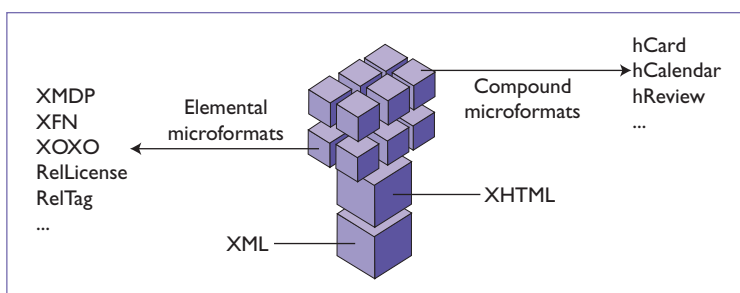| Relationship	Valid Values	Constraints
Friendship	Contact, acquaintance, friend	Pick one
Physical	Met	Presumed symmetric
Professional	Coworker, colleague	
Geographic	Coresident, neighbor	
Family	Child, parent, sibling, spouse, kin	Pick one
Romantic	Muse, crush, date, sweetheart	Not always symmetric!
Identity	Me	Excludes all other types

Figure 2. Standards on which microformats build. Microformats build on successive layers of existing standards. Just as XHTML depends on XML, the compound microformats for semistructured data reuse the simpler elemental microformats.

provides the foundation for microformats. It's more of a human-readable help file than a machine-readable set of rules for automating parsing and validation. Again, by providing this sort of "80/20" solution (focusing on the smallest subset of the problem that benefits the most users) microformats are making headway as a simple authoring solution while more complete Semantic Web description languages remain less widely adopted.

As the microformats community grew, Technorati hired a few of the initial advocates. This led to an unfortunate conflation between the general concept and a single start-up company's specific interests. CommerceNet.org, a long-standing nonprofit organization that promotes electronic commerce on the Internet, then helped sponsor the transition to a neutral home at Microformats.org in June 2004 and continues to help promote the technology and support the community in various ways.

An important test of the new regime came earlier this year with the rapid development of hReview, a format for publishing reviews. (See http://cnlabs.commerce.net/~rohit/hReview-in-Review/ for the overview that I presented at the

2005 World Wide Web conference.) Whether of books, movies, restaurants, or any other items, reviews are a common idiom in blog posting, and tool developers wanted a common way to share them with search engines that could aggregate community opinions.

Jointly authored by individuals from AOL, Microsoft, Yahoo, and Six Apart, among others, hReview was a watershed because it doesn't appeal to the authority of prior art.

Rather than translating one existing specification into XHTML, it builds on widely divergent standards, from the Platform for Internet Content Rating Services (PICS)[12] to the layout idioms of e-commerce sites. The critical break from past standardization attempts was making hReview independent of the items being reviewed; it contains nothing to limit it to only books, movies, restaurants, or so on. Instead, the `item` property is merely a formatted name, link, image, or hCard (for reviewing a person or corporation). This sidesteps all of the item-specific complexity of mandating bibliographic data, menus, or track lists; it's just enough metadata to build a better search engine for humans to find and read reviews.

It's important to note that neither CommerceNet nor Microformats.org is a standards body. The microformats community is an open wiki, mailing list, and Internet relay chat (IRC) channel that has proven remarkably scalable and accommodating. The only hard-and-fast rule for participation is that copyrights and patents on resulting specifications must be openly published and entirely royalty-free, respectively. Above and beyond that, the culture values research into existing standards, which helps dampen the tendency to promote too many narrow innovations. "Ruthless self-criticism" is actually one of the community's stated values.

Random Mutations

Microformats are not the only alternative to the intelligent design of the Semantic Web. The original XML vision is also adapting to the blog environment — this time, under the banner of *structured blogging*. (See www.structuredblogging.org for a good overview.)

Rather than treating XHTML content as authoritative and weaving metadata around it, structured blogging embeds plain XML within `<script>` tags, and then uses additional Dynamic HTML techniques to present it to end users via browsers.

This brings the advantage of complete decen-

tralization for creating new vocabularies, with an additional level of schema constraints for enumerated values and other basic data types. The disadvantage is that the semistructured information is invisible to ordinary browsers and screen readers for the disabled, and it creates less incentive for modularization and sharing of common vocabularies.

A subtler consequence is that structured blogging forces all the structured information on a Web page into a single "ghetto" — an island of XML within the larger HTML document. In contrast, we can add microformats to more complex HTML structures, such as table layouts for agenda grids or formatted presentation of bibliographic records.

In a classic joke, an inventor is showing off his latest gadget to a scientist who says, "Yes, yes, it works in practice — but does it work in theory?" Comparisons between the nascent microformats movement and the Semantic Web tend to raise the same question.

On "developers' day" at WWW 2005, the organizers held a panel discussion with advocates from both communities, and at least two distinct responses emerged. On one hand, microformatted HTML might be a case of Richard Gabriel's dictum "worse is better" by precluding "complete" solutions such as FOAF in favor of XFN. On the other hand, it might be the breakthrough on-ramp that the Semantic Web needs — a seedbed of common personal information that grows alongside 600,000-word ontologies for oncology.

Evolution, whether guided intelligently or randomly, will eventually churn out answers through the process of (technological) selection. In the meantime, it might be more evocative to end by considering what microformats are not:

> "Microformats are not a new language; infinitely extensible and open-ended; an attempt to get everyone to change their behavior and rewrite their tools; a whole new approach that throws away what already works today; nor a panacea for all taxonomies, ontologies, and other such abstractions."
>
> —*Tantek Çelik*

By avoiding as much as possible any pretension to designing a comprehensive, "intelligent" solution to the challenge of authoring information for the Semantic Web, microformats may yet take hold in their ecological niche as an appropriately incre-

mental evolution of existing technologies that makes the Web more amenable to automated analysis without infringing on authors' authority to present that data as they wish.

References

1. D. Brickley and L. Miller, "FOAF Vocabulary Specification," xmlns.com, 27 July 2005; www.xmlns.com/foaf/0.1/.
2. "vCalendar: The Electronic Calendaring and Scheduling Exchange Format," v. 1.0, versit Consortium, 18 Sept. 1996; www.imc.org/pdi/vcal-10.doc.
3. R. King, "Microformats: What the Hell Are They and Why Should I Care?," Webzine 2005, Sept. 2005; www.theryanking.com/presentations/2005/webzine/.
4. G. Klyne and C. Newman, *Date and Time on the Internet: Timestamps*, IETF RFC 3339, July 2002; www.ietf.org/rfc/rfc3339.txt.
5. F. Dawson and T. Howes, *vCard MIME Directory Profile*, IETF RFC 2426, Sept. 1998; www.ietf.org/rfc/rfc2426.txt.
6. D. Connolly and L. Miller, "RDF Calendar: An Application of the Resource Description Framework to iCalendar Data," W3C Interest Group Note, 29 Sept. 2005; www.w3.org/TR/2005/NOTE-rdfcal-20050929/.
7. F. Dawson and D. Stenerson, *Internet Calendaring and Scheduling Core Object Specification (iCalendar)*, IETF RFC 2445, Nov. 1998; www.ietf.org/rfc/rfc2445.txt.
8. "What's the Next Big Thing on the Web? It May Be a Small, Simple Thing — Microformats," Knowledge@Wharton, 27 July 2005; http://knowledge.wharton.upenn.edu/index.cfm?fa=printArticle&ID=1247.
9. D. Sifry, "State of the Blogosphere — Part 3: Tags," Technorati.com, Aug. 2005; www.technorati.com/weblog/2005/08/37.html.
10. T. Çelik and E. Meyer, "XHTML Friends Network" *Proc. ACM Hypertext 04*, ACM Press, 2004; www.gmpg.org/xfn/intro.
11. T. Çelik, E. Meyer, and M. Mullenweg, "XHTML Meta Data Profiles," *Proc. Int'l World Wide Web Conf.* (WWW 05), ACM Press, 2005; www.gmpg.org/xmdp/description.
12. T. Krauskopf et al., *PICS Label Distribution Label Syntax and Communication Protocols*, v. 1.1, W3C Recommendation, 31 Oct. 1996; www.w3.org/TR/REC-PICS-labels-961031.

Rohit Khare directs CommerceNet Labs, which is investigating decentralized electronic commerce. His research interests are in the fields of Internet protocols and decentralized systems. Khare has a PhD in software engineering from the University of California, Irvine. He was the founder of KnowNow and 4K Associates, previously worked on Internet standards development with MCI's Internet Architecture Group and the W3C, and served as editor of the *World Wide Web Journal* (O'Reilly & Associates). Contact him at rohit@khare.org.