

Do Domain Names Matter?

Francis Hwang

sera@fhwang.net

!

Is it just me, or are we paying less attention to the Domain Name System than we used to? Seems like only a few years ago that the tech-culture world was attuned to every new angle in the ongoing struggle over the DNS' management. You couldn't read the front page of Slashdot without catching one heavily commented-upon story on alternate registries, trademark disputes, or the latest ICANN board meeting.

But today? Hardly a peep. Not because the problems have magically solved themselves: The MPAA, for example, just sent a cease-and-desist letter to a blogger with the domain name www.ratednc-17.com. But a story like this won't draw the same attention it would have before. And by the way, what ever happened to those new top-level domains, like .biz, .info, and .name? Some of those are two years old and wide open for business — homesteads desperate for homesteaders.

This could be simply a temporary development. If the economy picks up we might see an uptick in the number of dot-coms suing hapless webmasters, and our outrage might rise accordingly. Or maybe we've just been exceptionally distracted of late: ICANN pales in comparison to the new crop of acronyms — MPAA, RIAA, DMCA, TIA, USA PATRIOT — menacing us today.

But perhaps these trends obscure a deeper shift. At the beginning of the boom, the vast quantity of people and organizations online outstripped our ability to find them, and we pressed the DNS into service to help fill that gap. But this usage of the centralized, permanent DNS conflicted with the common-sense methods that people use to name things in their everyday lives, and as the internet continues to decentralize this dissonance only grows stronger. The conflict is being alleviated not by technical or political reform at the center of the network, but by innovation at its edges. As end-user applications mature, they increasingly allow individuals to develop and share their own naming systems — not to destroy the DNS, but to render it irrelevant.

Just another pyramid scheme?

The reasons that the DNS started to crumble under the pressure of commercialization have already been well documented. Writing in 1998, Ted Byfield noted that the DNS was never designed for that pressure in the first place:

DNS was built around the structurally conservative assumptions of a particular social stratum: government agencies, the military, universities, and their hybrid organizations — in other words, hierarchical institutions subject to little or no

competition. These assumptions were built into DNS in theory, and they guide domain-name policy in practice to this day — even though the commercialization of the Net has turned many if not most of these assumptions upside down.

One of the assumptions Byfield is referring to is the notion that name collisions could be greatly reduced by dividing the namespace into top-level domains and trusting that everybody would calmly accept their place in that hierarchy. But as domain names became associated with trademarks, common usage flattened these tidy divisions into one undifferentiated sprawl. Corporations saw the web as one more front in the battle for marketing and public relations — like television, only with a keyboard — and accordingly they didn't care much for quaint rules written by computer scientists. So when, say, Archie Comics sued a California man for registering veronica.org in honor of his daughter, arguments that the .org TLD didn't belong to companies fell on deaf ears. (Online protest eventually succeeded where quoting RFCs had failed, although today veronica.org redirects to SamsDirect.)

Many reformers aimed for a political solution, appealing to ICANN to keep the DNS safe for bit players. They felt that in kowtowing to the corporations, ICANN was bastardizing the simplicity of the system that Jon Postel had managed until handing it off in 1998. (When the final draft proposal for ICANN was finished, Wired called Postel “the Internet's own Obi-Wan Kenobi” — a phrase that would attain an eerie resonance with the Slashdot crowd when Postel passed away a month later and ICANN revealed itself to be a bit of an Evil Empire.)

But would even a perfectly managed DNS have functioned in accordance with its earlier hierarchical vision? The hierarchy made sense to the users of the early internet, but the noisier commercialized internet would have fit much less comfortably into such a scheme. Even if you could've assigned every person, place, and thing its proper slot, most people would not have bothered to learn what went where.

Take, for example, the .museum TLD, which has been open since 2001. Most prominent museums have avoided using this TLD; the Whitney, the Guggenheim, and the Museum of Modern Art all place their primary domain names under .org. Conceptually, .museum muddies the waters because it's not mutually exclusive with .org. And when it comes to marketing, .museum is a disaster since it only serves to distract the user — who ever heard of a six-letter TLD? — without making her life any easier.

Or to take a more high-profile example, look at the recent lawsuit that forced the World Wrestling Federation to change its name to World Wrestling Entertainment. The World Wildlife Fund had sued the Federation for breaking the terms of a 1994 contract dictating who could use the initials WWF, in what media, and how prominently. Now, strictly speaking this wasn't solely an issue of domain names. The Fund's spokeswoman attributed the suit to an “explosion” of the acronym's use in three media: online, satellite TV, and cable TV. But one of the major grievances in the Fund's suit was the Federation's registration of the domain name www.wwf.com in 1997.

According to the neighborly rules of the pre-boom internet, this should not have been a problem: The Federation got wwf.com and the Fund got wwf.org. This solution works if you care about those tidy hierarchical divisions. Most people don't. This is one of the reasons that the new TLDs have been so underwhelming: People don't see the world as cleanly divided into

discrete categories, with the corporations in this corner and the non-profits in that corner. It's all one namespace to them.

Mutually exclusive hierarchies are convenient, but they only work on a small group of items. Once that group gets too big and diverse — a comic book artist here, an airplane-parts manufacturer there — any hierarchy that might reasonably hold that group becomes too cumbersome for people to use. When people want to organize large groups of items they often find it easier to use overlapping sets instead. That's why filesystems have symlinks. That's why many of Apple's OS X programs, such as iTunes and Address Book, let you drag-and-drop your MP3s or contacts into as many groups as you want. Why bother fretting over whether you should put Christine from work in your Friends group or your Coworkers group? Put her in both and get on with your life.

Six billion naming systems and just one internet

A hierarchical, precise DNS is a perfect system for computers. Human beings, however, prefer to rely on systems that make use of their own technical strengths, such as the ability to adapt their language to the preconceptions of your audience, and the ability to adapt their own conception of the world to accommodate new knowledge. Common sense, in other words. If, in the days when World Wrestling Entertainment was still a federation, you used the initials WWF in a conversation, chances are your listener would be able to figure out which one you were referring to. Humans do this by drawing on the context of the conversation to make the correct match. Are you talking about panda bears, or Stone Cold Steve Austin?

In real life, people have almost no problem resolving name collisions — a good thing, considering how often they happen. There are two types of Dove bars you can buy in a supermarket: One is chocolate and the other is soap. There are three famous Dres in hip-hop: Dr. Dre of NWA, Dre of Outkast, and Dre of Dre and Ed Lover. Hip-hop fans know how to tell them apart.

It happens on a personal level, too. In my freshman year of college, my dorm floor had four Mikes and four Daves. We resolved these name collisions by settling on nicknames for everybody: Big Dave, Sophomore Dave, Asshole Dave, etc. People who lived outside our floor didn't know who was who, but they didn't need that system anyway. We did, and it worked fine for us.

What we didn't do, however, was make use of last names, even though they offer a more global, permanent method of differentiation. Last names were less memorable to us than the jokey, college-guy associations we could invent on our own. Clay Shirky wrote that the aims of the DNS are to be memorable, global, and non-political. "Pick two", he said, but in fact most of the time people only care about the first: As long as names are memorable, people don't mind that they're local and highly subjective. Techies are an exception, since they spend much of their time crafting language for machines, and as such are accustomed to treating language as a brittle, precise tool. But most people like their language loose and contextual, thank you very much, and the hierarchies of the DNS demanded a rigor that never seemed worth the trouble to them.

The mnemonic and the meaningful

Another source of pressure on the DNS was the system's shifting role from one that was primarily mnemonic to one that was meaningful as well. The difference is subtle, but important. Consider the phrase "Every good boy deserves fudge", which music students sometimes learn to help them memorize what notes correspond to the lines of the treble clef. The phrase is helpful, but its content — boys deserving fudge — has nothing to do with music. It's mnemonic, but not meaningful.

The two can co-exist, and originally the DNS was a mix of both. A domain name like `gandalf.cs.columbia.edu` could give you important information — namely, that this domain is administered by somebody in Columbia's computer science department — but then, what does this domain do? Is it a mail server? A MUD? Knowing that somebody in Columbia's CS department likes *Lord of the Rings* is almost redundant.

Originally the purpose of a domain name was to be an address that was easier to remember than an IP address. This changed during the boom, as users and companies developed the notion that the function of a domain names was to serve as a self-explanatory pointer to a discrete real-life entity — a writer, perhaps, or a corporation or a museum or a hacker's collective. Of course, this was never fully realized, and it made little sense if you weren't swinging for the big leagues of global name recognition. If your site was niche enough that you could make use of an odd URL like `http://c2.com/cgi/wiki` (the first wiki, hosted on Ward Cunningham's web server), then those awful domain name disputes were somebody else's problem.

Today, internet services are becoming cheaper, more specialized, and easier to use, with the result that every day more people and organizations create a persistent online presence. And as the internet takes shape as — to borrow David Weinberger's phrase — small pieces loosely joined, the use of the DNS as a meaningful system is in further decline.

In the commercial world, companies ranging from small retailers to leading credit card providers use third-party services to manage online bill payment or e-commerce checkout. In doing so, they happily give up part or all of their domain-name branding in exchange for technical convenience.

Among individuals, of course, the most significant relevant trend is blogging. By some estimates there are already more than a million bloggers, and Lord only knows what those numbers will be like after AOL rolls out its blogging product later this year. Many of these bloggers don't have their own domain names. Instead, they're contained in subdomains (`http://jwz.livejournal.com/`), directories (`http://weblogs.mozillazine.org/hyatt/`), and CGI arguments (`http://www.xanga.com/home.aspx?user=sweetly_forgotten`).

If you can click on it, it's software

The meaningful DNS simply can't cope with a world of, say, 10 million bloggers, but luckily we have other ways to make sense of the internet. In the last five years, we've gained a number of powerful navigational tools, and these allow the DNS to pull back to a less high-profile role. The most obvious example is Google, which has done more than any other dot-com to make it easy to find your way around the internet.

But Google is still a centralized service, and as such there are limits to how much it can help. There is more promise at the edges of the network, where end-user software makes it easier for individuals to name, remember, and share URLs. Some preliminary examples:

- Subscribe to a blog's feed in your RSS aggregator and you might never have to type that blog's URL again.
- Blogging tools decrease the amount of manual work that bloggers have to do to pass links along. The beta version of Google's browser toolbar even has a BlogThis button.
- Apple's web browser Safari integrates with its Address Book to automatically bookmark the websites of your contacts.
- Almost all email clients and chat clients will automatically turn URLs into clickable links, relieving you of the need to even cut-and-paste.

None of these innovations are groundbreaking. But taken together they add up to an environment where users delegate to computers the dirty work of handling URLs. Consider, for perspective, this 1999 article by usability author Joe Clark:

A long URL works poorly in stationery, in articles in the print medium, and in advertising (e.g., on TV, with its low resolution and the short time a URL can be shown, or on radio, where it must be read out loud), and is a usability disaster in one-to-one conversations. (That's conversations, as in voice, as in getting together or talking on the phone.)

Today, this is still sensible advice if you're the webmaster for a Fortune 500 company or a popular dot-com. But for an increasing number of people, keeping URLs short isn't as important as it used to be. To take myself as an example: I'm relatively tech-savvy, and I create some sort of online content for a small, tech-savvy audience, so a short URL is much less important to me than it was only four years ago.

Stationery? I write in my Handspring Visor more than I write on paper. Print articles? My site isn't mass-market enough, and my target audience probably reads most of its news online anyhow. Television advertising? I'd love to have that problem.

And what about communicating a URL through speech? Personally, I find that this happens much less often than it used to. Maybe somebody will speak a URL out loud if she's referring to an easy domain name like half.com. But just as often she'll offer to email you that link, or IM it.

A URL can be both text and a software component. You can write it out longhand, but if you put it in an email client or a chat client, it's as much a software function as the Undo command: Click on it, and your computer responds. It's functionality that can be serialized into text if that makes it easier to transmit. And if you and your social circle are never far from computers with persistent broadband connections, then it's simpler to treat that URL as functionality rather than text: Rather than spell it out over the phone, email it or IM it.

Not that you should go making your URLs 400 characters long now. Shorter URLs are still better, or else why would we have those services that let you create a short URL to redirect to a longer URL of your choice? Notice, however, that the main purpose of these services is to

facilitate the machine transfer of URLs, since some email clients get confused when handling a long URL.

In fact, many of these services making the URL itself less meaningful, since they don't let you choose which key to assign to your long URL. Is <http://tinyurl.com/6a2> a map to your friend's party? That PDA your girlfriend is considering buying? The CIA World Factbook's entry on Afghanistan? These short URLs — and, of course, the domain names they contain — tell you absolutely nothing about what they point to. You'll have to rely on context to figure that out. Your friend writes you an email, says "Here's that restaurant where we're meeting for dinner on Thursday," and includes a short URL below. The URL itself means nothing. It takes its entire meaning from the conversation it's imbedded in.

Mere user customization is loosed upon the world

If the DNS is fading in importance, it won't be a surprise to everybody. Byfield, for one, wrote that "DNS's level of abstraction is sinking relative to its surroundings." A year later, in 1999, Jakob Nielsen predicted the same, and with pretty good timing to boot.

It is likely that domain names only have 3-5 years left as a major way of finding sites on the Web. In the long term, it is not appropriate to require unique words to identify every single entity in the world. That's not how human language works.

Today, in 2003, this is what the future of the domain name looks like: For the major players, the system will remain more or less unchanged. There will always be a small cast of large organizations and companies who will have domain names with household recognition: ebay.com, fbi.gov, etc.

But for the rest of us, we can increasingly rely on the fact that software is allowing users to build their own naming systems around their desktops, and then sharing and cross-pollinating those systems within their social circle. If you use the OS X Address Book, you can browse through your Safari bookmarks to find the link to, say, David Johnson's website. Which David Johnson? The one you care about.

So as decentralization continues, we can largely ignore the frustrating world of the DNS and focus our efforts on other ways to make connections. We can work on establishing our own roles in communities that are intimate and deep, not broad and shallow. And we can think less about marketing, and get back to just communicating.

!

[Francis Hwang (<http://fhwang.net/>) is the Director of Technology at Rhizome.org (<http://rhizome.org/>). His most recent artwork is firmament.to (<http://firmament.to/>), which uses a Perl script and the Google Web API to turn any HTML page into a free-associative index for the rest of the web. He has written about technology and culture for Spin, Wired, ArtByte, and Feed.]