ROBERT G. FERRELL

/dev/random

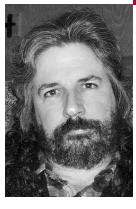

Robert G. Ferrell is an information security geek who enjoys surfing (the Internet), sashimi (it makes great bait), and long walks (to the coffee machine and back).

rgferrell@gmail.com

THIS WHOLE DOMAIN NAME SERVICE business has got me to frettin'. Folks have been predicting the imminent catastrophic downfall of the Internet for almost as long as the masses have been aware there was such a thing, but if that vile prognostication ever does come to pass, my money's on DNS being the culprit. Despite some of the vitriolic rhetoric I've seen concerning BIND and other DNS clients, I don't think it's fair to blame them, either, any more than it's fair to blame the jet stream for the untimely demise of your rhododendrons. The real point of failure here is the IP-address-to-name mapping scheme itself.

In the early incarnations of the public Internet, people discovered one by one that if they typed "foowidgets.com" into their browser they often found themselves on the Web page for the Foowidgets company, like as not replete with rotating animated gifs and blinking text (after the advent of Mosaic, anyway). Almost overnight, domain names became precious commodities; people fell all over themselves and anyone in their way in the rush to register names that pertained to their company, product, service, or person.

Once search engines came along and wormed their way into our daily routine, the point of this exercise began to slide inexorably toward mootness (mootitude?). That's not to say that the competition for domain names isn't as fierce as ever, or that people don't continue rightly to despise the bottom-feeding slime creatures who buy up names for which they themselves have no legitimate use, in the hope of reselling them for outrageous profits. They do. It *is* to say, however, that, quite frankly, the DNS system as it currently exists really isn't necessary. (No, that wasn't a misprint: I just had a lot of unused commas lying about that I bought on eBay in a moment of weakness.)

By "really isn't necessary," I mean to say it really isn't necessary, as in, not essential to the proper functioning of the Internet. "How," you may well ask with a hint of incredulity in your voice, "is that possible?" "Why," you indeed will in all probability now further inquire, "would we keep such an elaborate, high-maintenance, high-risk mechanism in place if it weren't vitally important?" The answer to both these questions lies embedded inextricably in the fundamental architecture of the Internet itself.

Pundits like to explain that DNS exists because remembering dotted quad IP addresses is too cumbersome and counterintuitive for humans. Here's why that's just a flat-out dumb statement: IPv4 addresses are 12 digits (of which 4 *must* be <= 2). U.S. domestic phone numbers are 10 digits. I know people who have dozens of phone numbers memorized. How many of you (other than network admins) have a dozen or more IP addresses committed to memory? The fact of the matter is that IP addresses are no more difficult to remember than phone numbers, but no one's ever written a hit song featuring one. IP address assignment and telephone number assignment are very much analogous processes: both involve locating an individual node in a complex multitiered network environment. Using the directory feature on your phone is equivalent to performing a DNS lookup in that it matches a human-friendly name with a network address. This is convenient, yes, but is it vital to contacting that person via phone? Not while we have both paper and electronic phone directories to help us out. In many ways, these tools are like manual search engines for telephone numbers, except you can't zap the pop-up ads as easily.

Search engine algorithms have grown so powerful and the tendrils of their crawler bots so far-reaching that the index page of virtually any domain you could possibly want to visit is already cached in multiple databases totally removed from those used for RFC 1034 (et al.) lookups. It wouldn't take a great deal of tweaking to make these distributed search engine databases authoritative for domain translation. We could also greatly diminish the competition for scarce names and in the process ruin the business models of domain squatters by employing Wikipedia-esque disambiguation pages that allow a large number of people to share one domain name. Forget TLDs: They are instruments of the diabolic and exist primarily to increase the revenue of domain registrars (another business model that will go belly-up). Besides, I'm tired of waiting for ICANN to authorize ".duh" and ".wtf."

How do you go about "registering" a Web presence in this brave new world? Just put it up with the appropriate tags and labels and tell the major search engines about it. As soon as their crawlers index the site, you're on the map. Anyone who types in a search term related to your business will eventually find you, the same way they do now. The only real difference will be that these search terms will be linked directly to an IP address, rather than to a domain name that then must be translated via DNS. You can find a fair amount of this sort of thing already in existence, since not everyone with a Web page has a domain name. You might argue that this gives the search engine companies enormous control over online commerce, but guess what? That's already the case. If you don't believe me, take this simple test:

1. Think of any product or service offered by a company whose name or URL you don't have memorized.
2. Go online and try, without the use of any search engine, to locate someone selling it.
3. Ingest some rich, creamy, artificially flavored trans fat as balm for your humiliating failure.

Your eConsumerism—nay, your very iExistence—is under the thumb of Google, and there ain't much you're willing to do about it, is there?

If the preceding proposal is too radical for you, here's a kinder, gentler alternative. Right now the DNS root zone system is, to borrow Dan Geer's word, *monolithic*, in that it relies on a baker's dozen more or less identical root name servers to point TLD queries in the general direction of the machine authoritative for a given zone. If these root servers are knocked offline or otherwise become unreachable, DNS lookups grind to a screeching halt. But

what if we distributed the functions of DNS *at every level* using a BitTorrent-type model, with many, many copies spread across the Internet? The odds against damaging or bringing down all of them at once are considerably higher than in the current failure mode, and as a bonus no single country could claim sole sovereignty over "the Internet." The powers that be claim that's the case now, but I think we all know deep down that it isn't in practice.

Oh, heck: /etc/hosts, anyone?