by MICHAEL O'BRIEN

"*We just register names. If you've got trademarks you get to fight it out.*"
– Network Solutions Inc.

"*When the installation process asks for the address of a name server, type in these magic numbers.*"
– Internet service provider

"*Get yer new domains here!*" – The new kids on the block

Mr. Protocol Revisits Names

Q: *This is amazing. Everything is* `something-or-other.com` *these days. I mean, there are a few* `.gov`*, and a fair number of* `.edu`*, hosts out there. But most everything anybody would want to deal with ends in* `.com`*. It's gotten to the point that all you have to do is type a company name in a Web browser's window, and it automatically sticks* `http://www.` *at the front of it and* `.com` *at the end, and there it is. Even sites in other countries have entries in* `.com`*! But other sites in the same countries have host names in their own country's domain. This is all a big mess. What's going on here? Doesn't anybody run this stuff any more?*

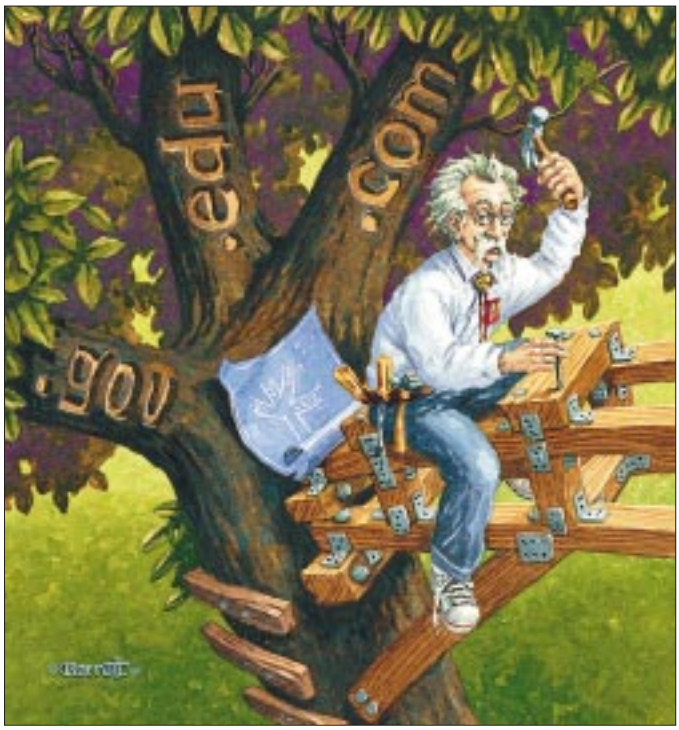

TOM BARRETT

A: Yes. No. Maybe. Sort of. It depends on whom you ask. It's a messy situation, but it might be cleared up soon. Or maybe not. See, it's like this…

Mr. Protocol has maundered on before (please catch that guy who's running out of the room screaming) about how the Domain Name System got started. In the Good Old Days when the Net was far smaller and less interesting than it is today, Stanford Research Institute ran the Host Table service. The master host table for the entire ARPANET was kept in a file on a well-known SRI machine, available for anonymous FTP. You fetched and installed the host table, ran whatever magic local program you needed to, and presto! You had the address of every host on the Net.

This was fine so long as there was only the one Net around to be updated. All connections to the ARPANET had to be centrally approved, so the host table was guaranteed to be reasonably up-to-date. On those occasions when it wasn't, you could edit it manually, or run a script to fix up long-standing inaccuracies. The Net was small enough that you could keep up with its intricacies, or at least, those intricacies that affected other hosts that you tended to talk to most often.

Usually, fixing up and installing the host table was a weekly ritual. Those who took pride in a smooth-running operation would install a new one daily. In the latter days of the ARPANET, daily updates became a practical necessity, as the transition to Internet technology began to heat up.

It was evident that as the Net became distributed, the name service also had to be converted to a distributed service. Paul Mockapetris, then of the Information Sciences Institute, designed the Domain Name System in response.

The DNS was a necessary development. Without it, the Internet could never have succeeded. No one can keep a host table of all of the millions of machines on the Net. Even if it could be kept, it would be far too large to be distributed, and even if it could somehow be distributed, it would be far too large for the majority of hosts on the Net to use.

Under the DNS, the only thing known to the Net as a whole are the IP addresses of a number of "root" name servers. To decode a fully qualified host address, such as `www.nostril.nose.com`, a host starting only with the

knowledge of a root name server would feed the name to the root server, which would return a response saying, in effect, "I have no idea what the IP address of that host is, old mangelwurtzel old thing, but I can tell you the address of a name server or three who know about the .com domain."

Armed with this new knowledge, our host would then feed the name to a server for the .com domain and would receive a response saying, "Sorry, never heard of that host, but here's where you can find out about the nose.com domain."

The server for nose.com, in turn, would return a pointer to the nostril.nose.com domain. Of course, it's possible that the name servers for nose.com and nostril.nose.com are the same, in which case, our host would get the IP address one step early.

It's the tree-structured nature of this beast that makes it work. Nobody is responsible for more than their own part of the tree. The servers for the higher-level domains in each domain name "delegate" the authority for the lower-level domains. This means that they don't actually contain information about those domains. Instead, they point at other servers that do. So, the only thing that the server for .com contains regarding the nose.com domain is a pointer to the server that actually does contain information about nose.com.

This separation of the layers, where the higher layers know only enough to get by (i.e., where the lower-level information starts), has made the Internet possible. The only thing that's wrong with it now that the Internet is part of everyday life is that it doesn't work.

What do you mean it doesn't work? Mr. Protocol is glad you asked.

The .com Problem

The problem is not with the idea of the DNS per se. The problem is the way people are using it.

Consider the way in which you typically go looking for a company Web site. You do it just as we mentioned in the question up at the top of the column: You take the company name, stick http://www. at the front and .com at the end, and pray. Now, no matter what your religious convictions may be [and it should be noted here that Mr. Protocol thinks "religious convictions" means "you were tortured by the Inquisition after being found guilty by an ecclesiastical court," which just goes to show why I'm writing this column and he's not], the canonical, standards-based method of discovering the correct name of a Web site should not be prayer. Yes, I know what "canonical" means, so shut up.

The problem is that the way people are using the .com domain, we might as well all be using host tables. There is no real difference between naming all the Web hosts www.whoever.com, and just naming them whoever. It still results in just one huge host table, only now it's the main server(s) for the .com domain that have to maintain it and serve it out, instead of SRI. This just doesn't work. Everybody knows it doesn't work, or at least, it won't work for much longer, at the rate the .com domain is growing.

Taking a dispassionate view of this, the real problem is one of directory services. The Internet has never had any. The only thing close to it that we've got are the Web search engines. These work by name only, and cough up anything that even remotely mentions the company you might be looking for. Real directory services would, ideally, take the name of a company, and bring up whatever resource you were looking for at the time, from the primary Web page for a Web browser to the email address of the chairman of the board if you were feeling peevish.

But directory services have we none. This isn't for lack of trying, mind you. All sorts of working groups have done large amounts of work on defining directory services, but nothing has ever been widely deployed. The DNS itself doesn't have a failing so much as a real need for an entirely different kind of service: directory services.

It's ironic that the ISO world, whose protocol suite was displaced from wide deployment by the rampaging success of the Internet protocol suite, actually has very good standards for directory services. It's anybody's guess why these have never been modified for the Internet and deployed (by which we mean that everyone knows, only everyone knows a different reason). For now, though, we stagger along without directory services.

Mark Andreeson, who developed the basic technology of the World Wide Web, never intended for Uniform Resource Locators (URLs) to be visible to human beings. These oddball collections of characters were to remain invisible. Humans were only supposed to see Uniform Resource Names…only these, representing a type of directory service, never really got hammered out and deployed before the invention of the Mosaic browser caused the Web to take off like a brushfire, and the Internet with it. It's as if the entire Internet were using raw IP addresses instead of host names.

Well, we do use host names, but thanks to the fact that everyone wants them to be easy to remember, since there are no real directory services to look them up, there's some strain showing. Eventually, there will be so many entries of the form www.foozle.com that the servers for the .com domain won't be able to take care of all the requests, because they won't be able to delegate any of them.

Since it doesn't look like directory services will be deployed any time soon, people are casting about for other solutions. This is where things get really interesting. Up till now, any time the Internet started to show strain, it was the usual cast of Internet Engineering Task Force suspects who came along and stuck their fingers in the dyke. But the Internet is big business, so we have the unprecedented appearance of entrepreneurs who are trying to a) solve the problem, and b) get rich at the same time.

Under contract to the U.S. government, a company called Network Solutions Inc. provides the domain registration service for all of the subdomains within .com, .edu, .gov, .org, .net and .us. Separate contracts support the operations of the six root name servers, which, in turn, generally also contain the servers for these same domains. The only thing a host has to know to bootstrap itself into using the DNS is the IP address of at least one of

these root name servers. Those servers which, in turn, serve the .com domain get their DNS tables from Network Solutions.

But who said there can only be these six root name servers?

The entrepreneurs look at it this way: The .com domain can't continue the way it has been because it will collapse under its own weight. If there were a much smaller number of subdomains within .com, and each of these gradually trickled down through a number of subdomains stacked one on top of the other before finally getting down to a real host name, the way the DNS is supposed to work, then all would be well. But everyone is trying to cram into the level just below .com, so that as many people as possible will be able to guess the correct name of their corporate Web server…or at least, remember it once they see it.

So why can't there be more top-level domains? Mr. Protocol is *really* glad you asked *that* one. There is no technical reason. It's all politics. Politics and human nature: same thing.

New Servers Are the Name of the Game

Let's say that you want to create a new top-level domain, say .biz, and start registering people in it. You're the Foofle Corporation, and someone's already grabbed www.foofle.com? No problem. Just go to "the other folks," and register as www.foofle.biz and be done with it.

And that's just what people are doing. The entrepreneurs have started an effort to create a whole new suite of root name servers. These root name servers will point to the existing name servers for the existing top-level domains, but will also point to other name servers for new top-level domains created out of whole cloth by the entrepreneurs.

So far so good, maybe. If you use one of these new name servers, you're getting a superset of the "traditional" Internet name space. Not being stupid, these people won't put up competing registrations for an existing top-level domain. That would be a quick and obvious shortcut to chaos. Instead, they only create new top-level domains and start accepting registrations for them.

The problem is that these new name servers are only useful if they're used…and used by everybody. Let's say you buy off on this, and get a nice shiny new domain name in one of these new domains. Well, presumably you point to the new name server so you have no problem getting at hosts in your own or anyone else's domain. But what happens when you send email to someone? If that someone doesn't also point to the new root name server, they're going to have no idea how to resolve that startling new host name in the "From:" line of your message. There's no way they can get an MX record or an IP address for delivery without also consulting the new root name server, of which they may never have heard.

Most people don't consciously make a decision. Certainly customers of CompuServe, America Online, or any of the other major on-line services do not choose which name server they're going to use. Their provider makes all of these decisions. In most cases, the providers give setup instructions that point customers at the provider's own name servers. These name servers, which are authoritative only for the provider's own domain, nevertheless serve as the DNS resolvers for all of the provider's customers. Customers who are trying to send mail or Web surf send host names to the provider's DNS server, which, in turn, consults the root name servers to begin resolving the host names. And there are not that many providers who are willing to hang the DNS usage of their entire customer base from an entrepreneurial root server, as opposed to one of the tried-and-true root servers that everyone else uses.

In a way, it's like that company you hear about on the radio, the International Star Registry. You send money and the company will name a star after your Aunt Tilly, and send you a nice certificate for you to give her on her birthday. The company will even send you a sky map pointing out which little speck of light is now Tilly Major. And it will duly record these names in its star chart registry books, which (it says) are kept in a bank vault in Switzerland.

All well and good. The company *does* name the star after your Aunt Tilly. The problem is, no one else cares. Certainly, the scientific community doesn't. They have their own naming authority, and it's not susceptible to you or your aunt. And no astronomer looks up star names in the registry of the International Star Registry, Swiss bank vault or no. (Swiss bank vaults sound impressive. Must be to keep hordes of disgusted astronomers from burning the books.)

So, if Aunt Tilly spent enough time fooling around with Linux, she could start her own root DNS server. And anybody who pointed to it would never know the difference, as long as a) they only sent out host names that other people could resolve without reference to Aunt Tilly, and b) Aunt Tilly is a good enough systems administrator that her server isn't down all the time. Point (a) is a bit of a killer. What good is having a brand-new domain if you can't tell anybody about it because they can't reliably resolve names in that domain? As for point (b), well, we can only hope for Aunt Tilly's sake that she's got a good operations staff.

The fact of the matter is that none of the "majors," by which we mean people like UUNET, MCI, Sprint and other large Internet service providers, are pointing to the new name servers. It's a rump movement with no large constituency, though they will argue that point. How large is "large"? Depends on whom you ask.

Nevertheless, .com is gonna break, and it's gonna break soon. What's to do? The Internet Assigned Number Authority, into whose lap this whole thing falls, at least theoretically, has done a very IETF-ish thing. It has formed the IAHC, whose impressive acronym means the International Ad Hoc Committee, which means, "We picked these people because they come from all around the world and it's mostly gonna be a problem they have to face so let's face it now rather than later and rope them in to help us." The immediate charter of the committee is to come up with a mechanism that will be recognized as legitimate by everybody for coming up with new

top-level domains. This mechanism, in turn, will be used to take the heat off of the .com domain.

Will it work? Almost certainly. Although cracks are showing in the facade, the IETF still speaks more or less with one voice. The entrepreneurs have a real chance, though. They have the option of approaching the committee, and saying, "We have established these domains, this technical mechanism, and this administrative mechanism, and it's fair, and equable, and easy to administer, and it solves the problem. And it's up and running now. We suggest you go with it. And it'll work under just about any conceivable future load, for these reasons, and it doesn't lend an unfair advantage to any nationality or group, for these reasons, and it's self-supporting and we stand to make a buck off it without its being onerous, for these reasons."

The entrepreneurs have some good ideas, too. Their plan is to run many root name servers, on the order of 64, rather than just six. They also plan to run separate servers for the top-level domains, so that, for example, the .biz servers aren't the same servers as their root servers. The current "mainstream" servers actually still have a single server serving root requests and top-level domain requests. Obviously, that will have to change.

The watchword of the IETF has always been "rough consensus and working code." The entrepreneurial group is fast approaching the "working code" state. The most distressing part about all of this is that the main spokespeople for this group in the IETF mailing lists have a tendency to argumentation and occasional flaming, which sabotages their arguments. It will be fascinating to see what relations they forge with the IAHC. From the other side, it will be interesting to see to what extent the IAHC can accept technical input while allowing personalities to slide off. Currently, there are enough angry people on the other side to claim that what we are dealing with here is "no consensus and bogus servers."

Is all this a take-charge solution in the finest American tradition? Or is it the biggest Ponzi scheme yet to hit the network? Only time will tell.

Mike O'Brien has been noodling around the UNIX world for far too long a time. He knows he started out with UNIX Research Version 5 (not System V, he hastens to point out), but forgets the year. He thinks it was around 1975 or so.

He founded and ran the first nationwide UNIX Users Group Software Distribution Center. He worked at Rand during the glory days of the Rand editor and the MH mail system, helped build CSNET (first at Rand and later at BBN Labs Inc.) and is now working at an aerospace research corporation.

Mr. Protocol refuses to divulge his qualifications and may, in fact, have none whatsoever. His email address is amp@cpg.com.