

More TLDs: why and how

This is a joint posting; John Levine posted it to [his blog](#) and Paul Hoffman posted it to [his blog](#).

[Susan Crawford](#), a new member of the ICANN board, asked about [auctions and lotteries for new gTLDs](#). Lots of people responded in the comments, and then the two of us kind of took over. We have now stopped, and are posting here.

The two of us agree on some things, and disagree on others. We agree that:

- The best thing for the internet is to release the next batch of gTLDs all at once, and to have that batch be about 50 new gTLDs. That way, no one gets much advantage of having a new TLD other than the semantics of the TLD's name.
- Anyone even considered being awarded a gTLD needs to prove in advance that they have the technical ability and skills to run a gTLD. This means existing DNS nameservers in geographically and topologically diverse areas, lots of bandwidth to those nameservers, and the ability to maintain a large number of queries on each one. Fortunately, these numbers can be agreed to ahead of time, and simple tests can be created to test the applicants. This is like the qualification process for a CLEC (competitive local phone company) in the US. In many cases, we expect applicants would simply contract with someone that already runs other TLDs to provide the technical infrastructure.
- ICANN should not profit from releasing the new gTLDs. If ICANN thinks it needs a bigger budget, it should go through a budget process to justify that budget, not just get windfall money from the gTLD process.
- The winners of the new gTLDs get no exclusivity over any other names issued by ICANN. That is, the first round might include *.rugs* and *.carpets*, or the first round might include *.tooth* and the next round might include *.teeth*; the winner of *.rugs* cannot prevent *.carpets* from being created.
- Generic (non-country) TLDs have failed at being directories, and that failure is getting worse over time. Users will always prefer to use search engines to find which sites relate to a particular topic than to assume that only domain names with a TLD that is semantically linked to that topic are of interest.
- There are probably only two plausible routes to a successful TLD: *.com* clones and certification. The clones are just like *.com*, only less crowded; we already have two of those, and a few more could probably be useful. The other route is certification. The reason that *.edu* is a success while *.coop* and *.museum* aren't is that people care whether something is an actual degree granting institution, while few people care about "real" museums and "real" co-ops. ("What a fool I was, they said they were a co-op but really they were only a producer's collaborative.") For a certified TLD to be useful, it has to cover an area that is relevant to many Internet users and be managed by an organization that will ensure that only bona fide SLDs are issued.
- Some of the most useful domains like *.edu* are not particularly large or lucrative. None of the proposed schemes (including our own) are likely to find them. If there is a place for creativity to be focused, it should be on figuring out which new TLDs are the most useful to typical Internet users and make sure those TLDs exist and are well-managed.

The two of us disagree on the best way to make these bunches of 50 gTLDs

appear.

John sees two routes to selecting TLDs. For TLDs intended to make money, the best approach is an auction, with the N highest bids getting to pick their N favorite domain strings, and the money given away to a suitable worthy cause, not ICANN. Other people have made more detailed proposals to deal with the obvious trademark issues, e.g., only IBM can pick *.ibm* but they still need a winning bid to do so. As Paul notes below, ICANN's beauty contest has picked losers, and a lottery tends to turn into auctions where the lottery winners keep the auction proceeds. Possible approaches include a separate lottery for five or ten names for which only non-profits can apply, giving virtuous bidders funny money they can use in the auction, as was tried in the PCS frequency auctions in the US. John doesn't have any great confidence that these will work, but if the auction process can be made simple and predictable enough, it should be possible to try one approach this year, another next year, and so on until one turns out to work.

Paul believes that there is no way to predict which TLDs might be "best". The track record so far [is abysmal](#). Having an auction might get people to think harder about which gTLDs would work best, but it is completely unclear who should profit from the auction. Instead, a lottery based on the desires of the organizations who qualify to be gTLD owners could be designed to get a wide variety of TLDs, with some organizations becoming big winners and the rest having ones that don't cost much to run. A lottery would prevent ICANN from making unnecessary money on the system, and would open the market to many companies who might otherwise be locked out.

Both of us agree that once the 50 gTLDs are assigned, there will be a lot of buying and selling of assets, regardless of what the rules for the auction or lottery say. Just live with it; that is how big business works. But the values of the new gTLDs will be much lower than might be expected because there are so many of them, with maybe another 50 or 100 a year later.

Reflections on new TLDs (by John Levine)

While pondering the renewal prospects for the three sponsored TLDs, *.aero*, *.museum*, and *.coop*, I went back and looked at the original applications for those and also for the unsponsored TLDs approved at the same time, *.BIZ*, *.INFO*, *.NAME*, and *.PRO*. Two lessons leapt out at me

If you build it, they won't care

All of the new TLDs offered a range of estimates of how many registrations they'd get, at 10%, 50%, and 90% confidence level. Their confidence was misplaced, since none of them have come close to their least optimistic 90% confidence level. The closest is *.INFO* which predicted between 3 and 11 million registrations and actually has 2.38M. The farthest is *.PRO* which predicted between 3 and 15 million and actually has 4900, but it's a special case due to being stuck in limbo for several years when its sponsor ran out of money.

You can fairly say that *.INFO* and *.BIZ* (with 1.28M registrations vs. predicted 3.8 to 30M) are successes, but pretty pale versions of success since each is in practice no more than a clone of *.COM* where people go if there's nothing good available in *.COM*. It's hard to think of a well known domain in either other than MTA.*INFO*, the New York City transit authority. Neither has come close to catching up with *.ORG* with 4.4M or *.NET* with 6.8M. The next biggest is *.NAME* where I counted 142,000 domains against a predicted

7 to 58M. Although my numbers may be a bit low since they sell individual names, it's unlikely they're anywhere close to recovering the \$7.5M the proposal said they were investing. The three sponsored domains, as I've noted before, each have only 5% of the minimum they projected.

The comparison between .AERO and .TRAVEL is most interesting. The new .TRAVEL domain has an aggressive campaign to sell domains to members of trade organizations, and they say they already have over 10,000 registrations, twice what .AERO has, even though they haven't formally launched yet. A little spot checking showed an awful lot of .travel domains landing on a registrar's parking page, suggesting they're defensive trademark registrations, but we'll see soon enough.

In any event, it's clear that with the exception of .COM clones, selling new TLD domains to prospective registrants is somewhere between hard and impossible.

Domains are not directories

The .MUSEUM and .AERO domains both made attempts to build a directory structure into the domain. The .MUSEUM domain has a lot of geographic names and keywords you can string together to guess a name, like `strong.museum.history.museum`, and a wildcard record so if you guess wrong, they show you a web page with suggestions of names you might be looking for. It's not awful, but nobody I know is likely to type `getty.museum` into their browser address bar rather than `getty museum` into their Google search box.

The airline industry already has a lot of widely used short names, including two-letter codes for airlines and three letter codes for airports. About 25,000 two- and three-letter codes are reserved for airlines and airports that might want them. Again this is not an inherently awful idea, but again that's not how people look for airlines and airports. Only 41 out of nearly a thousand two-letter airline codes are claimed, and under 500 of 9500 three-letter airport codes. [Bret Fausett noted](#) the Little Rock Ark. airport went to the effort of registering `littlerockairport.aero` but didn't bother to claim their reserved `lit.aero`.

The .TRAVEL domain seems to have recognized this, too, and plans to provide a searchable directory in which .TRAVEL registrants are invited to list themselves. I still doubt that it'll supplant Google, but it's easier to make a case for an industry-specific directory than to expect people to remember the naming rules for a directory-like TLD.