Article development led by **acmqueue**
queue.acm.org

Technology business plans that assume no competition—ever.

BY PAUL VIXIE

Arrogance in Business Planning

IN THE INTERNET addressing and naming market there is a great deal of competition, margins are thin, and the premiums on good planning and good execution are nowhere higher. To survive, investors and entrepreneurs must be bold. Some entrepreneurs, however, go beyond “bold” and enter the territory of “arrogant” by making the wild assumption that they will have no competitors if they create a new and profitable niche. So it is with those who would unilaterally supplant or redraw the existing Internet resource governance or allocation systems. Because alternative Domain Name System (DNS) roots provide such a well-proven and understood example of this kind of arrogance, this article begins with a short slog through that swamp before discussing the more

current and topical matter of alternative numbering Whois.

The DNS *root* is the dictionary of top-level domain names such as .COM or .US. It is managed cooperatively and transparently by a community that includes the Internet Activities Board (IAB), which designates and recognizes the Internet Assigned Number Authority (IANA); the Department of Commerce (U.S. DoC), which contracts for IANA services; and Internet Corporation for Assigned Names and Numbers (ICANN), which operates the IANA functions under that contract. The IANA functions contract includes among other things the job of editing the DNS *root zone* to add new top-level domain names such as .XXX. Each of these entities (IAB, U.S. DoC, ICANN) is itself a multistakeholder body that engages with the community to gather input to the decisions it makes about DNS. This governance model is imperfect, but it has worked for a long time and continues to evolve.

Technically speaking, every Internet device using DNS to look things up assumes there is a universal name space with a root zone to describe the top-level domain names, and there are some well-known root name servers to publish this root zone. To be universal in this context means that every name has a specific identity and will always mean the same thing no matter where you are on the Internet when you look that name up. The Internet Engineering Task Force (IETF) periodically revises the DNS protocol to add new capabilities, but this is always done in a backward-compatible way because of the installed base of hundreds of millions of connected devices. So while we could discuss a possible future in which new devices are connected to the Internet having a broader or somehow multiplicitous view of the DNS name space, as of today the only reliable way to treat this name space is as universal.

Given the high visibility and economic value of a new top-level domain name, DNS has been under considerable pressure to add more such names

NET

.NL
.CZ
.PL
.JP
.CN
.INFO
.SE
.FR
.EE
.ORG

.EU
.SK
.LV
.BR

.ES
.BZ
.AR
.TW
.CH
.CA
.TV
.UA
.NO
.ZA
.VN
.LT
.BE
.RO
.UK
.CC
.IN
.DE
.RU
.HU
.AT
.ME
.DK
.MX
.KR
.US

ever since the Internet climbed down from its academic ivory tower and became a world-changing dominating commercial and social apparatus. Prior work in this area includes adding a handful of new top-level names (.INFO, .MUSEUM, .BIZ, .XXX, and so on), and current work involves throwing the doors open to hundreds or thousands of new top-level domains (.APPLE or .MICROSOFT could soon exist). In addition to that, several bold (or dare I say, “arrogant”) entrepreneurs have tried to enter the market unilaterally.

Here is how this kind of unilateralism goes: first you create your own root zone, usually by copying the IANA root zone at some point in time; and then you try to get ISPs to use your root name servers instead of the IANA root name servers. If you succeed at this, then you try to sell name registrations in your alternative name space, where your new names will be visible only to the ISPs you have convinced to subscribe to your system. No such alternative root zone has really taken off, since this value proposition is pretty shaky—there is no way to manage the risk of conflict between an alternative name and some future real name in the IANA system. There is also no good way to align the interests of the people publishing the alternative names with the interests of some population who might want to look up such names.

What’s arrogant here isn’t the willingness to charge ahead in spite of the shaky value proposition; it’s the assumption that there will be only one alternative DNS name space, even if it is a financial success. Does anyone really think that other investors and entrepreneurs would not follow almost immediately, that other teams looking for their next opportunity would say, “Well, one is enough,” or even, “Being a late entrant into that market will be too difficult”? I cannot think of a single supporting example; success breeds copycats, in all times and all places.

It’s a marvel why the investors in today’s alternative DNS systems didn’t ask about copycatting. This is a pretty standard investment question. A bunch of copycats who pull various ISPs into competing alternative DNS systems could all sell the same names to different DNS operators, and there

It’s a marvel why the investors in today’s alternative DNS systems didn’t ask about copycatting. This is a pretty standard investment question.

would be no way for customers to tell the difference. Being first would count for nothing.

This spotlights a good test for whether some technology is a candidate for Internet governance infrastructure: Does it have to be done cooperatively, or do the physics allow for competition?

Alternative Numbering Whois

So far I’ve discussed the governance and economics of domain names, but there is another kind of Internet resource that has some superficial similarities to DNS: Internet numbering resources. Every network and every connected Internet device needs a number. This article focuses on Internet Protocol version 4 (IPv4) addresses, which are usually written as four numbers separated by three dots (e.g., 192.5.5.241 or 192.168.1.1). Some of these numbers are private and can be used only for local communication—for example, the address 192.168.1.1 is used by almost every cable or DSL router in every home in the world. Hosts connected to private networks rely on their routers to translate their private addresses into public addresses, a process known as NAT (network address translation). For the purpose of this article, the discussion is limited to public IPv4 addresses that are globally unique and used without NAT.

Before the commercialization and privatization of the Internet in the 1990s, the U.S. government assigned blocks of IP addresses without fee or contract. This befits the original purpose of the Internet, which was to be an interconnection mechanism for the government and its contractors. When commercialization and privatization began, the IP address-allocation function was moved out of government hands and into an regional Internet registry (RIR) system, which now consists of five registries serving the regions of North America and the Caribbean, Africa, Europe, Asia/Pacific, and Latin America. Each RIR is a non-profit association serving a community of network operators including both service providers and end users. Allocation policy is set in each region by a public policy development process, and resource allocations are governed by agreements that clearly describe

the allocation as being based on “demonstrated need” for network growth. These agreements also declare that number resources are not property.

Legacy numbering allocations made in the decades before the RIR system was put in place were very large because of the technical limitations of the time. The effect of this today is about half of all allocated numbers are of the legacy type even though most allocations are of the RIR type. Now that the Internet is running short of new IPv4 numbers for network growth, many network operators are looking for ways to acquire the rights to as many IPv4 numbers as possible so they can continue to grow their networks while the Internet converts from IPv4 to IPv6. This makes the older and larger legacy numbers very attractive, since the allocations were larger and are often held by older companies and universities whose needs may be modest by current standards. The holders of legacy numbers have no contractually explicit rights concerning those numbers unless they have sought safe harbor by entering into an RIR contract, but as a practical matter anyone who is using legacy addresses received in the pre-RIR era can safely continue to do so.

The RIR system permits designated transfers between address holders. The goal of the RIR transfer regime is to bring more IPv4 addresses into active use to facilitate network growth during the IPv6 transition. Any network operator who can demonstrate near-term operational need for number resources and who can negotiate a transfer with the current holder of those resources can simply sign an RIR contract and receive rights to the resources. Because this transfer regime was developed through a public policy development process, which is therefore bottom up rather than top down in nature, these rules are literally what the community of network operators asked for—such rules cannot be imposed by any government. Some interested parties, however, may not be able to demonstrate an immediate operational need and thus will not qualify as number-resource recipients. One class of such parties is the network operator who desires a long-term forward reserve. Another class is speculators who will never have need for the numbering

resources in their own names but who would like to hold the resources for later monetization (for example, rental or trading in futures).

It’s necessary to digest all of this background information to understand that not all interested parties are qualified recipients by the current transfer policies and not all transferable resources are under an explicit contract. The oft-stated concern is that these resources will be traded outside the system and that the RIR records (called Whois) will become useless. Since network operators use the RIR records every day to manage and diagnose their networks, these records should be complete and accurate. One proposal often heard in this context is that RIRs should not regulate transfers in any way and should simply record any transfer brought to them by a co-operating seller and buyer. A supporting argument for this proposal is that Whois can be run by anybody and if the RIRs won’t run an accurate Whois system (which is to say, a permissive system accepting the results of any and all transfers without limitation), then somebody else will do so. This argument breeds arrogance.

A strong advantage of the RIR Whois system in the eyes of network operators is that it is universal. There is only one entry for any given netblock and, therefore, effectively only one Whois system even though each RIR independently runs its part of that system. Let’s assume for the purposes of argument, however, that an alternative Whois system is created and enough network operators trust it that this alternative system becomes operationally relevant and that a non-RIR resource transfer regime becomes practical. Does anybody really believe that there would be only one alternative Whois system—no copycatting? Or as in the case of alternative DNS described earlier, would not the number of potential alternative Whois systems be limited only by available capital?

It would be technically possible to maintain a list of all alternative Whois systems and to query them all in parallel whenever network operations require knowing the details about a block of IP addresses. Inevitably, however, the same network would appear to be registered to different operators

in different Whois systems since freedom from transfer limitations is the stated reason for the very existence of the alternative systems. While anybody can start a new Whois system at any time, the operational usefulness and therefore the relevance of a Whois system depends on coherence and cooperation—two properties that an alternative Whois system and the alternative transfer market it supports would not have.

Conclusion

Any proposal for a competing Whois registry model is as doomed by design and destiny as every alternative DNS system. Even if it succeeds at first, it would fail after copycatting occurred. Participants in RIR public policy development would do well to remember this when evaluating dire warnings of RIR Whois irrelevancy because of an RIR transfer regime having a requirement of near-term demonstrated operational need. Speculators who want to monetize future need and network operators who want a forward reserve might still find ways to act outside the system, but resources will have to come into the system when their ultimate recipients qualify to receive the resources due to then-immediate operational need. The RIR system has no power to govern such private actions, but it need not and should not cede authority over the transfer policy and Whois registry—because that’s in the physics. ■

Related articles on queue.acm.org

DNS Complexity

Paul Vixie

<http://queue.acm.org/detail.cfm?id=1242499>

What DNS Is Not

Paul Vixie

<http://queue.acm.org/detail.cfm?id=1647302>

Successful Strategies for IPv6 Rollouts. Really.

Thomas A. Limoncelli, Vinton G. Cerf

<http://queue.acm.org/detail.cfm?id=1959015>

Paul Vixie is president of Internet Systems Consortium (ISC), a nonprofit company that operates the DNS F root name server and publishes the BIND software used by 80% of the Internet for DNS publication. He is also chairman of American Registry for Internet Numbers (ARIN), a nonprofit company that allocates Internet number resources in the North America and Caribbean region.