

DNS – A Failed Idea

This is a response to a [post](#) on Dave Farber's IP mailing list.

ICANN's does have a complementary role to the IETF in that it makes the IP address work well enough so you can argue it works modulo the IETF. The problem is that today's Internet protocols have no way to define stable relationships that are independent of the provided network.

The simple example I use is the problem of defining the relationship between a light switch and a light fixture. You simply can't define the relationship. You can't use the IP address because it's not stable (putting aside NATs for the moment). You can't use the DNS name because you can't look up the identifier if you're not fully connected with the rest of the Internet. Perhaps you can cache one if you were lucky enough to anticipate a problem but that's just playing the odds.

Even if you can assume you are connected to the root we still have a problem. The DNS name is not a stable identifier because you lease it (and agree to abide by social and commercial policies). You lease it because meaningful tokens are a scarce resource. That's because we attempt to treat the DNS as Procrustes' own version of trademark. And it fails in that role since it doesn't take into account any social aspect of naming and discovery. It's also a failure because you better not trust it – in many fonts, for example `rn` (`r n`) and `m` look identical so you can't trust what you see and a single typo goes far off and you can't use John Smith. So why do we compound the problem by encouraging more use of the DNS as a human-meaningful handle. There are far better ways to provide humans a way to express intent without creating the same degree of hazard.

XNS did use the MAC address as a stable identifier which was a step in the right direction though you still needed to get the MAC address from a central source of prefixes and chip supplier. That still leaves the problem of mapping the MAC address into the path. The bigger problem with the MAC address is that it is associated with physical objects and making them end points makes no more sense if we're talking about abstract protocols. It makes even less sense when the end point is a process running in the cloud that can float independently of the iron.

NATs are a feature in that they allow individuals to extend the reach of connectivity without having to get permission. The failure is in the need to translate addresses because we

use attempt to the IP address as more than a transient identifier.

You can just say fine – just do what Skype does and create your own identifiers on top – but why limit oneself to layering on existing protocols? In fact weren't the IP protocols just that – a layer on top of the previous Ethernet and telecom protocols in support of end-to-end solutions? So why patch over it again when we can revisit the defining principles.

There's also the issue of the IP address making routing unnecessarily difficult but that's another topic ... I'll just say that the network exists as a service to assist application in doing networking and should not be seen as a layer. It's a peer capability.

One point I make in <http://frankston.com/?name=OFI> is that the importance of the Internet is in the end-to-end argument – the idea that we can develop solutions without being tied to the network or, for that matter, any network at all.

I'm not surprised that the IETF is complicit in continuing business as usual though I am disappointed that it is not doing more to put an end to the confusion between "broadband" and "The Internet".

It's OK to patch up the existing protocols but we need to acknowledge the failures rather than pretending the bugs are features. Otherwise we will stay mired in the past which we insist on calling the future.