1

Defining internet law

A friend of mine set up an email server. He was really relieved at how easy it had been. It only took a few online searches, downloading some free software and following the steps. Now all he wanted to do was check to see if it was RFC compliant.

RFCs are usually described as technical standards for computer networking. RFCs establish protocols, procedures and conventions used in or by the internet. RFC 2026 says they are designed to help facilitate best practice in terms of:

- technical excellence;
- prior implementation and testing;
- clear, concise and easily understood documentation;
- openness and fairness; and
- timeliness.[1]

RFCs date back to the development of ARPANET in 1969 and the decision to document the process of designing networking applications. It is partly due to the success of these networking applications, and the take up of these early documentation practices, that now over thirty years later, it is relatively easy to set up your own email server from your home computer, if you want to. But that the internet is easy to access and easy for users to develop their own nodes or networks is only partly due to the successful distribution of good technical standards.

RFC actually means 'request for comments'. This terminology signals something of the ethos of those involved in the original developments.

Stephen Crocker, who wrote RFC 1, says in RFC 1000:

> The precise usage of the ARPANET was not spelled out in advance, and the research community could be counted on to take some initiative . . . Most of us were graduate students and we expected that a professional crew would show up eventually to take over the problems we were dealing with . . . we found ourselves talking to people whose first concern was how to get bits to

> flow quickly and reliably but hadn't – of course – spent any time considering the thirty or forty layers of protocol above the link level.
>
> . . . it became clear to us that we had better start writing down our discussions . . . I remember having great fear that we would offend whomever the official protocol designers were, and I spent a sleepless night composing humble words for our notes. The basic ground rules were that anyone could say anything and that nothing was official. And to emphasise the point, I labelled the notes 'Request for Comments'. I never dreamed these notes would be distributed through the very medium we were discussing in these notes. Talk about Sorcerer's Apprentice![2]

The 'just get on with it' approach to innovation and an unassuming attitude toward authority characterises much of the early development of internet standards.

This attitude morphed into a broader cyberlibertarian ethos that characterised the discussion of the internet of the 1990s. Richard Barbrook and Andy Cameron described it as a 'Californian Ideology':

> By integrating different technologies around common protocols, something is being created which is more than the sum of its parts. When the ability to produce and receive unlimited amounts of information in any form is combined with the reach of the global telephone networks, existing forms of work and leisure can be fundamentally transformed. New industries will be born and current stock market favourites will be swept away. At such moments of profound social change, anyone who can offer a simple explanation of what is happening will be listened to with great interest. At this crucial juncture, a loose alliance of writers, hackers, capitalists and artists from the West Coast of the USA have succeeded in defining a heterogeneous orthodoxy for the coming information age: the Californian Ideology.[3]

The success of early network applications was not simply because they worked. What carried them forward and contributed to their successful adoption was the promotion of a culture associated with the use of the technology.

Discussion of the internet has been strongly influenced by this early history. The conventional wisdom is that we have since moved on from those humble, innocent days. Increasingly internet architecture has become less 'open' and surveillance and control have assumed a definitive role. The professional crew that Crocker feared would turn up and take things over did eventually arrive. But they didn't turn out to be a bevy of like-minded

computer scientists interested in efficient applications. They were an army of intellectual property and contracts lawyers, practising a far more arcane science and with no interest in how quickly or reliably the information flowed. And then the lawyers got together with some of the computer scientists, which could mean the death of the internet as a commons, so the story goes.

There is something in these now familiar popularist histories of the internet and the politics of its changing cultures that is missing.

To understand what that is, we need to return to the discussion of RFCs and the peoples who develop and use them. Why do little nodes, such as my friend's small email server which is attached to a car club site, need to be RFC compliant? As the origin of the name infers, RFCs begin as humble suggestions about good design. They say no more than 'hey, adopting these standards might be a good idea'. There is no law of the internet that says your site or network must be compliant. If you are happy with how your little piece of the network works, why bother?

There is more to standards than functionality

The usual explanation for seeking compliance with RFC standards is to achieve the level of functionality that comes with adopting a tried and tested 'best practice'. But there are lots of organisations, some voluntary, others that are commercial operations, testing for RFC compliance. These organisations compile listings of non-compliant addresses and relays, and if you are not compliant, you are open to being blacklisted. These lists are very frequently updated and emailed to service providers and other intermediaries, who then deny service to blacklisted parties or locations. The point of blacklisting is to make the work of spammers, virus and worm distributors and other anti-social members of the internet community – people whose activities capitalise on sloppy network design – as difficult as possible.

The design issues RFCs address are real ones. And viruses, worms and spam are a pain that makes the technology difficult to use and a major expense for individuals and organisations to combat. As everyone who uses computers knows well, they pose a threat that can compromise the entire communications medium.

When you have a system that allows for points of entry into the network for people with relatively low levels of technical skill, you can't always be sure they realise the risks they are creating with the technical things they do, that can affect everyone else who uses the medium. That's why there is a need for RFCs and for bodies that test RFC compliance. These people don't tidy up

after poor designers, but by setting standards and by creating blacklists, they try to minimise the risks of damage. And that there is no official overseer, and that there are many that are running different compliance checks, is in keeping with the original engineering philosophy.

But this openness to authority structures can itself create problems.

Recently a blacklist operated by 'RFC-ignorant.org' listed the whole of the '.com.au' domain space. RFC-ignorant is:

> the clearinghouse for sites who think the rules of the internet don't apply to *them*. We maintain a number of lists . . . which contain domains or IP networks whose administrators choose not to obey the RFCs, the building block 'rules' of the net. It is important to note that NOTHING requires ANYONE to comply with an RFC (pedantically a 'Request for Comments'), however, the 'cooperative interoperability' the net has enjoyed is based upon everyone having the same 'rule book' and following it . . . RFC-ignorant.org does not block anyone. We document who has chosen not to implement certain protocols described in the RFCs, and provide a means for allowing people to determine for themselves if they wish to communicate with non-compliant systems.[4]

A consequence of the listing by RFC-ignorant was that email to and from some '.com.au' addresses started to bounce. This caused serious concern amongst affected Australian businesses. They felt that their access to customers had been severely compromised, because of a decision made by unknown persons, from somewhere else, about a technical matter that they had little understanding of. And in any case, they had no ability to remedy a problem that affected an entire top level domain. There was nothing that they could do themselves about getting a blacklisting of '.com.au' lifted.

Australia was not the only country affected. RFC-ignorant also blacklisted '.cl' and '.pl' domains, affecting email addresses in Chile and Poland.

The most practical solution offered to Australian businesses was to try to buy a US '.com' domain name, and get the word out to their customers of the change of address. Of course it is most unlikely that were the same compliance problem to affect the '.com' top level domain that RFC-ignorant would presume to deal with it in the same way. And even if they did, it is hard to believe administrators would avail themselves of this clearing house's services any longer. Can you imagine the consequences of administrators of email servers choosing to exclude the entire '.com' domain space?

The reason given for the blacklisting of the '.com' domain was non-compliance with RFC 954 NICNAME/WHOIS, October 1985. This is a protocol about the listing of necessary contact and ownership details of registered domains.

From the beginning of the development of ARPANET it was recognised that there would be a need to be able to learn about the origin of the resources distributed over the network. To service this information need, a directory service run by the Network Information Center (NIC) was established at Stanford Research Institute. Each domain registration record usually has basic details about the registrant and the administrative and technical contact. The registrant 'owns' the domain. The other contacts may provide day-to-day maintenance of the domain details, such as where the domain is hosted on a server. Whilst a website often also has a contact email address listed on it, there is always some uncertainty about who this party actually is, and what kind and level of responsibility they have for the site.

To locate the appropriate persons responsible for the domain it is common to conduct an online search in a 'who-is' database. There are numerous who-is databases on the internet dealing with a wide range of top level and sub-domains. The information a who-is search may provide about a domain registration could simply be an email address and a NIC-name. The term NICNAME has come to refer to an alphanumeric code created at the time of registration, that provides a short cut to a fuller record of contact details of the domain registration. The full records may include a postal address, phone number, fax, and email addresses.

Some Australian domain name registrars, in accordance with the policy of the licensing and accrediting body, the Australian Domain authority (auDA), had removed public access to much of the contact information for registered domains, citing privacy concerns. There have been cases where, for example, personal contact details of a celebrity author were obtained from a 'who-is' database search conducted by an earnest fan. The author's website had been set up with an administrator as the email contact, precisely to try to keep unsolicited contact with her fan base at a distance.

Whereas 'open architecture' could be described as the internet culture of the 1980s and 1990s, with the pervasive spread of quite intrusive communications technologies privacy has emerged as a major concern. Privacy and open architecture are not necessarily at odds. There are also many RFCs dealing with aspects of privacy. However unfortunately RFC 954 NICNAME/WHOIS does not address privacy. This and related problems with the protocol has led to requests to replace it, including some by administrators of

country code top level domains who have implemented privacy policies and then found their domains listed as non-compliant.

The role of drafting and publishing RFCs is now undertaken by the Internet Engineering Task Force (IETF).[5] The organisation describes itself as 'a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual.' An IETF Working Group is looking at the drafting of a replacement protocol, recognising the lack of a privacy policy as one of the weaknesses of RFC 954.[6]

That the issue is being addressed attests to the open consultative processes of the IETF, and their well-developed sense of responsibility toward best facilitating the operation of the network as a whole, and attention to interoperability within it. However IETF only sets standards, and has no interest in hearing disputes about ways these are interpreted and used downstream. That is not their game. But given the way the network has evolved to date, nor is it anyone else's.

RFC-ignorant also values the 'co-operative interoperability' of internet architecture. However they see this dynamism articulated as rules established by a respected community of technical peers. Their services draw attention to those they believe are compromising the spirit of the enterprise in not conforming to the expert's rule book. They only point the finger, of course. But to creators, owners and administrators of domains, at whatever networking level, this subtle detail is hardly a relevant point.

There are many kinds of lists, but finding out who subscribes to what services, and why to those, is very difficult to determine. It is not something systems administrators are keen to openly discuss, fearing the discussion would in itself create some vulnerability. Administrators take advantage of various kinds of listings hoping to keep the spaces they bear responsibility for operating as their clients expect – with minimal intrusions and maximum network efficiency. Systems are continually tested for myriads of weaknesses, and no system can ever be presumed as secure, reliable or invulnerable. Blocking is not a simple exercise, and spammers and the like can be constantly moving targets. It is well recognised that one of the consequences of using lists may well be the blocking of legitimate users from time to time and place to place.

Responsibility for 'collateral damage' produced by the network is always deferred. 'Co-operative interoperatibility' is partly a messy war of attrition, and those fraternising with poor practice are cast as co-conspirators in a flagrant theft of resources. Network resources are assumed to really only be

the natural inheritance of those that can be tested and proven as technically worthy.

As the provider of one aggressive Mail Abuse Prevention System says:

> The Mail Abuse Prevention System's Realtime Blackhole List (MAPS RBL) can be used by any interested party in the configuration of their own network or mail relay, toward the goal of limiting theft of resources by spammers. This step must not be taken lightly – the MAPS RBL creates intentional loss of connectivity for anyone who chooses to use it. While we try to limit that connectivity loss to only networks which are friendly or neutral toward spam, sometimes a spammer hides in and amongst nonspammers so as to share a more positive fate with those nonspammers. What actually happens is that the nonspammers share an unpleasant and negative fate with spammers in that case. In other words, if you are not willing to occasionally throw out a baby with the bathwater (figuratively speaking, of course), then the MAPS RBL is not for you.[7]

The hope is that eventually the nonspammers, finding themselves thrown out, will force their service providers to take the action required to rectify whatever is deemed to be the cause of the potential spamming problem. And if the provider doesn't act, their disconnected clients are exhorted to go elsewhere.

Despite the salient advice that may be offered, not only administrators that share a hardline approach will necessarily subscribe to these kinds of lists. There is a strong pragmatic interest in using whatever tools are available to help with the job – so long as it shows results. The desired situation is to draw neither spammers nor anti-spam vigilantes to your system. A successful system is one that in technical terms goes unnoticed. To help with this objective, increasingly email software comes with settings already determined to gather various kinds of default listings that facilitate the spread of preferred technical 'solutions'.

Whilst businesses struggling with bouncing email will be fuming, these people are also likely to share an interest in opening their email and finding the inbox filled with messages actually relevant to the business or their other lives. The real source of anger is frustration at finding oneself unceremoniously dumped from a network you have come to depend upon. Some of the disconnected may well have a gut-expectation that modern notions of administrative fairness and accountability in decision-making would apply before draconian steps were taken. But these particular legal ideals are alien to this diffuse, global network.

The interface between law and technology is difficult because despite little in the way of formal authority structures, there are clearly forces that operate very much like laws. However the culture that has grown up around the internet uses references to the voluntariness of the 'protocols', the virtues of 'decentralisation' and 'openness', and 'choice' about compliance, about subscribing to lists and taking action. This language, used especially by those technicians and managers whose actions power the system, deflects any address toward the reality of decision-making structures existing. It also destroys any practical expectation of formal responsibility being taken for those hit with the friendly fire produced by the maintenance of smooth operations.

The other kinds of laws

We are accustomed to thinking of laws in terms of formal mechanisms of governance, of centralised rules, run by hierarchically organised bureaucracies, with courts and other legal personnel that serve the public. And the problem is that this level also permeates the network at various points. For example, to act constructively against RFC-ignorant it would be useful to know who is using their lists, and thereby be able to contact these parties directly with news of the fallout from the recent initiative, and tell them about the privacy problems associated with complying with RFC 954. However the client list of RFC-ignorant is commercially valuable information. Quite ironically, in these circumstances this means that the privacy of RFC-ignorant and their clientele could be protected under trade secret or confidential information laws. If a third party were to find a means of accessing these details, the courts could be asked to prevent the disclosure and use of that information, because disclosure could harm the commercial interests of the clearing house and its clients.

The information contained in who-is databases might also be regulated by more specific privacy legislation. This area of jurisprudence is relatively new and strongly influenced by the European Union (EU) directive on the protection of personal data adopted in July 1995. This directive regulates transfer of data to non-EU countries, who are required to provide an adequate level of protection for personal information as a condition of the free flow of information from all EU states. Out of a concern for the privacy issues raised by new technologies, and for the trade implications of not having acceptable privacy laws, many EU trading partners have found it essential to develop them. Australia, for example, has national and state legislation

that creates obligations concerning the collection and use of personal information obtained by government departments, agencies and many private sector organisations. There are principles governing how personal data is to be stored, made secure, the conditions under which it may be disclosed, and the rights of the individual to access the information. auDA's *WHOIS Policy No: 2002–06* says: '4.2 In order to comply with Australian privacy legislation, the street address, telephone and facsimile numbers of registrants will not be disclosed.' It is the application of this policy, which is in accordance with national law, that created the non-compliance issue noticed by RFC-ignorant.

Whilst the openness of the network emphasises global interoperability and global technical norms that operate like 'rules', there is also a need to attend to the layers of domestic laws that could apply. Clearly these laws can affect the operations of all parties resident within the state or country, and parties are not free to ignore them without risking legal consequences.

It was commonly thought that online operations need only concern themselves with complying with domestic laws where their business has a significant presence, such as an office and assets. However this has proven not necessarily to be the case, as Dow Jones, the US-based publisher of the financial publication *Barron's Magazine*, found out. In 2001 Dow Jones was sued under an Australian state defamation law over comments made concerning the dealings of an Australian businessman, Joseph Gutnick. The article was made available to subscribers of the Barron's news service on Dow Jones' website. A small number of subscribers downloaded the publication in Victoria where Mr Gutnick is a resident and where he sought to defend his reputation.

The Australian High Court considered the question of local jurisdiction and legal responsibility in a global communications medium in 2002.[8] In addressing these issues Justices Gleeson CJ, McHugh, Gummow and Hayne noted:

> Dow Jones submitted that it was preferable that the publisher of material on the World Wide Web be able to govern its conduct according only to the law of the place where it maintained its web servers, unless that place was merely adventitious or opportunistic . . . The alternative, so the argument went, was that a publisher would be bound to take account of the law of every country on earth, for there were no boundaries which a publisher could effectively draw to prevent anyone, anywhere, downloading the information it put on its web server.

However as their Honours saw it:

certainty does not necessarily mean singularity. What is important is that publishers can act with confidence, not that they be able to act according to a single legal system, even if that system might, in some sense, be described as their 'home' legal system. Activities that have effects beyond the jurisdiction in which they are done may properly be the concern of the legal systems in each place.

Justice Kirby said:

The genius of the common law derives from its capacity to adapt the principles of past decisions, by analogical reasoning, to the resolution of entirely new and unforeseen problems. When the new problem is as novel, complex and global as that presented by the Internet in this appeal, a greater sense of legal imagination may be required than is ordinarily called for.

. . . If the place of uploading were adopted as the place of publication which also governs the choice of applicable law, the consequence would often be, effectively, that the law would assign the place of the wrong for the tort of defamation to the United States. Because of the vastly disproportionate location of webservers in the United States when compared to virtually all other countries (including Australia) this would necessarily have the result, in many cases, of extending the application of a law of the United States (and possibly the jurisdiction and forum of its courts) to defamation proceedings brought by Australian and other foreign citizens in respect of local damage to their reputations by publication on the Internet. Because the purpose of the tort of defamation (as much in the United States as in Australia) is to provide vindication to redress the injury done to a person's reputation, it would be small comfort to the person wronged to subject him or her to the law (and possibly the jurisdiction of the courts) of a place of uploading, when any decision so made would depend upon a law reflecting different values and applied in courts unable to afford vindication in the place where it matters most. At least in the case of the publication of materials potentially damaging to the reputation and honour of an individual, it does not seem unreasonable, in principle, to oblige a publisher to consider the law of the jurisdiction of that person's habitual residence.

Justice Callinan remarked:

what the appellant seeks to do, is to impose upon Australian residents for the purposes of this and many other cases, an American legal hegemony in relation to Internet publications. The consequence . . . would be to confer upon one