

July 2012
Geoff Huston

Occam's ITRs

Some Further Reflections on revising the International Telecommunications Regulations

It's been a quarter of a century since the world's governments convened to draft up a common set of regulations about the conduct of international telecommunications. In December of 2012 the world's governments will convene to reconsider these regulations, to hopefully sign an updated set of regulations. This time around, this activity is generating considerable levels of public interest. Congressional hearings in the United States have been held, and various pronouncements of intent from various governmental, regional, and industry groups have been made. The level of interest in international telecommunications is high, and the diversity of views about what should be expressed in a revised set of regulations is also evident.

Rather than adding specific measures, conditions, or constraints, it may be prudent to consider a set of regulations that says far less and encompasses our common aspirations in the area of international telecommunications, rather than attempting to arbitrate among a diverse set of often conflicting specific demands.

To provide some context on the current set of regulations, the International Telecommunications Regulations (ITRs) of 1988 define telecommunications in very broad terms:

Telecommunication: Any transmission, emission or reception of signs, signals, writing, images and sounds or intelligence of any nature by wire, radio, optical or other electromagnetic systems.

International telecommunication service: The offering of a telecommunication capability between telecommunication offices or stations of any nature that are in or belong to different countries.

While this definition is highly inclusive, the ITRs were in reality referencing only telephony and related telephony-based telecommunications networks. At the time, all wide-area, national, and international data networks were built on the margins of oversupply of existing telephony-based infrastructure. It's therefore unsurprising that the effort to generalize the concepts of international telecommunications was naturally defined, and limited, by that telephony-based communications paradigm and its concepts of technology, tariffs, and inter-provider interaction. While the regulatory language was generic, the concepts described by these regulations matched quite precisely the technological, operational, and business profiles of telephony.

By 1988 the data network industry was flourishing, and Ethernet was a nearly universal substrate for local data communications.

It's interesting to observe just how revolutionary Ethernet was at the time. The prevailing data communications technologies were point-to-point technologies, that connected two computers together. Ethernet was a common bus technology that allowed hundreds of computers to connect to this common bus, allowing any computer to transmit a data packet that was addressed to any other connected computer. But it was not just the transformation from point-to-point to common bus that was so radical here. It was also the transfer of functionality from the network to the computer. Point-to-point circuits were typically "network-clocked," so that the network provided the timing signal for the transmission. The network clock had to be stable over periods of months or even years, which implied a centralised clocking model with high precision and very stable clocks and the propagation of the clock signal across the carriage domain. Obviously this was not cheap. Ethernet used "self-clocked" packets. While each computer's etherclock was a high speed oscillator (20Mhz), it only had to maintain sync for around 12,000 cycles, allowing for a far cheaper time function to support the data transmission. This allowed Ethernet to be a purely passive cabling system, where the entire functionality of the data communications service was defined by the connected computers, with no active network component. This was indeed a revolutionary shift for data communications.

But such specialized data technologies operated over a dimension of a campus at best. Any communications over longer distances, in particular internationally, required the transformation of the data into a stream that matched the characteristics of the voice carriage hierarchy, injecting it into the voice network, and performing a complementary extraction of the data stream from that network at the other end. In 1988 it was still possible to assume that at the carriage level, the general entirety of global telecommunications was in the form of telephony.

A lot has changed during the last 24 years. Large-scale data networks are constructed using a digital data transmission network, which is architected from the photon all the way up to the packet. These systems are commonly engineered to carry IP, the Internet Protocol, and comprise the modern Internet. The massive array of applications, ranging from traditional data streams all the way to today's social networking environment, have been constructed and layered above this common carriage substrate.

On the Internet, voice, mail, messaging, and television all group together as instances of application categories on this common data carriage network. There is no enduring need to architect service bundles, such as the once popular "triple-play". There are now meta-applications that blur the distinction between particular communication transactions. Social network applications have achieved huge success due to their agility in bringing together what were considered to be discrete application environments in novel forms. The combination of these varied media and communications models have given rise to an immersive peer-to-peer environment of interaction among literally hundreds of millions of people.

The "datagram", which is the very heart of the technology of the Internet, is what makes this all possible. What this represents is a stripping down of the functions performed by a network to the most simple and basic operation. Every "transaction" is a single packet, and there are no guarantees about each transaction. It's up to the devices at either end, outside of the network, to construct all other aspects of the communication. The network is as minimal and efficient as possible.

If anything about the Internet could be termed "revolutionary," it is this concept of removing functionality from the network itself and shifting that functionality to the computers that are collectively the sources and termination points of all these packets. But what particularly characterises IP is the extent of this shift. Compared with other very generic packet switching protocols of the 1970's and 1980's, including OSI, X.25, Frame Relay and ATM, the position taken by IP was extreme. It was extreme in terms of removing so much functionality from the network function that what was left was incapable of making any guarantees at all about packet transmission. No only could individual IP packets be dropped in transit, but a stream of IP packets could be delayed by variable periods, re-ordered, or even fragmented on the fly. The result of this network architecture was an extremely simple network that was incredibly cheap. The switching units for IP packets had no need to operate with expensive clocks, and no need to maintain local state for every active traffic flow passing through the switch. There were no network admission functions, or transaction setups, or even transactional states generated within the network. This shift of functionality was one that shifted the onus of service support from the network to the devices at the edge of the network, and even further, a shift of making a service an outcome of the operation of the network to a service that is an outcome of an application that is run on end users' systems. This is at the heart of why the Internet heralded such a significant, even revolutionary, change in the telecommunications environment. It really did turn everything we thought we understood about networking inside out.

This stripping out of network functionality in the Internet has profound implications in terms of the business models of interaction of network operators. In telephony the "transaction" was visible to the network operator as a "call", which in turn created a resource reservation state in the network. These "calls" were asymmetric, with a "calling party" at one end of this virtual circuit, and a "called party" at the other. Calls also have duration that is visible to the network operators. These characteristics formed the basis of the tariff models for the telephone network's subscribers, where callers were charged for the duration of calls, at a rate that reflected the distance between the two parties. This common tariff structure was reflected in the financial arrangements made between telephone operators, which formed the basis of inter-provider settlement of caller payments via the call accounting financial settlement models that are described in the ITRs.

In a stateless datagram transmission network, a "call" is an abstract concept managed by an application, and there is no counterpart in terms of a transaction that is visible at the network level. At the level of IP packets, the level where the networks themselves operate, there are no calls, no calling party, and there is no duration or any clear concept of distance. There is not even any knowledge of the application that generated the packet. A generic tariff concept such as "caller pays" simply has no analogy at the network level of the datagram Internet. This implies that the concept of inter-operator call accounting at a network level also has no counterpart in Internet networks. The larger significance is that some concepts previously thought to be generic concepts, which spanned all telecommunications media in 1988, were in fact concepts that were implicitly tied to a particular family of network architectures.

In revisiting the ITRs today, and attempting to update them, it may be tempting to try to integrate the Internet into the ITRs through an editorial process. Proposals to this end have been presented, simply inserting "and the Internet" to the provisions of the existing document, while other proposals advocate extension of existing definitions in the document to include Internet concepts which appear to deserve

similar treatment. There are also proposals to add specific references to particular Internet-related activities, such as packet routing or data caching. Some proposals also attempt to bring some of the more poorly defined concepts that have not enjoyed any widespread deployment in the public Internet, such as Quality of Service mechanism concepts, into the regulatory domain. Such efforts at incremental editing run the risk of producing a regulatory framework that is a sequence of compromises that is internally inconsistent and largely disconnected from today's reality, let alone being applicable to tomorrow's reality.

Aside from the question of merging a fundamentally new environment with a framework built on one that is undeniably outdated, there are other difficulties appearing in the current early-stage negotiations on the ITRs. The most visible of these is an apparent lack of common motivation and common perspective regarding the ITRs, the Internet, and telecommunications regulation more generally. Some nations that have experienced a long-term erosion of revenue streams from international telephony financial settlements may understandably seek some relief in a new set of ITRs. However there are also national telecommunications environments that, via a process of progressive liberalization of their domestic telecommunications markets, have experienced large-scale adoption of these new computer-mediated communications services. These services are layered upon a common data substrate of the Internet, and these national economies have realized large-scale economic benefits from such changes in their telecommunications environment. Understandably, such nations may be very reluctant to impose additional regulatory-inspired overhead or inefficiencies onto what they regard as a highly beneficial environment. If no commonality of purpose can be found in this diversity of national interests, then it seems a challenging objective to define common "solutions" through changes to a single common instrument, that is, the ITRs.

It may be prudent to consider how to avoid a committee outcome that demonstrates potentially the poorest outcomes of such a process: a pastiche of compromises that neither satisfies nor offends anybody, but at the same time contains such a volume of internal inconsistencies that its utility is completely compromised.

To avoid such an outcome, one possible approach is to maintain a perspective of the ITRs as an aspirational document that expresses common expectations and desires about worldwide telecommunications in a succinct manner. Perhaps it would be better to aim for a much smaller body of text that simply states a common desire to promote the development of telecommunications services and their most efficient and beneficial operation, while harmonizing the continued development of facilities that enable worldwide telecommunications.

P.S. And if you did not quite understand the reference in the title of this article to the writings of William of Ockham (c. 1288 – c. 1348), an English Franciscan friar and scholastic philosopher, then http://en.wikipedia.org/wiki/Occam's_razor is as good a reference as any as to the merits of economy and siccinctness. Less is indeed often better.

Disclaimer

The views expressed are the author's and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.

About the Author

Geoff Huston B.Sc., M.Sc., has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and has been active in the Internet Engineering Task Force for many years.

www.potaroo.net