The ISP Column
A monthly column on things Internet

Geoff Huston
October 2015

Some Thoughts on the Open Internet

I'm sure we've all heard about "the Open Internet." The expression builds upon a rich pedigree of term "open" in various contexts. For example, "open government" is the governing doctrine which holds that citizens have the right to access the documents and proceedings of the government to allow for effective public oversight, a concept that appears to be able to trace its antecedents back to the age of enlightenment in 17th century Europe. There is the concept of "open society," a theme that was developed in the mid 20th century by the Austrian philosopher Karl Popper. And of course in the area of technology there was the Open Systems Interconnection model of communications protocols that was prominent in the 1980's. And lets not forget "Open Source," which today is an extremely powerful force in technology innovation. So we seem to have this connotation that "open" is some positive attribute, and when we use the expression of the "Open Internet" it seems that we are lauding it in some way. But in what way?

So let's ask the question: What does the "Open Internet" mean?

> The Federal Communications Commission of the United States has published its views on this question:
>
> 'The "Open Internet" is the Internet as we know it. It's open because it uses free, publicly available standards that anyone can access and build to, and it treats all traffic that flows across the network in roughly the same way. The principle of the Open Internet is sometimes referred to as "net neutrality." Under this principle, consumers can make their own choices about what applications and services to use and are free to decide what lawful content they want to access, create, or share with others. This openness promotes competition and enables investment and innovation.
>
> 'The Open Internet also makes it possible for anyone, anywhere to easily launch innovative applications and services, revolutionizing the way people communicate, participate, create, and do business—think of email, blogs, voice and video conferencing, streaming video, and online shopping. Once you're online, you don't have to ask permission or pay tolls to broadband providers to reach others on the network. If you develop an innovative new website, you don't have to get permission to share it with the world.'
>
> *http://www.fcc.gov/openinternet*

The FCC's view of an "Open Internet" appears to be closely bound to the concept of "Net Neutrality," a concept that attempts to preclude a carriage service provider from explicitly favouring (or disrupting) particular services over and above any other.

Wikipedia offer a slightly broader interpretation of this term that reaches beyond carriage neutrality and touches upon the exercise of technological control and power.

> "The idea of an open internet is the idea that the full resources of the internet and means to operate on it are easily accessible to all individuals and companies. This often includes ideas such as net neutrality, open standards, transparency, lack of internet censorship, and low barriers to entry. The concept of the open internet is sometimes expressed as an expectation of decentralized technological power, and is seen by some as closely related to open-source software."
>
> *http://en.wikipedia.org/wiki/Net_neutrality#Open_Internet*

In this essay I'd like to expand upon this theme of openness and extend it to include considerations of coherence within the Internet and also consider fragmentary pressures. I'd like to see if we can provide a considered response to the question: Is today's Internet truly "Open?"

What does the "Open Internet" mean?

Let's examine the attributes of an "Open Internet" through the lens of its component technologies. The questions being addressed here for each of these major technology activities that support the Internet are: What would be the expectations of an "Open Internet"? What would a truly open and coherent Internet look like?

The technology model used here is an adaption of the earlier Open Systems Interconnection reference model (ISO/IEC 7498-1), where each layer of this reference model uses services provided by the layer immediately below it, and provides services to the layer immediately above it. It is a conventional taxonomy for networking technologies. An Internet-specific protocol reference model is shown in Figure1, based on the technology model used in RFC1122 ("Host Requirements" RFC)

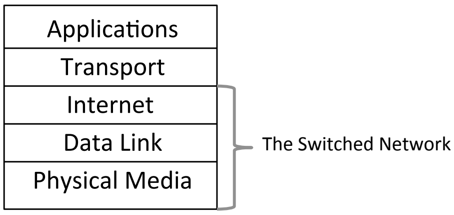

Figure 1 – A Protocol Reference Model for the Internet (after RFC 1122)

The concept of "Openness" as applied to networks carries an obvious connotation of accessibility. This is not the same as a free service, but it is a service where access to the service is not limited or restricted in arbitrary ways. An open network is an accessible network. This concept of general accessibility encompasses more than accessibility as a potential consumer of the network's service. It also implies that there are no inherent restrictions or arbitrary inhibitions for anyone to provide services, whether it's services as a provider of transmission capacity, switching, last-mile access, mobility, names, applications, of any of the other individual components that make up the Internet. This concept of openness also extends to the consequent marketplace of services that exists within this networked environment. Consumers can make their own choices about the applications and services that they choose to use in such an open network. The environment promotes competition in the supply of goods and services, and stimulates investment in innovation and development that provides evolutionary pressure to expand and diversify the ways in which we make use of this common network.

Such outcomes are the result of the application of the same set of overall principles into each of the areas of technology that form the essential components of the Internet.

An Open Switched Network

The theoretical model of an open and coherent network is a re-statement of an interpretation of the end-to-end principle in packet-switched networks, where the network's intended role is strictly limited to the carriage of individual packets from source to destination, and all users, and all functions and services that populate the network, are provided by devices that sit outside of the network itself. These devices communicate between themselves in a manner that is largely opaque to the packet-switched network. Furthermore, edge devices are not expected to communicate with packet switching devices within the network, and equally, packet switching devices within the network do not directly communicate with edge devices (with the one exception of the generation of Packet Control messages (ICMP message in the context of the Internet Protocol). Network consistency implies that all active (packet switching) elements within the network that perform switching functions in IP packets use a consistent single interpretation of the contents of an IP packet, supporting precisely the same IP protocol specification.

> The seminal work on the end-to-end principle is the 1981 paper "End-to-End Arguments in System Design" by J.H. Saltzer, D. P. Reed, and D. D. Clark, published in Proceedings of the Second International Conference on Distributed Computing Systems. Paris, France. April 8–10, 1981. IEEE Computer Society, pp. 509-512.
>
> A simple restatement of the end-to-end principle is that the network should not replicate the functions that can be performed by communicating end systems.
>
> A further paper on the topic is "Tussels in Cyberspace: defining Tomorrow's Internet" by D. D. Clark, K. R. Sollins, J. Wroclawski and R. Braden, published in SIGCOMM'02, August 19-23, 2002.
>
> A restatement of this paper's thesis is that in an unbundled environment each actor attempts to maximize their role and value, and when applied to networks and applications they may create conflicting situations, which undermines a pure end-to-end network design.

The Internet Protocol has chosen a particular form of packet switching which is a "stateless" switching function. Within each active switching element each IP packet is forwarded towards its intended destination without reference to any preceding or following packets, and without reference to any pre-configured state within the switching element. This implies that every IP packet contains a destination address that is not relative to any assumed network state or topology, and that this address has an identical unique interpretation across the entire network domain. In this way IP addresses are not relative to any locality, network or scope, and each IP address value is required in this architecture to be unique across the entire Internet.

> This is a claim more honored these days as an exception rather than general practice. Network operators have eschewed passing all responsibility for packet transmission to end points and have responded by constructing internally segmented networks that rely on various forms of virtual state within the network. This extends from the extensive use of VLANs in ether-switched data services to the almost ubiquitous use of MPLS in wide network networks. The current enthusiasm SDN is no exception to this bias towards the use of virtual circuits within networks.

A Open Consistent Address Space

In an open and consistent Internet every destination on the Internet is reachable from any location on the Internet. The way this is achieved is the universal ability to send a packet to any destination, and this implies that all such destinations require an IP address that everyone else may use. These IP addresses must be allocated and administered such that each address is uniquely associated with a single attached network and with a single attached device within that network. The network itself cannot resolve the inconsistency of address clashes where two or more devices are using the same address, so the responsibility for ensuring that all addresses are used in a manner that is unique is left to the bodies who administer address allocation and registration.

> This has been an evolutionary process. The original address administration and registry function was managed through the US research agencies, and the evolution of this model has lead to the creation of five "Regional Internet Registries" each of which serve the address allocation and registry function needs of regional communities. The administration of the central pool of unallocated addresses is part of the IANA function. The policies that govern the administration of the distribution and registration functions within each of these regional registries are determined by the regional communities themselves, in a so-called "bottom-up" self regulatory manner.

The practices relating to access to address space through allocation and assignment are based on policies developed by the respective address communities in each region. The general theme of these address distribution policies is one of "demonstrated need" where addresses are available to applicants on the proviso that the applicant can demonstrate their need for these addresses within their intended service infrastructure.

Open End-to-End Transport

The service model of a stateless packet switched network is one of unreliable datagram delivery. This service model is inadequate for most useful network services. The Internet has commonly adopted a single end-to-end stream protocol, the Transmission Control Protocol (TCP), that is conventionally used by communicating end systems to transform the network's unreliable datagram delivery service into a reliable lossless byte stream delivery service.

This is not the only end-to-end transport protocol in common use. Another protocol, the User Datagram Protocol (UDP), is a minimal abstraction of the underlying IP datagram behavior, commonly used by simple query/response applications, such as the DNS resolution protocol.

While many other transport protocols have been defined, common convention in the Internet has settled on TCP and UDP as the two "universal" end-to-end transport protocols, and all connected systems in an open coherent network would be expected to be able to communicate using these protocols. The uniform adoption of end-to-end transport protocol behaviors is a feature of such an open network, in that any two endpoints that both support the same transport protocol should be able to communicate using that protocol. In this open network model, the operation of these end-to-end protocols is completely opaque to the packet-switched network, as it concerns only the communication signaling between the two end systems.

This perspective of the end-to-end protocols in use in the Internet also makes a critical assumption about the nature of the flow control processes. This model assumes that TCP is the predominate protocol used by end hosts and, most critically, that the flow control algorithm, used by all TCP implementations, behaves in very similar ways. This model assumes that there is no central method of allocation or governance of network resource allocation to individual end-to-end conversation flows, and instead the model relies on the aggregate outcome of the TCP flow control protocols to provide a fair share allocation of common network resources where an approximately equal proportion of network resources is utilized by each active conversation flow.

> The conventional flow control process is one of additive increase in flow rates (slow) and multiplicative decrease (fast), or "AIMD". TCP sessions have no arbitrary speed settings, and each TCP session will both impose pressure on other concurrent sessions and respond to pressure from other concurrent sessions to try and reach a meta-stable equilibrium point where the network's bandwidth is, to some level of approximation, equally shared across the concurrent active flows.
>
> Packet loss is the signal of over-pressure, so a flow will gradually increase its sending rate to the point of onset of packet loss, and at that point it will immediately halve its sending rate and once more gradually probe increased rates until the next packet loss event.
>
> TCP implementations that use a different flow control algorithm normally fare worse, as their efforts to place greater flow pressure on concurrent flows often results in higher packet loss rates in their own flows. However, there has been a significant body of research into flow control algorithms and there are TCP flow control algorithms that can secure a greater relative share of the network over a conventional AIMD flow control algorithm without the element of self-damage. These algorithms are capable of exerting "unfair" pressure on other concurrent TCP flows, and can consume a greater proportion of network resources as a result.
>
> One aspect of the "network neutrality" debates is the assumption of a relatively passive network where the network's resources will be equitably allocated due to the general fair-shared outcome that is achieved by the uniform use of particular TCP flow control behaviour. The TCP ecosystem is changing with entrants such as Akami's use of FAST, Google's use of QUIC with Chrome and some Linux distributions using CUBIC, and these assumptions about the general equity of outcome of competing end-to-end streaming sessions are now an increasingly approximate set of assumptions.
>
> *"TCP Protocol Wars", http://ipj.dreamhosters.com/wp-content/uploads/2015/07/ipj18.2.pdf*

A Open Consistent Name Space

This open and coherent model of the Internet is not limited to the network packet switching and end-to-end transport functions. A critical component of the Internet is implemented as a distributed application that sits alongside clients and servers at the "edge" of the network rather than within the network's direct purview. This is the Internet's symbolic name space, the Domain Name System (DNS).

This name space is the combination of a name structure and a name resolution function that allows a user level discourse using familiar symbols and terms to refer to service points connected to the Internet that are identified by IP addresses and transport protocol port numbers.

While it is conceivable to think about many diverse name spaces and even many diverse name resolution protocols, and the Internet as such would not necessarily prevent such an outcome, a coherent view of the Internet requires that the mapping of symbols to IP addresses follows a uniform and consistent convention across the entire network. Irrespective of where and how a DNS query is generated, the response should reflect the current state of the authentic information published in the DNS. The implication here is that an open and consistent DNS uses the hierarchical name space derived from an single and unique root zone, and that all name resolvers perform the resolution of a name query using a search within this same uniquely rooted name space. This is the essential element of a consistent name space for all of the Internet.

Open Applications

The context and content of individual conversations in this open coherent network model is also the subject of a number of common conventions, where certain commonly defined application level protocols are defined for common services. For example, applications wishing to pass email messages are expected to use the SMTP protocol, the retrieval of web pages to use the HTTP protocol, and so on.

This implies that the protocols used to support network-wide functions, including for example data transfer, electronic mail, instant messaging, and presence notification, all require the adoption of openly available protocol specifications to support that application, and that these specifications are openly implementable and not encumbered by restrictive claims of control or ownership.

Much of today's environment also relies heavily on the concept of "open source" technologies. The Unix operating system, originally developed at AT&T Bell Labs in the 1970's and distributed as open source, is now the mainstay of the much of today's environment. The implementation of the TCP/IP protocol suite by the Computer Systems Research Group at the University of Berkeley in the 1980's was made available as open source and the ready availability of this software package was part of the reasons behind the rapid adoption of this protocol as the common computer networking protocol in the 1990's. Subsequent "open" implementations of popular applications, such as sendmail for Mail, BIND for the DNS, Apache for Web servers, added further momentum to this use of open source, and these days the concepts of open source is fundamental to much of the technology base of not only the Internet but to the entire information technology world.

Open Security

Security functions include both the open and unrestricted ability for communicating end users to invoke protection from third party eavesdropping and the ability for these end users to verify the identity of the remote party with whom they are communicating, and to authenticate that the communication as received is an authentic and precise copy of the communication as sent. This is useful in many contents, such as for example in open communications environments using the radio spectrum, or in environments that trade goods and services where authentication and non-repudiation is vitally important.

To allow such functions to be openly available to all users requires the use of unencumbered cryptographic algorithms that are generally considered to be adequately robust and uncompromised, and the associated availability of implementations of these algorithms on similar terms of open and unencumbered availability.

An Open Internet

One view of an open network is a consistent network, in that the same actions by a user will produce the same response from the networked environment, irrespective of the user's location and their choice of service provider. In other words, the interactions between the application on the user's device and the application that serves the referenced content should not be altered by the network in any way, and users should see identical outcomes for identical inputs across the entire network.

These considerations of the prerequisites of an open coherent Internet do not imply the requirement for an Internet that is operated by a single operator, or one where services are provided via a single service delivery channel or technology provided through a single channel. While the Internet is an amalgam of tens of thousands of component networks, populated by millions of services, and services by thousand of suppliers of services and technologies it is still feasible that this collection of service providers are individually motivated follow common conventions, and to operate their component services in a fashion that is consistent with all other providers. The property of coherence in an open internet is an outcome of individual interests to maximize their effectiveness and opportunities by conforming to the common norms of the environment in which they operate.

An Open Internet is not one where open access equates to costless access. The considerations of openness in such a model of an open network relate to the absence of arbitrary barriers and impositions being placed on activities.

What these considerations imply is the ability to evolve the Internet through incremental construction. A novel application need not require the construction of a new operating system platform, or a new network. It should not require the invention and adoption of a new network protocol or a new transport protocol. Novel applications can be constructed upon the foundation of existing tools, services, standards and protocols. This model creates obvious efficiencies in the process of evolution of the Internet.

The second part of the evolutionary process is that if a novel application uses existing specifications and services then all users can access the application and avail themselves of its benefits if they so choose. Such an open unified environment supports highly efficient processes of incremental evolution that leverage the existing technology base to support further innovation. The process of evolution is continual, so it is no surprise that the Internet of the early 1990s is unrecognizable from today's perspective. But at the same time today's Internet still uses the same technology components from that time, including the IP protocol, the TCP and UDP end-to-end transport protocols, the same DNS system, and even many of the same application protocols. Each innovation in service delivery in the Internet has not had to reinvent the entire networked environment in order to be deployed and adopted.

Much of the Internet today operates in a way that is consistent with common convention and is consistent with this model of an open, unified and accessible public resource. But that does not mean that all of the Internet environment operates in this manner all of the time, and there are many fragmentary pressures.

Such pressures appear to have increased as the Internet itself has expanded. These fragmentary pressures exist across the entire spectrum of technologies and functions that together make up the internet.

Some of these fragmentary pressures are based in technology considerations, such as the use of the Internet in mobile environments, or the desire to make efficient use of high capacity transmission systems. Other pressures are an outcome of inexorable growth, such as the pressures to transition the Internet Protocol itself to IPv6 to accommodate the future requirements of the Internet of Things.

There are pressures to increase the robustness of the Internet and improve its ability to defend itself against various forms of abuse and attack.

How these pressures are addressed will be critical to the future of the concept of a coherent open Internet in the future. Our ability to transform responses to such pressures into commonly accepted conventions that are accessible to all will preserve the essential attributes of a common open Internet. If instead we deploy responses that differentiate between users and uses, and construct barriers and impediments to the open use of the essential technologies of the Internet then not only will the open Internet be threatened, but the value of the digital economy and the open flow of digital goods and services will be similarly impaired.

The Where and How of "Internet Fragmentation"

In defining what is meant by "Internet Fragmentation" it is useful to briefly describe what is meant by its opposite, an "Open and Coherent Internet". As we've explored in the previous section, "coherence" implies that each of the elements of the Internet are orchestrated to work together to produce a seamless Internet which does not expose the boundaries between discrete elements. Coherence also implies consistency, in that the same trigger actions by a user produce the same response from the networked environment, irrespective of the user's location and their choice of service provider. Openness also implies the ability to integrate and build upon existing tools, technologies and services to create new technologies and services, and in turn allow others to further evolve the technology and service environment.

"Fragmentation" on the other hand encompasses the appearance of diverse pressures in the networked environment that leads to diverse outcomes that are no longer coherent or consistent. In the context of the Internet, fragmentation also encompasses various ways in which openness is impaired, and also can include consideration of critical elements of service and the fragility of such arrangements when the supply of such services is left to a very small number of providers.

This section contains some notes on where and how there are fragmentary pressures that are driving apart aspects of the Internet and create various "islands" of differentiated functionality and connectedness. It concentrates on the technical aspects of these pressures for fragmentation and does not attempt to analyse public policy implications.

IP level Fragmentation

The issues around address exhaustion in IPv4 and the transition IPv6 deserve attention in relation to any discussion of potential Internet fragmentation.

The transition to IPv6 is still a process without clear coherence or assured outcomes. It is possible that the work undertaken already by a relatively small number of retail Internet access providers, including notably large ones such as AT&T, Comcast, Deutsche Telekom and KDDI, will generate sufficient impetus in the market to pull both content providers and other ISPs along with them in embarking on IPv6 services. This is, however, by no means an assured outcome, and the continued expansion of Network Address Translators (NATS) in IPv4 Internet appears to have no immediate end in sight. The market signals are as yet unclear and the public policy actions have not yet provided adequate impetus, with the result being that the general response from the majority of players has been insufficient to make any real progress in trying to shut down the use of IPv4 in the Internet.

Due to the address exhaustion of IPv4, increased use of Carrier Grade NATs is being made to share this scarce address resource across a greater number of users In other words, address exhaustion for IPv4 is creating larger and larger networks of "semi-opaque" connectedness within the public network. IPv4 addresses used in conjunction with NATs no longer have a clear association with a single end user and the most probable outcome is that parts of the net will "go dark" in the sense that user's actions within this "dark" network are effectively untraceable. These devices also compromise other aspects of robustness in the engineering of the Internet at this level of operation. The requirement to pass all traffic to and from an external site through the same address translation unit impairs some forms of robust network operation that uses diverse points of interconnection and diverse connectivity, and instead this form of state-based middleware creates critical signal points of failure. Given the critical importance of content delivery in many networks, the presence of CGNs creates incentives to place selected content distribution functions on the "inside" of the CGN. This runs risks of the network discriminating between various content delivery systems through this ability to position some content in an advantaged position as compared to others. The longer-term pressures are difficult to discern at this stage, but the longer this hiatus in addresses lasts the greater the levels of address pressure. The greater the address pressure on the IPv4 network the greater the fragility and complexity of networks using address sharing.

Another side effect of IPv4 address exhaustion is address trading. This market has appeared organically and there is growing evidence that transferred IPv4 addresses are not all being registered in the established address registries. Some of this is evidently due to address "leasing" where the lessee is not registered at the current beneficial user of the address, but also some times due to a reluctance of the address holder to enter the address into the address registry because of concerns over address title policies or similar concerns for the parties involved. The larger the pool of unregistered addresses the greater the pressure to fracture the address space. There is no clear way back when or if the space fractures in this manner.

With the exhaustion of the address allocation framework for IPv4 and the established common belief that addresses are plentiful in IPv6, then much of the original rational for the regional address registry structure is weakened.

> Much of the original rationale for the regional internet address distribution framework lay in the perceptions of scarcity in supply of addresses in the IPv4 address plan, and the need to perform a complex rationing operation. The clearly finite pool of addresses and the larger visions of the Internet's future implied that it was not possible to simply allocate an adequate pool of addresses to each network operator to meet perceived needs, and instead each regional registry devised a rationing scheme based around the principle of "demonstrated need". The original objective of this process was to ration the consumption of addresses in IPv4 until such time as IPv6 was prevalent, and there was no further need for IPv4 addresses. Without the need for further rationing and its associated administrative overhead, and a reversion to a potentially far simpler registry model then the case for regional fragmentation of the registry function is an open question.

However, not all of the pressures in this space are directed towards aggregation of the registry function into a single operation. When coupled with a cyber security perspective that its "good to know where every address is in a country" its reasonable to anticipate further pressure to further fracture the regional structures into national structures. In the Asia-Pacific region, APNIC already has China, India, Indonesia, Korea, Japan, Taiwan and Vietnam all operating such national address registries, and in Latin America there are comparable structures in Brazil and Mexico. It is an open question whether this will spread in response to these pressures of national security and the effective end of the conservative address allocation function.

Routing Fragmentation

The routing system is intended to ensure that every switching element is loaded with consistent information such that every attached device on the Internet is reachable by any other device. The Internet uses a two level protocol routing hierarchy. The set of local routing domains (or "Autonomous Systems" (AS's)) use a variety of routing protocols. As they do not directly interact with each other this is not an issue at all. The second routing domain (or the "Inter-Domain" space) uses a single routing protocol called the Border Gateway Protocol (BGP).

The protocol BGP, and the broader Internet routing space, is under various pressures.

The AS identification field was defined as a 16 bit number field. The Internet community is close to exhausting this identifier space, and needs to move to a larger 32 bit field. Over the past 20 years the problem has been identified, technical standards produced, software has been deployed by vendors, the transition strategy defined, and the process has been started. In Europe, the process is well under way, while in North America (Canada and United States) the process has stalled and almost no 32 bit ASs are in use. The subtle difference is AS-specific communities appears to be the issue here. The Canadian and United States ISPs appear to make use of these AS-specific communities for routing policy, and are reluctant to use 32 bit AS numbers for this reason. The European ISPs appear to make more use of routing registries to describe routing policies, and these registries are largely agnostic over the size of the AS number field. It is unclear how the North American ISPs are going to resolve their issues given that the 2 byte AS number pool will be exhausted in the coming months.

The routing system is under constant pressure from false routing advertisements. Some of these are local scope advertisements intended to comply with national directives to filter certain IP addresses as part of a national content filtering directive. Some are the result of mistakes in router configuration. Others are deliberately malicious advertisements designed to redirect traffic in some manner, or to disrupt the genuine service. Efforts to improve the security of the routing system are being explored by the Internet Engineering Task Force (IETF), but the measures being contemplated imply additional overheads in routing, increased complexity and increased brittleness. The security is most effective when the entirety of the routing space adopts the technology, and balancing the local costs against the common benefit that is contingent on universal adoption is a significant issue for this approach.

The routing space is not a uniform space, and different address blocks are visible in different parts of the Internet, and there is often no clear reason why. There are "ghost routes" where the original withdrawal has not successfully promulgated across the entire network and some networks are still carrying reachability information for the old route. There are "islands" of more specific routes, which are blocked from universal promulgation by various prefix length filters being applied by various routers. There is the selective absence of routing information because some routing domain use 'Route Flap Damping' and others do not. Local routing policies may apply differential treatment to various routes, such as is seen in the distinction between transit, peering and customer relationships as implemented in the routing space. The result is that there is no clear consensus as to what constitutes "the Internet" in a routing sense. Each AS appears to see its own routing view of the Internet and there are invariably some number of subtle distinctions between each of these local views. The result is that it is not assured that every single end point on the Internet can send a packet to any other connected end point at any time. In some cases the routing system simply does not contain the information to support this universal form of connectivity. Though for most means and purposes the overwhelming majority of all end points can be reached.

ISP Peering and Transit

The Internet architecture does not expose user level transactions to the network and inter-network arrangements are not based on transaction accounting. At its heart, every user of the Internet pays for his or her own access. In turn, their ISP undertakes to provide connectivity to the rest of the Internet. It cannot do this without the support of other ISPs. Two ISPs that interconnect and exchange traffic typically do so under one of two broad models. One model is the transit relationship, where one party pays the other, and in return is provided with all of the routes available on and via the other network. The transit model is used in open competitive markets when the interconnection is perceived as being asymmetric and one party has significant assets, and the other party is wanting to access those network assets. There is no particular assurance from this model that a customer of a transit provider necessarily sees the entirety of the Internet.

The typical transit arrangement is that the customer is given access to the route set controlled the transit service provider that the ISP cannot obtain as efficiently by any other means. The other broad model is a peering model where neither party pays the other, and each party learns only the customer routes of the other. Through the use of peering, ISPs can reduce their transit costs, as they do not need to purchase transit for that traffic. To save interconnection costs ISPs establish or make use of Internet Exchange Points (IXPs), where they can peer with multiple networks at the same time. The peering model is often seen in open competitive market situations where the two providers bring, in each party's perception, approximately equal assets to the connection, so neither party believes that there is undue leverage of the investments and assets of the other. Peering arrangements are at times challenging to sustain. Some networks grow and want to change their position to a seller of transit connectivity. They may then opt to de-peer some networks in order to force them to become customers. There are also various hybrid approaches that combine peering of customer networks with the option of also purchasing a transit service. For example, Hurricane Electric has an open peering policy, while at the same time selling an optional transit service to the networks it peers with.

The market-based approach to connectivity as represented by this model of interconnection is efficient, and relatively flexible, and it embraces entirety large proportion of the inter-provider relationships in the Internet. Its divergence from the model supported by telephony is still a source of continuing tension in certain international circles. Efforts by certain countries to assert some form of paid relationship by virtue of their exclusive role of access to a national user base have, in general, been relatively self-harming in terms of a consequence of limited external visibility on the part of the national user community that was used in such negotiations. Nonetheless, where commercial negotiations do take place and in the absence of sufficient competition, one player may leverage their position to endeavor to extract higher rents from others. In those instances, and the reason why the Internet has become so successful in competitive markets, ISPs then have the option to bypass each other using transit if they find that more economical.

Other tensions have appeared when the two parties bring entirely different assets to a connection, as is the case with Content Distribution Networks connecting with Internet Access Providers. One party is bringing content that presumably is valued by the users, the other party is bringing access to users that is vital for the content distribution function. Whether or not a party can leverage a termination monopoly in such situations depends on the competitive market situation of the location it operates in. For example Free in France for a while demanded paid peering from Google and would not upgrade saturated interconnects, but in 2015 upgraded its peering with Google without receiving payment.

Name Space Fragmentation

The name space has been under continuous fragmentation pressure since its inception.

The original public name space has been complemented by various locally scoped name spaces for many years. As long as the public name space used a static list of top level domains the private name was able to occupy unused top level name spaces without serious side effects. The expansion of the gTLD space challenges this assumption, and the collision of public and private name spaces leads to the possibility of information leakage.

Other pressures have, from time to time, taken the form of augmenting the public root with additional TLD name spaces through the circulation of "alternate" root servers. These alternate roots generate a fractured name space with the potential for collision, and as such have not, in general, been sustaining. The problem with these alternate systems is that a name that refers to a particular location and service in one domain may refer to an entirely different location and service in another. The "use model" of the Internet's user interface is based on the uniqueness of a domain name, and in fact based on the graphical representation of a domain name on a user's device. So if a user enters a network realm of an alternate root name space where a previously known and trusted domain name is mapped to a different location, then this can be exploited in many ways to compromise the user and the services they use. The confidence of users and the trust that is placed in their use of the Internet is based on a number of premises, and one of the more critical premises is that the Domain Name space is consistent, unfragmented and coherent. This premise is broken when alternate root systems are deployed.

As well as these fragmentary pressures driven by an objective to augment the name space in some fashion, there are also pressures to filter the name space, so that users are unable to access the name information for certain domain names. In some cases these are specific names, while in other cases it has been reported that entire TLD name spaces are filtered.

> It has been observed that the TLD for Israel, .il, is filtered in Iran, such that no user in Iran is able to resolve a DNS name under the .il TLD.
>
> *http://www.potaroo.net/reports/2015-07-GTLD-Universal-Acceptance-report-v1.pdf*

The resolution process of the DNS is also under pressure. The abuse of the DNS to launch hostile attacks has generated pressures to shut down open resolvers, to filter DNS query patterns and in certain cases to alter a resolver's responses to force the query to use TCP rather than DNS.

This abuse has also highlighted another aspect of the DNS, namely that for many service providers the operation of a DNS resolution service is a cost centre rather than a generator of revenue. With the advent of a high quality high performance external DNS resolution services in the form of Google's Public DNS, the Open DNS resolver system and Level 3's long standing Open DNS Resolver service, many users and even various ISPs have decided to direct their DNS queries to these open resolver servers. Such a response has mitigated the effectiveness of local name filtering, and at the same time allowed these open providers to gain substantial market share of the Internet's DNS activity.

> Google recently noted that "Overall, Google Public DNS resolvers serve 400 billion responses per day."
>
> *http://googlewebmastercentral.blogspot.fr/2014/12/google-public-dns-and-location.html*

Not only is the DNS protocol enlisted to launch attacks; the DNS itself is under attack. Such attacks are intended to prevent users from obtaining genuine answers to certain queries, and instead substituting a deliberately false answer. The DNS itself does not use a "protected" protocol, and such substitution of "false" answers is often challenging to detect. This property of the DNS has been used by both attackers and state actors when implementing various forms of DNS name blocking. Efforts to alter the DNS protocol to introduce channel security have been contemplated from time to time and have some role in certain contexts (such as primary to secondary zone transfers) but they have not been overly effective in the area of resolver queries. Another approach is to allow the receiver of a response to validate that the received data is authentic, and this is the approach behind DNSSEC. Like many security protocols, DNSSEC is most effective when universally adopted, in that at that point any attempts to alter DNS resolution responses would be detectable. With the current piecemeal level of adoption, where a relatively small number of DNS zones are signed (and even where the zones are signed, DNSSEC uptake at the domain name level is vanishingly small – even amongst banks, large-scale ecommerce providers), then the value of this security approach is significantly smaller than would be the case with general adoption.

The contents of the DNS are also under some pressure for change and there are various ways that applications have chosen to handle them. This is evident in the introduction of scripts other than ASCII (so-called "Internationalized Domain Names," or "IDNS"). Due to a concern that DNS implementations were not necessarily 8-bit clean, the introduction of DNS names using characters drawn from the Unicode character space required the application to perform a transform of the original unicode string to generate an encoded string in the ASCII character set that strictly obeyed the "letter, digit, hyphen" rule. Similarly the application is required to map back from this encoded name to a displayed name. The integrity of the name system with these IDN components is critically dependent on every application using precisely the same set of mappings in their applications. While this is desirable, it is not an assured outcome. A second issue concerns the "normalisation" of name strings. The ascii DNS is case-insensitive, so that query strings are "normalized" to monocase when searching for the name in the DNS. The issue of normalised characters from non-ascii scripts presents some issues of common use equivalence by communities of users of a particular language, and what may be regarded as equivalent characters by one community of users of a given language may not be equivalently regarded by users of the same language. In recent years, engineers and linguists within the ICANN community have been working towards a common set of label generation rules, and have been making real progress. This is a particularly complex issue in the case of Arabic script which has many character variants (even within individual languages) and uses some characters which are not visible to humans (zero-width joiners). The DNS is incapable of handling such forms of localisation of script use. Despite being available for 15 years, IDNs are still not working seamlessly in every context or application in which an ASCII domain is used. This issue is called "universal acceptance". Although it applies equally to other new gTLDs, it is far more complex and challenging to overcome in IDNs. Examples include IDN email addresses – while Google announced last year that Gmail will support IDN addresses, this only applies when both the sender and receiver both have Gmail accounts. IDN email addresses are not supported by any of the major application providers in the creation of user accounts (which often use email addresses as the user's unique account identifier), nor in digital certificates, DNS policy data or even many web browsers.

> In a test involving some 304 new generic top level domains a problem was observed with punycode-encoded IDNs where a combination of Adobe's Flash engine, the Microsoft Windows platform and either the Internet Explorer or Firefox browsers was incapable of performing a scripted fetch of an IDN. The problem illustrates the care needed in application handling where entirely distinct internal strings (in this case the ascii punycode and the Unicode equivalent) refer to the same object.
>
> *http://www.potaroo.net/reports/2015-07-GTLD-Universal-Acceptance-report-v1.pdf*

The fragmentation risk is that the next billions of Internet users who are not Latin-script literate – for example, 80% of India's 1.2 billion population is unable to speak English – will not be able to benefit from the memorability and human-usability of the domain name system. The problem has been masked to some extent by the demographics of Internet uptake to date, but is likely to become more apparent as the next billion comes online. Another possibility is that such populations will simply not use the domain name system. Uptake of domain name registrations (both ASCII and IDN) in Arab States and Islamic Republic of Iran is extremely low, and stands in stark contrast to the enthusiastic uptake of social network platforms (Egypt has 13 million Facebook users; Saudi's Twitter usage grew by 128% in 2013, to 1.8 million.

> Another issue with the use of IDNs concerns the "homograph" issue, where different characters drawn from different scripts use precisely the same display character glyph on users' screens. The risk here is of "passing off" where a domain name is registered with a deliberate choice of script and characters that will be displayed using the same character glyphs as a target name. This has led to different applications behaving differently when handling exactly the same IDN domain name. Some applications may choose to display the Unicode string, while others may elect to display the ascii encoding (Punycode) and not display the intended non-ascii string Unicode equivalent.
>
> *http://en.wikipedia.org/wiki/IDN_homograph_attack*
> *https://wiki.mozilla.org/IDN_Display_Algorithm*

Another quite different precedent has been created by the IETF when they opened up the "Special Use" domain name registry (RFC6761). There was a common convention that names that looked like DNS names were in fact DNS names, and were resolvable by performing a DNS query. Other name forms, and other non-DNS forms of name resolution, were identified principally through the URI name scheme, which used the convention of a scheme identifier and a scheme-defined value. The registration of the top level name ".onion" is anomalous in this respect in that the names under .onion are not DNS labels, and are not resolvable using a conventional DNS query. Similar considerations apply to names under the ".local" top level domain, which are intended to be local scope names resolved by multicast DNS queries. Such implicit scope creep of the DNS label space to encompass names that are not resolvable by the public DNS, yet otherwise resemble conventional DNS names, offers greater levels of potential confusion.

Application Level Fragmentation

Similar considerations apply at the application level, and there are tensions between maximizing interoperability with other implementations of the same application and a desire to advantage the users of a particular application.

> Apple's iMessage application used a similar framework to other chat applications, but elected to use encryption with private key material that allowed it to exchange messages with other Apple chat applications, but no other.

There is an aspect of fragmentation developing in the common network applications, typically over behaviors relating to mutual trust, credentials and security. For example, the Mail domain is a heavily fractured domain, predominately because of the efforts of the 'mainstream' mail community to filter connections from those mail agents and domains used by mail spammers. It is no longer the case that any mail agent can exchange mail with any other mail domain, and in many cases the mail agent needs to demonstrate its credentials and satisfy the other agent that it is not promulgating spam mail.

Similar pressures are emerging to place the entirety of application level interactions behind secure socket and channel encryption using the service's domain name as the key. It is possible that unsecured services will be increasingly be treated as untrustable, and there may be a visible line of fracture between applications that use security as a matter of course and those that operate in the clear.

Another potential aspect of fragmentation concerns the demarcation between the application, the host operation system, the local network and the ISP environment. The conventional view of the Internet sees this as a supply chain that invokes trust dependencies. The local network collects IP addresses and DNS resolver addresses from its ISP. The local network uses the ISP's DNS resolver and relies on the ISP to announce the local network's address to the Internet. Devices attached to the local network are assigned IP addresses as a local network function. They may also use the local network gate was a DNS resolver. Applications running on these devices use the network services operated by the system running on the device to resolve domain names, and establish and maintain network connections. Increasing awareness of the value of protection of personal privacy, coupled within increasing use of these devices in every aspect of users' lives and increasing use of these devices to underpin many societal functions implies increasing levels of concern about information containment and protection. Should a local network trust the information provided through the ISP, or should it use some other provider of service, such as a VPN provider, or a DNS resolution service? Should an attached device trust the local network and the other attached devices? Common vectors for computer viral infection leverage these level of trust within local networks, and services such as shared storage system and similar can be vectors for malware infections. Similarly, should an application trust its host? Would a cautious application operate with a far greater level of assured integrity were it able to validate DNS responses using DNSSEC directly within the application? How can the application protect both itself and the privacy of the end user unless it adopts a suitably cautious stance about the external environment in which it operates.

As computing capability becomes ever more ubiquitous and the cost, size and power of complex computing falls, there is less incentive to create application models that build upon external inputs, and instead incentives appear to draw those inputs back into the explicit control of the application, and use explicit validation of external interactions rather than accepting them on trust. This does not necessarily fragment the resultant environment, but it does make the interactions of the various components in this environment one that is far more cautious in nature.

Security Fragmentation

It has often been observed that security was an afterthought in the evolution of the Internet protocol suite, and there is much evidence to support this view. Many of the protocols and applications we use are overly naive in their trust models, and are ill-equipped to discriminate between authentic transactions and various forms of passing off and deception. None of the original protocols operated in a "private" mode, allowing eavesdroppers to have a clear view of network transactions. The routing system itself forms a large mutual trust environment where rogue actors are often difficult to detect. This is coupled with an environment where edge devices are similarly vulnerable to subversion, and the combination of the two has proved to be exceptionally challenging.

Today's situation appears to be that that the very openness of the network and its protocols is being turned against the Internet, and the level of cohesion of how to improve the situation is one that is still evolving and incomplete. Instead, it is evident that there is a certain level of fragmentation and diversity in the current endeavours in network security.

The domain name security framework is a good example of this. When the requirement emerged to be able to associate a cryptographic key pair with a service that was delivered from a given domain name, then the logical answer would have been to bind the public key to the domain name in the same fashion as an address was bound to a domain name, namely through the use of a DNS resource record and publication within the DNS. However, this was not possible at the outset due to the lack of mechanisms in the DNS to permit validation of a DNS response and the lack of a standard way of representing a public key in the DNS. The interim measure was the enrolling of a set of third party certification authorities who generated certificates that associated a given domain name to a particular public key. This set has expanded to many hundreds of these third party Certification Authorities (CAs), all of whom are trusted by end users for their security needs.

The central problem now is that the user does not know in advance which certification authority has issued a public key certificate for which name, so the user, and the subject of an issued certificate are both forced to trust the entire set of CAs. If any CA is compromised then it can be coerced into issuing a fake certificate for any domain name, compromising the privacy and integrity of the service offered by this domain name. This is not a framework that naturally induces high quality and integrity. The system is only as strong as the weakest CA, and the risks inherent in this framework does not lie in the choice of a particular CA to certify a domain name, but in the level of integrity of the entire collection of CAs.

> Google posted this entry on their online security blog in March 2015:
>
> "On Friday, March 20th, we became aware of unauthorized digital certificates for several Google domains. The certificates were issued by an intermediate certificate authority apparently held by a company called MCS Holdings. This intermediate certificate was issued by CNNIC.
>
> "CNNIC is included in all major root stores and so the misissued certificates would be trusted by almost all browsers and operating systems. Chrome on Windows, OS X, and Linux, ChromeOS, and Firefox 33 and greater would have rejected these certificates because of public-key pinning, although misissued certificates for other sites likely exist.
>
> "We promptly alerted CNNIC and other major browsers about the incident, and we blocked the MCS Holdings certificate in Chrome with a CRLSet push. CNNIC responded on the 22nd to explain that they had contracted with MCS Holdings on the basis that MCS would only issue certificates for domains that they had registered. However, rather than keep the private key in a suitable HSM, MCS installed it in a man-in-the-middle proxy. These devices intercept secure connections by masquerading as the intended destination and are sometimes used by companies to intercept their employees' secure traffic for monitoring or legal reasons. The employees' computers normally have to be configured to trust a proxy for it to be able to do this. However, in this case, the presumed proxy was given the full authority of a public CA, which is a serious breach of the CA system. This situation is similar to a failure by ANSSI in 2013.
>
> "This explanation is congruent with the facts. However, CNNIC still delegated their substantial authority to an organization that was not fit to hold it.
> …
> "As a result of a joint investigation of the events surrounding this incident by Google and CNNIC, we have decided that the CNNIC Root and EV CAs will no longer be recognized in Google products."
>
> *https://googleonlinesecurity.blogspot.ca/2015/03/maintaining-digital-certificate-security.html*

Improving this situation has proved to be exceptionally challenging. Adding digital credentials into the DNS to allow DNS responses to be validated can provide a robust mechanism to place public keys into the DNS, but the cost of such a measure is increased fragility of the DNS, increased complexity in zone registration and administration, increased time to perform a DNS query and much larger DNS responses. In addition, the chosen form of DNS security is one that interlinks parent and child zones, so that piecemeal adoption of this form of security has limited benefit. All of these considerations, coupled with the incumbency of a thriving CA industry, have proved to be inhibitory factors in adopting a more robust form of associating domain names with public keys to improve the integrity of secure communication.

Similar issues have been encountered in the efforts to retrofit secure credentials to allow authentication into other protocols. The Internet's inter-domain routing protocol, the Border Gateway Protocol (BGP), is mutual trust environment where lies, whether deliberate or inadvertent, readily propagate across the Internet, causing disruption through the diversion of traffic to unintended destinations. Efforts to improve the situation by using public key cryptography to provide a framework that allows routing information to be validated against a comprehensive set of digital credentials involve considerable complexity, add a potentially large computational overhead to the routing function, and contribute further fragility to a system that is intended to be robust. Such systems can only authenticate as valid information that is already well signed. False information is invariably indistinguishable from unsigned information, so the ability of the system to detect all attempts of abuse is predicated on universal adoption of the technology. This provides little incentive for early adopters and such retro-fitted security systems are forced to compromise between a desire for complete integrity and a system that can provide incremental benefits in scenarios of partial adoption.

Similar problems have been encountered in email systems and the scourge of spam mail. The original open trust model of email has been comprehensively abused for some decades and efforts to add digital credentials into mail have also been unable to gather a critical mass of adoption. The weaknesses in the security model of email has induced many end users to subscribe to the free email services provided by the very largest of mail services, such as Gmail and Yahoo, simply because of their ongoing investment in spam filters, at the cost of a level of digital privacy.

Insecurity does not only occur at the level of application protocols that sit above the transport services provided by the IP protocol suite. We are seeing these underlying protocols also becoming the subject of abuse. The User Datagram Protocol (UDP) is now a major contributor to Distributed Denial of Service **(**DDOS) attacks. There are few clear practical responses that would mitigate or even prevent this. "Just block UDP" is tempting, but two of the more absolutely critical services on the Internet, the DNS and the network time protocol, are feasible only with this lightweight query/response interaction, and it is the DNS and NTP which are being used in these DDOS attacks. This raises the question of how this lightweight efficient query response protocol can be used moving forward if its abuse rates are unacceptably high to the extent that they overwhelm the Internet itself.

The outlook is not good. When the overall environment becomes toxic the motivation of individual actors is to spend resources to defend themselves, rather than attempting to eliminate the source of insecurity. This increase in defensive capability induces ever larger and more toxic attacks, and the ensuring escalation ensures that all other actors who cannot afford such large budgets to defend their online presence in the face of such continual attack are placed in a precarious position. As with electronic mail, the business of content hosting is now shifting into a role performed only by a small number of highly capable and large scale hosting providers. The provision of DNS services is undergoing a similar shift, where the activity is only viable when undertaken by a large scale incumbent operator who has the wherewithal to protect their service from this continual onslaught of attack. This is no longer an open and accessible market for the provision of such services.

Local Filtering and Blocking

The public policy objectives in the area of content filtering and blocking space are intended to fulfil certain public policy objectives by preventing users within a country from accessing certain online content. The motives for such public policies vary from a desire to uphold societal values through to concessions made to copyright holders to deter the circulation of unauthorised redistribution of content. This chapter will not evaluate motives for content filtering, but look at the technology that can be used to support filtering and the potential side effects of such approaches.

Content filtering, or preventing users accessing certain online content, can be achieved in a number of ways, including routing filtering, DNS name resolution filtering and traffic interception

Route Filtering

Route Filtering takes the IP address of the service to the filtered and creates specific routing forwarding rules to treat all packets directed to this address in a manner that prevents the packets reaching their intended destination.

This can be achieved in a number of ways. One way is for all ISPs and transit operators to use a list of routes to be filtered and use this as a filter list to be used to filter routing information learned by routers. This is not altogether effective, in so far as a set of addresses may be encompassed within an aggregate route object. Removing the entire route object may have unintended consequences for third parties whose services are addressed in the span of addresses covered by the aggregate announcement, and leaving the aggregate route objects in place may also support continued access to the addresses to be filtered. This approach is only effective if uniformly applied, of course.

Another approach to blocking at the routing level is a technique of advertising specially crafted false route attributes for the IP addresses as "specific" routes. Local routers pick up the filtering route advertisement and due to the falsified attributes of the advertisement will prefer this route over the route to the original content service location. Aggregate routes that span this filtered route will still be learned by the local router, so access to third parties remains unaltered. The new route for the filtered addresses can lead to a "null route" which will cause the router to discard the traffic, or it can redirect the traffic to a traffic interceptor, as described below. A concern with this form of use of the routing protocol itself to convey the addresses to be blocked is that the propagation of the false route itself needs to be carefully contained. A further concern is that this filtering route is indistinguishable from a routing attack, in so far as a third party is attempting to inject a route into the network that causes a drop in reachability to a particular destination. Where service providers already deploy routing systems that are intended to detect and drop efforts to inject bogus routes into the network these "negative" routes will be filtered and dropped unless each ISPs has specific measures in place to detect and accept these particular routes.

> The Pakistan YouTube Incident.
>
> In early 2008 engineers at Pakistan Telecom created a couple of false routes that were intended to prevent Internet users within Pakistan from access the YouTube video sharing social media web site. The technique used was a common one, namely the creation of "more specific" route advertisements, so that local routers would prefer these synthetic routes over the genuine routes advertised by YouTube itself.
>
> As soon as the routes were created they quickly permeated beyond Pakistan and were heard through the much of the Internet, causing YouTube to be inaccessible for these users.
>
> The ease with which the false routes were created and the speed at which they were propagated well outside of their intended scope has caused others to respond by revitalising the work on securing the routing infrastructure such that router would be capable of believing such synthetic route advertisements.
>
> *http://www.cnet.com/news/how-pakistan-knocked-youtube-offline-and-how-to-make-sure-it-never-happens-again/*

The ongoing issues with the exhaustion of the supply of IPv4 addresses has meant that there is no longer a clear association of an IP address with a particular service, a particular content publisher, or particular content. Web hosting providers use a technique of "virtual name hosting" to place the content from many different sites behind a service portal which uses a single public IP address. Blocking that IP address at the routing level not only blocks the intended service point, but also blocks the entire set of sites that exist behind the same IP address. Such collateral damage limits the efficacy of address-based content filtering.

> The Australian Melbourne Free University Incident
>
> An Australian federal government agency, the Australian Securities and Investments Commission, used powers under the Australian Telecommunications Act to require that Australian ISPs block access to a particular IP address in February 2013. The Commission told a Senate hearing that "Recently, in targeting a scam investment website, we received information that another site [Melbourne Free University] that shared the same internet address was also blocked. ASIC was un-aware the IP address was shared by other websites." Evidently more than 1,200 different web sites were hosted on the same IP address at the time.
>
> *http://www.theweeklyreviewmelbournetimes.com.au/story/1565198/asic-admits-to-melbourne-free-university-block-that-took-more-than-1000-websites-offline/?cs=12*

Not only is there increasing use of shared IP addresses for hosting content, content itself is increasingly agile across IP addresses by adding further hosts for their content. IP level blocked sites can readily circumvent such IP-level interception mechanisms by shifting their content to other hosting agencies. When content is passed into a widely distributed Content Distribution Network the content is no longer associated with a set of IP addresses, but often is served from the CDN provider's IP addresses, along with all the other content hosted by the CDN. Such measures negate the effectiveness of content filtering by IP address by removing the stable relationship between content and address. At the same time end users can readily circumvent localised IP routing filtering by using public Virtual Private Network (VPN) services to perform a "virtual relocation" of their point of interconnection to a location where the local route blocking no longer applies.

Name Filtering

The DNS is also used as a means of enforcing filtering of content. In the simplest form of name filtering a list of proscribed DNS names is circulated to internet Service Providers, and this list is used to configure their user-facing DNS resolvers, so that queries directed to these resolvers for the filtered names result in a synthetic response.

If the user uses the configuration settings as provided by default, then their DNS resolution function will direct their queries to the ISP provided resolvers who, in turn, will apply the filter to the proscribed names list. There are a number of potential DNS responses, and operational practice varies. Some providers elect to return a response to the name query, but provide a private (unrouted) address in response to the query. Some elect to provide an address that points to a resource that describes why the name has been redirected. Others elect to send a DNS result code that the name does not exist.

Such name filtering operations are readily circumvented, and many users appear to learn of the availability of unfiltered open DNS resolvers, such as those operated by Google (Google's Public DNS), OpenDNS or Level 3. By replacing the reference to the ISP's resolver with a reference to one of more of these open resolvers in their devices, the user effectively restores a complete view of the Internet's name space and bypasses the locally imposed name filter.

However, such local efforts of remediation are not without their own downsides. It can be that the use of these non-local resolvers reduces the overall perception of performance of the Internet, particularly when the non-local resolver is located far from the user's device, as the DNS transaction to resolver the domain name takes longer for more distant resolvers.

The use of non-local resolvers also impacts on content distribution systems and content localisation. A number of content systems direct the user to different content depending on the supposed location of the user, and the way this can be done is by assuming that the user is located in the same locale as their DNS resolver. By providing a particular response based on the locale of the resolver that is asking the DNS query, non-local use of DNS resolvers effectively teleports the user into the locale of the remote DNS resolver. Some times this is undertaken deliberately by the user as a mediation against content blocking, and in other cases it may be the cause of user complaint about inadvertent exposure to inappropriate content.

The use of non-local DNS resolvers also leads to information leakage. A DNS query is the precursor to almost all forms of transactions on the Internet, and knowledge of the sequence of DNS queries being made by a user can be analysed to provide insights into user's behaviour, both in the aggregate and, with appropriate data volume and analysis, data profiles that can potentially close in on individual users. While national regulatory frameworks may safeguard the collection, storage and use of data as it relates to data about a country's citizens, the safeguards that relate to non-local users may not be present.

It may also be that the name filtering function includes traffic inspection and blocking attempts to use non-local DNS resolvers. Again this is readily circumvented. The most obvious form of circumvention is to place the IP address of the blocked site in the user's local hosts.txt configuration file. In this case the user's applications can then perform the name to address translation without using the DNS and thereby circumvent the DNS block. More general circumventions against DNS blocking encompass techniques that encapsulate DNS queries within other commonly used protocols to bypass the traffic inspection and blocking function (HTTP secure access, TCP port 443, is a common circumvention method). This form of circumvention is less commonly used at present, as there is a slightly higher technical barrier to using such DNS tunnelling solutions, but as with other circumvention methods, the more widespread the blocking the more widespread the tools and techniques for circumvention. Tunnelling DNS queries through the traffic filter makes the user's DNS queries completely opaque to the local network operator, and hides the user's behaviour behind a widely used conventional form of payload encryption. Providing further incentives for users to turn to deliberate obfuscation of their online behaviours may be seen in a positive light by some interests, but as a highly retrograde step by others.

These DNS-based name filtering interventions also rely on changing the result of the DNS query and providing the user with a synthetic result that is not the original DNS data. The adoption of DNSSEC, the security system for the DNS can prevent this synthetic data being accepted by the user. If the blocked name is signed using the DNSSEC technology, and the user is performing validation of DNS results, then the attempt to substitute a synthetic response would be flagged by DNSSEC validation as an attempt to subvert the DNS, and the response without be withheld from the application. The user is then aware that the name has been filtered by an intermediary, and the substitute response is not accepted by the user's DNSSEC validating resolver.

The diverse and uncoordinated nature of the application of name filters make a subsequent task of undoing the name filters extremely challenging. If a name has been hijacked and used for purposes that trigger the imposition of name filters, then once the name is restored the subsequent task of identifying if and where such filters have been applied is extremely daunting.

Name filtering can be an expeditious and efficient way of blocking access to inappropriate content. However it is readily circumvented, and the mechanisms for circumvention often lead to an outcome where users are further motivated to adopt technologies that are intended to hide their existence and their actions from the local network. Conventional tasks such as attribution, localisation, and customisation are impeded by such steps.

Facebook and China

A example of name filtering can be found in use in China, where there is widespread use of DNS interception and synthetic responses to DNS queries, providing unrouted addresses in response to queries to resolve blocked web sites (law.harvard.edu/filtering/china/appendix-tech.html).

There have been instances where the redirection via these DNS interceptors has used a third party's IP address. In one reported case it lead to a denial of service attack, and in another it lead to a form of cross-border information leakage.

http://www.theguardian.com/technology/2015/jan/23/what-happens-when-most-of-china-visits-your-website-it-dies-a-horrible-death
http://www.potaroo.net/presentations/2013-10-15-facebook.pdf

Traffic Interception

The technique of route redirection can be coupled with traffic interception in order to address some of the shortfalls of IP address filtering. This approach uses some form of routing level interception to direct the traffic to some form of traffic interceptor. At this point the interception agent functions like a conventional web proxy unit, undertaking the initial protocol with the end user as a proxy, then receiving the URL being sought by the user. At this point the interceptor can determine if the URL is part of some blocked list, in which case the connection can be terminated by the agent, or whether the proxy can forward the fetch request to the intended destination as a conventional proxy.

This form of traffic interception has fewer side effects. When a single IP address is used by multiple web sites, the original routing redirection operates at the level of an individual address, then the traffic interceptor works at the application level and applies the content access policy at the particular web sites that are the subject of the policy.

Many web sites use SSL (secure socket layer) tunnelling to prevent certain forms of traffic interception and eavesdropping. Using an appropriate traffic cipher algorithm the interactions of the application are impervious to application level interception, as the steps taken to encrypt the application session happen at the level of the initial protocol connection (immediately following the TCP handshake) rather than as a function performed within the application level interaction. However this does not prevent all forms of interception, it just prevents interception of the application level interaction. As the server's certificate name is transmitted in the clear to the client, an interception engine can proxy the initial setup of the SSL session and then read the server name, as described in the server certificate, and decide whether or not to terminate the session at that point. Once the session proceeds past the initial crypto handshake visibility into the session is lost. This is equivalent in functionality to the name filtering approach examined above.

Cleanfeed: UK Traffic Interception

A hybrid approach of routing and traffic interception is used by the 'Cleanfeed" system used within the United Kingdom. Route redirection is used to redirect the client traffic destined to IP addresses that are associated with filtered content to a web proxy unit. This web proxy completes the initial protocol handshake and session setup up to the point of the HTTP protocol GET command, which contains the named URL that the user is attempting to fetch. If the named URL is on the list of filtered content, then the proxy session is terminated towards the user by responding with an HTTP 404 response (Page Not Found). Otherwise, the proxy performs as a conventional proxy and initiates a new session with the requested content, relays the content request and passes the result back to the user.

There have been reports of a number of issues and shortcomings with this approach, including that the blocking filter operates only on the Domain Name part of the URL and not the trailing locator string, which has meant that large content sites (such as Wikipedia) have been filtered by Cleanfeed in the past and Cleanfeed reportedly was unable to distinguish served content beyond the simple domain name. Another reported limitation is that Cleanfeed is unable to proxy SSL connections, so sites that are being a non-Server Name Indication SSL connection cannot be filtered in this manner by Cleanfeed, as the GET commend is performed. If, however, the site uses SSL with Server Name Indication, then the intended DNS part of the URL is exposed in the clear prior to the start of the channel encryption process.

https://wiki.openrightsgroup.org/wiki/Cleanfeed
http://thenextweb.com/insider/2012/01/04/as-the-sopa-debate-rages-a-look-at-the-messy-state-of-website-blocking-in-europe/

Locality and Interdependence

The Internet was not originally designed as a single network that serviced much of the world's digital communications requirements. Its design was sufficiently flexible that it could be used in many contexts, including that of small network domains that were not connected to any other domain, through to large diverse systems with many tens of thousands of individual network elements. If that is indeed the case then why is it that when networks wish to isolate themselves from the Internet, or when a natural calamity effectively isolates a network, the result is that the isolated network is often non-functional. Where is the inter-dependence in the Internet that binds each network component into the whole and what efforts are being made to reduce this level of interdependence?

Locality and GeoLocation of Names and Addresses

Names and addresses on the Internet are not intrinsically tied to any particular physical location, nor any country or region.

While a class of domain names are associated with individual countries, namely the per country Top level Domains (ccTLDs), it is a matter of policy by the administrators of these TLDs whether the registration of subdomains within these ccTLDs is limited in any way to entities that reside in the associated country, of whether the services named within a ccTLD name space is restricted in some fashion to refer to a service located in that country. Domain names can be see as symbolic names that refer to attachment interfaces to a network, through the mapping of a domain name to an IP address, and it is the IP address, and via this address, the location of the device that is using this address that is the actual geolocation of the address and, ultimately, the address of the DNS name.

An IP address does not contain any internal structure that identifies a country or particular locale, nor does it intrinsically identify a network operator. There is no single authoritative maintained database that maps IP addresses to the geographic location of where that address is being used.

The address allocation process used by the Regional Address Registries administers a registry that records allocated addresses and the country where the address holder is located. Of course that is a slightly different definition of location than the country where the device that uses an address is located, but in most cases these are the same.

Private interests have constructed more detailed databases that attempt to provide location of addresses to finer levels (*https://www.maxmind.com/en/geoip2-services-and-databases*), and similar studies have been undertaken by academics with public inputs (*http://internetatlas.org*).

These approaches cannot map the entire address space of course. Shared addresses used in private contexts (such as the widely used 192.168.0.0/24) have no particular location. Addresses used in satellite-based access systems similarly have no fixed locale on the earth's surface. Addresses used by operators of mobile data services generally can be mapped to countries, but resolution to finer levels may depend on the address management practices used by the mobile operator. If the access network operator uses address sharing, such as a Carrier Grade NAT (CGN) then the physical location of the address may not be clearly established, and it may be only possible to resolve the location to country or region depending on the internal structure of the network and its operational practices in address sharing.

With a sufficiently coarse level of resolution, such as that of location to countries, a number databases exist, both public and private, that map addresses to countries with a reasonable level of accuracy. Finer levels of resolution are available in certain countries and for certain address ranges.

This locality information, that associates IP addresses to locations, is used to inform much of the operational work when implementing policies about locality of online services and infrastructure.

Locality in Routing and Traffic

Locality in traffic flows can be a critical issue for users. The shorter the network path between communicating end points then the lower the time taken to send a packet and receive a response (a so called "round trip time"). The lower the round trip time the faster a flow controlled protocol can sense the equilibrium flow rate for a traffic stream, and the greater the Transmission Control Protocol (TCP) carriage capacity, all other things being equal. If shorter paths across the network produce better outcomes for users in terms of perceived performance of network transactions, and such shorter paths allow the transport protocols to make more efficient utilisation of the network by the traffic flows triggered by these transactions, then how are this localised network paths realised?

The behaviour of the routing system is the first place where there is a natural bias towards finding the shortest possible paths in the network. Most automatic routing protocols take an arbitrary interconnection graph between switching elements and select candidate paths that represent the shortest path through the network (If every switch-to-switch link is assigned a "cost`", the these routing protocols reach a converged state by computing the shortest possible path through the network, where the path cost is simply the sum of the individual link costs). Such link based routing protocols are used within individual network domains. Between such routing domains the internet uses the Border Gateway Protocol (BGP) as its interdomain routing protocol. Rather than looking at network paths at a level of detail that considers each path as a sequence of individual point-to-point links, BGPO looks at the network as a set of interconnections between component networks, and each path through the aggregate network is defined as a sequence of networks. BGP is a shortest path selection routing protocol, and it selects candidate paths to use that transit the fewest possible networks.

Of course a routing protocol cannot create new inter-domain connections, and if the shortest possible path between two networks that are not locally interconnected is a path that traverses an arbitrary number of external networks, the BGP will select such a path.

One way to assist the routing protocols to select localised paths is to enrich the local network interconnection at the physical layer. One of the more efficient ways to achieve this is through use of a local traffic Internet Exchange (or "IX"). An exchange can be seen as a switching mesh: an individual network that connects with a single connection to an exchange can create virtual point-to-point connections to all other parties who are also present at the same exchange. In this scenario all parties connected to an exchange point can directly exchange traffic with all other parties at the exchange without the traffic traversing third party networks. Many exchanges also permit selective forms of interconnection, where each pair of networks represented at the exchange can determine whether they interconnect, ands the commercial terms of such an interconnection independently of any other connections that they may have set up at the exchange. Exchanges are commonly operated at a city level, or a national level. A smaller number of exchanges operate with large number of local and international providers, essentially operating in a role of being a regional connection hub, and examples of these are at London (LINX), Amsterdam (AMSIX) and Frankfurt (DECIX).

Exchanges operate at a number of levels. They allow access providers to interconnect customer traffic without the use of a transit provider acting as a middleman. They allow access providers to "see" a range of competing access providers, bring competitive pressures to bear upon the transit role. Given this rich collection of connectivity, exchanges are also powerful attractors for content distributors. A content distribution network located at an exchange can directly access a far larger number of access networks and their end user population, and often do so at a lower cost than remote access that is arbitrated by a transit network's services.

Exchanges can assist in keeping network traffic between local end points local by facilitating rich interconnectivity between local access providers and transit and content providers.

Localisation may also be an outcome of regulatory actions, where specific regulatory measures prevent the "exporting" of data beyond national boundaries.

There are also aspects of tensions between localisation and various commercial pressures. Conventionally, a structure of dense localisation of interconnection would be expected to be an outcome that reduce costs for all parties, but at times local circuits may incur higher coists for a network operator, or the use of local sender keep all commercial interconnection arrangements may expose one service provider's network assets to its competitors without due financial compensation. In such cases the outcome of such environmental conditions is local fracturing of the network, and consequently in these cases external connections and services are an essential component of ensuring local inter-connectivity.

It may also be the case that true locality of traffic flows is not the same as the IP-centric view of the locality of a traffic flow. While at the IP level it may be that the start and end points, and even the IP intermediate switching points may be able to be mapped to geographical locations that sit within the bounds of some locale, such an IP-centric view of the network may not necessarily reveal intermediate encapsulations, such as Multi-Protocol Label Switching (MPLS) or other forms of IP tunnelling, including IPv6-in-IPv4 transition tunnelling mechanisms. When these location of the basic paths of the carriage system and the carriage level switching equipment is revealed it is always possible that the physical path of the traffic flow does not match the logical view provided at the IP level.

Locality and the DNS

While there is a desire to restrain traffic flows so as to ensure that Internet traffic between users in a particular scope or locale remains within the realms of a certain locale, country or region, there is a similar consideration relating to DNS queries.

The DNS is an intrinsic part of almost every form of user transaction on the Internet, and the choice of resolver used to perform DNS resolution is one aspect of locality and interdependence. Users may use DNS resolution tools that make use of non-local DNS resolvers. The issues of the interaction of these measures with DNS-based approaches to content filtering is a consideration here, as such DNS-blocking measures intended to support local content filters are circumvented by this form of use of non-local DNS resolvers. The use of such non-local DNS resolvers also creates an external dependence that would be broken if the local network were to lose a route to the non-local resolver.

There is also the consideration of the flows of meta-information. DNS queries can be analysed to provide near real time information about the online behaviour of users. The use of non-local DNS resolvers potentially results in this information being passed across national borders into regimes that may operate with different frameworks concerning personal data concerning individuals who are not citizens of that country.

> The issue here is that the DNS was not designed with privacy in mind at the outset, and the DNS is notorious in leaking information to third parties.
>
> The problem lies in the conventional practice of passing the full domain name of the query to parent zones when the desired response is the name servers of the child zone. It is slightly less efficient, but far more secure, to minimise the query string when resolving a name. This is the topic of current study by the DNS Operations Working Group of the IETF.
>
> *https://tools.ietf.org/html/draft-bortzmeyer-dns-qname-minimisation-02*

Sample-based measurements show that the extent of this use of non-local DNS resolvers appears to range between 1-2% of users in some countries to in excess of 60% of users in other countries.

http://www.potaroo.net/ispcol/2014-11/resolvers.html

Another aspect of the DNS concerns interdependence. The DNS name space is a rooted hierarchy, and the DNS resolution protocol makes the assumption when resolving a name that critical aspects of the DNS are available answering queries. In particular this concerns the availability of the DNS root zone servers (but also extends to the authoritative servers for other popularly queried TLDs). The widespread use of anycast servers for the DNS root zone has improved the performance of the DNS in terms of the time taken to resolve a DNS name, but these resolvers need to regularly refresh the content of their local cache with the content of the primary server. The implication here is that isolation of local DNS resolvers from the authoritative servers that serve the root zone of the DNS will eventually mean that the local servers will cease to answer queries in an extended hiatus of connectivity.

"Anycast" is a deliberate approach to deploying multiple servers using the same IP address at diverse locations. As long as each anycast instance responds in precisely the same manner to queries, then anycast is an efficient method to improve the performance and robustness of a service. With anycast a user's query is directed to the "closest" instance of an anycast server constellation. What is "close" in this context is the outcome of the routing protocols selection of shortest path. If a server in the anycast constellation fails then as long as the individual anycast routye is withdrawn, then the local traffic that would normally be directed to this instance is directed elsewhere. Anycastr also assists in DDOS mitigation. If the attack originates from a single source, or a small set of related sources then the attack will be directed to a single instance of the anycast server constellation and all other servers will operate without interruption. A broad scale distributed source attack will be spread across the anycast constellation, so that any individual server will experience only a small fraction of the total DDOS volume.

Locality and Content

Locality of network extends beyond access and transit networks and the mechanisms of controlling traffic flows to the aspect of what content is accessible by which users.

While the intended operation of the Internet was to blur geography and create a network where location was not a visible part of the user experience per se, the questions relating to the local of users and content remain. Can content be customized to locales? What information is available to generate this locale tagging?

The externalities of the commercial arrangements relating to content distribution may impose a limit on access to certain content to users located in certain locales. This can take the form of IP address geo-location, where the locale of the user is guessed by consulting a geographic location database with the remote side IP address making the connection to the content.

Alternatively, this can be performed by the DNS name resolution, where the IP address of the resolver querying for the DNS name of the content is used to lookup a geographic location database. The response provided by the authoritative name servers to the query is based on the assumed location of the user, as the assumption here is that the user and their DNS resolver are closely located. With the increasing popularity of common open DNS resolver services, such as Google's Public DNS and the Open DNS project, such assumptions are no longer always the case. The assumed location of the resolver that puts a query to an authoritative name server may not be closely located to the original user who made the query.

Such approaches to limit content to defined geographic locales generates a response from some users to want to bypass such geolocation blocking of content. In an open environment such as the Internet there are providers of services that are intended to address precisely this problem and virtually relocate the user's device to a network location within the desired geographic region. Such approaches typically involve the use of secure VPN technology, where all the user's traffic is carried inside an encrypted IP tunnel to emerge into the internet within the desired locality. The side effect of this increased adoption of such advanced tunnelling solutions to access local content implies an increased level of use of encryption for user traffic.

It is evident that there is no clear cut absolute solution for a content provider to unambiguously determine precisely where an end user is physically located, and efforts by a content provider to enforce geographical based differential outcomes for content distribution creates a secondary market in responding to these measures. The increasing sophistication of the measures and their responses in effect fuels demand for increased levels of user anonymity. This, in turn, frustrates many of the conventional measures used by law enforcement agencies, as it fuels the creation of an online environment where individual actions by users are effectively anonymous and cannot be readily mapped back to a physical identity or location.

Security

Any form of public communications network necessarily exposes some information about the identity and activity of the user's of its services. The extent to which such exposure of information can be subverted and used in ways that are in stark opposition to the users' individual interests forms part of the motivation on the part of many users to reduce such open exposure to an absolute minimum. The tensions between a desire to protect the user through increasing the level of opacity of network transactions to third party surveillance, and the need to expose some level of basic information to support the functions of a network lies at the heart of many of the security issues in today's Internet.

Security and Authenticity Requirements

The public sector is as acutely aware of the need for accessible online security as the private sector. Internet users need to be assured that the material they retrieve from an online service is genuine, and any data that they provide in response is also handled with appropriate levels of confidentiality and care.

The security frameworks used on the Internet are devolved frameworks as distinct from tightly interdependent and mutually constrained frameworks. The algorithms used to encrypt and decrypt data, and the algorithms used to generate and validate digital signatures are open algorithms that can be readily implemented, and an application wishing to use encryption can chose a mutually agreed algorithm. The manner in which security is applied to a network transaction can vary according to the motivation to use a secure solution, such as the desire to prevent third party inspection of the network traffic, or the desire to support a transaction that cannot be tampered with buy third parties, or where the identities of the parties are assured. The selection of trust points to allow parties to validate the authenticity of digital signatures is an independent consideration, as is the processes used to support awareness of revocation of otherwise valid credentials. At the heart of the implementation of secure systems is a software library of security tools and procedures, and the current mainstay of many of the online security systems is the open source OpenSSL library.

> The "heartbleed" bug of early 2014 that affected OpenSSL hosts illustrated the extent of use of Open SSL in today's world.
>
> "A serious overrun vulnerability in the OpenSSL cryptographic library affects around 17% of SSL web servers which use certificates issued by trusted certificate authorities….17.5% of SSL sites, accounting for around half a million certificates issued by trusted certificate authorities."
>
> *http://news.netcraft.com/archives/2014/04/08/half-a-million-widely-trusted-websites-vulnerable-to-heartbleed-bug.html*

The trust and confidence such security mechanisms engender underpins much of the trust in moving a large diversity of economic activity into the digital realm. Online banking, business-to-business transactions, public services provided to citizens, all rely on these mechanisms. These mechanisms are under continual pressure, and the points of vulnerability are exploited from time to time.

Address and Routing Security

The integrity of the Internet's IP address and routing systems are essential to the operation of the Internet. Two of the fundamental assumptions of the Internet's architecture are that every point of attachment of a device to the network has a unique IP address, and that every routing element in the network has a mutually consistent view of how to direct a packet towards its addressed destination. When an IP packet is passed into the network, the packet's destination address should be a unique and unambiguous value, and irrespective of where the packet is passed into the network it should be forwarded to the same intended destination.

> The Internet uses a "self learning" process of understanding the location of attached devices. This self-learning process is achieved via the operation of routing protocols. The language of routing systems includes the announcement (or "advertisement") of an address to a network, and the propagation of this announcement to all other networks across the Internet. The security question here concerns the ability to insert of false information into this self-learning system.

The question is how the integrity of this system is managed. What is to prevent a malicious party from announcing someone else's address into the routing system? Indeed that is the very nature of the IP address filtering process, where deliberately false forwarding directions are inserted into parts of the network. What is to prevent an attached device presenting a packet into the network with a deliberately falsified source IP address?

> This is a more subtle form of exploitation of weaknesses in address and routing security. Every packet has in its header the intended destination of the packet (the "destination address") and the identity of the device that created the packet (the "source address"). Why would any party want to lie about its source address? There are some use cases that relate to IP mobility, but the overwhelming issue here is the use of this technique to mount hostile attacks. If the attacker is able to pass in simple query packets using the UDP protocol to a conventional server, then by using a packet that contains the source address of the intended victim, the server sends its response to the victim. By replicating this query across many servers the attack volumes that can be bought to bear on the victim can be measured in the 100's of gigabits per second.
>
> *http://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack/*

The approach used by the Internet has been largely one of conventions of behaviour. The common interest is to ensure that packets reach their intended destination, which aligns with local interest. The original models of addressing and routing were based largely on mutual trust. Each network operator was expected to announce only those addresses that were assigned to that network, and each network was expected to propagate routing information that accurately reflected the local connectivity and routing policies used by the local network. The system as a whole operated correctly if each component network operated according to these simple principles. While these conventions continue today (such as the "Routing Manifesto" at https://www.routingmanifesto.org/manrs/) the scope and diversity of today's network means that such conventions can be abused from time to time.

The onus is placed on each network to defend itself from believing incorrect route advertisements. One approach here has been in the use of "route registries", where each network records in a common registry its local addresses that it originates, and the local connections to adjacent networks, and the routing policies that it applies to each such adjacent network. If every network operator diligently maintained such data in a routing registry then other operators could use this information to generate a local filter to apply to incoming routing advertisements. The contents of the routing information would be constrained by the information in the route registry, and any additional information could be rejected as a bogus route.

Some local communities have used the route registry approach, and it has proved useful in minimizing certain forms of routing attack, but these approaches have not translated into the entirety of the Internet, and more recent efforts have been directed towards a more rigorous form of processing routing information that will distinguish between receiving genuine routing information and synthetic (or false) information.

The more recent approach has borrowed the conventions used in Public Key Infrastructure (PKI) models, and enrolled the address registry operators as Certification Authorities (CA). Allocations of addresses are accompanied by a certificate, where the address registry attests that the entity who holds a certain public/private key pair is the current holder of a particular collection of IP addresses. The address holder can use their private key to sign various attestations about the use of their addresses, and third parties can validate these attestations using the collection of published certificates. The routing system is a form of circulation of such implicit attestations, where networks are advertising reachability to addresses, and the secure routing model calls for digitally signatures to be attached to these route advertisements, making these implicit attestations of reachability to be explicit and testable as to their validity. Receivers of such routing updates that describe the reachability of an address block can validate the authenticity of the routing update by validating the digital signatures that are associated with the update.

There are two significant issues with this approach, which warrant noting in the context of an open and coherent Internet.

The first is that digital signatures and the associated set of credentials that permit such signatures to be validated do not intrinsically identify erroneous material. Such signatures can only validate that the received information is authentic and has not been tampered with in flight. The use of this approach in a manner that would reliably identify instances where addresses and route advertisements are used in an unauthorised manner is only possible in an environment where every network and every address holder generates, maintains and publishes their signed attestations and associated certificate-based credentials. The issue here is that this approach to address and routing security has a prerequisite condition of universal adoption in order to be comprehensively effective.

> Any scheme that uses positive attestations can only identify what's "good." If everyone who can generate positive attestations does so, then attestations that cannot be validated can be safely assumed to be "bad." But in an environment where not everything has an associated attestation, or digital signature in this case, then the situation is not so clearly defined. In the case of digitally signed route attestations partial deployment infers that you have a category of routes that are "not signed". These routes are not provably authentic and not provably not authentic as they are simply not signed at all. But if validation fails for a route attestation then this is semantically equivalent to "not signed". It is not possible in such a partial adoption space to distinguish between falsified information and unsigned information.

Orchestrating an outcome of universal adoption of any technology is extremely challenging in the diverse and loosely coupled environment of the Internet, and particularly more so when the motivation for individual adoption of a technology is based on perceptions of the value of risk mitigation. Differing risk appetites and differing perceptions of liability and loss lead to fragmented piecemeal adoption of such a technology, and in the case where the major benefit of the technology is only attainable if there is universal adoption. The prospects of achieving a critical mass of local adoption that would influence the remaining parties to also adopt the technology are challenging, particularly when there is no public policy framework that encourages individual actors.

Domain Name Based Security

Much of the current system of security is that used by the Web, and in particular the system of the issuance and validation of domain name certificates used by the Secure Sockets Layer protocol (SSL) for channel encryption. This form of security is based on the attestation of a third party that a particular operational entity controls a given domain name. Lists of collections of third parties who are trusted to publish such attestations are packaged with popular browsers. These lists are similar, but can deviate from time to time between browsers. This deviation can cause user confusion, where a site will report itself as "secure" when using some browsers and generate a security exception alert when using other browsers.

The second issue with this approach is that browsers are generally unaware which third party Certification Authority (CA) has actually attested which domain name. This leads to the vulnerability that is a trusted third party is compromised then any fake attestations it may subsequently generate would be trusted by users' browsers.

> The Diginotar CA Compromise
>
> In 2011 a European Certification Authority, Diginotar, had its online certification system hacked, and 344 false certificates were minted for a number of popular domain names, with both the public and the private key of these false certificates published as part of the compromise.
>
> Because the private key of these fake certificates was published, and because many popular browsers were willing to trust all certificates that were issued by the Diginotar CA, then it was possible for a party who was in a position to intercept user traffic to any of the services named by these domain names to substitute the fake certificate and then successfully dupe the user's browser and masquerade as the intended service.
>
> A number of parties were affected, including the Dutch tax authority who used Diginotar as the CA for the certificates used by its online taxation service. According to a report by Fox-IT commissioned after the incident, the false certificates appeared to have been used to spy on a large number of users in the Islamic Republic of Iran by performing a man-in-the-middle attack.
>
> *https://www.rijksoverheid.nl/bestanden/documenten-en-publicaties/rapporten/2012/08/13/black-tulip-update/black-tulip-update.pdf*

The underlying issue with this model is that there are hundreds of Certification Authorities (CA) that are trusted by browser vendors (see https://wiki.mozilla.org/CA:IncludedCAs for one such list, also https://www.eff.org/files/colour_map_of_cas.pdf), and any CA can issue a certificate for any domain name. The implication here is that this is an asymmetric trust model. While a certificate subject is trusting on the integrity of a single CA to correctly mint a certificate for the domain name in question, the consumer of the certificate, namely the end user of the service, is trusting the entire collection of CAs. The user is taking as a matter of implicit trust without explicit validation that none of these CAs, and none of their registration agents (RAs) have been compromised and none of them have issued a certificate for the same domain name under false premises. The CA system is only as strong as the most vulnerable CA and has only as much integrity as the CA that performs the minimum of integrity checks. The standards of the entire CA system are only as strong as the lowest individual set of standards from any single CA.

This form of domain name certification is a form of third party commentary on the actions of others. The certificate issuer for such certificates is neither the domain name holder nor the domain name registrar, and the integrity of the entire system depends on the robustness of the checks a Certification Authority and its Registration Agents perform to ensure that the domain name certificate is been issued to the correct party, namely the party that "owns" the associated domain name. Over time the system has been subject to various forms of abuse by applicants and erosion of integrity of validation by some CAs. In response, the CAs issued a variant of the domain name certificate, termed an "Extended Validation" Certificate that was intended to demonstrate that the CA had performed their validation of the applicant with some further vigour. The competitive pressure for CAs, and the inability of the CA system to create a sustaining market for certificate based on integrity of the product, is a common weakness for this system.

> As the certificate is an undistinguished product, then it makes sense to use the cheapest CA to obtain the certificate. But the cheapest CA can become the cheapest CA by reducing its expenses by reducing the number and efficacy of the tests it undertakes to confirm that the applicant does indeed "own" the domain name. Competitive market pressures between CAs create pressures that erode the integrity of the product they produce.

Alternate approaches being considered by the technical community are directed towards removing the concept of a third party commentary on Domain Name registrations and instead use DNS security (DNSSEC) and place the domain name public keys directly into the DNS alongside existing data and rely on DNSSEC to provide this key information to users in a robust and secure manner. Such a framework for name-based security would remove the shared fate risks of the current CA model, and mitigate the broad consequences of compromise of an individual CA's operation.

However, such an approach places even more pressure on the DNS. This approach relies on the widespread adoption of DNSSEC, both in signing zones and in validating DNS responses, and progress toward this objective goal is not thought to be overly impressive to date. Studies of validation indicate that some 1 in 7 users, or some 13% of the Internet's total user population pass their DNS queries via resolvers that use DNSSEC to validate the authenticity of the signed response that they receive (http://stats.labs.apnic.net/dnssec). Within individual countries the proposition of users who use DNSSEC-validating resolvers varies from 70% (Sweden) to 2% (Republic of Korea).

As with the considerations of the vulnerabilities associated with a single trust anchor for the PKI that is proposed to be used for the address and routing infrastructure, a similar consideration applies to the approach used by DNSSEC, and a brief comparison of the existing third party certification model and a DNSSEC-based model is useful here. A distributed third party certification model appears to offer a robust approach, in so far as the system is not reliant on the actions of any individual CA, and the provision of security services is subject to competitive pressures as any CA can certify essentially any domain name. There is no requirement for universal adoption, and incremental adoption creates incremental benefits to both the service provider and the consumer of those services. Unfortunately this is not quite the entire story, and, as already pointed out, compromise of an individual CA can lead to compromise of the integrity of any domain name certificate, and the robustness of the entire system is in fact critically reliant on the robustness of each and every CA, and failure of one CA leads to potential failure of other elements of the system. The system creates strong interdependencies, and has no mechanism for limiting potential damage. The DNSSEC model is strongly hierarchical, and at each point in the name delegation hierarchy the administrator of the zone file is in a defacto monopoly position of control over both the zone and all the subzones from that point in the name hierarchy. The root of the name space is also the root of the trust model of DNSSEC, and this root zone and it's associated key signing key represents a single point of vulnerability for the entire framework. The accountability of the body administering the root zone of the DNS to apply the best possible operating practices to ensure the absolute integrity and availability of the keys that are used to sign the root zone of the DNS is a matter of legitimate public interest.

DNSSEC and DANE

DNSSEC allows an application to validate that the responses it receives from DNS calls are authentic, and precisely match the data that is held in the authoritative name servers from the zone. The DNSSEC signatures ensure that no third party can tamper with the response in any way without this tampering being clearly evident.

DANE is a technology that allows a DNS zone administrator to provide information related to the SSL encryption key value used in conjunction with that DNS name to be placed into the DNS.

The combination of DANE and DNSSEC allows this secure channel bootstrap procedure to operate at a far greater level of robustness than is the case with the model of third party CAs. The service provider has a number of options as to how to insert the key information into the DNS, but the result is similar, in so far as rogue CAs are in no position to mislead the application with a false key information.

There is one side effect from this structure that impinges on the opening up of more generic top level domain names (gTLDs) in the DNS. Validation of a DNS response require that the client performs a "signature chase" to the key of the root zone. This means that to validate the signed zone "service.example.com", then the zone and key signing keys for "example.com" also need to be validated, as to the zone and key signing keys for "com", and these need to be validated against the root key. If the zone management of either "example.com" or "com" were to fail to maintain correct key signing state then "service.example.com also fails validation. The closer a domain name is to the root of the DNS the fewer the number of intermediaries, or the fewer the number of external dependencies for ensuring name validation. Also the fewer the number of intermediate zones the faster the entire DNS validation process as undertaken by the application.

In an Internet that makes extensive use of both DNSSEC and DANE, and in an Internet that relies on DANE to securely transmit session encryption keys to the client application, then it would be anticipated that were the gTLD space to be available for use by service providers, then there are clear incentives for such providers of service who which to use secure channels a manner that is as reliable and robust as possible to make use of name spaces that are located in the root zone itself, namely as a gTLD.

Block Chain Security and Bitcoin

A conventional view of security is that trust relationships are explicitly visible, and the task of validating an attestation is to find a transitive trust path from a trust anchor to the material being tested. The implicit trust relationship here is that trust in party A implies trust in all parties trusted by A, and so on. In a large and diverse environment such as the Internet there is a critical reliance on such trust relationships, either as an explicit hierarchy with all trust ultimately being placed in the operator of the apex point of the hierarchy, or as a common pool of trust with trust being placed in a collection of entities, few (or often none) of whom have any direct relationship with the user and the user has no rational grounds to invest them with any degree of trust.

An alternative model of trust was originally developed in concepts of the "web of trust" where trust was established by a process of engaging others to sign across an attestation, and the larger the pool of such third party signings the greater the level of trust could be inferred in the original attestation. This avoids the use of hierarchies and single trust anchors for an entire system.

Bitcoin uses a similar approach in its blockchain, where records of individual transactions are widely distributed, and meshed into an interconnected chain linked by cryptographic signatures. Validation of a ledger entry is effectively a process of consultation with the collective wisdom of a set of holders of these blockchains and there is no need for distinguished points of the origination of trust relationships.

To what extent the increased complexity of such blockchain models obfuscates inherent vulnerabilities of such an approach is of course open to further consideration and debate, but it does represent a secure system of trustable attestations that does not require the imposition of trusted points of authority to seed the entire system.

Denial of Service Attacks

It is a common assumption in many aspects of networking that the transactions that are presented to the network are authentic. This means, for example, that the packet contents in each IP packet are genuine, or that the initial steps of a protocol handshake will be followed by the subsequent steps, or that the payloads presented as part of an application level interaction represent an effort to work within the framework of the intended scope of the application. This trust in the authenticity of what is presented to a network could be said to be part of the Robustness Principle (RFC791) that protocol implementations should "be liberal in what [they] accept from others".

Later experience with the diversity of motivations on any large public network lead to a refinement of this principle that developers should "assume that the network is filled with malevolent entities that will send in packets designed to "have the worst possible effect"" (RFC1122). This was written in 1989, but it has been extremely prophetic.

The Internet has proved to be a surprisingly hostile environment, and almost every aspect of both application, operating system, protocol and network behavior has been exhaustively probed, and vulnerabilities and weaknesses have been ruthlessly exploited. Much of the intent of this exploitation is not to gain unauthorized access to an IT system, but to ensure that other legitimate users of an online service cannot gain access, or, in other words to deny the service to others.

The internet has a rich history of exploitation of various weaknesses, including, for example, TCP SYN attacks designed to starve a server of active connection slots, preventing legitimate users from accessing the service, or the injection of TCP RESET commands into long-held TCP sessions to disrupt their operation, such as is the case for the BGP routing protocol.

Some of the more insidious attacks involve the combination of deliberately falsified source addresses in IP packets, the UDP protocol, and applications where the response is far larger than the query. This is the case for the DNS protocol and certain forms of command and control in the NTP network time protocol. Denial of service attacks have been seen that generate 100's of gigabits per second (https://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack/). These attacks do not explicitly require that systems be infected with a virus, or otherwise enlisted into a bot army to mount the attack. The nature of the attack actually requires that system operate entirely as normal, and even the servers that unwittingly coopted into being part of the attack are assumed to be functioning perfectly normally and are simply responding to what they believe are perfectly normal queries. What allows such attacks to take place is the ability to inject packets into the network that use an incorrect (or "spoofed") IP source address.

The initial technical response to this form of source address spoofing is a document that was published some 15 years ago in 2000 (https://tools.ietf.org/html/bcp38). All the evidence suggests a general reluctance by network operators to equip their networks with additional points of control and filtering. Given that this has proved challenging the next steps have been to look at the DNS protocol itself to see if it is possible to prevent the DNS for being used as an unwitting vehicle for this form of attack. Again, this remains an extremely challenging exercise, and changes to infrastructure protocols are not one that can be made and adopted with any speed, so other forms of response have been used in the face of such attacks.

The most pragmatic response so far has been to equip service points and networks themselves with sufficient capacity to absorb attack volumes and maintain the service. Such systems are engineered for peak attack loads rather than peak service load, and are able to maintain a service consistency and quality even in the face of an attack.

While this is a pragmatic response, it does have its own consequences. If online services and facilities need to be provisioned not just to meet anticipated service loads, but provisioned to the extent to meet peak abuse loads then the scale and associated cost of mounting an online service rises. This escalation of cost implies that the barriers to entry for such services rises and established service providers who operate significant footprints within the network are in an advantaged position. The rise of the Content Distribution service business is partially a result of the escalation in the requirements for online services as a result of these form of service abuse. This represents a subtle change in the picture of the online world. The model of the service provider directly using information technology solutions to connect with customers over the network is evolving to a model that includes an intermediate party who takes content from the service provider and operates a business of standing that content online in such a manner that it is resilient to most forms of hostile attack and attempts to disrupt the service. Consumers of the service still interact with the original content, but do so via the content distribution system and the distribution provider, rather than more directly with the original service provider. There is the potential for concerns here about the extent of alternative supply and competition in this aggregated world of content distribution networks, and the extent to which such intermediaries can exercise control and influence over both consumers and users and the providers of services over the network.

Going Forward

There is no doubt that effective and efficient security in the Internet is of paramount importance to the digital economy, and the public interest is served by having a widely accepted mechanism for making trustable digital attestations that support this function.

It is also in the same public interest to ensure that this system operates with integrity, and that vulnerabilities are not exploited to the extent that they erode all trust. There are aspects of the current security framework in both the address and routing infrastructure and in the naming infrastructure that could be changed, and it is possible to propose secure systems that set a higher standard of security than what we use today. However, this exposes further questions that relate to public policy. Who should underwrite the costs associated with such a change? Should certain levels of adoption of secure practices in online services be mandated? Should certain security standards in consumer equipment and services be enforced? Is it even feasible for countries to set their own domestic standards for online security without reference to the larger Internet? Is the level of interdependence in the Internet now at such a level that such national determination of codes of practice in security and standards for online products neutralised by the globalization of this open digital economy?

Although the topic of competition reform in the communications services sector absorbs a huge amount of attention by policy makers and public sector regulators. The studies on the evolution of the mobile sector, with the introduction of mobile virtual network operators, the concepts of spectrum sharing and the issues of legacy dedicated voice services and voice over data and similar topics dominant much of the discussion in this sector. Similarly there is much activity in the study of broadband access economics, and grappling with the issues of investment in fibre optic-based last mile infrastructure, including issues of public and private sector investment, access sharing, and retail overlays again absorbs much attention at present.

You'd think that with all this effort that we be able to continue the impetus of competitive pressure in this sector, and to continue to invite new entrants to invest both their capital and their new ideas. You would think that we would be able to generate yet further pressures on historically exorbitant margins in this sector and bring this business back to that of other public sector utility offerings. But you would be mistaken.

Open competitive markets depend on common platforms for these markets that support abundant access to many resources, and in terms of todays communications networks abundant access to protocol addresses are a fundamental requirement. But over the past decade, or longer, we have consistently ignored the warnings from the technology folk that the addresses were in short supply and exhaustion was a distinct possibility. We needed to change protocol platforms or we would encounter some serious distortions in the network.

We have ignored these warnings. The abundant feed of IP addresses across most of the Internet has already stopped, and we are running the network on empty. The efforts to transform the Internet to use a different protocol sputter. A small number of providers in a small number of countries continually demonstrate that the technical task is achievable, affordable, and effective. But overall the uptake of this new protocol continues to languish at levels that are less than 3% of the total user population.

The ramifications of this are, in general, completely invisible so far to the consumer. Web pages still appear to work, and we can all shop online from our mobile devices. But the entrance of new players, and the competitive pressure that place on the market is drying up. The lack of protocol addresses is an extremely fundamental barrier to entry. Only the incumbents remain.

Shutting down access to the Internet to all but existing incumbents should be sending chilling messages to regulators and public policy makers about the futility of generating competitive infrastructure in copper and in radio spectrum if the same cannot be maintained in the level of provision of Internet access and online services.

This is not a comfortable situation and continued inaction is its own decision. Sitting on our hands only exacerbates the issue and todays situation is assuming a momentum that seats incumbents firmly in control of the Internet's future. This is truly an horrendous outcome. Its not "open". Whatever "open" may mean, this is the polar opposite!

The Open Internet?

Today we just don't have an “Open” Internet.

The massive proliferation of network-based middleware has resulted in an internet that has few remaining open apertures. Most of the time the packet you send is not precisely the packet I receive, and all too often if you deviate from a very narrowly set of technical constraints within this packet, then the packet you send is the packet I will never receive. The shortage of addresses has meant that the rigors of scarcity has replaced the largesse of abundance and with this has come the elevation of what used to be thought of as basic utility, including privacy and security in online services into the category of luxury goods only accessible at a considerable price premium. Our technology base is being warped and distorted to cope with an inadequate supply of addresses and the ramifications extend out from the basic domain of the internet protocol upwards into the area of online services and their provisioning. From the crowding out of open technology by encroaching IPR claims, to the problems of the mass of our legacy base restricting where and how we can innovate and change, and the rigors of scarcity of addresses, the picture of the technology of the Internet is now far from “open.”

Maybe the “open” Internet is something entirely different. Maybe it’s about the policy environment, and the competitive landscape. Maybe it’s about the attributes of having no barriers to entry in the supply of goods and service using the Internet. This could be deregulation of the carriage and/or access regime, allowing competitive in packet transport. Or the ability to deliver content and services without requiring the incumbents’ permission and without extortionate price gouging on the part of providers of critical bottleneck resources. Maybe in the “open” Internet we are talking about the benefits from low barriers to entry, innovation, entrepreneurialism and competition in the provision of goods and services over the Internet platform.

But this is not "open" either. The fact that we’ve exhausted our stock of IP addresses impinges on this considerations of markets for the provision of goods and services on the Internet, and their open operation. Without your own pool of IPv4 addresses you cannot set up a packet pushing business, so that's no longer "open". And without your own pool of IPv4 addresses you cannot set up secure services as a content service provider. So as long as you are willing to offer goods and services over an open, insecure, untrusted channel, and as long are you are willing to put the fate of your enterprise in the hands of your virtual neighbours with whom you are sharing IP addresses and hosting platforms, and so low as the price of access to these shared address resources is not in itself a barrier to entry, then perhaps this niche is still accessible. But its not what we intended it to be. It’s not “open”.

Perhaps the “open” Internet, in the sense of being an "open" platform that can carry the hopes and aspirations of a socially transformative power of a post-industrial digital economy, is now fading into an ephemeral afterglow.

Maybe its not too late, and maybe we can salvage this. But there are many moving parts here, and they all deserve attention. We need to use an open common technology platform that offers abundant pools of protocol addresses. Yes, I’m referring to IPv6. But that’s by no means the end of this list. We need continuing access to software tools and techniques. We need open software. We need open technology. We need open devices and open access to security. We need to open competitive access to the access infrastructure of wires and access to the radio spectrum. We need open markets that do not place any private or public enterprise in overarching positions of market dominance. We need an open governance structure that does not place any single nation state in a uniquely privileged position. We need open dialogues that enfranchise stakeholders to directly participate in conversations that matter to them. Indirect representation is just not good enough. We need all of these inputs and more. And each of them are critical, in as much as we are aware from centuries of experience that failure in any of these individual aspects translates to catastrophic failure of the entire effort.

Yes, this is asking a lot from all of us. But, in particular, its asking a lot from our policy makers and regulators. The mantra that deregulated markets will naturally lead to these forms of beneficial outcomes that enrich the public good ignores a rich history of market distortions, manipulations and outright failures. An "open" Internet is not a policy free zone where market inputs are the sole constraint. Markets aggregate, monopolies form, and incumbents naturally want to set forth constraints and conditions that define the terms of any form of future competition. And in this space of market behaviours our only residual control point lies in the judicious use of considered regulatory frameworks that encourages beneficial public good outcomes.

At best, I would label the "open" Internet an aspirational objective right now. We don't have one. It would be good if we had one, and perhaps, in time, we might get one. But the current prospects are not all that good, and talking about today's Internet as if it already has achieved all of these "open" aspirations is perhaps, of all of the choices before us, the worst thing we could do.

Today's Internet is many things, but it's certainly not an "open" Internet. It could be, but to get there it's not just going to happen by itself. It's going to need our help.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He 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 was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

www.potaroo.net