The Internet Invariants: The properties are constant, even as the Internet is changing

Leslie Daigle
Ldaigle@thinkingcat.com

May 16, 2019

Executive Summary

The Internet, itself in constant innovation since its inception, has historically supported unprecedented innovation across the globe, driving considerable growth in technology and commerce.This paper reviews a set of properties set out by Internet experts in 2012, which aimed to capture the unvarying properties that defined the Internet ("the Invariants"). The paper's exploration of the different levels of considering the Internet runs along the divide of "networking" versus "use" of the Internet, with a close look at the "application infrastructure" that couples them.

An early realization was that that the Invariants not only capture an ideal form of the Internet, they describe a generative platform — a platform capable of continuous growth and fostering the expansive development of new things upon itself. In 2012, the Internet had been developed and experienced as a generative platform for global communications activities and services. In reviewing the Invariants themselves, and examining today's Internet using the Invariants as a lens, it becomes apparent that the Invariants are still as relevant as ever, in describing the fundamentals of the Internet as a generative innovation platform, and, they are a constructive framework for enabling discussion of technology and policy choices, in terms of whether those choices would bring the Internet closer to, or push it further away from, the ideal of the Invariants.

Notably, several technologies being developed and deployed in today's Internet don't conform to those Invariants, and thus are not laying the foundation for similar innovations in the future. With the Invariants in hand, however, we have a tool to evaluate the state of the Internet and any proposed changes that would impact it, and support discussion between and among technologists and policy makers to help ensure that future choices foster a better Internet, aligned with the ideal expressed in the Internet Invariants.

This is "climate change" of the Internet ecosystem: absent concrete action to address the departure of the application infrastructure of the Internet from the ideal outlined in the Invariants, the experience of the Internet going forward will not feature such a rich diversity of solutions to the needs of the world's population.

Table of Contents

Introduction

That the Internet is changing and facing new pressures internally and externally is nothing new. Its flexibility and responsiveness to embrace changing circumstances are, in a large part, why it has firmly established itself as the global infrastructure of our daily lives. At some point, however, technical, commercial, and policy pressures may force adaptations that irrevocably break the Internet as we know it. Are we there now? How would we know? What can we do to avert catastrophe? What can we do to effect change to address real world challenges without causing catastrophe? How can we, collectively and constructively, discuss tradeoffs?

In 2012, a group of Internet experts discussed and enumerated the characteristics that described and seemed to define the Internet, as it had existed to that time. These were the features that made the network "the Internet", a platform that had seen incredible growth and fostered unprecedented innovation in the uses to which it was put. They were perceived as "unvarying" features, important to preserve in the Internet going forward. They were captured in the Internet Society's publication "Internet Invariants: What Really Matters"[1].

In the years since then, with various pressures brought to bear on the Internet and its services (social, economic, business, policy, political), one question is whether those properties were, indeed, unvarying: do they still describe the Internet as we know it today?

Initial review, in response to that question, highlighted that the Invariants actually capture the essence of the Internet "as a generative platform" — that is, a platform capable of increasing growth of itself and new and novel applications upon it. They remain a crisp articulation of the unvarying properties of a generative platform, and the question becomes: is the Internet still a generative platform?

This paper will review the origins of each of the Invariant characteristics in the context of the Internet, and then look at the state of the Internet today to review whether it still is well aligned with the characteristic. The purpose is simple: to determine whether or not the Internet of today remains a generative platform.

This review will be carried out at two levels: the baseline Networking technology, and

[1] https://www.internetsociety.org/internet-invariants-what-really-matters/

the Application Infrastructure that has been developed through open standards and the same development and deployment ethos as the network itself. The message is simple: if we want the Internet to remain successful as a growing platform supporting unprecedented innovation, these are the characteristics to preserve at those levels, in the face of technical, business and policy evolution.

Both the Networking and Application Infrastructure standards were developed in the Internet's early years, where collaboration was necessary in order to build anything that would scale to a globe-spanning solution. By contrast, much of today's experience of the Internet comes from using the commercial services that have driven the success of the Internet. These vertical solutions have grown to the scale where they can effect globe-spanning solutions single-handedly. The second part of this paper will explore this general class of User Experience in the light of the Invariants. While these applications and services were never intended to be anything other than their own solution, understanding the degree to which they do not, collectively, form a generative platform is important to understanding why the network and application infrastructure layers cannot succeed if they evolve to be less aligned with the Invariants, and more commercially owned.

Finally, the original purpose in describing and publishing the Invariants was to provide a tool for analysing prospective technical and policy issues to weigh the pros and cons of possible outcomes, in terms of their impact in bringing the Internet closer to (or pushing it further away from) the Invariants. Several case studies will be reviewed in the last section, to illustrate this application.

The Internet — in 2019

There is a considerable difference between what network engineers mean when they refer to "the Internet", and what the general public experiences as the Internet. The diagrams below describe the different layers of technology that combine to form each view.

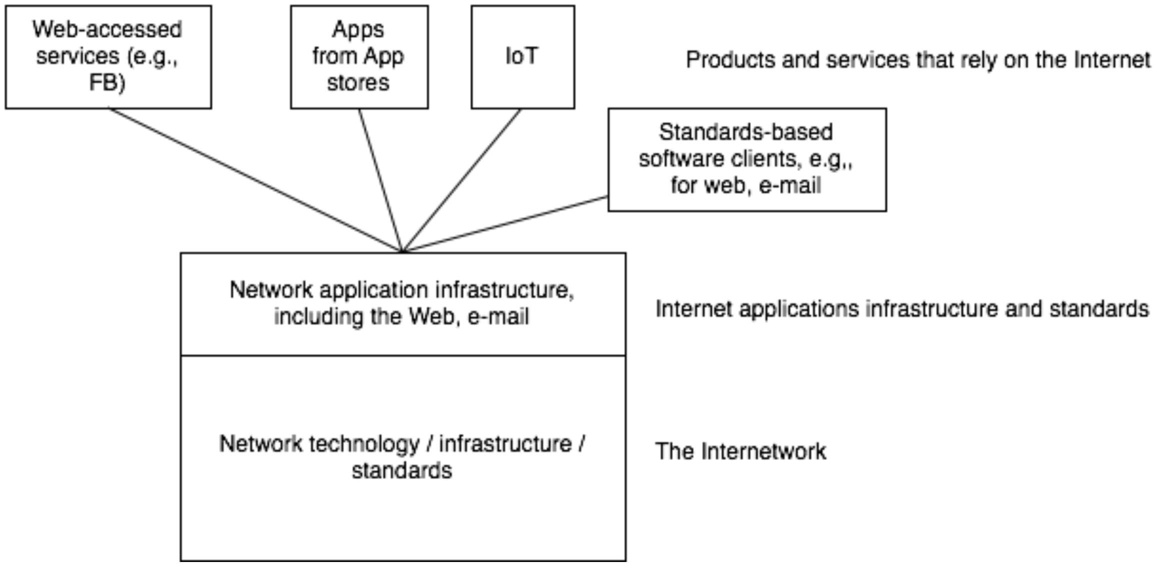

In the strictest sense, “the Internet” is the lowest layer in the diagram – the collection of technologies that combine to move packets about the network and across the globe. A slightly more generous interpretation includes the set of equally-openly-defined application services and infrastructure that run on top of that network technology – e.g., e-mail, the World Wide Web, etc. Throughout the rest of this document, we’ll refer to the former as “Networking” and the latter as “Application Infrastructure”.

Users’ perception, or experience, of the Internet is shaped by their use of products and services that, in turn, leverage the open technologies of the Internet (hereafter, “User Experience”). Facebook is a private corporate service that uses (primarily) WWW technology to deliver a particular experience to users. App store applications generally use Internet communications technology to connect to servers (including cloud services) to carry out their particular function.

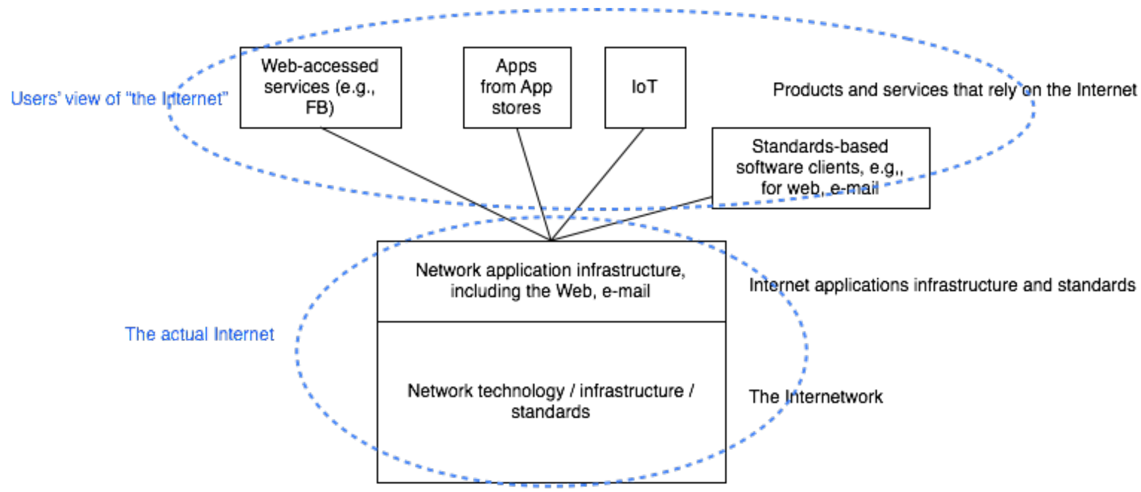

Both the engineering and the user's perception of the Internet are important for reviewing the state of the Internet and making choices that impact its future. Engineering choices will impact the technology and deployment of the Internet. Users' perceptions will affect corporate plans for future services and products, and will also be a major driver of local and global policymaking interest. Policies, in turn, shape the use and potential scope of the Internet and all the services that run on top of it.

However, only the Internet layers (Networking and Application Infrastructure) were built in open technologies that provided a "generative platform" — constantly evolving themselves, and allowing for the unfettered growth of innovative tools, products and services that run on top of them.[2] What most people experience as "the Internet" — the User Experience — is a collection of largely proprietary products, services, and focused ecosystems of technology. This collection evolved on its own, through disparate agencies, and not as a generative platform. Where the Internet continues to support the generation of new forms of User Experience, the latter technologies do *not*, as a general rule, support further evolution and innovation as a layer beyond. The User Experience can grow wider, not taller.

This paper will revisit the properties of the Internet that capture its ability to provide a

[2] It is also important to note that there are networks that use the Internet Protocol suite, and some variations of its technology, but which are deliberately controlled for a particular purpose. At various times, this has included mobile phone backhaul networks (which have since largely migrated to use the open Internet), large-scale networks operated by individual corporations (e.g., Amazon and Google), and national or regional networks that are heavily blocked at the network and/or application and content level. These networks may connect to the Internet in some fashion, but are not truly the Internet, nor do they espouse the properties reviewed in the following sections. For those reasons, they are not considered in scope for this document.

generative platform, reviewing the Networking and Application Infrastructure layers to see how well they align with the properties in 2019. It will also review the User Experience of the Internet through the same set of properties, to illustrate the degree to which it is not, and cannot be, a platform to generate a whole new wave of innovation. This distinction is important not just for its own sake, but also to understand how to foster more innovation (create alignment with the properties), and the dangers associated with allowing existing infrastructure to become subsumed in proprietary systems.

Part I

Review of the 2012 Invariants

In 2012, a group of Internet experts articulated a set of 8 properties that the Internet demonstrated; properties that appeared "unvarying" over time, and necessary for the Internet to maintain its primary essence. In retrospect, those properties describe a "generative platform" — a framework that supported not only its own continued growth and development, but also served as a basis for the creation of new and novel frameworks upon it. From its inception until the experts' discussion in 2012, the Internet (Networking and Application Infrastructure) was demonstrably a generative platform.

Those unvarying properties ("Invariants") were:

- Global reach, integrity
- General Purpose
- Supports innovation without permission (by anyone)
- Accessible
- Based on interoperability and mutual agreement
- Collaboration
- Technology — reusable building blocks
- There are no permanent favourites

At first glance, these may not sound as familiar as several technical principles that have been bandied about in policy discussions for years. The reality is, the developers of the Internet's technologies held a variety of technical principles to heart. Many of these continue to be held up as defining requirements of successful Internet technology, even as operational reality has undermined them. For example, the strict definition of "endpoint" has been unclear for a couple of decades (since the introduction of load-balancers for servers) – making the "end-to-end" principle a little hard to argue. Or, the call for a "dumb network, with smart edges" still holds a lot of merit, but the reality is that networks have challenged that on more than one front.

Thus, in articulating the 2012 Invariants, a different approach to characterizing the key facets of the Internet was taken, by articulating a set of features that had, to that point, been unvarying. These appeared to be the properties worth defending, in order to preserve the Internet as it was known.

These properties are reviewed below, in terms of today's Internet reality, to determine whether the Networking and/or Application Infrastructure are still aligned with the Invariants, and thus providing a generative platform for onward development. The collection of Invariants is considered, in the light of whether it should be updated (added to, reduced in number, modified).

Later sections will consider the User Experience, the components of which were never built with the intention of providing a generative platform, through the lens of the Invariants, to understand what to expect from future development at that level.

Finally, the updated Invariants will be used as a tool for reviewing current major technical and policy questions.

Global reach, integrity

> As written in 2012: "Any endpoint of the Internet can address any other endpoint, and the information received at one endpoint is as intended by the sender, wherever the receiver connects to the Internet. Implicit in this is the requirement of global, managed addressing and naming services."

Historical perspective

Networking technology was established to connect nodes, or "endpoints". Inter-networking was conceived to hook together networks into one global network connecting nodes. Application protocols were designed to connect node to node, using the same process for finding destination node addresses (initially "host tables", and then DNS lookup) and route to that address (hop by hop through neighbouring networks) irrespective of the physical or network location of either the source or the destination endpoints. These processes were part of the "network stack" in the operating system that promoted the initial deployment of the Internet – Unix.

This global reach meant that there was wholeness, or "integrity" of the network, which was further bolstered by the implicit expectation that networks would not inspect or modify packets they carried, beyond whatever was necessary to route the packet as advertised.

This global reach and integrity was so much part of the ethos of early design that applications (ultimately, Application Infrastructure) were designed to be similarly uniform even when operated in distributed instances across the globe — DNS, gopher and then the World Wide Web, and so on.

Networking, 2019

At the network layer, it is still mostly true that the Internet has global reach: in principle, it is still true that any Internet endpoint can send packets to any other endpoint. In practice, however, there are networks that intentionally prevent some connections; connection through a proxy or a gateway changes the Internet boundary; and some networks use Internet protocols but are deliberately constrained to the point of not being part of the Internet.

Connection prevention can happen at the network level by dropping packets for traffic using particular protocols (e.g., ICMP, UDP), or for particular destination ports (e.g., port 25). The destination endpoint simply never gets the packet. The motivation for implementing this sort of blocking often stems from an attempt to effect some kind of security policy in the network (preventing ICMP reflection attacks; reducing the opportunities for producing and sending spam e-mail), but the impact is nonetheless an erosion of the global reach and integrity of the Internet.

Proxies and gateways are well known in the context of home and enterprise network connections. Home gateways typically manage sharing of a single public IP address among connections from all devices within the home network. That is, while all devices in the home can access the Internet, the whole home network appears as one location on the Internet. Also, they typically include "firewalls", with built in rules to prevent network traffic from reaching a device in the home network, except for replies to communication initiated by the home device. That fits well with the conception of a home network that supports computers for individual family members, but as IoT has brought more lightweight devices into the home network, it has meant a need for external services in order for those devices to be configured or used from outside the network. This is both a complexity introduced to overcome the proxy blocking, and an opportunity for more security issues for users (who must rely on a multiplicity of device vendor services "in the cloud" to connect and manage their home devices).[3]

Proxies are perhaps less recognized in the context of "smart phones". With the introduction of Android and Apple devices, most modern day smartphones give every appearance of being "on the Internet", when in fact they are making use of mobile data networks to connect to a gateway that allows them access to the global Internet (but not vice versa). Even in the case of a computer that is "tethered" to the Internet through a mobile phone's access to its data network, the real point of connection to the Internet is somewhere in the physical hardware belonging to the mobile network operator. The computer is tethered to a proxy, at some level.

The invariant describes the need for a coherent, global addressing and naming services. IPv4 has been the standard addressing system for Internet connections since 1983. However, with only ~4 billion addresses in the entire space, it is inadequate in the face of the popularity and importance of the Internet. There are fewer IPv4 addresses than people on the planet, and there are more connected devices on the planet than people. For each device to be part of a globally coherent network, more addresses are needed. This has been a known problem for decades, and the successor technology (and address space) has been defined: IPv6. Continued

[3] https://www.cigionline.org/sites/default/files/gcig_no33web.pdf

deployment of IPv6, and eventual deprecation of IPv4, is imperative to ensure global reach and integrity of the Internet.

Application Infrastructure, 2019

There is some erosion of users' experience of the application infrastructure in terms of global reach and integrity.

With networks blocking individual application ports (e.g., port 25 for sending e-mail), users could struggle with applications that work in some networks but not others. In fact, users do sometimes see this in networks with known significant constraints (e.g., "My phone can't send e-mail from the corporate network"). To get around this, most applications and services are now designed to run on a very few well-known ports (e.g., port 80, which is for HTTP — web traffic). And, where "network neutrality" principles are not recognized or upheld there can be significant deprecation in the user's experience using applications in different networks, undermining the integrity of the global network.

Within the context of open standard applications themselves, the current status is somewhat mixed. E-mail more-or-less works globally, except for inconsistent implementation of internationalization standards (in domain names, and in e-mail addresses), and in overzealous implementation of spam filtering rules.

Generally speaking, the World Wide Web still has global reach — any host can reach any public website, globally. However, an individual user's experience of services reached over the web may be highly dependent on where the user sits. While a user can generally access a global service like Google or Amazon from any Internet connection in the world,[4] the actual endpoint they reach will be different for reasons of load balancing, optimization of traffic routing, and efficiency of delivery of content. The user's experience of the service may be deliberately tailored to their locale — as will be discussed in a later section.

Summary

[4] For the purposes of this discussion, the state of network/access filtering in China is such that it is arguably not part of the global Internet

	Networking, 2019	*Application Infrastructure, 2019*
Global reach, integrity	The network level of the Internet still features global reach, with reasonable integrity	Mixed bag — traditional, standards-based applications have global reach and integrity, but some are blocked by network choices, in the name of reducing unwanted traffic

Global reach and integrity remain core to the health of the Internet. It is, in so many words, the purpose of the Internet.

There is no inherent architectural reason that networks are not end-to-end addressable globally. Where there are strong enough motivations, the blocks that have been put in place can be removed. That said, the expectation of global reach and integrity has been greatly reduced and anyone building a service today is likely to build it on top of something that is known to work. As a result: so many things run over HTTP.

General Purpose

As written in 2012: "The Internet is capable of supporting a wide range of demands for its use. While some networks within it may be optimized for certain traffic patterns or expected uses, the technology does not place inherent limitations on the applications or services that make use of it."

Historical Perspective

The Internet was designed as a "dumb network" — any state (data) related to the communication activity was intended to be stored at one or both ends of the communication, and choices about managing the communication flow were also a responsibility of "the edge". This is distinct from circuit-switched telephony networks, where the endpoints are "dumb", and the network does the heavy lifting of managing the communication flow (reserving the circuit, monitoring performance, etc). One of the implications of this approach is that the network itself is not *a priori* designed for any particular application type that might run over it. This also led to the focus, in the 1990's, to extend the application naming and location approach introduced in the World Wide Web (Uniform Resource Locators, etc) — with application-independent naming and identification, the Application Infrastructure would continue to support a wide variety of future applications (I.e., was general purpose).

Networking, 2019

In principle, the Internet is still a general-purpose network. Efforts to develop and deploy network-shaping protocols have been more successful within dedicated networks than across the networks of the Internet.

In practice, however, a number of network optimizations and management practices have been deployed that have been based on Internet uses that were historically valid, as opposed to necessarily keeping the network as general as it might be. Early home broadband networks provisioned users with asynchronous links, allowing much greater download speeds than upload rates (and then were surprised when home users then started uploading bandwidth intensive data such as videos, *en masse*). Less visibly, NATs and other middleboxes have played a role in driving the Internet to become less general insofar as they are all built based on a retrospective look at network uses: they are designed and provisioned based on the way users *have* used the Internet, not the possible uses of the future. Efforts to deploy new network technologies that should have worked in a general purpose Internet based on open standards have been stymied by these middleboxes (e.g., quic, TLS1.3).

Application Infrastructure, 2019

Virtually all applications run over ports 80 and 443, as the least-commonly-blocked application ports. It is now easier to develop new standards, or revise old ones, to work within the confines of port 80 (and, often HTTP itself) than to continue to develop and expect to deploy entirely new standards for applications and services. That represents a significant level of constraint in terms of supporting entirely new uses of the Internet.

Also, the general purpose nature of the Internet has meant that hasn't always been fit-for-purpose for particular commercial uses. Commercial (I.e., proprietary) Content Delivery

Networks (CDNs) appeared first as a means to ensure that corporate website content was reliably and efficiently accessible for consuming users, whether across the continent or across the globe. Today, CDNs have become synonymous with delivery of video content. While the general purpose Internet can deliver all the packets necessary for playing videos, the transmission realities in the best-effort Internet packet delivery are such that the resulting delivery is not necessarily adequate for reliable and smooth video display.[5] Also, by positioning copies of video content close to consumers of the video, the Internet backbone isn't clogged with continuously and repeatedly shipping the same popular video content to multiple end users. So, arguably, CDN services have become an important part of allowing the Internet to deliver video. However, by addressing the efficiency of content delivery through proprietary technology, not open standards and services, the Internet's base technology level is not improved, and the advantages of efficient data delivery are not available to all applications and services.

Summary

	Networking, 2019	*Application Infrastructure, 2019*
General purpose	The Internet is still reasonably general purpose at the network level	Constrained by past uses and purposes (e.g., port blocking)

The Internet remains reasonably "general purpose", and needs to do so in order to continue to support the diverse range of purposes for which it is used across the globe. As more demanding uses are identified and become common (as did video delivery), the challenge is to ensure that the Internet's technology can support those uses, without requiring the development of specialization or proprietary overlays.

Supports innovation with requiring permission (by anyone)

[5] See graph of YouTube delivery improvement with CDNs in https://blog.apnic.net/2019/01/24/five-years-at-the-edge-recording-the-evolution-of-web-usage-from-an-isp/

As written in 2012: "Any person or organization can set up a new service, that abides by the existing standards and best practices, and make it available to the rest of the Internet, without requiring special permission. The best example of this is the World Wide Web – which was created by a researcher in Switzerland, who made his software available for others to run, and the rest, as they say, is history. Or, consider Facebook – if there was a business approval board for new Internet services, would it have correctly assessed Facebook's potential and given it a green light?"

Historical Perspective

When the original Internet networking technology was being defined, the challenge wasn't figuring out how to keep activities *off* the Internet, it was figuring out how to allow many diverse networks to join together to *make* the Internet, supporting the uses to which it was being put. It was not a question of "inclusion at all cost", and permission-free innovation did not extend an invitation to set up traffic to run rampant on the Internet or otherwise abuse its resources. Then, as now, inclusion was open to those hosts, networks, and applications that were willing to abide by the shared rules and agreements of the network ecosystem.

Networking, 2019

At the network layer, it is still largely true that Internet technology supports innovation without permission. Anyone can, in principle, send any type of packets, in whatever form works for the innovation. What happens to the packets in flight will be determined by the existing operational standards and practices.

The practical challenge is that, as the network operators have adopted various common practices (e.g., port blocking, as noted earlier), truly innovative network technologies can be stifled, and new applications must conform to use existing common ports if they don't want to be blocked by network practices.

Many of these operational choices, and even some development ones, are implicitly based on the assumption that the way protocols or parameters are used today is, in fact, how they are meant to be used forever. If a protocol was initially defined to use a fixed set of (say) five possible values for a parameter, software will be built that checks for those values, and *only* those values. If a certain manufacturer's hardware is known to ignore or flag an error whenever something other than those fixed values is encountered, it is practically impossible to update or extend the standard to include other parameter values. This is a significant enough concern that some protocols developers have opted to populate all values, requiring software to be more generous in its interpretation from the outset.[6]

The term "ossification" has been used to describe the Internet's technology — in the current era of stabilized network hardware and configurations (for so-called commercial grade performance and to manage capital and operational expenditures), the trap is "that abides by the existing standards and best practices". It's updating those standards and practices that is particularly difficult, absent the early research years' mentality of continuous update and adaptation.

Application Infrastructure, 2019

It is still largely true that new services and products can be developed and deployed on the Internet as a network, at will. This is especially true in terms of web-based services. There is nothing technological that stops someone from setting up new and innovative offerings via the web. It's not even particularly expensive to do so, in most parts of the world.

A person (or, more likely, an organization) can still stand up an information service anywhere, any time, and build new websites and new services from the ground up. This has allowed Facebook to have competitors, for instance (though, none are anywhere near the size or scope of Facebook at a global level). From new voice services, to chat services, to websites and online activities, new entrants are possible. They are not necessarily probable, and they are likely either to fail to gain sufficient notice to survive, or to be bought out by larger players in a defensive corporate move, but the operative point here is that the Internet does not *a priori* prevent it.

It is, nowadays, more challenging to deploy new services and offerings in user experience if the application runs afoul of the typical blocking of ports, as outlined above, or when the offering has data delivery requirements that are extreme, by current general purpose network standards. This sort of extreme data delivery requirement includes online gaming, where slight differences of perceived realtime can variously advantage or disadvantage players, or large quantities of data that must be transferred smoothly.

[6] For example, https://tools.ietf.org/html/draft-davidben-tls-grease-00 "Applying GREASE to TLS Extensibility", Benjamin, D.

Summary

	Networking, 2019	*Application Infrastructure, 2019*
Supports innovation without requiring permission (by anyone)	It is still possible to add new functionality to the Internet without requiring permission	Largely still true, though there is significant constraint to run over existing application ports and protocols (to avoid blocking)

Permission-free innovation remains core to the Internet's continued growth and evolution — at all levels.

Concerns have been expressed that permission-free innovation is a license to do anything with or to the Internet. As outlined in the full text of the invariant's description, that is not the case. The boundaries on "anything" include all applicable standards, best practices and regulations.

While perspectives certainly can, and do, vary as to whether innovations are necessarily always positive, the overall strength that this property gives the Internet is that the source of innovation is not limited. The next generation of Internet is not limited by the scope of the imagination and capabilities of any single set of experts. By not having to predict where the next improvement will come from, or what it might address, the Internet is open to the full scope of human imagination and endeavour.

Accessible

> As written in 2012: "It's possible to connect to it, build new parts of it, and study it overall: Anyone can "get on" the Internet – not just to consume content from others, but also to contribute content on existing services, put up a server (Internet node), and attach new networks."

Historical Perspective

The very purpose of the original Internet technology development was to figure out how to enable networks to connect. It was not devised so that *specific* networks could connect, but rather that any network willing to abide by the rules (standards and practices) could connect and become part of the whole.

Historical application protocol design is often criticized for its apparent disregard for security considerations. However, in the early days of the Internet, accessibility, traceability and transparency of protocols improved monitoring and debugging options, which were perhaps more immediate concerns at the time.

Networking, 2019

At the network layer, this is still largely true. It is still possible to "get on" the Internet, pretty much anywhere. You can set up a server (a reachable device capable of responding to incoming requests on chosen protocols). Anyone can still add networks to the Internet, by connecting with neighbouring networks through agreed practices and standard technology.

However, some networks are restricted in permitted use – consumer access networks have long stipulated that customers not run servers on their connections, and some ports are actively blocked. Coupled with the practice of dynamic assignment of IP addresses (and, possibly, port ranges in the case of so-called carrier grade NAT), home network users are relegated to the role of consumer of services from the network, not participants. Mobile data networks are entirely restricted, as noted earlier.[7]

For commercial access, it may be expensive to find a suitable connection point — finding a "neighbour network" means different things in different contexts, whether it's enterprises buying access from commercial access providers, or remote regions wiring up and connecting to distant partners.

This invariant property speaks also to the question of examinability and transparency of the network. Much network connectivity information is visible today – at least from the standpoint of routing announcements and registration of resources. Things that are opaque (and largely always have been) include backhaul and backup links, and some level of peering information — customer relationships are always sensitive information. For other information, various services are still heavily used by network operators and motivated independent engineers — such as Route Views,[8] and the RIPE Atlas[9] probe network.

[7] As a concrete test: tether two computers to the Internet via two separate phones on the same data service, and attempt to have one computer reach the other in some way, as if they were on the same network. Unlike the home network, this sort of direct host-to-host connection is not likely to work.
[8] From the Route Views homepage, http://www.routeviews.org/routeviews/ :
"The University's Route Views project was originally conceived as a tool for Internet operators

Application Infrastructure, 2019

It is still largely true that new services and products can be stood up, with the caveats mentioned above about some commercial restrictions applied to different types of access networks' use. This is especially true in terms of web-based services. There is nothing technological that stops someone from setting up new websites or other forms of communication. It's not even particularly expensive to do so, in most parts of the world.

In some parts of the world, this is definitely not true, as peoples' experience of the Internet is limited to particular applications (e.g., zero-rated Facebook or WhatsApp). In those contexts, users do not have the ability to fully engage in contributing to the Internet.

Summary

	Networking, 2019	*Application Infrastructure, 2019*
Accessible	Internet networks are largely still accessible.	The Internet is still largely accessible at the Application Infrastructure level where the network is still general purpose. Definitely not accessible in parts of the world where single specific apps are financially preferred (e.g., zero-rated)

---

to obtain real-time BGP information about the global routing system from the perspectives of several different backbones and locations around the Internet. Although other tools handle related tasks, such as the various Looking Glass Collections (see e.g. TRACEROUTE.ORG), they typically either provide only a constrained view of the routing system (e.g., either a single provider, or the route server) or they do not provide real-time access to routing data.
While the Route Views project was originally motivated by interest on the part of operators in determining how the global routing system viewed *their* prefixes and/or AS space, there have been many other interesting uses of this Route Views data. For example, NLANR has used Route Views data for AS path visualization and to study IPv4 address space utilization (archive). Others have used Route Views data to map IP addresses to origin AS for various topological studies. CAIDA has used it in conjunction with the NetGeo database in generating geographic locations for hosts, functionality that both CoralReef and the Skitter project support."

[9] https://atlas.ripe.net/

Accessibility remains an important feature of the Internet — for the historical reasons of fostering growth of the network itself, and to allow onward development of the technologies themselves.

Apart from the challenges to accessibility at the level of Application Infrastructure outlined above, threats to accessibility at the Internet level include nationalized service requirements, and extraterritorial legislation,[10] which set up barriers to entry for new services (too complex for a small entity to figure out compliance across multiple jurisdictions) and drive consolidation of services (to the large players that can take care of complex compliance). Other illustrations of potential challenges to accessibility include proposals to address recent concerns about China establishing network points of presence in North America to divert (and presumably monitor) traffic outside of its national boundaries. Some of the proposals include suggestions for denying companies the right to establish network points of presence in North America because they are Chinese[11]. There is no question that diverting traffic is wrong. But establishing (licensing) requirements for setting up networks would be wrong, insofar as it would undermine the accessibility of the Internet.

Based on interoperability and mutual agreement

> As written in 2012: "The key to enabling inter-networking is to define the context for interoperation – through open standards for the technologies, and mutual agreements between operators of autonomous pieces of the Internet."

Historical perspective

Rather than attempting to define and mandate a uniform network substrate within every network, the technical approach to defining the Internet focused on developing standards for interoperation *between* networks. These standards were not imposed, but rather they were recognized as norms to be adopted voluntarily and followed by mutual agreement. As a natural extension of the "mutual agreement" nature, it followed that the entities that would be agreeing

[10] https://www.internetsociety.org/blog/2018/10/splintering-the-internet-the-unintended-consequence-of-regulation/
[11] https://www.theglobeandmail.com/politics/article-china-telecom-hijacked-internet-traffic-in-us-and-canada-report/

to connect could participate in the definition of the standards by which they would connect - and, thus, the Internet is defined in specifications developed in an open standards process.

In this way, diversity was supported, while coherence was preserved. The emphasis on interoperability meant that diverse implementations could be supported (and were encouraged), while providing coherence in the system by specifying how different systems were meant to interoperate.

Similarly, the World Wide Web defined a uniform means of accessing and referring to content, rather than relying on proprietary software. Before that, there were Gopher, and Anonymous FTP. There were open e-mail relays, in the days before spam overran e-mail. All of these early pieces of Application Infrastructure were designed for interoperation and deployed in the expectation of making the Internet work through mutual agreement.

Networking, 2019

At the Networking level, this invariant largely still holds — it's pretty difficult to ship packets across heterogeneous networks, there does need to be interoperability between them.

The focus on interoperation, rather than uniformity of networking approach within and across networks, has meant that individual network operators are free to do whatever they want within the privacy of their own network (subject to any commercial or legal service requirements). While there are networking best practices, there is huge variability in network architecture, resource allocation, and traffic handling between networks. Some of that variability stems from the purpose of the network (e.g., transit network versus consumer access network), the geography and span of the network, and the vendor of the hardware that supports the network. Nevertheless, in spite of these variations within networks, the Internet as a whole continues to function because the inter-network focus is on interoperation defined within open standards.

Application Infrastructure, 2019

This is still mostly true in terms of Application Infrastructure technologies defined through an open standard process (e.g., at the IETF or the W3C). E-mail services are still designed to connect to other e-mail services the world over. The World Wide Web continues to be a distributed information service, with individual organizations managing their own resources and linking to others as necessary. Software clients are designed to work with the World Wide

Web — I.e., interoperable with any WWW server, no matter the service or organization providing the information. (This is where a distinction can be made between accessing a service over the WWW (e.g., Facebook or LinkedIn) versus using an app that has been purpose-built for the service. The apps are not interoperable — your Facebook app cannot be used in place of the dedicated apps for Twitter or LinkedIn to access those services).

And, yet, often in the name of managing security risks and exposure, service blocking does occur. For example, mail services have maintained and blocked connections from "known" spam sources for years. This has collateral damage insofar as legitimate e-mail users at those sources are then unable to send e-mail to some recipients. As a handful of e-mail providers emerge as the leaders of outsourced e-mail services, it becomes increasingly difficult to operate an independent e-mail server. Independent (small) e-mail services are less likely to be recognized as safe, and more likely to be blocked as a "security measure", undermining e-mail's global reach and utility as an interoperable infrastructure.

Summary

	Networking, 2019	*Application Infrastructure, 2019*
Based on interoperability and mutual agreement	The Internet still functions on mutual agreement and interoperability between networks	While still true, largely, at the level of Application Infrastructure, consolidation and the "winner takes all" mentality make it harder to operate application infrastructure independently

It is fair to say that commercial interests keep interoperation to a minimum: there's little business mileage in sharing with potential competitors. However, the network layer has not had much choice about interoperation to date, as there still is no reasonable possibility of a complete, globe-spanning, full mesh network from one provider.

As an illustration of the phase change at the application level, consider “single sign on” and “federated identity”. While there was active work on defining standards for identity management and federated identity services in the early 2000’s (Liberty Alliance), the concept of accessible, standards-based identity management servers operated by neutral third party services has not taken off. Instead, outside the context of an enterprise, most of the world knows “single sign on” as “using Facebook credentials to sign into other services” (thereby granting Facebook that much more access to their information habits).

Collaboration

> As written in 2012: “Overall, a spirit of collaboration is required – beyond the initial basis of interoperation and bi-lateral agreements, the best solutions to new issues that arise stem from willing collaboration between stakeholders. These are sometimes competitive business interests, and sometimes different stakeholders altogether (e.g., technology and policy).”

Historical perspective

Historically, the very process by which Internet technology was developed and advanced was highly collaborative. That collaborative spirit was (and remains) an important driving force for its development: the engagement of the network engineers developing hardware and operating networks in the development process is necessary to ensure that existing problems are tackled and pursued to find pragmatic solutions. Open standards processes were established to support and foster the collaboration needed to achieve those outcomes.

The networking technology and application infrastructure developed for the Internet was also based on collaboration. This was a conscious strategy, in an era when it wasn’t possible for a single entity to define and deploy a globe-spanning network. To span the geography of interested research networks, collaboration of multiple (diverse) participating networks would have to be harnessed.

And then, when numbering resources (IPv4 addresses) became scarce, due to the unexpected popularity and growth of the Internet, the natural step was to create open, bottom-up policy development fora and organizations to manage the allocation of resources based on those policies. This yielded the Regional Internet Registry (RIR) system, which serves to allocate numbering resources to this day.

For clarity, "interoperation and mutual agreement" focuses on operations of existing systems, and the *status quo*. However, to address new challenges and resolve problems, it is necessary for involved parties to step out of regular operations existing practices, and work together to find new methods, technologies, and means.

Networking, 2019

It is still true that the globe is spanned by a coherent Internet due to the collaboration of participant networks — even competing commercial interests. And, Network Operator Groups (e.g., NANOG, RIPE) are still popular platforms for sharing experience and expertise. When there are network-wide issues to address, competitive organizations can come together to find and implement solutions — e.g., the World IPv6 Launch in 2012,[12] and the Mutually Agreed Norms for Routing Security (MANRS)[13] today.

Application Infrastructure, 2019

Application Infrastructure continues to be developed, through open standards processes. This has been most evident as updates to existing infrastructure — DNS, WWW, and e-mail. Some steps have been taken — e.g., WebRTC, and SIP for voice over IP. Much implementation of Application Infrastructure software is done collaboratively — through Open Source Software (OSS). For example, the Apache Software Foundation[14] has been responsible for the pre-eminent web server software for two decades, and has done so by focusing on merit and collaboration.

And, yet, the level of collaborative development of application infrastructure has slowed considerably since the early days of the World Wide Web. As the interest in operating application infrastructure in distributed services, operated based on Interoperability and Mutual Agreement has declined, so has the impetus to develop such services, collaboratively. (See earlier remarks about efforts in federated identity management standards and services). What's harder to do now than it was in the early days is set up any kind of an "infrastructure" service – something where collaboration is required. Arguably, DNS couldn't be deployed today as it was built, because there isn't a successful operational business model that supports collaborative

[12] https://www.worldipv6launch.org/
[13] https://www.manrs.org/
[14] https://en.wikipedia.org/wiki/The_Apache_Software_Foundation

infrastructure operation, and few entities are accepted as neutral parties to operate infrastructure.

Summary

	Networking, 2019	*Application Infrastructure, 2019*
Collaboration	As the network is still heavily dependent on interoperability, network operators still engage in collaborative activities to address mutual, network-spanning issues.	While there is still some collaborative development in Application Infrastructure, consolidation and the "winner takes all" mentality make it harder to operate application infrastructure independently

By leveraging the combined skills and resources of a diverse collective, collaboration has addressed major Internet challenges from the outset — leading to successes well beyond what might have been imagined. As an ideal, it remains as applicable today as ever.

One of the challenging aspects of collaboration, however, is that it is very difficult to plan, map, or direct. For example, in 2007, several uninvolved Internet network experts collaborated to track down and resolve the 2007 Estonia cyber attack.[15] That was a normal network engineer reaction — see a problem, collaborate to fix the problem. However, for companies and governments looking to measure or at least describe the level of security of increasingly critical networks, "we'll figure it out" isn't a crisp enough answer.

Technology — reusable building blocks

> As written in 2012: "Technologies have been built and deployed on the Internet for one purpose, only to be used at a later date to support some other important function. This isn't possible with vertically integrated, closed solutions. And, operational restrictions on the generalized functionality of technologies as originally designed have an impact on their viability as building blocks for future solutions."

[15] https://www.wired.com/2007/08/ff-estonia/

Historical perspective

Standards for the Internet's underlying technologies have traditionally been developed and specified in terms of specific messages exchanged, expected ordering of messages in an exchange, and intended impacts. To a technologist, that means that the framework for re-use and/or building on top of the specification is clear. As a result, technologies designed for one purpose have been re-used as the basis for successive generations, as well as to support other innovations by other technologists (e.g., DNS has been used for telephone number lookup (ENUM)). Without this building block approach, each new innovation would have to be developed from the ground up, re-inventing solutions that had, effectively, already been addressed elsewhere (and perhaps not as successfully as the earlier iterations).

Networking, 2019

Internetworking is still largely carried out through standardized building block technology. Best practices are developed to articulate effective approaches to achieve desired networking goals.

And, general purpose networking technology has been used to create very specific networks — whether IP-based delivery of content within access providers, or other forms of specialized networking.

Application Infrastructure, 2019

Those applications that have been supported by the development of open standards (e.g., e-mail, WWW) do still feature this building-block technology, and they are extended and re-purposed continually.

However, to the extent that application infrastructure services have become commodity services, and in some cases dominant players have emerged in the resulting markets, these tools become less flexible and less open to collaborative development and reuse — I.e., less like building blocks. As large providers of e-mail services increasingly block connections from independent servers, e-mail becomes more like a proprietary application service than infrastructure, for example. This tension is not new: concerns over the business handling of domain name registrations caused concerns of unfair monopolies and failure to adequately respect trademark law. The community concern and responses drove the development of the DNS registry/registrar role split in the late 1990's, and the creation of ICANN to provide an environment for ensuring the ongoing operation of DNS registrations in the interest of the

Internet as a whole. (Note that is scoped to obtaining (registering) domain names, and resolution at the root and top level, and provides little or no constraint on local resolution of domain names in enterprises, at ISPs and at open resolvers such as Google's "8.8.8.8").

Summary

	Networking, 2019	*Application Infrastructure, 2019*
Technology — reusable building blocks	The Internet's network technology is still largely based on (open standard) technology building blocks	Application infrastructure is still built with building-block technology, but it is increasingly offered through proprietary services.

The historical perspective, and thus the invariant, is still applicable in today's Internet. Where things have gone sideways is at the level of things built on top of the network.

There are no permanent favourites

As written in 2012: "While some technologies, companies and regions have flourished, their continued success depends on continued relevance and utility, not strictly some favoured status. AltaVista emerged as the pre-eminent search service in the 1990's, but has long-since been forgotten. Good ideas are overtaken by better ideas; to hold on to one technology or remove competition from operators is to stand in the way of the Internet's natural evolution."

Historical perspective

The Internet was developed in a research and academic environment, where the pursuit of better ideas was an important motivator. In order to advance from "working with regular interruptions and challenges", to "functions reliably and expected", change was necessary and accepted. There was no sense that the Internet's technology was "done", rather that successive iterations of development and deployment were necessary and appropriate.

One other facet of the historical perspective that is important here is that the network landscape presented a very heterogeneous reality in the early days. The purpose of the *inter*network was to hook these disparate networks together for common purposes. The diversity of the networks turned out to be an advantage — while one network technology or use might be predominant at a given moment, another would be in development. Hard cutovers were rare, and challenging to engineer. Support for this diversity meant, among other things, they were not always necessary.

Networking, 2019

At all levels, this is less true now than 20 years ago. Consider it part of the "ossification" or stabilization of the Internet.

At the physical level of networking, it is hard to lay out new physical infrastructure to compete with incumbents, whether it's for consumer access, or trans-oceanic transit, because access to right-of-ways may be restricted, and the cost of deploying enough physical infrastructure to reach the size of incumbents is prohibitive. Competitors may be able to get a toehold where there is a change in technology medium — e.g., when fixed wireless becomes a serious competitor with fibre to the home — but otherwise there is little room for new entrants in regions of developed networks.

The regulatory environment plays a role in shaping the extent to which new entrants can get a start. Wireless spectrum and nation-spanning fibre backbones are natural monopolies. On the other hand, a regulatory regime that 'unbundles' provision of wholesale connectivity from residential and enterprise 'retail' network provision can help to ensure some competition in the marketplace.

Corporate realities are such that network operators can't be changing out their equipment and technology on a regular basis. Not only is it prohibitive in terms of capital expense, reaching all corners of the network can be challenging. A clear illustration of this was provided by access network operators gearing up for the World IPv6 Launch in 2012. Not only did the access network operators have to change equipment throughout their networks, in all neighbourhoods they served, but there was often a need for CPE updates — I.e., changing the "customer premises equipment", or "CPE" or "modem". If the access provider didn't own the equipment, they had little leverage to get customers to change it.

Apart from the capital expense and challenges of physically changing out equipment throughout a network, even updates to existing hardware can be challenging. Updates have to be tested and thoroughly vetted, rolled out through a managed process throughout a network (and rolled back, if issues are discovered). Finally, a change in technology (hardware or software) means a change in operational practices, which further means updated training and changed procedures, adding to operational costs.

Application Infrastructure, 2019

At the level of Application Infrastructure, the question of whether there are still “no permanent favourites” yields a mixed bag of results. While it is true that no single application infrastructure is being unduly propped up to prevent competition, it is also true that there is less evidence of “better” overtaking “good” in this space. Mostly, there are updates to existing application infrastructure protocols, and some attrition of services (Jabber for instant messaging has been taken over by proprietary platforms, for example).

Summary

	Networking, 2019	*Application Infrastructure, 2019*
There are no permanent favourites	Networks are becoming more stable / ossified, and it’s harder to make significant changes	There is nothing inherent in technology that makes it less true at the Application Infrastructure level, but scale and business practices are making it hard to dislodge incumbents.

The historical perspective that led to this invariant still applies: the Internet is not “done”, insofar as there are still problems to solve, innovations that could improve it at both the network and user experience level. This invariant is still applicable, even as the existing Internet is slipping away from realizing it.

Diversity remains an important feature of the Internet. As in many ecosystems, supporting or being supported by a broad base of different systems makes it harder to harm the ecosystem in a permanent way.

Some conclusions from the review of the 2012 Invariants

This section provides some conclusions about the collection of Invariants as an expression of "the ideal Internet", as well as some final reflections on how the current Internet aligns with that ideal.

Briefly, the Invariants appear still to capture important aspects of the Internet's design, development and deployment. As such, they remain a relevant characterization of the "ideal" Internet, as a generative platform. However, the degree to which the Internet still aligns with the Invariants does depend on whether one examines the Networking layers of the Internet, or focuses on the Application Infrastructure. The latter is seeming less aligned with the generative platform ideal than it was in 2012.

The Invariants as an expression of the Ideal Internet

The review of each of the individual 2012 Invariant properties of the Internet considered the historical motivation for the property, and included some reflections on the applicability of the Invariant as one considers the Internet, today. Before reviewing the collective assessment of the current Internet as modeled against the Invariants' ideal, a first question to consider is whether the ideal described by the Invariants is still relevant.

From the review in the preceding sections, it does seem that the Invariant properties, with some modest update to text descriptions, do seem still to capture the essence of the Internet in ideal form, as a generative platform for innovation.

	Relevance and applicability of the ideal	*Expression, updated as necessary*
Global reach, integrity	Global reach and integrity remain core to the health of the Internet. It is, in so many words, the purpose of the Internet	Any endpoint of the Internet can address any other endpoint, and the information received at one endpoint is as intended by the sender, wherever the receiver connects to the Internet. Implicit in this is the requirement of global,

		managed addressing and naming services.
General Purpose	The general purpose nature of the Internet remains important — if it became highly specialized, or if parts of it were highly specialized, its overall utility would be lessened.	The Internet is capable of supporting a wide range of demands for its use. While some networks within it may be optimized for certain traffic patterns or expected uses, the technology does not place inherent limitations on the applications or services that make use of it.
Supports Innovation without Permission (by anyone)	It is still important that the next generation of Internet is not limited by the scope of the imagination and capabilities of any single set of experts. By not having to predict where the next improvement will come from, or what it might address, the Internet is open to the full scope of human imagination and endeavour.	Any person or organization can set up a new service, that abides by the existing standards and best practices, and make it available to the rest of the Internet, without requiring special permission. The best example of this is the World Wide Web – which was created by a researcher in Switzerland, who made his software available for others to run, and the rest, as they say, is history.
Accessible	Accessibility remains a key tenet of the Internet — at the level of Networking and Application Infrastructure.	It's possible to connect to it, build new parts of it, and study it overall: Anyone can "get on" the Internet – not just to consume content from others, but also to contribute content on existing services, put up a server (Internet

		node), and attach new networks.
Based on interoperability and mutual agreement	The invariant text still stands as a description of Internet-level reality. As later sections of this document will explore, there are hard questions to ask whether there is the desire to see User Experiences of the Internet mimic the openness and tractability of the Internet layer (I.e., through a similar ideal)	The key to enabling inter-networking is to define the context for interoperation – through open standards for the technologies, and mutual agreements between operators of autonomous pieces of the Internet
Collaboration	By leveraging the combined skills and resources of a diverse collective, collaboration has addressed major Internet challenges from the outset — leading to successes well beyond what might have been imagined. As an ideal, it remains as applicable today as ever, even as it is challenging to plan, map or direct.	Overall, a spirit of collaboration is required – beyond the initial basis of interoperation and bi-lateral agreements, the best solutions to new issues that arise stem from willing collaboration between stakeholders. These are sometimes competitive business interests, and sometimes different stakeholders altogether (e.g., technology and policy).
Technology — reusable building blocks	Technologies designed for one purpose have been re-used as the basis for successive generations, as well as to support other innovations by other technologists. Without this building block approach, each new innovation would have to be	Technologies have been built and deployed on the Internet for one purpose, only to be used at a later date to support some other important function. This isn't possible with vertically integrated, closed solutions. And, operational restrictions on the generalized

	developed from the ground up, re-inventing solutions that had, effectively, already been addressed elsewhere (and perhaps not as successfully as the earlier iterations).	functionality of technologies as originally designed have an impact on their viability as building blocks for future solutions.
There are no permanent favourites	The Internet is not "done", insofar as there are still problems to solve, innovations that could improve it at both the network and user experience level. Diversity is also an important feature of the Internet. As in many ecosystems, supporting or being supported by a broad base of different systems makes it harder to harm the ecosystem in a permanent way.	A healthy Internet ecosystem relies on diversity — of technologies, companies, environments that foster those that become or continue to be relevant and useful. Good ideas are overtaken by better ideas; to hold on to one technology or remove competition from operators is to stand in the way of the Internet's natural evolution.

Summary report card of the Invariants, 2019

Having concluded that the Invariants provide a useful model of the Internet as a generative platform, let's consider the summary of the individual Invariant reviews from earlier sections.

Invariant	**Networking, 2019**	**Application Infrastructure, 2019**
Global reach, integrity	The network level of the Internet still features global reach, with reasonable integrity	Mixed bag — traditional, standards-based applications have global reach and integrity, but some are blocked by network choices, in the name of reducing unwanted traffic
General purpose	The Internet is still reasonably general purpose at the network level	Constrained by past uses and purposes (e.g., port blocking)
Supports innovation without requiring permission (by anyone)	It is still possible to add new functionality to the Internet without requiring permission	Largely still true, though there is significant constraint to run over existing application ports and protocols (to avoid blocking)
Accessible	Internet networks are largely still accessible.	The Internet is still largely accessible at the

		Application Infrastructure level where the network is still general purpose. Definitely not accessible in parts of the world where single specific apps are financially preferred (e.g., zero-rated)
Based on interoperabilitiy and mutual agreement	The Internet still functions on mutual agreement and interoperability between networks	While still true, largely, at the level of Application Infrastructure, consolidation and the "winner takes all" mentality make it harder to operate application infrastructure independently
Collaboration	As the network is still heavily dependent on interoperability, network operators still engage in collaborative activities to address mutual, network-spanning issues.	While there is still some collaborative development in Application Infrastructure, consolidation and the "winner takes all" mentality make it harder to operate application infrastructure independently
Technology — reusable	The Internet's network	Application infrastructure

building blocks	technology is still largely based on (open standard) technology building blocks	is still built with building-block technology, but it is increasingly offered through proprietary services.
There are no permanent favourites	Networks are becoming more stable / ossified, and it's harder to make significant changes	There is nothing inherent in technology that makes it less true at the Application Infrastructure level, but scale and business practices are making it hard to dislodge incumbents.

Reflections on the Invariants' review

The Invariants-based report card for the current Internet tells a pretty clear story. The Networking layer itself is pretty stable, if ossifying. However, where the Application Infrastructure layer was originally designed with the same generative platform principles in mind, it is slipping further from that ideal expressed in the Invariants. T

This disconnect is not the fault of any individual actor, but is rather the outcome of decisions being made outside the framework of collaboration that fostered the Internet's development. The more proprietary solutions are built and deployed instead of collaborative open standards-based ones, the less the Internet survives as a platform for future innovation.

Such proprietary solutions are feasible now in ways that were impossible two decades ago. Corporate scale and computing power are such that a single organization can own and operate global-scale information services and networks. The easier path, today, is for individual companies or consortia to create a product or service and then work to become the dominant (I.e., winning) solution. At least, that's easier than the effort required to engage in the cross-community, open, collaborative processes that yielded the Internet.

The real loss, however, is that the resultant collective Internet technologies, which don't conform to the generative platform ideal of the Invariants, will not continue to grow and show the kind of innovation that has been the Internet's hallmark to date. It's not possible to tap into the global pool of expertise to foster new developments when the User Experience level is entirely built of non-interoperable proprietary technology.

This is "climate change" of the Internet ecosystem.

Part II

The User’s Experience of the Internet

Previous sections have reviewed the Internet’s Networking and Application Infrastructure, which were developed and deployed in a manner consistent with providing a platform for growth and continued innovation. Collectively, they have enabled the seemingly-endless stream of applications and services that have appeared on the Internet in the past twenty years — Facebook, Google, Amazon, Facetime, app stores. At the same time, websites, e-mail and other application infrastructure services became more critical to business activity. As this criticality increased, websites, e-mail services and all the related infrastructure became more feature-rich (complex) and challenging to manage for operational reliability, scale and security. “Outsourcing” became a thing, and specialized support services became commercially viable (e.g., CDNs, Wordpress for website and blogging platforms). Even though each of these products and services was developed independently, as proprietary systems, it was the openness of the Internet’s core Networking and Application Infrastructure made them possible.

When people talk about “the Internet” in 2019, they are including those built up applications and services, or maybe even thinking of them exclusively — that’s where and how they use the Internet, as the days of configuring one’s own network connections are long past. Consequently, much discussion of new technical directions and policy issues for the Internet focus on the applications and services. But, unlike the Networking and Application Infrastructure discussed so far, this “User Experience” of the Internet is based on (largely) proprietary vertical services which, necessarily, have not been built using the same ethos and collaborative approach needed to create a platform that aligns with the Invariants. This is a sign of success — it is a testimony to the power of the Networking and Application Infrastructure in supporting myriad new things.

However, it is important to understand certain implications of the differences brought to bear by the fact that the User Experience of the Internet is different than the generative platform discussed in the previous sections:

1. The collected applications and services that make up the User Experience of the Internet today do not individually or collectively provide a generative platform for the kind of expansive growth and innovation that has been the hallmark of the Internet to date.
2. As application and service operators grow to have global span, and take on some of the

services that have been operated as part of Application Infrastructure, the consequence is that pieces of application infrastructure become more like proprietary services than part of the Application Infrastructure platform, supporting the Invariant properties.

3. In seeking to modify and address technical and policy issues related to the User Experience of the Internet — arguably, where most of today's issues lie — it is important to ensure that proposed paths forward don't reach further than the applications and services themselves to impact the Networking and Application Infrastructure's inherent Invariant characteristics.

To explore that further, the next sections review the User Experience of the Internet through the lens of the Invariants — even though these applications and services were not built as a collective effort, it is important to understand in some detail how they differ from the open Internet systems that have provided a generative platform for growth and innovation.

Global reach, integrity

As noted in reviewing Application Infrastructure in 2019, an individual user's experience of services reached over the web may be highly dependent on where the user sits. While a user can generally access a global service like Google or Amazon from any Internet connection in the world,[16] their experience of the service is deliberately tailored based on the service's detection of locale. Google defaults to the language of the local country, Amazon directs to the local "store" which carries products available for delivery in, and pricing that is applicable for, the particular geographic region or country. Typically, a user can access service in another region (e.g., a user in the US can see Google search results from France, using http://www.google.fr), although some offerings may not be available (e.g., Amazon sellers may not ship to the user's location). These are choices made by the company offering the service, and while they may mean the user experience is not *uniform* across the globe, they do not undermine its global reach or integrity because users can access other regions' content. Less consistent are services offering content with geographically restricted intellectual property rights. For legal reasons, the Netflix library varies from region to region. If a user can see the content, but not access it when out of its region, we might say that's an issue outside the scope of the Internet. If the user can't even see the existence of the content outside of a given region, one could argue it is less supportive of a global network with integrity.

As a general rule, any Internet-accessible application or service can be used globally, unless specifically blocked by a network or geofenced by the application/service provider itself. Being operated by a single entity, a service has an element of integrity ensured, even if it is not uniform.

Looking at the span of Internet applications and services offered individually, it can also be noted that they don't collectively form an application platform with global reach and integrity. The simple fact that users' access credentials vary from service to service, and that it is difficult for users to share their own data between services, belies a lack of integrity in the universe of applications and services.

[16] For the purposes of this discussion, the state of network/access filtering in China is such that it is arguably not part of the global Internet

As more Application Infrastructure services are outsourced and consolidated, they become part of this vertical silo world. Increasingly as e-mail providers consolidate, the likelihood of “spam prevention measures” causing e-mail from smaller independent sites to be blocked is also on the rise.

If a user’s experience of the digital world is uniquely through applications available for their mobile devices, the reality is starker still. The range of applications available through “app stores” varies from one region to another. Typically, a user is bound to the app store of their home region, so they may not even be able to access an application that is scoped to the region in which they are traveling.

Summary

	Invariant	*User Experience, 2019*
Global reach, integrity	“Any endpoint of the Internet can address any other endpoint, and the information received at one endpoint is as intended by the sender, wherever the receiver connects to the Internet. Implicit in this is the requirement of global, managed addressing and naming services.”	Application vertical silos are growing, undermining the integrity of User Experience across the globe.

General Purpose

There are so many applications and services available over the Internet, it’s hard not to think of it as “general purpose” at the User Experience level. However, unlike the Networking level, which can literally be put to any use that packet sharing can achieve, a user's experience is limited to the products and services that are made available. Users are consumers.

Importantly, the set of applications and services that are made available today are not themselves a platform for new uses. While people may make use of Facebook for unexpected things (raising awareness for campaigns, community news services, cat videos), that's not the same thing as building a voice communication service that undermines a global industry (as Skype did with telephony, building on the Networking and Application Infrastructure of the Internet).

So, rather than saying the User Experience of the Internet supports a "general purpose" reality, it's fair to say that it's "multi-purpose", but finite.

Summary

	Invariant	*User Experience, 2019*
General purpose	The Internet is capable of supporting a wide range of demands for its use. While some networks within it may be optimized for certain traffic patterns or expected uses, the technology does not place inherent limitations on the applications or services that make use of it	Limited to the (admittedly vast) set of applications and services that have been built to run on the Internet.

Supports innovation with requiring permission (by anyone)

Coming back to the fact that many people experience the Internet uniquely through a set of proprietary services — whether web-based services (e.g., Facebook) or application platforms (e.g., iPhone, Android) or now IoT support services — it is important to note that there is very little room for innovation and development *upon* these proprietary frameworks. That is, while the open standards of Internet application infrastructure make it an ideal substrate for the development of new technologies and innovations, proprietary services are not. Even services that offer "APIs" to allow other services to access them (e.g., Twitter, Facebook) challenge those other services' innovation because of the API-offerer's natural control of changes to the API or even its existence. As far as business practices are concerned, there is nothing wrong with application platforms controlling their APIs. From the standpoint of User Experience defining a perspective on the Internet, these proprietary APIs fail to support innovation without permission (from anyone).

Summary

	Invariant	*User Experience, 2019*
Supports innovation without requiring permission (by anyone)	Any person or organization can set up a new service, that abides by the existing standards and best practices, and make it available to the rest of the Internet, without requiring special permission. The best example of this is the World Wide Web – which was created by a researcher in Switzerland, who made his software available for others to run, and the rest, as they say, is history.	Pretty much non-existent

Accessible

At the level of applications and (commercial) services, the Internet is not particularly accessible — it's generally not possible to "pop the hood" and peer inside them, which you can do with the general Internet. Much discussion has been made of services' "algorithms", the proprietary approach to using data to provide their services. These "algorithms", and even most of the data sources upon which they operate, are closely held corporate secrets.

Also, whether or not you can get to a given app or service is dependent on where you are, and what you are willing to pay. While Internet access itself is a paid service, much effort was put in at the outset of the Internet's public success to endeavour to ensure that access is affordable. Net net, no one can get on all apps and services, and everyone is consuming rather than building with many of them.

Summary

	Invariant	*User Experience, 2019*
Accessible	It's possible to connect to it, build new parts of it, and study it overall: Anyone can "get on" the Internet – not just to consume content from others, but also to contribute content on existing services, put up a server (Internet node), and attach new networks.	Mixed bag — the User Experience of the Internet is opaque. However, much of it is still accessible in terms of consumption.

Based on interoperability and mutual agreement

Where Internet Application Infrastructure was built to support global applications at a time when it was not feasible to have a single entity providing a monopoly global application or service, the last two decades have seen the emergence of just such services. We have individual services that have come a lot closer to providing the globe-spanning information system than was possible in the late 1980's.

In today's application and service world, apart from areas where cooperation is absolutely needed (e.g., certificates), the name of the game is not interoperation. Everything is winner-take-all. For example, Facebook and Twitter initially provided public "APIs" that allowed third party developers to create apps that would read and/or update users' Facebook and Twitter feeds. However, as those companies developed their own apps, they shut down the APIs' functionality so that the third party apps were not competitive with the platforms' own offerings.

Summary

	Invariant	*User Experience, 2019*
Based on interoperability and mutual agreement	The key to enabling inter-networking is to define the context for interoperation – through open standards for the technologies, and mutual agreements between operators of autonomous pieces of the Internet	While it exists in specific instances of mutual collaboration, the general reality is that interoperability and mutual agreement are almost non-existent between applications and services that make up the User Experience of the Internet.

Collaboration

In the world of proprietary applications, services and app stores, height matters: verticals win. Commercial emphasis is on building (and/or buying) the most complete independent silo, for everything from CDNs to app stores to social networking. Commonly, new, small competitive services are bought to be integrated or killed.

There is very little collaborative spirit at the level of Internet application software and services offered commercially. Existing technologies based on application infrastructure standards interoperate, but new developments have been largely undertaken by closed environments — consortia, or individual proprietary companies.

Summary

	Invariant	*User Experience, 2019.*
Collaboration	Overall, a spirit of collaboration is required – beyond the initial basis of interoperation and bi-lateral agreements, the best solutions to new issues that arise stem from willing collaboration between stakeholders. These are sometimes competitive business interests, and sometimes different stakeholders altogether (e.g., technology and policy).	Very little perceived need for, and thus actual existence of, collaboration amongst applications developers and service operators

Technology — reusable building blocks

The typical User Experience of the Internet, beyond e-mail and general web surfing, is through specific social media services and various device "app stores" and their apps.

Proprietary applications are far less likely to feature this ability to be repurposed, because they are purpose-built and defined by the needs of the creator. New innovations in the commercial application space do have to reinvent everything from the ground up, and not always as effectively (or efficiently, or securely) as previous systems. Thus, these products and services are not the building blocks for another layer of technology beyond their current purpose.

A significant concern is that there have been aspects of application infrastructure that should (or, at least, could) have been developed through open standards and adoption, but have not been. Identity management is one such case. While standards were developed (cf Liberty

Alliance),[17] they didn't achieve adequate uptake, and instead of open standards-based federated identity management, we have much more significant use of platforms' credentials being used across the board (e.g., Facebook or Google logins). Apart from any potential issues with that approach, it does mean that these technologies are not available for extension or improvement in any next turn of the development crank.

Looking forward, as various trends are pushing for more software, less hardware implementation of network functions (Software Defined Networking (SDN), network function virtualization (nfv)), it will be important to see if the Internet level trends similarly away from reusable technology building blocks.[18]

.

Summary

	Invariant	*User Experience, 2019.*
Technology — reusable building blocks	Technologies have been built and deployed on the Internet for one purpose, only to be used at a later date to support some other important function. This isn't possible with vertically integrated, closed solutions. And, operational restrictions on the generalized functionality of technologies as originally designed have an impact on their viability as building	The applications and services that make up the User Experience of the Internet are typically vertical, integrated, closed solutions, which do not permit their reuse as building blocks.

[17] http://www.projectliberty.org/

[18] See also the Internet Architecture Board workshop on "Semantic Interoperability" (https://www.iab.org/activities/workshops/iotsi/) for reflections on the state of Internet of Things and interoperability. Also, see the case study below.

blocks for future solutions.	

There are no permanent favourites

It is still true that a company failing is not going to end the Internet. There might be significant vacuum and disruption in peoples' day-to-day lives, and corporate functioning, however. While one can believe that Facebook (or Instagram or Twitter or any other social media platform) could cease operations, the impact would be a scattering of contacts, as different people head to different alternative platforms. But, the vacuum would be filled.

It is, however, harder for new entrants to overtake existing services, the way Facebook and Google did. That's partly by business design of those organizations, but also in part because it's very hard to get noticed in today's world of silo'ed applications, and large corporations buying out perceived competition before it has a chance to develop a threatening following. As a concrete example, Google doesn't seem poised to lose ground in the search market, any time soon.[19]

Summary

	Invariant	*User Experience, 2019*
There are no permanent favourites	A healthy Internet ecosystem relies on diversity — of technologies, companies, environments that foster those that become or continue to be relevant and useful. Good ideas are	There is nothing that props up any given product or service that make up the User Experience, and new products and services do get built on the open platform of the Internet and its Application Infrastructure. However, it is difficult for new

[19] https://sparktoro.com/blog/2018-search-market-share-myths-vs-realities-of-google-bing-amazon-facebook-duckduckgo-more/

overtaken by better ideas; to hold on to one technology or remove competition from operators is to stand in the way of the Internet's natural evolution.	entrants in existing services to get noticed in the shadow of major players. (So, it's "easier" for Uber or AirBnB to establish themselves and grow than for a new social media platform to unseat Facebook, for example).

Reflecting on Internet User Experience through the lens of the Invariants

As noted earlier, the products and services that provide the primary User Experience of the Internet today were never designed with the intention and expectation of creating a generative platform. As a result, it is hardly surprising that today's User Experience doesn't align well with the Invariants.

However, even if there is a clear demarcation in the stylized diagram of Internet levels provided earlier, in reality there is no bright line between "Application Infrastructure" and "User Experience" services. The distinction is increasingly blurred as traditional Application Infrastructure elements (e.g., DNS and email services) are offered commercially, based on proprietary systems offering scale and efficiencies — and a little less openness, transparency, and support of the Invariants.

Where the Application Infrastructure in 2012 helped give rise to the characterization of the Internet through the Invariants, it is far less conformant today. In this paper's review of the Internet's alignment with the Invariants at all 3 levels, it has become clearer that the Application Infrastructure level of Internet is trending to be more like the (non-generative) User Experience level.

This trend bodes ill for future innovations of products and services running on the Internet — the less the Application Infrastructure is developed and extended through the open and collaborative principles that created it, and the less it presents a generative platform for building up on it, the fewer options there are for innovation at the User Experience level.

Of course, not everyone has given up on development of new Application Infrastructure to address current issues — Sir Tim Berners-Lee launched "Solid":[20]

> Solid (derived from "social linked data") is a proposed set of conventions and tools for building decentralized social applications based on Linked Data principles. Solid is modular and extensible and it relies as much as possible on existing W3C standards and protocols.

The remaining question is whether Solid, and other Application Infrastructure efforts, can find the right factors for uptake in the modern Internet realm.

[20] https://solid.mit.edu/

Collecting all the individual assessments, this is how the User Experience of the Internet stacks up against the properties that characterize the underlying network of the Internet.

Invariant	**Networking, 2019**	**Application Infrastructure, 2019**	**User Experience, 2019**
Global reach, integrity	The network level of the Internet still features global reach, with reasonable integrity	Mixed bag — traditional, standards-based applications have global reach and integrity, but some are blocked by network choices, in the name of reducing unwanted traffic	Application vertical silos are growing, undermining the integrity of User Experience across the globe.
General purpose	The Internet is still reasonably general purpose at the network level	Constrained by past uses and purposes (e.g., port blocking)	Limited to the (admittedly vast) set of applications and services that have been built to run on the Internet.
Supports innovation without requiring permission (by anyone)	It is still possible to add new functionality to the Internet without requiring permission	Largely still true, though there is significant constraint to run over existing application ports and protocols (to avoid blocking)	Pretty much non-existent
Accessible	Internet networks are largely still accessible.	The Internet is still largely accessible at the	Mixed bag — the User Experience of the Internet

		Application Infrastructure level where the network is still general purpose. Definitely not accessible in parts of the world where single specific apps are financially preferred (e.g., zero-rated)	is opaque. However, much of it is still accessible in terms of consumption.
Based on interoperabilitiy and mutual agreement	The Internet still functions on mutual agreement and interoperability between networks	While still true, largely, at the level of Application Infrastructure, consolidation and the "winner takes all" mentality make it harder to operate application infrastructure independently	While it exists in specific instances of mutual collaboration, the general reality is that interoperability and mutual agreement are almost non-existent between applications and services that make up the User Experience of the Internet.
Collaboration	As the network is still heavily dependent on interoperability, network operators still engage in collaborative activities to address mutual, network-spanning issues.	While there is still some collaborative development in Application Infrastructure, consolidation and the "winner takes all" mentality make it harder	Very little perceived need for, and thus actual existence of, collaboration amongst applications developers and service operators

		to operate application infrastructure independently	
Technology — reusable building blocks	The Internet's network technology is still largely based on (open standard) technology building blocks	Application infrastructure is still built with building-block technology, but it is increasingly offered through proprietary services.	The applications and services that make up the User Experience of the Internet are typically vertical, integrated, closed solutions, which do not permit their reuse as building blocks.
There are no permanent favourites	Networks are becoming more stable / ossified, and it's harder to make significant changes	There is nothing inherent in technology that makes it less true at the Application Infrastructure level, but scale and business practices are making it hard to dislodge incumbents.	There is nothing that props up any given product or service that make up the User Experience, and new products and services do get built on the open platform of the Internet and its Application Infrastructure. However, it is difficult for new entrants in existing services to get noticed in the shadow of major players. (So, it's "easier" for Uber or AirBnB to establish themselves and grow than for a new social media platform to unseat Facebook, for example).

Part III

Using the Invariants as a Framework for Discussion

Given the vastness of the Internet, it is challenging not to fall into the mistakes outlined in the parable of the blind men and the elephant when we come together to discuss the state of the Internet, proposed policy, or technical changes. In focusing on specific features or segments of the Internet, it is easy to lose sight of the whole. It can also be difficult to discuss changes to one aspect of the Internet and have an understanding of the relative impact on other areas. Rather than focusing on one aspect, or trying to fit the whole Internet into a box, this paper reviews a framework for considering the Internet as a whole, cutting across many different areas of the Internet and providing a tool for supporting broad Internet discussions.

Those discussions are important because the Internet is growing and changing all the time, as well as provoking change in the rest of the world. The Internet is known as a technological and social platform that has itself featured continuous innovation and growth, as well as launching change of unprecedented scope across the globe. This continuous change has simultaneously been heralded as its greatest feature or bug, and has raised concerns about consequences (intended or otherwise) for personal and civil realities. So, apart from any issues accruing from its own growth and maturity, as the Internet becomes increasingly part of daily personal and commercial life, there are also increased tensions between fostering its continued growth and expansion, and trying to stop some or all aspects of the Internet and its use that are impacting the daily life of ordinary citizens, at the very least in order to provide safeguards.

However, each identified issue is like a single part of the elephant. One person's obvious next step to improving a technology is another's threat to the Internet's very fabric. It can be difficult, and misleading, to evaluate individual questions without the context of the Internet as a whole.

Having reviewed and established the Invariants as an articulation of what the Internet should be (at both the level of the Internet itself, and users' experience of it), we can use them as a framework for considering the impact of specific technology developments and policy changes on the Internet as a whole. To the extent we wish to see the Internet – as a network, or as the world experiences it – continue to be an open platform for innovation and openness, those impacts must be weighed.

By way of example, the following pages consider several technical and policy contexts, through the lens of the Invariants. Through that lens, we can understand how the contexts — whether technology developments or policy choices — would push the Internet closer to, or farther away from the ideal express in the Invariants. Detail is given for any impacted Invariant.

These case studies are drawn from real situations. As noted earlier, change is necessary and constant in the world, and thus, the Internet. There is no intention of judgment about the fact that these changes have arisen, but with the framework of the Invariants we can see where alternative approaches might be better, in the interest of ensuring the Internet continues to be successful.

Each of the case studies is also set at a particular moment in time. A fairly negative review of the state of one technology space (e.g., Internet of Things) does not mean the technology is doomed for all time, but rather highlights the current set of challenges to address.

Technology and policy developments will continue to arise and impact the Internet, well beyond the scope of this handful of case studies. If they successfully illustrate the use of the Invariants to have "whole elephant" discussions of impacts on the Internet from any technology or policy development, they will have served their purpose.

Case Study: Internet of Things

As originally conceived, the Internet connects network “hosts” — full server machines, each running a full suite of network service servers (DNS, e-mail, finger, telnet/ssh, etc). That model hasn’t represented reality for years, or even decades. With home computers connecting to the Internet via dial-up access providers, running web and e-mail servers was actively discouraged. Since then, the size and complexity of devices connecting to the network has only reduced. Although an Apple watch has more processing power than the original Cray supercomputers, and thus is computationally capable of running the services outlined for host computers of the early Internet, it does not. And, yet, we say it is “on the Internet”. Colloquially, anything that is on WiFi, and thus passes packets to the Internet, is “on the net”.

As the “Internet of Things” gathers momentum, and a surprising array of devices in industry, commercial buildings, and private homes become WiFi enabled, the range and scope of devices “on the net” changes significantly.[21]

An important aspect of these devices is that, irrespective of the size of physical object they are tied to (from individual LED bulbs to refrigerators, for example), the networking components are simple and generally designed to be hands-off. This is the nature of embedded systems — build, ship, forget.

Effectively, this means that within a decade, over half the devices connected to the Internet will be designed to get themselves onto the Internet with minimal user intervention, sending traffic and (in some instances) taking instructions from a set of monitoring and control services.[22]

As the marketplace has blossomed, those services have been built and operated in a proprietary fashion. There are, for example, competitive markets for home automation in

[21] From https://iot-analytics.com/state-of-the-iot-update-q1-q2-2018-number-of-iot-devices-now-7b/ — by 2021, the number of IoT devices “on the net” is expected to meet the number of non-IoT devices; thereafter, there will be more IoT than not, with a projection of over 21B IoT devices by 2025

[22] Recently, RFC8520, “Manufacturer Usage Description (MUD)” (https://datatracker.ietf.org/doc/rfc8520/) has been published as a proposed standard for IoT devices to share manufacturer’s usage expectations within destination networks. This is important work in improving the technology around handling IoT devices in networks. Although MUD does not address the specific issues highlighted here, it highlights that work is ongoing to address known issues.

general, and thermostat control in particular. You can't just have a "smart thermostat", you have to have an account with the thermostat's operator to control it. And that may or may not be the same operator as the service for your house "smart locks". It is almost assuredly different than the system that is monitoring and controlling your solar power, and has nothing to do with the controller of your automated pet feeder. The challenges are only magnified when taken to a corporate or industrial level.[23]

While people have become adept at managing a large number of accounts with a variety of online providers, there are two major threats that this explosion of proprietary services offer today:

- Security
- Longevity

In late 2016, Dyn[24] suffered the consequences of a major DDoS attack mounted by an army of IoT devices that had been taken over by a malicious party (the Mirai botnet). The fault was largely tagged to IoT devices that are so simplified that they have no security controls, or default user names and passwords that can't be changed.

Requiring better security practices in devices is clearly necessary. But, operational security issues also come in the shape of exposing one's network and premises to the vagaries of a remote service. Even if individuals keep their passwords well-chosen and hidden, if a remote service is malicious, or is hacked, the device owner may leak information, or lose control of their devices, without even being aware of it.

The longevity of these embedded systems is also important to consider. What functionality does an IoT device have if the company operating the monitoring and control service elects to shut it down, change the terms of use, or finds itself out of business? While consumers have grown accustomed to regular replacement of phones and TVs, at least in some parts of the world, larger appliances such as fridges and washing machines are expected to last for a decade or more. And where a conventional appliance will continue to carry out its function until it physically breaks, and consumers are left high and dry because they can't get replacement parts from a defunct company, an IoT-driven appliance for which the function is entirely dependent on its monitoring and control service will be useless the day that service stops.

[23] https://www.networkworld.com/article/3243928/internet-of-things/what-is-the-industrial-iot-and-why-the-stakes-are-so-high.html
[24] https://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/

Before dismissing that concern as out of place in a discussion of Internet health, consider *why* these devices will age prematurely: they are dependent on proprietary services, and competition is fierce to own market spaces. And, given the fact that IoT devices are going to be the largest representative of Internet user experience within a decade, it's relevant to review this situation in the light of the Invariants.

IoT and the Invariants

To a certain extent, IoT technology has been the victim of failure of integrity (*Global reach, Integrity*) and reduction of the *General Purpose* nature of the Internet. With NATs in home gateways, and expectations of blocked ports, new devices have a very constrained Internet to work with.

The specific issues considered here are:

- Reliance of IoT devices on remote services
- The network conditions that have prompted IoT development choices (blocked ports, etc)
- The development environment for IoT devices

	Networking	*IoT user experience*
Global reach, integrity	Failure of IoT developers to universally take up IPv6, the current standard of the Internet, is harmful to IoT ecosystems and the Internet as a whole	IoT devices cannot be directly-addressable; there are more of them than global IPv4 addresses
General Purpose	Blocked ports are a challenge to more general IoT services	
Supports Innovation without Permission		Within IoT ecosystems innovation may be supported, but not

		proprietary services
Accessible		Individual components of the technology for access may be (they are built on open standards) — but the services that monitor and control them are impenetrable and untransferable
Based on interoperability and mutual agreement		Apart from IETF standards, most specs are developed in industry consortia — not accessible, often competitive
Collaboration		No — it's a race to own the space
Technology — reusable building blocks		The parts that are necessary in order to build devices and ecosystems; the control and monitoring services are not
There are no permanent favourites		If you knock off a service, its whole IoT ecosystem is impacted, irreparably

On the whole, we can conclude that "the Internet of Things", as currently implemented, does not provide a generative platform for Internet innovation.

Case Study: Monoculture

The Internet was created to deal with diversity of requirements, geographical location, local network technology, and so on. The Internet spread rapidly because it could support a broad range of technical and user realities, in a global and integral whole.

As the Internet and the industries that support it mature, some aspects of the Internet have become less diverse. This is happening both at the network level (as companies that provide networking and its services are consolidating) and the user experience level, as “Internet access” becomes synonymous with one particular application or service.

Consolidation can happen in different ways, including through typical commercial activities — market maturation, individual market dominance, and mergers and acquisitions. As markets mature, the scale of operation necessary to disrupt an incumbent is significant. While this consolidation can be considered “simplifying”, simplicity is not always the best outcome. Some of this simplicity takes the form of the kind of control for which commercial anti-monopoly policies are in place: an incumbent can implement changes that are proprietary, or resist adhering to updates developed and proposed by others.

User experience “monoculture” occurs when people can only access the Internet through one application service. Different applications are dominant in different parts of the world — WhatsApp is all over India, for example, while Facebook is predominant in the Philippines. At issue is the situation when users effectively have no choice - they cannot use other applications, or it would be cost-prohibitive to do so (e.g., when certain applications are “zero-rated”, or not charged against the user’s data usage). In terms of information experience, this means that the reality that is promulgated on those platforms cannot be reasonably challenged by fact-checking elsewhere. Several countries have reported civil issues due to this monoculture.[25]

Monoculture and the Invariants

We can look at the issue of monoculture through the lens of the Invariants, specifically focusing on:

- Large players in network operations
- Consolidation of application services
- Application access limited to specific platforms

[25] See http://time.com/5505458/facebook-maria-ressa-philippines/ and https://www.wired.com/story/how-whatsapp-fuels-fake-news-and-violence-in-india/

	Networking	*User Experience*
Global reach, integrity	Integrity of the network can suffer if the largest incumbents elect to play (only) by their own rules.[26]	Global reach and integrity is completely undermined as users can only see what's in their application's bubble
General Purpose	Excessive monoculture in the network means it supports existing purposes; may stifle growth to address new purposes	Limited to the purpose of the monoculture app
Supports Innovation without Permission	Can be gated by the dominant players.	Very limited — scoped to the use of the app
Accessible	Only parts that are not proprietary (e.g., can't see into FB's proprietary transport protocol)	No
Based on interoperability and mutual agreement	Less necessary	No
Collaboration	Less necessary	No — it's a race to own the space
Technology — reusable building blocks	No — rise of proprietary technology	No — there are no building blocks here
There are no permanent favourites	No — dominant players are difficult to unseat	No — the permanent favourite is the only reality

[26] E.g., Facebook deployed its own transport protocol, unshared with anyone else. See https://smartdata.polito.it/five-years-at-the-edge-watching-internet-from-the-isp-network/

It is fair to say that monoculture is bad for ensuring a generative platform at any level.

Case Study: Specialized Services

Where the monoculture discussed earlier considers an existing service or experience being taken over by a single player, a related issue is specialized services. These are services that are built and deployed to address particular network services more effectively or efficiently than can be achieved through distributed systems that might be operated by more than one entity. Examples of this include content delivery networks (CDNs), and cloud services such as Amazon's Web Services (AWS).

These specialized services are possible because of the general purpose nature of the Internet, and its accessibility. And, from a practical standpoint, they provide important services to resolve commercial problems.

However, from the standpoint of considering the Internet as a generative platform, solving these problems through proprietary solutions, rather than collaborative open standards, undermines the Internet's future. CDNs, as implemented today, rely on DNS hacks and other infrastructure assumptions (sometimes called trickery), which undermines the Internet's global reach and integrity. To the extent that AWS is successful, it makes customers reliant on a commercial product that is not the Internet, and locks them in.

Specialized Services and the Invariants

Looking at specialized services, we can focus on:

- Networks that are purpose-built for improved performance for a particular application or type of application
- Proprietary application services that support particular application issues (e.g., "world wide *wait*" for CDNs a decade ago)

	Networking	*Application Infrastructure*
Global reach, integrity	Integrity of the network can suffer if the largest incumbents elect to play	This is undermined as the application will function differently depending on

	(only) by their own rules, building specialized networking services	the specialized service
General Purpose	Specialization in the network may mean it supports existing purposes; may stifle growth to address new purposes	Limited to the service of the specialized app — e.g,. Videos or web content that have paid for the service
Supports Innovation without Permission	Can be gated by the dominant players.	Very limited — scoped to the use of the app
Accessible	Only parts that are not proprietary (e.g., can't see into FB's proprietary transport protocol)	No
Based on interoperability and mutual agreement	Less necessary	No
Collaboration	Less necessary	No — it's a race to own the space
Technology — reusable building blocks	No — rise of proprietary technology	No — there are no building blocks here
There are no permanent favourites	No — dominant players are difficult to unseat	No — the permanent favourite is the only reality

None of the above suggests that such specialized services improve the fabric for future innovations.

Nevertheless, with all that said, it is interesting to note that the industry itself does bump up against the limitations of closed technologies, and coalitions form to address them. In particular, consider the "Open Caching"[27] work being done at the Streaming Video Alliance, which builds on the IETF's CDNI (CDN Interconnection)[28] work.

Case Study: SOPA/PIPA that wasn't

In 2012, the United States was considering legislation targeting the operation of DNS, a core Internet infrastructure technology. Intended to help curb issues of intellectual property and copyright infringement, particularly on websites of organizations outside the United States' jurisdiction, SOPA (Stop Online Piracy Act) and PIPA (PROTECT IP Act) incited protests by network services such as Google and Wikipedia.[29] Following the protests and reversal of opinion of key supporters, the proposed laws were withdrawn "indefinitely".

The proposed laws may have been withdrawn because of industry pressure stemming from concerns of business impacts, but the technical impacts of the implementation of the laws would have been profound. The proposed laws would have required US network operators to block access to domains identified as harbouring copyright-infringing material, by failing to provide DNS answers for those domains. That is, operators would be required not to provide addresses for the named domains.

Concerns were expressed about the laws themselves. Some people were concerned that it would be too easy for domains to be added to the list of sites to block. Others expressed fears that the language was open to broad interpretation. Additionally, blocking an entire site was an overly inclusive requirement, and it would have considerable collateral damage.

The technical community was also concerned about the impact on operations and on a core Internet technology, DNS. Requiring selective blocking of domain names at the time of resolution would mean that DNS was no longer a global system. Worse, this approach would also have undermined the deployment and use of the DNSSEC protocol to cryptographically certify DNS responses.

SOPA/PIPA and the Invariants

We can look at this in more detail through the lens of the Invariants, focusing specifically on the issues of:

[27] https://www.streamingvideoalliance.org/technical-work/working-groups/open-caching/
[28] https://datatracker.ietf.org/wg/cdni/about/
[29] https://www.theverge.com/2012/1/18/2715300/sopa-blackout-wikipedia-reddit-mozilla-google-protest

- DNS as an open standard application technology
- Impact on the overall network reachability

	Application Infrastructure	*User Experience*
Global reach, integrity	No longer global and integral — the DNS would provide different responses in different areas	No longer an integral whole — users in some countries would be able to resolve affected domain names, but not in the US.
General Purpose		Added collateral damage: blocking a domain affects more than its web server, but also e-mail and other services
Supports Innovation without Permission		Chilling effect on content — some impacted sites might be perfectly legal in the jurisdiction in which they operate (just not in the US)
Accessible		
Based on interoperability and mutual agreement	No — operational impact by being required to break DNS services	Undermined — affected sites would be unilaterally removed in the US
Collaboration	No	No
Technology — reusable building blocks	Broken — DNS not operating as specified, and things built on it (e.g.,	

	DNSSEC) impacted	
There are no permanent favourites		

From the above, it's clear that SOPA/PIPA would have pulled the Application Infrastructure and User Experience further away from the ideal of a generative platform.

Case Study: National boundaries on networks

As outlined in "On the Nature of the Internet",[30] a contribution to the Global Commission on Internet Governance, the constant push to align Internet network architecture with geopolitical boundaries runs counter to the ideal expressed in the Invariants.

The impetus is understandable — as governments are concerned with anything that affects their citizens' well-being and proper behaviour, this extends to a desire to have more control over the network infrastructure that impacts so much of their personal and commercial lives. For example, that paper describes:

> In 2013, revelations of US government data collection practices caused other countries' governments to consider how much of their citizens' traffic flows through the United States, whether or not it is destined for any user or service there. These realizations have led to calls to reroute major Internet links to avoid having traffic transiting US networks. Changing network connections (and, thus, routes) is a common and ongoing occurrence, but it is usually driven by needs for network efficiency and resiliency. Attempting to re-architect the Internet so that citizens' traffic remains within certain geopolitical boundaries is at odds with responding to the global Internet's needs, and may well lead to less diversity and resiliency in (national) networks. A look at global connectivity maps provides some surprising information — Internet connections do not naturally align with political boundaries. For example, Canada has an immense geography and a modest population. Population centres (and, therefore, obvious locations for networking hubs) are generally spread apart. Since the Internet's routing technology is designed to pick efficient steps between origin and endpoint, it is not surprising that it is sometimes

[30] Daigle, Leslie, "On the Nature of the Internet", GCIG, 2015. https://www.cigionline.org/publications/nature-internet

cheaper, easier and faster to route Internet traffic from one end of Canada to its middle via a connection point in the (much more densely populated) United States, Canada's neighbour to the south. So, traffic from Canadian cities Vancouver to Toronto might reasonably bounce through US cities Seattle and/or Chicago. Similarly, many international connections out of countries in Latin America terminate in Miami. Miami terminates important data links from other continents. Rather than building individual links between every country in South America to every other continent (or country), it has been most effective and efficient to build large-capacity links to Miami from South America, and have South American traffic transit Miami on the way to or from countries in Europe. "Cheaper," in the context of interconnections, can mean more than a slight savings for companies involved.

However, as that paper notes, such network realignment is at odds with the Invariants — especially in terms of its Global Reach, Integrity and in terms of diversity of infrastructure and connections.

National Boundaries and the Invariants

Specific points of interest for review:

- Changing the shape of networking connections, which was designed to support connections wherever they made sense physically and/or commercially

	Networking	*Application Infrastructure*
Global reach, integrity	Imposing national borders would reduce opportunities for multiple interconnections and alternative routes	Facilitates the possibility to constrain the infrastructure (e.g., SOPA/PIPA)
General Purpose		
Supports Innovation without Permission		

Accessible	Constrained	
Based on interoperability and mutual agreement	Constrained — forced to limited opportunities within jurisdiction	
Collaboration		
Technology — reusable building blocks		
There are no permanent favourites	Sets the stage for incumbents and natural monopolies	

Thus — national boundaries on the Internet's networking would clearly be antithetical to the Internet's ideal at its most basic level.

Case Study: Extraterritorial legislation

If the Internet can't be made to reside within jurisdictional boundaries, recent trends have been to push legislation to have extra-territorial reach.

Early examples included requirements for user information not to be exported from that user's country, irrespective of the base of operations of the company providing the service. The implication would be that every (major) web service would have to establish data stores in each country of operation, and then ensure that data did not inadvertently migrate. This is not only complex for multinational organizations (and a significant barrier to entry for smaller organizations), it also raises privacy concerns for those who fear their government's own actions.

The General Data Protection Regulation (GDPR) has broadened the requirement for companies operating in the EU. The introduction of GDPR caused some services to stop doing business with customers in the EU, and block access based on IP address geolocation.[31] For

[31] The Chicago Tribune, among others. See https://www.marketwatch.com/story/chicago-tribune-la-times-go-dark-in-europe-after-gdpr-fail-2018-05-25

companies not based in the EU, the GDPR imposes a requirement to understand and be responsive to a legal requirement in a foreign regime. As if that's not complex enough, other jurisdictions are considering "similar" regulations; in future, companies endeavouring to offer services globally will have to navigate an increasing mire of potentially conflicting legal requirements to serve their customers.

Concerns include that this will variously drive smaller organizations out of business, provide increased barriers to entry for new services, and, at the very least, drive further consolidation of customer-supporting services. From Amazon Web Services:

> Having previously worked in an area where regulation required us to segregate user data by geography and abide by data sovereignty laws, I can attest to the complexity of running global workloads that need infrastructure deployed in multiple countries. Availability, performance, and failover all become a yak shave as you expand past your original data center. Customers have told us that they need to run in multiple regions, whether it is for availability, performance or regulation.[32]

Extraterritorial Legislation and the Invariants

Looking at this through the lens of the invariants, we can look specifically at:

- Impact of implementing responses to legislation through network filters
- Tactics for application services to comply with legislation
- Tactics for application services to (legally) avoid legislation (I.e., not serve customers in impacted regions)

	Networking	*User Experience*
Global reach, integrity	No longer global and integral — some IP addresses blocked for connection	No longer an integral whole, as services go dark in some jurisdictions rather than comply — often including unexplained unreachability

[32] https://aws.amazon.com/blogs/aws/new-aws-global-accelerator-for-availability-and-performance/

General Purpose		
Supports Innovation without Permission		Chilling effect on developing new services
Accessible		Undermined
Based on interoperability and mutual agreement	No — this is mandated	
Collaboration		No
Technology — reusable building blocks		
There are no permanent favourites		No — drives toward more consolidation of services

Concluding thoughts: the continued importance of the Invariants as a lens for reviewing and discussing the Internet

The basic conclusions from this review are that the Invariants are still as relevant as ever, in describing the fundamentals of the Internet as a generative innovation platform, and, they are a constructive framework for enabling discussion of technology and policy choices, in terms of whether those choices would bring the Internet closer to, or push it further away from, the ideal of the Invariants.

The initial review of the Invariants also highlights that the Internet's Application Infrastructure is moving significantly away from providing a generative platform for future innovation. There is no single thing that drives that condition, but it would take significant collaboration and work to foster the development and deployment of open standard solutions to address the problems that are currently more readily "owned" by proprietary interests.

And, if that effort is not made, the future Internet will not feature the kind of adaptability and innovation that has brought it to where it is today.

Appendix: The Internet Invariants, 2019

This is the collected updated text of the Internet Invariants.

Global reach, integrity	Any endpoint of the Internet can address any other endpoint, and the information received at one endpoint is as intended by the sender, wherever the receiver connects to the Internet. Implicit in this is the requirement of global, managed addressing and naming services.
General Purpose	The Internet is capable of supporting a wide range of demands for its use. While some networks within it may be optimized for certain traffic patterns or expected uses, the technology does not place inherent limitations on the applications or services that make use of it.
Supports Innovation without Permission (by anyone)	Any person or organization can set up a new service, that abides by the existing standards and best practices, and make it available to the rest of the Internet, without requiring special permission. The best example of this is the World Wide Web – which was created by

	a researcher in Switzerland, who made his software available for others to run, and the rest, as they say, is history.
Accessible	It's possible to connect to it, build new parts of it, and study it overall: Anyone can "get on" the Internet – not just to consume content from others, but also to contribute content on existing services, put up a server (Internet node), and attach new networks.
Based on interoperability and mutual agreement	The key to enabling inter-networking is to define the context for interoperation – through open standards for the technologies, and mutual agreements between operators of autonomous pieces of the Internet
Collaboration	Overall, a spirit of collaboration is required – beyond the initial basis of interoperation and bi-lateral agreements, the best solutions to new issues that arise stem from willing collaboration between stakeholders. These are sometimes competitive business interests, and sometimes different stakeholders altogether (e.g., technology and policy).

Technology — reusable building blocks	Technologies have been built and deployed on the Internet for one purpose, only to be used at a later date to support some other important function. This isn't possible with vertically integrated, closed solutions. And, operational restrictions on the generalized functionality of technologies as originally designed have an impact on their viability as building blocks for future solutions.
There are no permanent favourites	A healthy Internet ecosystem relies on diversity — of technologies, companies, environments that foster those that become or continue to be relevant and useful. Good ideas are overtaken by better ideas; to hold on to one technology or remove competition from operators is to stand in the way of the Internet's natural evolution.