

The Internet as Infrastructure FAQ

Introduction

This is a companion to the writings such as [Broadband to Infrastructure](#) and [Internet Native Policies](#).

I view the Internet more as infrastructure than as a utility. The word “utility” implies something that is provided and may be metered. The Internet is best considered a technique for using whatever facilities are available in the same way we take advantage of sidewalks when we take walks. We improve our ability to travel by providing sidewalks as infrastructure rather than requiring people wait for the next bus or trolley (light rail).

When I present that perspective, I find myself responding to many of the same questions hence this FAQ or Frequently Asked Questions.

I plan to update this as needed.

What is “The Internet

Less important than the dictionary definition of the Internet is the generative dynamic that comes from the strict separation of the packets of data from what we do with them.

This is embodied in end-to-end principle (<http://bit.ly/EndToEnd>) which says that we can only depend on what we can do at the end points and not inside the network. Note that this is just the opposite of what people mean when they use the term end-to-end to include every step along the path. This is understandable because we assume that someone along the way must do something to make it work.

To put it another way the Internet is not a network (as a service) but rather a way of networking (as something we do by working with others). Perhaps a better comparison is to think of it as social networking rather than arranged marriages.

The Generative Dynamic

Voice over IP (VoIP) is a good example. In theory it wasn't supposed to be feasible because voice requires very precise timing. This relied on a carrier to manage the complete path of the connection. In the 1990's software engineers at a small company, VocalTec, used clever programming to make voice work on their local networks. They were surprised to find their customers were using the technology to speak over the wider Internet. It turned out that the capacity created for the web increased the capacity of the Internet so that there was enough capacity for video.

Skype can offer video at no additional cost because it doesn't promise video so that it doesn't have to reserve capacity. If there isn't sufficient capacity it can reduce the resolution or fall back to audio or even texting. It is this resilience that drives the generative dynamic.

Rather than engineering a solution to a program, the fact that VoIP worked can be considered discovered as much as engineered. This is the opposite of the normal problem-solving paradigm. It is about creating and taking advantage of opportunities.

This is a very powerful dynamic but it's counter-intuitive since we are taught to solve problems. The opportunity dynamic is powerful but requires discovering what works rather than having a narrow definition of success.

One result is that we aren't dependent upon a provider making promises and assuring the entire path is reserved, we can interconnect our local facilities in the same way we drove across the country without depending on the interstate highways. The interstate highways are not reserved paths either. They are just optimizations that we can apply once we are supporting infrastructure at a national scale.

Another way to think about this is the networking is something we do ourselves without depending upon a provider doing it for us – akin to driving vs taking the train.

This brings us to the crux of the matter.

The Infrastructure Argument

It's one thing to recognize that we needn't depend on a provider for services such as phone calls, movies, games, apps, and television. That alone challenges the business model of the telecommunication industry.

But if we only need local facilities management (or, if you wish, local networking) then we don't need networking as a service. This means we must find a way to pay for the facilities, support and maintenance. The answer, of course, is to pay for a common infrastructure.

The question is not whether we need separate but neutral pipes, the question is whether we need a pipe provider at all!

The Tragedy of the Commons?

The term “Tragedy of the Commons” is used to refer to situations in which there is a limited resource with the us-

ers all using us much as they can without regard to the others. In practice we are very good at managing shared resources.

What is remarkable about the Internet is that it's basic design is based on protocols that foster decentralized cooperation. It has worked very well despite predictions of collapse. Using the shared facilities across city services will create financial incentives to invest and maintain facilities that the city relies on.

The real innovation is in taking advantage of the capacity available. The winner is not the one who can "hog" the most capacity. For example, video streaming (Netflix, Hulu, YouTube) requires far less capacity than traditional television.

But what about Potholes?

I'm surprised how often people tell me that cities can't maintain infrastructure because the streets often have potholes. Given that the cities can maintain complex infrastructure such as water systems and roads, they are able to maintain, or hire companies or people to maintain, vital infrastructure. By comparison, assuring connectivity is far simpler thanks to the resilient nature of Internet protocols and the wide availability of networking expertise the city can hire or contract for.

Transition: How do we get there?

The transition can be incremental.

The transition has already started to a small degree with your home network. You own the wires within your home and then pay for a broadband pipe which interconnects you to the rest of the world. With Dropbox, OneDrive and other such services you are a provider of content.

You can share your broadband connection with your neighbors by using a router which splits the connection into two independent home networks. Concerned about security? See Privacy, below.

When we share among multiple apartments and homes we shift from Do It Yourself (DIY) to Do It Ourselves (DIO). It is this approach of working together and interconnecting these connected zones that gave us the term inter-networking or, for short, the Internet.

An apartment house owner or condo board could pay for a single connection to the rest of the world and then make it available as a common resource. The cost, per apartment, is far lower than today's broadband prices which makes this a compelling option. Owning the physical infrastructure means no longer paying a monthly fee to a provider.

As we increase the number of participants the cost savings increase rapidly. That alone will drive a transition to an infrastructure model. Even without collateral benefits.

Privacy!

Privacy is a real concern when using the Internet. Having the city operating the facilities will make little difference. What will make a difference is the ongoing shift to encrypting all the traffic so no one, government or private, can read the traffic. Virus and Malware protection, password and other internal privacy measures are always good practice.

We already take responsibility for the privacy of the devices within our homes. We don't depend on a provider to protect us.

"Free-to-Use"

One reason that infrastructure creates opportunity is that it is free-to-use. Free-to-use is not free – it means we are paying for the facilities in such a way that we can make it freely available.

Sidewalks are a good example. They may be expensive to build but we don't charge you for each step nor collect a fee when you leave the house.

Common Carriage

This is a principle applied to public services such as railroads. While it is hard to apply this principle to the Internet as a technique it does make sense to apply it to the channels that telecom provides as a transport.

This is not about regulating simply about the use of the facilities owned by telecom operators as one means of inter-networking. It matters because they control the bulk of the capacity between local facilities.

Regulations, Network Neutrality and Competition

Today there is a debate over deregulation and network neutrality. The regulations exist for a number of reasons including addressing the inherent conflict of interest which exists when a provider's own services compete with others, such as Netflix, who depend on the provider. The regulatory system was created to address a dysfunction in the market in a marketplace with high capital costs and little differentiation among services

An infrastructure approach addresses this conflict of interest. With an infrastructure approach, we get companies vying for the community's business with people choosing what's best for themselves rather than what's best for the

provider. This is a market that is in balance with little need for regulation.

The market for services is no longer tied to particular cables or radio bands. This is already happening – not only with Netflix but also with HBO and other content providers.

The idea of the city as the broadband provider may seem to some as though it would eliminate competition in the market. Just the opposite is true. Broadband competition is carefully regulated and managed with providers committed to particular technologies. The city or municipality as vendor makes no difference, provided the vendor meets standards and is in compliance.

Collateral Benefits

The idea of “one Internet” is very powerful. That is the spirit of *Ambient Connectivity*. You’re simply connected. When you bring a device home you simply turn it on and it works. As we shift to an infrastructure approach, travelers will be able to assume connectivity. You will automatically get your email on any device (not just your phone) and browse the web without having to think about how to get connected again and again. The heart monitor on the wrist always be connected to your medical service if you want to be monitored.

Ambient Connectivity offers “one stop” connectivity, always on and attainable, instead of the mish-mosh of WiFi and local network services one must sign into in order to achieve connectivity. Today it’s complicated to connect devices to (the rest of) the Internet. You have to go through complex rituals to connect to your home WiFi and then do it all again if you change your router. Each stop along the way when you travel, or even as you make your way around your local environs, needs a different setup.

With mobile phones, you need to make sure your carrier has service where you are and that you have paid for the right plan for each device. Then you may need to pay for yet another account for connected cars.

Ambient connectivity is much simpler. Connectivity “just works.” That’s resilience. This becomes life-saving in an emergency when you can harness any available connectivity to get a message out.

This is important for day-to-day use and is essential in an emergency when connectivity can be quickly extended without worrying about whose connectivity it is. It is ours.

Sidewalks?

Like any analogy there are many ways the city assures that we have sidewalks (or other rights of way). It might pay for it out of general funds or it may make property owners

responsible. Sometimes people need to walk on the side of the road or find some alternative. What is important, in terms of policy, is that we assume we can walk around a city and we can accommodate the few exceptions.

In the same way, we would assume Ambient Connectivity. Sometimes the city provides it, sometimes property owners contribute to the whole. What is important is that we can assume connectivity without having to negotiate with providers again and again.

Network Layers

This is for the networking geeks. We are taught that networks are layered with layer zero being the physical layer all the way up the stack to the session and application layers.

That is just a story we tell or a model. It is not hard reality. Such layering or abstraction models can be useful working decompositions, but they are not the only decompositions. Such models presume that one layer is dependent upon the one below. The power or generative idea of the Internet is that we don’t have to depend on what layers do but can take advantage of any available opportunity as a resource.

While we’ve always been free to ignore the purpose of a piece of hardware or service and use it as we please (such as treating a table as a chair when we sit on it), we can now use software to implement our interpretation. In this case we can use existing telecommunications facilities as a best effort packet transport without depending on layered services such as reliable delivery.

There is a disconnect between communications engineers who presume that they are assuring our ability to communicate and the Internet in which communications in the sense of exchanging meaning is not dependent upon the network itself. All we ask is their best efforts to forward packets without guarantees.

Vocabulary

It’s difficult to have a conversation when we don’t have a shared understanding of the meaning of words. In the case of the transition from the telecommunications services to the Internet as infrastructure the change in meaning is at the heart of the shift.

The repeated theme is the economy of language in using a word to describe a set of concepts that seem to be part of a single whole. This makes it difficult to, well, communicate when the meanings are no longer aligned. This is the classic problem of a paradigm shift as our understanding and our vocabulary are no longer in alignment.

Communications / Communicating

A simple example of communicating is two people speaking to each other and arriving at a shared understanding. When we needed to communicate over a distance we would write a letter and the post office would take responsibility for making sure the message with its meaning was delivered intact. The same is true when we use the electronic version and tele-communicate.

There was no need to use a separate word for the technology because that would be redundant. Of course, the telegraph operators and, later, the phone companies communicated on our behalf by taking responsibility for carrying messages intact. It's expensive keep such promises.

This all changed with the Internet. We send just numbers (data packets) that have no inherent meaning. Only the sender and receiver understand the meaning. This is not entirely new because one used to be able to send coded messages that were only understood by the sender and receiver. What is different is not only that this is the norm, but we take advantage immediacy of the Internet to have a conversation. Instead of relying on a carrier to make sure the message is delivered intact, we can correct for lost packets by simply asking for them to be resent. It's the equivalent of saying "huh" and repeating the messages.

This is a sharp contrast with a game in which we whisper a message passing it from person to person and laugh at how distorted the message is at the end. The game is often called, appropriately enough, telephone game.

The carrier is not no longer carrying a meaningful message and thus, is not communicating on our behalf. We need to acutely aware that the business of communications engineers has changed and is now very different from professional communicators who use whatever means is available to get the meaning across.

Quality of Service

One concern with best efforts network is how to make promises like quality of service or QoS. But the whole idea of QoS is tied to the idea that the network operator is providing services (such as telephony). But with the Internet those services are created entirely outside the network and there isn't a single definition so how can a provider assure QoS? Furthermore, there is no relationship between the packets and their intent, so the provider doesn't even know the purpose of the packets. This means that if a provider see two packets at a router there is no way to say which is more important than the other.

More about this in [Purpose vs Discovery](#).

Further Reading

<https://rmf.vc/InternetNativePolicies> which explains some of the concepts in more detail.

<https://rmf.vc/ZeroRating>. Zero rating and the related concept of fast lanes and slow lanes are firmly in the world of telecommunications where they assume the services are defined inside the network. The concept of lanes simply does not exist in the Internet.

<https://rmf.vc/FurtherReading> for more pointers.