


## Systems, Assumptions, and the Origins of Internet Security<sup>1</sup>

### Or: So You Let Me Ramble and Babble At Y'all for A Bit One Evening

Obligatory Scholar's Caveat (When Talking More Casually):

This is – like any speech on events gone by – limited by the quality of sources and level of investigation. There is a LOT of early history to the Internet and this isn't anything approaching comprehensive or complete. It is merely a presentation of some basic notions, strung together. It is more a rationale, or even a theory!, of how some problems started – and persist – in the modern phenomena that is the Internet – as seen by an occasional systems analyst.

The person who drew this:



Probably would be somewhat staggered to consider all of the things we consider entirely normal to the smallest of the entities with a network link of any sort. That is actually somewhat amazing, because that image is the first map of what would become the Internet. What you can't tell from the scanned image is that it was drawn on a napkin in 1969<sup>2</sup> by an employee of a firm called BBN<sup>3</sup>. The Internet we take for granted – as with so many other great ideas (especially in Computing) – began over lunch.

Note the date.

This was not instant creation of the situation we know now.

Things had not settled down on ideas like 802.11 and Ethernet (now ubiquitous for a lot of interconnect needs) – all kinds of things were running around. There was no clear winner – rather, a plethora of options, all with suitably techie sounding names. X.25, DECnet, assorted things from OSI, offerings from IBM. Many of these were proprietary, some tightly linked to vendor specific platforms. (It wasn't

<sup>1</sup> I like footnotes; footnotes are cool. With apologies to Matt Smith as The Doctor. Note that he's correct about bowties. Also, fezzes.

<sup>2</sup> ARPA, for those interested is "Advanced Research Projects Agency" (now DARPA), a component of the US Department of Defense who funded a lot of the early work. More at: <http://www.vox.com/a/internet-maps>

<sup>3</sup> <https://www.nextnature.net/2008/08/the-first-map-of-the-internet/>

unusual for an institution to be “a DEC” or “an IBM” ‘shop’.) There are still lots of options, but at least some of the quirks and clunkier choices and (especially) the vendor related rigidity has gone away. (Mostly!)

Note, also, the complexity. A whopping 2 devices.

Yet for the utter simplicity of this notion – a host and a device to link you to other hosts, which wasn’t even entirely new(!), this napkin is a kind of crowbar.

It represents a wedge idea. The network that will become familiar to use as the internet is in many ways, present... A bunch of working notions is apparent:

- The idea of an “edge”
- Hardware is important... but protocol (“IMP”) is important as well. Eventually, the idea that protocol behavior need not be confined to one particular physical device becomes key.
- Core concepts of design and decision making (the lunch had more than one person!) would be determined through discussion and consensus. This will in part become realized in the idea of *interoperability* as an important characteristic. The lock in by vendor (or even excessive standardization) will come to be rejected for various key aspects.
- A lack of centralized, well, anything.
- (lunch, again) Cooperation, driven by needs and goals that are joint. Communication – moving data – solving the problem.

OK.

Deep breath.

Before we start doing the wave in the stadium, and spelling out Internet with cheers. There is a flip side to these decisions – *tradeoff* – an idea we’ll keep revisiting.

Also: remember – computing is based in large part on math. Math – as a community – is fond of proofs, which begin with assumptions. Good or bad, your proof (your solution) is OK... going from those assumptions.

Now, with the napkin, it was one step from running. Once it is running, it is. Corrections to a running system, even with awkward, messy or incorrect behavior are difficult. This idea, generally, also becomes an assumption.

In many ways, the Internet is well thought out. The primary standard was that it Work. Another assumption.

*Form followed function.*

One form that evolved was a notion that everyone had (or at least was responsible for getting) some sort of network connections within their own institution(s). In effect there was a need to provide a connection between not just hosts but between sites/groups/bodies/entities.

This changes the game a bit.

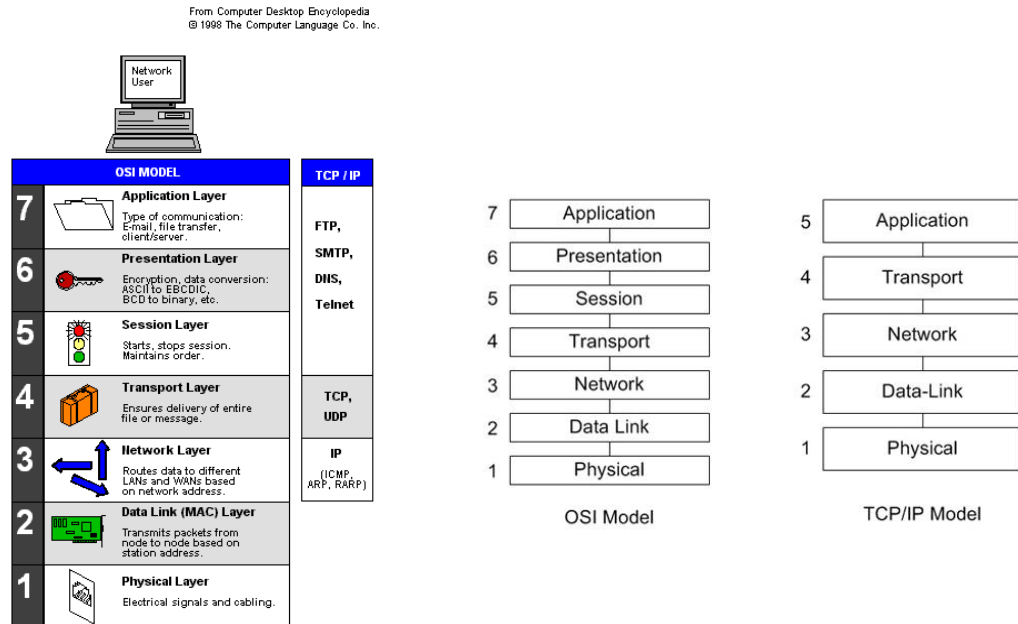
Any facility/college/business/university participating in/joining the network does so with a variety of things in place – among them a trust (as much as any computer person trusts their users) in their own local customers. These people/users were “known” and “authorized” by the people who employed this first generation of Internet builders and technicians. By linking the systems, there was in effect a decision to trust (to some degree) the rest of the participants. For the technical types, it also meant a trust in their counterparts elsewhere to make sure nothing spilled over – and that good practices were generally being followed. The tightness and “transitive” trust accepted by the original techie types was a Big if Unspoken consideration in helping all this work. (Yup, another assumption.)

This can’t (eventually, it didn’t!) scale forever.

Lots of things have to come into the equation as the system gets more complicated. But a few basic pieces remain pretty constant – partly because things are evolutionary.

The network admins, local developers and contributors, system administrators and sites all have something else – teaching or physics research or problem solving or all manner of other things – they’re doing. They are creating a network and tackling decisions sometimes out of need, sometimes out of curiosity, sometimes out of frustration – but that isn’t their primary goal, generally.

Even if the trust and good behavior mechanisms didn’t scale into the forever (and we’ll revisit ways they have opened the doors to problems)... some things did, largely because they worked. Or, at least, provided enough utility that replacement didn’t seem like a useful activity. One such is below:



These are the mental model for building the ‘protocol stack’ – familiar to any survivor of Networks. Both (TCP/IP and OSI) have their uses, but the TCP/IP one is a bit quicker to consider. At layers 3 and 4 in the TCP/IP model are Transmission Control Protocol and Internet Protocol – the basic items that give us a packet.

Also, this packet – and the model that goes with it - is (in many ways) what came out the practical winner<sup>4</sup>.

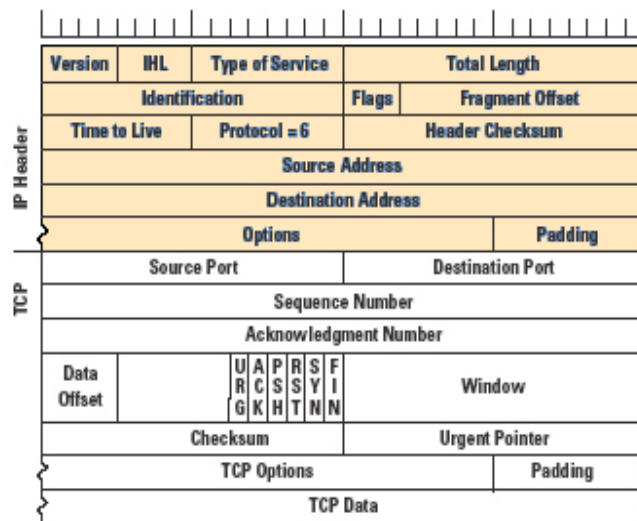
Winner how? It solved the problems of the authors of this setup... it was there, and able to connect things and the associated community was prepared to respond to practical needs. Most of the crew concerned was tasked with producing fixes to the needs of linking systems, or concerned with getting those links up and running. Big connections are expensive. Those who pay for them (in that era, the sciences, some big institutions, and the research community) are not always noted for their patience.

A minor miracle of record keeping and coordination also helped, largely performed by one person. Even this now-celebrated individual<sup>5</sup> had other things to do – local customers/users and an employer to keep happy.

All of this – the assumptions about appropriate use, the need to Get Things Done, meant a certain mindset. Some basic premises - *be conservative in what you do, be liberal in what you accept from others*<sup>6</sup> – are still touchstones in the decision making and operating rules for networks today.

Any Malevolence – bad behavior – was assumed to be derived from a poorly built ‘packet’ (transport/network layers in that model), less so from people.

The realization of that layering idea - and that packet (combining layers 3 and 4) looks like the blocks below.... Nothing more nor less than directions for how to read a set of bits coming cross fiber, wire, or radio.



<sup>4</sup> TCP/IP was, in fact, not universally adopted as many waited for OSI (a standards and technology group, still present, less US-centric / more international) to produce what was expected to be a Total Solution covering all manner of things. TCP/IP, if a bit scrappier, got implemented and deployed quicker....

<sup>5</sup> Jon Postel. In many ways one of the first who did ‘community service’ for the Internet. So much that one of the community awards is now named for him.

<sup>6</sup> “Postel’s Law”

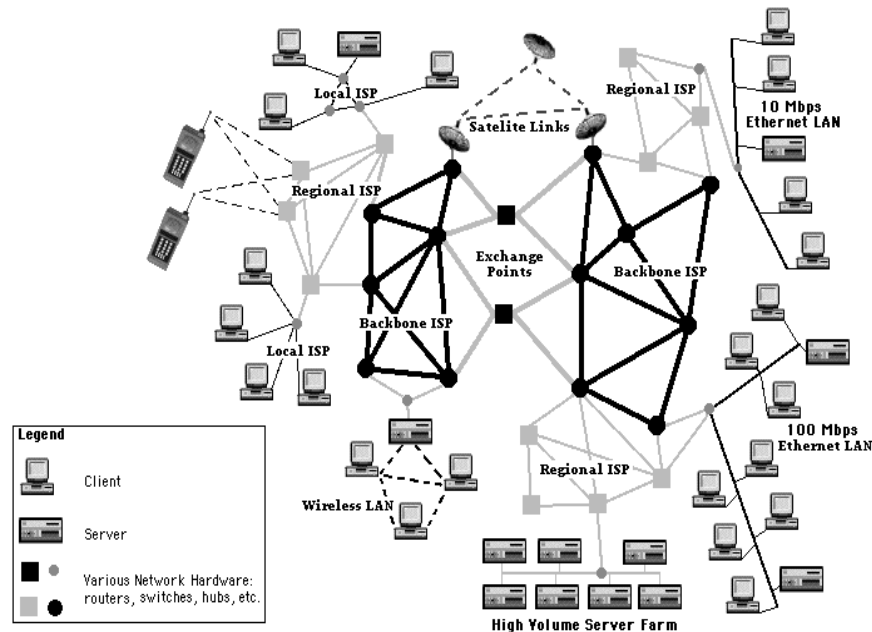
This familiar Networks item – the breakdown of the TCP and IP headers) has... astoundingly, no built in security features.

This is the basic unit of communication on the network. Without this, you don't have data moving really. No security feature....

In the primary implementation of data moving across this 'Internet' thing.

Nothing. Nada. Zip. Zilch.

And that was OK...to start. But it started show some cracks when we have made it to<sup>7</sup> this:



In terms of complexity of the Internet.

Here's the amazing part: even with the lack of deeply embedded, multi-layer (in terms of the model) security, that basic setup from 1969 \*still holds\*. Things have been added, but that common unit of communication has held up – even as the physical layer has changed, and routing has evolved, and all manner of crazy additions have been tacked on at the application layer, that packet is still it. It solves the basic needs of networked communication.

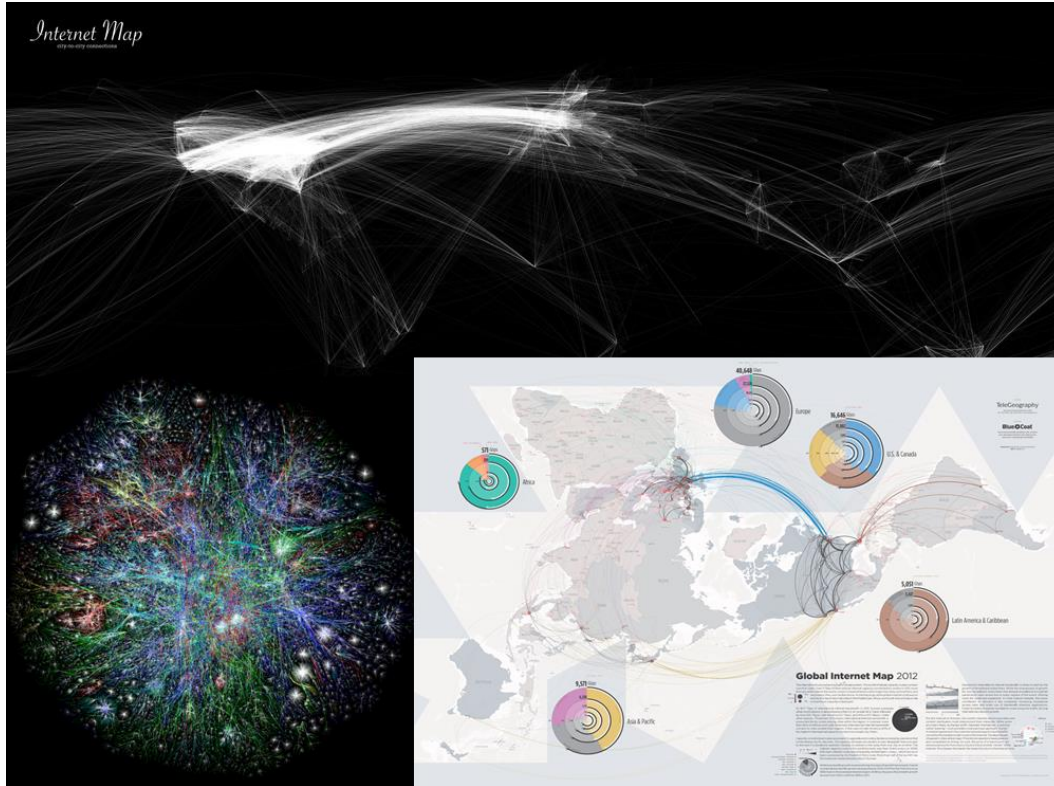
The “decomposition” (or “divide and conquer”<sup>8</sup>) of concerns across the different layers has stood the test of time (generally), so far. Only one design decision has proved profoundly problematic to the longevity of pragmatically brilliant work created then... The host address size. IPv4's numbers managers have run out of address space<sup>9</sup>. 1969 to 2015... not bad. And the problem doesn't stop it from continuing to operate.

That period of growth gets us to the modern configuration:

<sup>7</sup> <http://www.cs.utexas.edu/~mitra/csFall2012/cs329/lectures/webOverview.html>

<sup>8</sup> Yep, as in programming – same idea you've heard of before.

<sup>9</sup> [https://blog.apnic.net/2015/07/02/ipv4-address-exhaustion-comes-to-arin/](https://blog.apnic.net/2015/07/02/ipv4-address-exhaustion-comes-to-ar/)



Which takes several different stylized presentations to consider.... Some of which, like electric lighting (in some ways equally revolutionary) show the concentration of people on the Earth itself, along with a rough demonstration of the landmasses (North America, in particular!).

*Not bad for a napkin.*

--

*Interesting background on the early Internet:*

Exploring the Internet: A Technical Travelogue, By Carl Malamud

Various speeches by Vint Cerf – who helped created a bunch of this technology (and collected The Turing Award, the Nobel of Computing, in the process)

The RFC – Request for Comments - process, and Jon Postel's legacy, turn up in many sources. His obituary is at <https://www.rfc-editor.org/rfc/rfc2468.txt>

Other useful info: <http://www.internetsociety.org/internet/what-internet/history-internet/brief-history-internet>

Finally.... I think I've credited all the pictures used, apologies and no infringement/etc. intended!