

On NOT Baking in Special Services

The Internet is a work in progress. But when I see people wanting to give special treatment to voice or video it seems as if they've decided the work is finished and we only need to make what we have already work better by baking it in.

I am reminded of the apocryphal (that is, it didn't really happen) story of the patent office deciding it was going to close because everything had already been invented. Accepting today's Internet as finished means accepting the current hierarchical model of trust and all the naiveté associated with that assumption.

Building in preferential treatment will break the dynamic of the Internet. After all a key design requirement for today's Internet is to avoid having to build services into the network.

I've used the example of what would've happened if we'd made VoIP a special service in the 1980's. After all, if we didn't have the PSTN doing voice already then voice might have been considered a purpose for the Internet. The technology was available -- all we needed to do was reserve capacity along the path for those willing to pay a premium for a voice path.

After all, the idea that we could rely on best efforts for voice would've seemed silly back then because there would have been far too many dropped bits and far too little computational capacity to deal with them. Voice did work over some paths but you couldn't assume you'd have the right path.

In fact we did that experiment. If you read SS7 documents it is just like the Internet except designed with the presumption that you needed circuits as the basic stuff of which telecommunications was built and would reserve capacity along the path. (SS7 also assumes hierarchical trust). You can see how well that experiment turned out -- it failed and voice traffic now increasingly going over IP. But we wouldn't have known it failed because we wouldn't have had the success of the Internet for comparison. Instead bits would be held aside for high value services instead of contributing to a purposely limited pool of bits. And that is just what happens in today's telecom.

Over time, as the supply of bits increased, we discovered that we could use the available bits for VoIP. This was largely due to the Web which gave us a reason for having so many bits available. Demand actually increased to meet supply because no bits are special.

In practice things get more complicated as providers attempt to do us favors. One example was when we used modems to tunnel through the phone network by pretending to be voice calls. That meant we couldn't take advantage of all the capacity available "between" bits thus threatening the phone network.

A simple way to think about it is that reserving capacity is like reserving a highway lane before you set out to drive. Bits (or packets) are like individual cars that can share the lane or change lines to use the available capacity. Even better capacity is increasingly rapidly unlike a real highway that just gets more congested.

The solution was to handle bits without all the extra mechanism. This mismatch came up again with P2P and now with "wireless".

We go with what works rather than baking in requirements. This is counterintuitive. Remember that the Web started out on dialup and was usable and even exciting at 1Kbps. It was only as we started to be able to assume higher speed connections that the web site took advantage of the opportunity to add more glitzy bits. And then we discovered that if you repurposed a video distribution network it would be good for video.

But it is not good for the medical monitoring and services that can keep us alive. We can't assume that a connection is available because we've restricted access to billable paths. This is the price we pay for associating value with speed.

Just as we had the PSTN back then we still have a cable TV network. If we make video a special service on the Internet we will kill the dynamic once again. If you are forced to provide enough bits so video just happens to work you'll assure the abundant bits are available for all purposes. Solving problems by adding capacity also simplifies network management because you don't have to infest all network elements with special knowledge about video.

Even better, if you don't require the network to know about video then those selling video would be able to do so without having to pay all the providers along the path for premium services. This means anyone can offer video not just those who can make deals.

Perhaps we can build in video without a payment system -- but why just video and why video at all? What is the complexity of having to build in all the control plane mecha-

nisms that coordinate all the different services and pick winners and losers? And who says the bits are video anyway, especially as encryption becomes the norm in order to support the movement towards sending video over the top of IP? Do we prioritize "important" video or cute kittens?

All this comes back to the naive idea that the Internet is finished when, in fact, it has just begun.

The very idea of having an "Internet Inc." with a central source of IP addresses and global management violates the idea that we create services at the edge. One of those services is networking. The current implementation is a transitional stage.

As we reinvent the Internet from the edge we need to see how far we can go with the fewest requirements on the routers and other networking devices. We should assume nothing more than a best efforts requirement for moving the bits along. We also need to remember the Russian proverb, *trust but verify*. If a router says go west and I find myself staring into the sunrise I should be suspicious.

This would a first square mile (or kilometer) model in which communities, at scale, hire operators and buy gear with the best price/performance. It wouldn't be exclusive control so those offering additional capacity and bypasses would be contributing to the commons rather than competing with it.

This is a model that would scale far beyond what we can do with today's scheme that rewards scarcity. The differentiated services become more valuable by limiting the number of undifferentiated bits.

The idea that wireless is special and that Verizon should be able to use their capacity to sell us services in order to fund their network goes against this. While such concessions may be in Google's short term interest they are not in its long term interest.

It makes little sense to fight for network neutrality if some bits get special treatment either by protocols or by having a carrier hold capacity aside.

Perhaps we are betrayed by our own success. The wonders we can do with the portion of capacity we have now seems like so much to so many when it is really so little.

For more see:

- <http://rmf.vc/?n=IPVd> – why the we need to think about discovery rather than purpose in order to understand what the Internet is.

- <http://www.bricklin.com/qos.htm> for more on the problems with QoS as a general policy.