Is It Time To Make The Network Smarter?

Net Forecasts – Peter J. Sevcik
BCR Volume 34, Number 5
May 2004

When David Isenberg published "The Rise of the Stupid Network" in 1997, I was pleased to see that someone inside AT&T thought like the rest of us "Net-heads." I have long been an advocate of simple, fast network technology, having made the public case for the datagram over the circuit connection model two decades prior to Isenberg's paper.

The dumb network concept has become dogma. For example, in the February 2004 issue of *BCR* (pp. 14-15), my fellow columnist Dave Passmore took the stand that smart networks are a dumb idea.

However, the pendulum may have swung too far, such that networks are at risk of being too stupid. I think that the network community should at least re-examine the premise. And I see a wide range of opinions in the community.

Keeping The Internet Simple

There are two broad arguments for keeping the core of the Internet dumb, or a better term, simple: Speed and flexibility. As the 'Net grows in both users and applications, the traffic continues to increase at an alarming rate, meaning that the network's forwarding plane must stay as streamlined as possible.

Second, the Internet has been a wonderful platform for innovation due to its staunch adherence to the end-to-end principle, best described in RFC 1958 – "Architectural Principles of the Internet" – that states, "As a first principle, certain required end-to-end functions can only be performed correctly by the end-systems themselves."

Unconvincing Arguments for Change

Some arguments for making the 'Net smarter are simply not valid, I believe. The first is that this is a way to make money – i.e., by selling differentiated services or quality of service (QOS) on one network.

We have been hearing about the elusive profits in QOS for some time, but the business case for charging more to deliver the bits better is very limited (see *BCR*, March 1999, pp. 8-10). Scott Bradner of Harvard points out that, "No ISP has been able to convince users to spend more money all of the time to get better service some of the time."

In addition, any proposals for making the Internet smarter cannot force the core routers to perform deep packet inspection or other onerous tasks. Because of the continuing growth cited above, any added intelligence must be either at the service provider edge or in the networks' control plane/higher management plane (i.e., layers of intelligence separate from packet forwarding). This may help explain why Juniper recently acquired NetScreen, which is the leading deep-inspection *edge* device.

Finally, it is clear that the Internet of today, with all its shortcomings, is here to stay. Any proposals to fix the current Internet by changing or eliminating the best-effort service are a non-starter. The billion or so users of the current Internet are getting something done, so change had better not affect them negatively.

Making the Internet Incrementally Smarter

However, there are also three broad reasons for making some parts of the *network infrastructure* (as opposed to the Internet *per se*) smarter.

First, there is already a lot of end-system intelligence that makes the system work – unfortunately, this intelligence is so well distributed that some of it finds its way into the hands of unfriendly players. Even a few bad people can cause havoc, as witnessed by spam, viruses and denial-of-service attacks. Yes, the good people and machines on the edge can protect themselves, but the work is much harder if the network does not help at all.

To provide this help, the core or the service provider edge should ensure that the end-user edge systems are not operating maliciously. This can be accomplished by a smarter control plane or moving the network-edge boundary to keep some functions such as authentication out of the user's hands.

The second reason for a more intelligent transport is the increasing amount of real-time traffic. This is an issue primarily on the access link, rather than the Internet backbone.

A VOIP or IP-video call has two bi-directional edge access lines involved (my send/receive and your send/receive). Each of the four circuits needs proper traffic prioritization and policing, but I have control over only one-quarter of these resources (i.e., my send). More coordination is needed to ensure full end-to-end two-way quality. Vint Cerf of MCI (co-creator of IP), puts it well by asking for a better "interface that provides richer vocabulary for class of service inbound toward the network, and better status/performance information outbound from the network."

The third and final reason is that enterprises are using the Internet to connect to business partners and customers around the globe. At these distances, the speed of light can tremendously affect the performance of transactions. But the problem is even worse than the speed of light.

A true global enterprise application strategy requires that many users be connected to the application from many local ISPs that are then interconnected to ISPs that serve the enterprise. Peering points among ISPs are a chronic source of congestion because there is little business incentive to properly invest in these costly connections.

Thus, performance degrades as a function of ISP count and physical distance. Again, it would be easier for the end-to-end model to succeed if the end-systems could signal the transport system with a specification of performance they need and then be *assured* that all the ISPs in the path are adhering to the specification. This should be performed by a smarter control plane and better service interfaces between carriers.

The Grand Strategy

Traditional carriers have begun talking about building a separate, parallel, carrier-class IP network to provide business-class services apart from the best-effort Internet. The Internet as we know it would continue to serve consumers and many enterprises, and would remain open to everyone.

The new business-class Internet would be a parallel structure composed of many ISPs that interconnect to provide "better" global services. The key difference from the idea of a QOS-enabled Internet, as described above, is that you're fundamentally using a different network: You're not paying more to ride first class over coach on the same plane; you're paying more to fly in a plane rather than taking the train.

This new business-class Internet, or Infranet as Juniper calls it, would have all the services as the consumer-class Internet but stricter requirements. Businesses that connect to it would have to adhere to a higher standard of network hygiene so they would not pollute the commons, as has occurred in the current Internet.

In addition, there might be new and different services catering to the needs of businesses, which consumers would not need or be willing to pay for. Specifically, the business-class Internet could have assured end-to-end services, methods of identifying application needs and communicating them to the network, and better security.

In theory, a business-class Internet could be an economical alternative to private networks. The shift would reduce the carriers' revenues, but they might have to accept it anyway: As it is, they're likely to lose much of their existing high-margin private-network traffic, because enterprises are moving all their public-facing electronic interactions to the public Internet.

What's more, enterprises are also moving their back-office applications and networks onto the Internet. Enterprises are buying extra products and services to make the public Internet a better place to operate a business. In short, they've been forced by the end-to-end principle to make their own edges smarter.

So traditional carriers face a simple choice. Lose *all* of the high-margin private network business to a public Internet that the user improved on his own, or build a new, better Internet that does not need any performance enhancements. The margins will be lower than today's private network business, but that is still better than no revenue at all.

Dialogue and Standards Are Needed

We should constantly question old assumptions, even those that have served us well. We should open the dialogue rather than hide behind extremist views of pure "dumb" or "smart." At least two goals should be addressed: How to make the network smarter, and how to better define or move the border between user and network.

The vendors are already starting to add intelligence within an overall architecture. Cisco's Intelligent Information Network (IIN) and Juniper's Model for Integrated Network Transformation (MINT) are the beginnings of a smarter network. However, the two initiatives are not coordinated nor standards based, which is not good for any of the parties involved.

The smarter network should be discussed and defined in a public forum. There are some good candidate venues like the Internet Research Task Force (IRTF), Internet Engineering Task Force (IETF), and Dave Clark's new Knowledge Plane project (see www.isi.edu/know-plane/#intro).

Changing the fundamental assumptions of a system that works is risky. Detailed user-network interfaces and strict rules are what prevented the use of a modem on the Bell System. It took the Carterphone case to unlock a universe of new data applications on the old rigid voice network. Nevertheless, I think that some form of added network intelligence and better controls are essential in order to avoid network anarchy or solutions by corporate fiat.

We need to shift the dialogue. I find Fred Knight's remarks in the March 2004 *BCR* (see p. 2) a perfect example. He apologizes and laments having to take a free ride on a neighbor's wireless LAN when his DSL service failed because of the MyDoom virus.

In fact, he did just what the founding fathers of the Internet expected him to do: He made a smart user edge play. Instead of apologizing, he should be outraged that the 'Net was so bad that he had to do it at all! Serious enterprise applications that manage money, control health, or maintain national security can't operate on such a network.

Companies Mentioned
Cisco Systems (www.cisco.com)
Juniper Networks (www.juniper.net)

Peter Sevcik is president of NetForecast and is a leading authority on Internet traffic, performance and technology. Peter has contributed to the design of more than 100 networks, including the Internet, and holds the patent on application response-time prediction. He can be reached at peter@netforecast.com.