netWorker
THE CRAFT OF NETWORK COMPUTING

The Dawn of the Stupid Network

By David S. Isenberg

Originally published as the cover story of ACM Networker 2.1, February/March 1998, pp. 24-31.

In recent history, the basis of telephone company value has been the sharing of scarce resources -- wires, switches, etc. - to create premium-priced services. Over the last few years, glass fibers have gotten clearer, lasers are faster and cheaper, and processors have become many orders of magnitude more capable and available. In other words, the scarcity assumption has disappeared, which poses a challenge to the telcos' "Intelligent Network" model. A new type of open, flexible communications infrastructure, the "Stupid Network," is poised to deliver increased user control, more innovation, and greater value.

Telephone companies (telcos) have always pushed technology improvements that promote the smooth continuation of their basic business. They invented the stored program control switch in the 1970s, as a move toward cost reduction and reliability. Programmability also made possible certain call routing and billing services. In the 1980s, phone companies began marketing these services as the "Intelligent Network." Technology continued its trajectory of improvement, but because technology began to change the value proposition in ways that the old business could not assimilate, the telcos seemed to "fall asleep at the switch" at the core of their network. Meanwhile, the Stupid Network - based on abundant, high-performance elements that emphasized transmission over switching, as well as user control of the vast processing power at the network?s edges - was taking shape.

KEEP IT SIMPLE, STUPID

"Keep it simple, stupid," or KISS, is an engineering virtue. The Intelligent Network, however, is anything but simple; it is a marketing concept for scarce, complicated, high-priced services, surrounded by features like 800 service, call waiting, and automatic calling card validation. These are intertwined in the network architecture in a plethora of service adjuncts, each with its own systems for operations, provisioning and maintenance. This complicated, centrally controlled amalgam of systems is born of a single application - two-way real-time voice communication.

So what exactly is a Stupid Network? George Gilder observed more than five years ago, "In a world of dumb terminals and telephones, networks had to be smart. But in a world of smart terminals, networks have to be dumb." In a Stupid Network, control passes from the center to the edge, from the telco to users with an abundance of processing power at their fingertips. The center of the network is based on plentiful infrastructure - cheap bandwidth and switching - that is about as smart as a river. The water in a river, like a data object in a Stupid Network, gets to where it must go adaptively, with no intelligence and no features, using self-organizing engineering principles, at virtually no cost. Bits go in one end and come out the other. Data flows - like water - define the movements and channels within the system.

Eric Clemons, a professor at the University of Pennsylvania?s Wharton School of Business, makes the distinction between strategy and doctrine. "Strategy," he says, "is learning how to deal with dogs. Doctrine is about belief: "Dogs don?t do that." Telco doctrine, formed in the age of monopoly and scarce infrastructure, is rarely examined explicitly. When there was only one telephone company, what Ma Bell did defined how things were. So today, even though the new era of competition requires clear thinking and new beliefs, telco culture inextricably mixes doctrine and strategy.

THE CIRCUIT-SWITCHED LEGACY

The Intelligent Network concept has its roots in the first software-controlled switches in the 1970s. In those days, working with computers meant writing code to save a byte here and an instruction cycle there. Current software practices, such as object-oriented programming, were too inefficient for prime time and were relegated to the confines of academia.

Thus, telephone network equipment was designed in a climate of scarcity. Consider the local exchange, represented by the three digits of a telephone number that follow the area code (the nxx in the pattern nxn-nxx-xxxx). The local exchange "owns" the last four digits of a telephone number. Theoretically, a local exchange can serve up to 10,000 telephones, e.g., with numbers 762-0000 through 762-9999. The design assumption, though, is that only a certain percentage of these lines, maybe one in 10, are active at any one time. But should these assumptions change temporarily (e.g., an earthquake in California) or permanently (calls to AOL lasting several times longer than normal voice calls), the network hits its limit. Then, getting a dial tone becomes a matter of try, try again.

Even more assumptions have now changed permanently. Before 1996, long distance carriers like AT&T, MCI and Sprint cranked a precise set of capacity planning equations that told them which switches would need more circuits, which routes needed to be upgraded over the next year, and where to plan for new switches. Suddenly, increased Internet usage threw the telcos an unplanned 60% increase in data traffic. Suddenly, some points of their network hit capacity. The telco fallback position had always been to lease capacity from their rivals, but this wasn?t available either, because the other telcos were maxed out as well.

"INTELLIGENT" NETWORK SERVICES

In the late 1970s, telcos became fixated on their expensive investments in computer-controlled switching, and were intrigued by the prospect that they could do "intelligent" things with these investments. They reduced the cost of running the network and formed a platform for revenue producing services geared toward call set-up and billing. The concept of network control was extended to let digital switches communicate with databases (known as Service Control Points) and signal processing systems

(Intelligent Peripherals).

Intelligent Network (IN) specs were meant to encourage telecom equipment vendors to design their equipment to work in a multi-vendor environment, so telcos would not be locked into one supplier. In addition, IN equipment was designed to work with certain customer systems and databases. Some common Intelligent Network services include: routing calls to a number other than the one the caller originally dialed (the basis of 800 service); caller options ("press 1 for customer service," etc.); and supplying calling party numbers directly to customers for database lookup (which is why I must call my bank from my home phone when my new ATM card arrives in the mail).

STUPID IS BETTER

Stupid Networks have three basic advantages over Intelligent Networks - abundant infrastructure; underspecification; and a universal way of dealing with underlying network details, thanks to IP (Internet Protocol), which was designed as an "internetworking" protocol. Some key "two-fers" emerge from these basics: Users gain end-to-end control of interactions, which liberates large amounts of innovative energy; innovative applications are rapidly tested in the marketplace; and innovative companies attract more capital and bright people.

Abundant Infrastructure

In a Stupid Network, if you have congestion, you just add more connections, bandwidth, switching or processing power. If you want reliability, you add more routes or more redundancy. If you need more intelligence for features or services, you add it at the endpoints. As early packet network visionary Paul Baran points out, it is possible "to build extremely reliable communications links out of low-cost unreliable links, even links so unreliable as to be unusable in present networks."

Even as the costs of networks have dropped, capacity has improved manyfold. At the dawn of the digital transmission era, for example, you could run 1.5 megabits - 24 calls - on a coaxial cable as thick as your ankle. Today, network providers routinely put several tens of gigabits - a few hundred thousand calls - on a single glass fiber as thin as a human hair. Switching used to be scarce, too, but now it is equally abundant. Where a human operator could set up maybe 100 calls an hour, modern computer controlled switches, such as Lucent?s 4ESS, can now complete about 1 million calls in the same hour. Furthermore, if you consider that routing a single packet is equivalent to setting up a call, routers can now set up 3.6 trillion "calls" an hour. And the prices for these components have come way down: Today when you buy a Gigabit Ethernet switch, you get 1,000 chunks of 64 kbps throughput (each equivalent to a phone call) for every dollar.

This leads to two different models of capital investment. In the telco model, network expansion is a big decision that requires expert engineers, detailed Erlang models of congestion, and several consultants - and takes months, if not years, to implement. But if you’re running a Stupid Network, expansion is no problem. A few hundred gigabits? Put it on my credit card. More switching capacity? Take it out of petty cash.

Underspecification

The Intelligent Network is tightly specified for voice. All other data types require special leased access lines, or awful kludges like modems. The Stupid Network is underspecified - this means bits-in, bits-out. It is nothing special for underspecified networks to carry voice, music, bank balances, e-mail or TV on the same facilities. You stuff bits in one end of the network, and they find their way to the other end of the network. Packets carry their address with them, and out they come at the other end, right where you want them to be.

Underspecification also means that there is little thought for congestion control. So what if there is congestion, or even crashes? On the whole, the convenience of underspecification more than compensates for the occasional jam-up. And you can always add more "infrastructure."

Internetworking

Internet Protocol points the way to a key property of Stupid Networks. The foremost design goal of IP is to cross multiple, physically different networks. To IP, it doesn?t matter if the underlying transport is circuit, SONET, Ethernet, Bitnet, FDDI or smoke signals. An IP application works the same no matter what the underlying network technology. This makes the details of how a network works irrelevant (including how "intelligently" it is engineered).

IP neatly takes the provider of the physical network infrastructure out of the value proposition. No matter how intelligent a telco?s network might be, if it is running IP, its intelligence is reduced to commodity connectivity. Networks that run IP are left with one main source of distinction: how much connectivity they provide. Thus, the Internet that we know and love is a "virtual network" - a "network of networks" - that is independent of wires and transport protocols.

Because IP makes the details of the network irrelevant, all that matters is that the bits sent by your machine are received by my machine, and vice versa. In an IP communication application, users don?t care how the "call" is set up, or even if there is a telephone call that forms part of the communication path between endpoints.

User Control

This means that users are in control of their interactions. Suppose, for example, that two users want to bring a third party into an interaction; they just do it. An IP-connected user does not need to order special three-way connectivity service from the networking company. All that user needs to do is write (or install, or use) a program that sends packets to two different destinations and receives from both of them.

A BOOST TO INNOVATION

This ability to "just do it" liberates huge amounts of innovative energy. If I have a Stupid Network and I get an idea for a communications application, I just write it. Then I send it to my buddy, and my buddy can install it, too. If we both like it, we can send it to more people. If people really like it, then maybe we can charge for it - or even start our own company. Yahoo!

In contrast, the Intelligent Network impedes innovation. Existing features are integrally spaghetti-coded into the guts of the network, and new features must intertwine with the old. For example, until recently you could not get Caller ID for an incoming call when you were on the phone. To fix this, Bellcore had to invent a low-tech, low-functionality, high-complexity protocol called Analog Display Services Interface (ADSI). Call waiting with caller ID, however, would be a no-brainer under Internet telephony - a packet containing caller ID information could be treated just like any other packet arriving at my system, to be interpreted by the end-user IP telephony application.

New Internet capabilities are coming online that will fuel Stupid Network innovation. Internet Protocol Version 6 (IPv6), stabilized in 1995, is becoming available now. Its capabilities include expanded address space, real-time functionality, mobility management and carrier selection, hooks for authentication and data integrity, multicasting, and easy coexistence with and migration from the current IPv4 standard. IPv6 capabilities in the hands of innovators will foster whole new areas of applications.

A QWEST FOR BANDWIDTH

Perhaps the leading proponent of the Stupid Network is Qwest Communications, which is building a huge 16,000-mile-plus SONET based network that will reach 125 cities by next year. According to Qwest?s executive VP of products, Nayel Shafei, "Qwest is leveraging the greatest currency of all - our unlimited bandwidth - to shape the future of telecommunications." The company is about to go international, too; it will turn on its 1,400-mile Mexican backbone in mid-1998, and it recently announced a trans-Atlantic link.

The Qwest network will run native IP over SONET, according to CEO Joe Nacchio. This is an industry trend; Sprint, for example, recently went ATM-less on its SONET/IP backbone. This is possible because physical layer infrastructure is becoming more abundant and endpoints are becoming more capable.

Qwest is using its network to compete with established telcos on price. Qwest now offers long-distance telephone service over its IP backbone at 7.5 cents per minute, a 25% discount over the now industry-standard dime.

Shafei maintains that the sound of telephone service over Qwest+s network will not have the glitchy, scratchy quality that is associated with Internet telephony today. Calls will be mainlined into the jugular of Qwest?s unlimited bandwidth. Consumers will use Qwest?s IP long distance service by making a local call on a normal telephone, thus accessing a circuit-to-IP platform by Vienna Systems, a Newbridge Networks affiliate. The Vienna platform will not compress the voice; it will simply packetize the raw, 64 kbps signal and send it via IP. Shafei claims that the quality will be virtually as good as circuit-switched voice.

Several other entrepreneurial companies have tested the waters for Stupid Network innovations. Vocaltec, for instance, was the first company to commercialize the fact that you could use the Internet for voice communication. If I have Internet access, and you have Internet access, and we?re both running Vocaltec software, we can talk to each other for as long as we like, for no incremental cost, no matter where in the world we are. Though voice quality is still not ideal, and delay can interfere with the flow of a conversation, look for this new communications niche to merge with other Internet applications to create new value.

Placeware, a Xerox PARC spin-off for interactive multimedia meetings and conferences over the Internet, mixes Internet telephony with data sharing, presentation graphics, and a crude representation of the meeting space. Demos of the technology seem to add a lot to voice conferencing, and it gives a more participatory, less self-conscious feel than a video conference.

BEYOND QOS TO SIMPLE STUPIDITY

Intelligent Network advocates point out that networks need to treat different data types differently. Right now, they?re absolutely correct. There is a network for telephony, another network for TV, and proprietary leased-line networks for financial transactions - and none of these are ideal for public Internet traffic. You need to have low delay for voice telephony, the ability to handle megabit data streams with ease for TV, and low error rates and strong security for financial transactions.

Quality of Service (QOS) is an intermediate step in the journey from separate networks to a single, simple Stupid Network. QOS, in standard telco thinking, means a repertoire of different ways of handling each type of data on a single network. If the Stupid Network is to become a bona fide integrated service network, it will need to carry all kinds of data with different needs.

But suppose technology improves so much that the worst QOS is perfectly fine for all kinds of traffic, without a repertoire of different data handling techniques. Suppose, for example, that everyday normal latency becomes low enough to support voice telephony, while at the same time allowing enough capacity for video, plus data integrity strong enough for financial transactions. This would be a true Stupid Network - one treatment for all kinds of traffic.

Skeptics might say that there would have to be dramatic improvements in networking technology for this to happen. Well, we?re getting there. Routing switches from Madge and Foundry recently showed performance impressive enough to conclude that routing latency and jitter (variation in packet arrival time) may soon be a negligible issue. But these were lab tests, not field usage, and packet losses were as high as 1% under some conditions. So we are not there yet - but perhaps we will be soon.

PLAYERS IN THE NEW ORDER

Still other technologies of abundance, with the potential to break the telcos? foot-dragging hegemony, have attracted interest from entrepreneurial vendors. Here are some of the leading candidates:

- LMDS: This technology provides a wireless broadband last-mile path to the Stupid Network. The

FCC LMDS auction, completed on March 25, opens U.S. markets to deployment and service over the next couple of years. The two big bidders were Nextband and WNP Communications. Equipment manufacturers Hewlett-Packard, Stanford Telecom, Texas Instruments, Tadiran and others also will be beneficiaries.

- CDMA: Another wireless data access method; Qualcomm is still well positioned. Also watch Broadband CDMA, an emerging open standard that can deliver from fractional T-1 on up. Interdigital is one leading B-CDMA player.
- Gigabit Ethernet switching: This technology has moved from laboratory to the marketplace in a remarkably short time. While Ethernet has been synonymous with Local Area Networks in the past, "Neighborhood Networks" are replacing ATM as the vehicle of choice for campus nets. Can real neighborhoods be far behind? Gigabit Ethernet players include Bay Networks, Cisco, 3Com, Cabletron, Foundry, Extreme Networks, and many others.
- Cable Modems: Players include set-top box makers General Instruments and Scientific Atlanta, plus Motorola, Hybrid Networks and others. Cable provider Comcast is a good bet for cable modem service - Bill Gates thinks so, anyway (to the tune of $1 billion).
- The power companies: Also worth watching, following Nortel?s announcement last November of technology to deliver data to end users over power lines.

All these infrastructure improvements are rapidly making the telcos' Intelligent Network a distinctly second-rate choice. The bottom line, though, is not the infrastructure; it is the innovation that the Stupid Network unleashes. The Stupid Network assures the next paradigm-breaking, market-making "new thing." The only question is who will become the next Netscape, the next Microsoft - or the next Ma Bell. And that's not a stupid question.

The netWorker website that time forgot (v1.1 only) is http://www.acm.org/networker/default.htm

Posted on 6 June 1998