JUST DELIVER THE BITS, STUPID

John Nolan, First Mile Networks

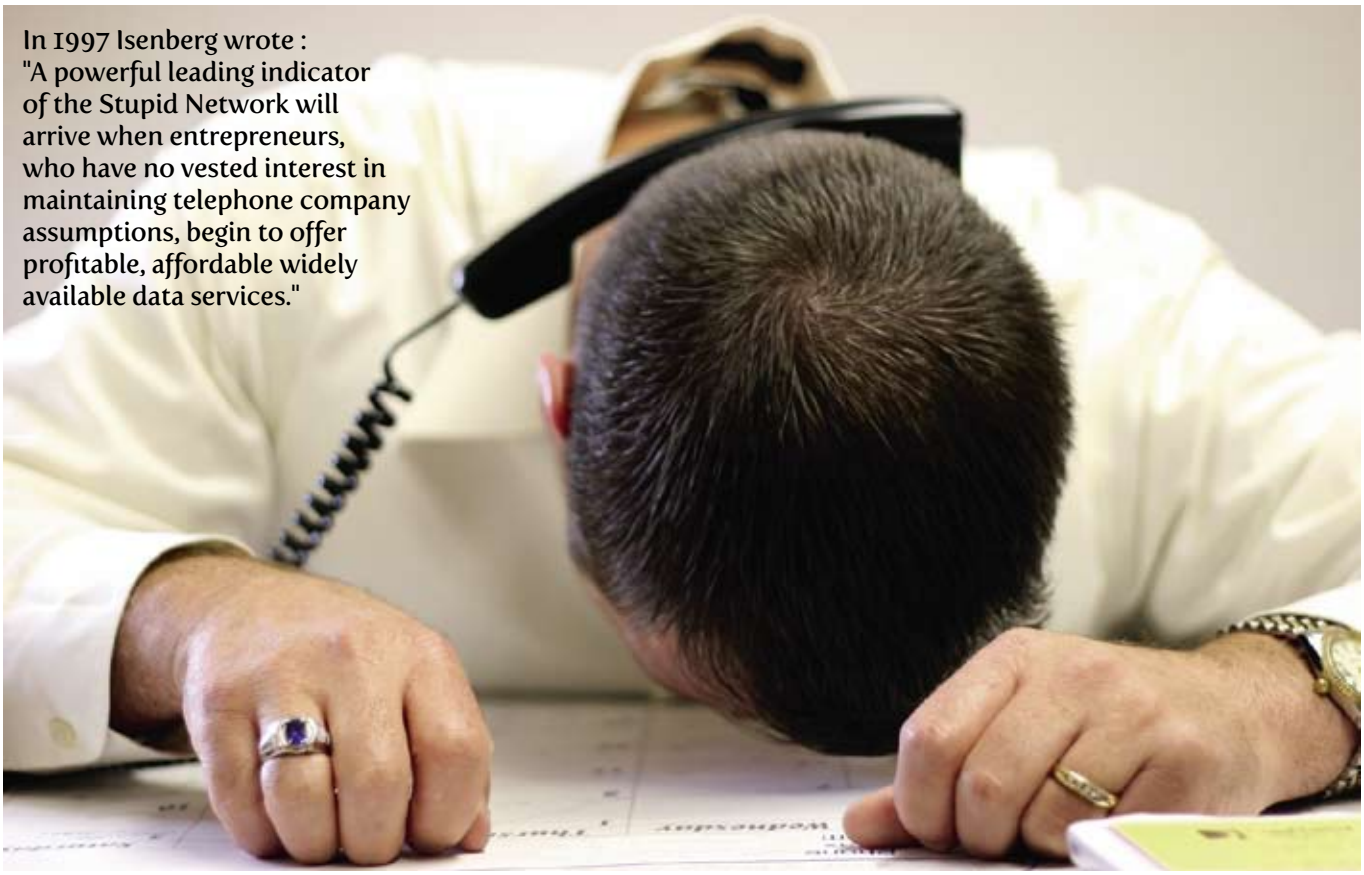

'Just deliver the bits stupid'. So wrote David Isenberg in his paper *The Rise of the Stupid Network*. John Nolan remembers it, re-introduces it and reviews it.

I can't remember if Isenberg's paper had the same effect on me as when I first heard *Up From the Skies* [1] way back in 1967, but it too made me sit up and take notice. Isenberg's paper is considered seminal in that, at the time, he was promoting an alternative approach to the existing orthodoxy in telecoms thinking and more recently, it was considered to be the eighth most important Internet paper ever published [2]. The paper is reproduced below with a view to encouraging discussion amongst ITP members. I will follow it with my own observations.

An historic note may be in order in that, like Isenberg, I too am a child of the all encompassing Telco as I started my working life in August 1971 in what was then Post Office Telecommunications, London City Telephone Area (and subsequently joined the predecessor of the ITP, the IPOEE very soon after). It should be noted that Isenberg's paper was first written in May 1997 and many of the 'Intelligent Network' concepts Isenberg describes are familiar to me. Whilst the paper did not have AT&T's endorsement, it was widely available on the Internet and was 'formally' published in Computer Telephony in August 1997. If I recall correctly, I was consulting to Enertel in the Netherlands at that time launching an indirect access product, but wasn't aware of the paper until early 2000. So here it is.

RISE OF THE STUPID NETWORK

Why the Intelligent Network was once a good idea, but isn't anymore. One telephone company nerd's odd perspective on the changing value proposition.

DAVID ISENBERG

OBSOLETE ASSUMPTIONS & ENDURING MENTAL MODELS

Design-by-assumption works as long as assumptions hold. Assumptions are shortcuts to useful efficiencies, provided they are not violated. The classic telephone company value proposition, embodied in today's telephone network, holds:

- That expensive, scarce infrastructure can be shared to offer premium priced services.
- That talk - the human voice - generates most of the traffic.
- That circuit-switched calls are the "communications technologies" that matter, and
- That the telephone company is in control of its network.
- Telephone companies still behave as if these assumptions hold despite:
- Up to several thousand-fold declines in key infrastructure costs over the last two decades,
- A 20-year double-digit annual growth rate in the volume of data traffic, so that the volume of data traffic is now overtaking the (also growing, but more slowly) volume of voice traffic,
- The many different data types that now travel over the telephone network (despite the fact that the network is not optimized for all these data types),
- The many different types of "communications technologies," from television to Ethernet, that are not part of telephone network architecture, and
- The Internet, which, because it makes the details of network operation irrelevant, is shifting control to the end user.
- The Intelligent Network is a straight-line extension of the four assumptions above - scarcity, voice, circuit switching, and control. Its primary design impetus was not customer service. Rather, the Intelligent Network was a telephone company attempt to engineer vendor independence, more automatic operation, and some "intelligent" new services into existing network architecture. However, even as it rolls out and matures, the Intelligent Network is being superseded by a Stupid Network,
- With nothing but dumb transport in the middle, and intelligent user-controlled endpoints,
- Whose design is guided by plenty, not scarcity,
- Where transport is guided by the needs of the data, not the design assumptions of the network.

The Stupid Network is not all here yet. It is in its infancy. It needs to get stronger and, well, a bit more coordinated.

Some telephone company people realize that things are changing, and must change. But they are hemmed in by conscious, deliberate, long established telephone company practices. Many are also hobbled by less conscious telephone company mental models of "communications," "technology," and "customer needs." While these people may realize that the old ways are becoming obsolete, they live in a world conditioned by an encompassing, arcane legacy that only remembers "rational," incremental change. (Note: here "telephone company" refers to large companies whose main business is to provide circuit switched voice calling service. In the US, most of these are the heirs of the Bell System legacy - but Sprint, MCI, GTE, SNET, and others might also try on this shoe, and if it fits...)

COMPUTERS AS SCARCE RESOURCES

It used to be more expensive to complete telephone calls than it is today. The operator-completed call gave way to call completion by electro-mechanical switch. Then, in the late 70s, the era of computer controlled electronic switching made placing calls even cheaper and more reliable.

In those days, computers, including those that controlled switching, were still considered expensive, scarce resources. When I worked in the nascent electronic toy industry in 1979, a single insight that eliminated six transistors paid my way. And the same factor - the need to save two expensive bytes of memory - laid the basis in this era for the Year 2000 Problem (stay tuned to the eleventh hour news for more on THIS story!).

Now computer circuits are thousands of times cheaper. Moore's Law is what we call the ongoing improvement in computing cost and power. But in the 70s it was not generally known to be a `law' - to most telecommunications engineers (and to humanity in general), it has become the most game-changing wild card played in recent times.

Telephone networks have been designed for optimal use of scarce resources. The local exchange in your city, which handles the last four digits of your telephone number, theoretically could handle up to 10,000 telephones, e.g., with numbers 510-547-0000, 0001, 0002, et cetera through 510-547-9999. But the switching office is not designed to handle 10,000 simultaneous calls. It is designed to handle far fewer, maybe one tenth of that, based on the assumption that even in the busiest time of the day, only a fraction of its telephones will be active at any one time.

The network works as long as engineering assumptions (e.g., the length of a call, number of call attempts, etc.) do not change. But let the assumptions change episodically (e.g., Rolling Stones tickets go on sale), or structurally (calls to Internet service providers last several times longer than voice calls), and the network hits its design limits -

> "Design-by-assumption works as long as assumptions hold."

"If you hate hanging on hold, you are part of a huge latent market – do you know anyone who doesn't."

completing a call becomes a matter of try, try again.

What if network design were based on another assumption - that computation and bandwidth were cheap and plentiful?

DOING 'INTELLIGENT' THINGS WITH PHONE CALLS

Once the telephone companies began doing digital switching, the idea that you could do "intelligent" things with calls was not far behind. The concept of network control was extended to let various centralized resources - digital switches, databases (Service Control Points) and signal processing systems (Intelligent Peripherals) - communicate among each other by extending the telephone network's control protocol (SS7).

As noted above, the main force motivating the Intelligent Network was a telephone company attempt at "vendor independence" so telephone companies could get better deals from their suppliers. Thus, Intelligent Network specs were meant to encourage vendors to design their equipment to work in a multi-vendor environment - to interoperate. As a side benefit, almost an afterthought, some of the newly specified equipment could also interoperate with the business systems of certain customers - but only via limited, cautiously designed interfaces. Virtually all of these services centre around call completion, automation, and billing. This, in a nutshell, is the concept marketed as the Intelligent Network. Some Intelligent Network service examples include:

- Routing calls to different numbers than the one that the caller originally dialed (this is the basis of e.g., 800 service).
- Giving caller choices before the call is completed ("push one for domestic reservations," etc.).
- Saying, "Calling Card, Collect, Third-Party, or Operator" to control payment options.
- Verifying that the calling card number is valid in "real time."
- Supplying calling party numbers directly to customers for database lookup (which is why I must verify from my home phone that I got my Citibank card in the mail).

Expensive computers, intertwined in central network operations, do this. Belief becomes reality. But wait! The telephone companies are now losing design hegemony - the news that "The Internet is here!" is beginning to penetrate the telephone company inner sanctum.

MEETING CUSTOMER NEEDS

The astute reader might by now suspect that the main beneficiaries of the Intelligent Network are the telephone companies themselves. Nevertheless, telephone companies propound a "philosophy" that the Intelligent Network makes it easy to introduce new services and new technologies, and to meet new customer needs.

New customer needs, when they are detected, filter into the telephone company slowly. Some needs, the ones with big, obvious, immediate payoffs, get attention from decision makers, who then request a business case, which must then get approved. The next step is the development plan, followed by the Operations, Administration, Maintenance, and Provisioning Plans. Then if all goes well, the telephone company might begin the process of implementation. This can take years, or even decades (witness ISDN).

If you hate hanging on hold, you are part of a huge latent market - do you know anybody who doesn't? Yet, telephone companies have yet to use Intelligent Network capabilities to effectively ameliorate this problem. Now, suppose Internet Telephony gets as good as telephone company telephony (see below), and some enterprising independent programmer wants to make a product that solves the problem of being on hold. They would simply write an end-user application and sell it from their web site. If it works, and people like it, they will sell lots of it. If not, they might try again. But they don't have to go through any long, bureaucratic economic justification, business planning, and technical development processes - they just do it. Internet Telephony, because the Internet Protocol works at the level that user software manages the session, takes the telephone company out of the value equation.

THE INTERNET DIS-INTERMEDIATES THE TELEPHONE NETWORK

The Internet breaks the telephone company model by passing control to the end user. It does this by taking the underlying network details out of the picture.

Let's look at how this works in the case of voice. To the telephone company, there is one main way of transmitting voice - sampled in 8-bit bytes, 8000 times a second, for an aggregate rate of 64 kbit/s. The entire telephone network is designed around this rate. But if you want to send voice on the Internet, you can encode it at any rate you want, and send it at any rate up to the one that the slowest underlying network link

supports. The recipient must have the right decoder running in her intelligent terminal, too.

The very name, Internet, denotes that it is designed to network networks. Y ou can use Internet Protocol on an Ethernet to communicate with an X.25 network, an FDDI network, or a modem - lower layer protocols are submerged, made irrelevant. So if you are on an (e.g., 10Mbit/s) Ethernet, and your endpoint application wants to send better quality 256 kbit/s voice, no problem. You can't do that with the telephone network.

Or, with a different application (on the same endpoint and network) you can send six different interwoven 10 kbit/s voice streams to six different destinations at the same time. And you don't have to tell your Stupid Network provider anything about it, or pay a premium to install anything special. The network provider becomes virtually irrelevant - the user controls the relevant capabilities.

TRUE VOICE, FALSE START

I contrast the flexibility of a Stupid Network with my experience as a member of AT&T's True Voice technical team. AT&T True Voice was a valiant attempt to improve circuit switched voice quality as much as possible in the context of current network architecture. If we had not been constrained by network architecture, the easiest way would have been to increase the sampling rate or change the coding algorithm. But to actually do this, we would have had to change every piece of the telephone network except the wires. So we had to work within the designed 64 kbit/s data rate.

An astute AT&T perceptual psychophysicist (and a friend of mine) determined that voice quality could be substantially improved by boosting the bass part of the signal, that part of the audio spectrum between 100 and 300 cycles per second. But as we set out to implement this conceptually simple improvement, we kept running into the problem that there were too many places in the network that had built in "intelligent" assumptions about the voice signal - echo cancellers, conference bridges, voice messaging systems, etc. - and too many devices that depended on these acoustic assumptions for their correct operation - modems, fax machines, and a surprising number of strange devices with proprietary analog protocols. After about two years of intense effort, we made a noticeable difference, one that most listeners preferred (if asked explicitly), but it was not as large as it could have been. There was too much "intelligence" intertwined with the basic transport.

> "With a Stupid Network, you'd get a different program, install it in your intelligent end user device and run it."

The True Voice experience led me to see the advantages of a network - a Stupid Network - that would let you stuff bits in one end and get them out the other without getting tangled up in cobwebs of legacy assumptions. Want a different voice quality? With a Stupid Network, you'd get a different program, install it in your intelligent end user device and run it.

A NETWORK ENGINEERED FOR USE

There is no longer first-order economic justification for a telephone company to engineer and control scarce, expensive, network resources - the basic conditions no longer obtain. The age of plentiful computing is here. I have a multi-color, three dimensional screen saver that uses the entire capacity of my 200 MHz Pentium. The designers of the Intelligent Network never imagined such "wasteful" use of processing "intelligence." The age of plentiful bandwidth is just around the corner, as several families of technologies (fiber, satellite, cable modems, xDSL, LMDS, and low power TV, to name just six) line up to break the local bandwidth bottleneck, and as the capacity of backbone fiber has risen from 2 to 6 to 10, 20 and 40 Gbit/s over just the last few years.

The age of centralized control is ending too, with the rise of the next generation of Internet - and especially the appearance of circuit-like Internet mechanisms, such as those in the latest version of Internet Protocol (IPv6), designed to tame delay and improve real-time two-way Internet voice.

JUST DELIVER THE BITS, STUPID

A new network "philosophy and architecture," is replacing the vision of an Intelligent Network. The vision is one in which the public communications network would be engineered for "always-on" use, not intermittence and scarcity. It would be engineered for intelligence at the end-user's device, not in the network. And the network would be engineered simply to "Deliver the Bits, Stupid," not for fancy network routing or "smart" number translation.

Fundamentally, it would be a Stupid Network.

In the Stupid Network, the data would tell the network where it needs to go. (In contrast, in a circuit network, the network tells the data where to go.) In a Stupid Network, the data on it would be the boss.

Instead of fancy "intelligent" network routing translation, in a Stupid Network, intelligent end-user devices would be connected to one or more high speed access networks - always listening for relevant information, for data addressed to their owner. Sometimes a "communication" might be a few bits, perhaps a short, pager-type message. Other times, it might be longer, like email. In the event of the need for two-way voice communication, an initial message might state the identity of the "caller," and/or inquire of the whereabouts of the owner. The intelligent end-user device could apply its knowledge of where its "owner" was, and who the caller was. Then, if it were programmed to do so, it could launch a message to its owner, telling of the call, the caller's identity, location, and any other information. It could also forward as much information as practical.

End user devices would be free to behave flexibly because, in the Stupid Network the data is boss, bits are essentially free, and there is no assumption that the data is of a single data rate or data type.

IDIOT SAVANT BEHAVIORS FOR DIFFERENT DATA TYPES

In the current telephone network, voice is the assumed data type, unless specially ordered, high cost services are ordered. But in the Stupid Network, because the data is the boss, it can tell the network, in real time, what kind of service it needs. And the Stupid Network would have a small repertoire of idiot-savant behaviors to treat different data types appropriately. If the data identified itself as financial data, the Stupid Network would deliver it accurately, no matter how many milliseconds of delay the error checking would take. If the data were two-way voice or video, the Stupid Network would provide low delay, even at the price of an occasional flipped bit. If the data were entertainment audio or video, the Stupid Network would provide wider bandwidth, but would not necessarily give low delay or absolute accuracy. And if there were a need for unique transmission characteristics, the data would tell the Stupid Network in more detail how to treat it, and the Stupid Network would do what it was told.

The Stupid Network would let you send mixed data types at will - limited only by the knowledge and imagination of the application programmer community. One-way voice messages, multi-way voice conferences, two-way video, email, documents, audio and/or video entertainment, whatever, could be mixed and interspersed at will, within and between sessions. You would not have to ask your Stupid Network provider for any special network modifications - its only function would be to, "Deliver the Bits, Stupid."

One thing about the Stupid Network is clear - the physical elements that comprise the network would be neither expensive nor scarce. There would be little profit margin in shipping dumb bits. There would be lots of high value Business Ideas supported by the Stupid Network, above and beyond transport.

LEADING INDICATORS

A rudimentary form of the Stupid Network - the Internet - is here today. The telephone companies are beginning to realize this. Fearing erosion of their control and, more importantly, their revenue stream, they have been quick to call for the banning of Internet Telephony, quick to call for the federal imposition of charges on Internet access, and slow to implement widely available, reasonably priced broadband services. This creates a chicken and egg problem - while the hungry wait for dinner and breakfast.

A powerful leading indicator of the Stupid Network will arrive when entrepreneurs who have no vested interest in maintaining telephone company assumptions begin to offer profitable, affordable, widely available data services. Watch Metricom's Ricochet modem service, an early entry in this market. Will entrepreneurial broadband service follow? There are several early efforts, for example, Sky Station International, which plans to launch self-propelled balloon-based transcievers over major cities to deliver personal 1.5 Mbit/s service. Meanwhile, we will see how advances in Internet Technology (such as IPv6 and the Internet II initiative of leading universities) evolve - here the ability of the Internet to offer low delay services, such as two-way voice, is the key indicator.

To counter these threats, the telephone companies are now speeding deployment of Intelligent Network services, much like sailing merchants responded to the threat of steam by inventing faster sailing ships in the mid 1800s. The beneficiaries of this accelerated Intelligent Network deployment are big businesses - who can offer cheaper help-desk type services with lower human labor costs. Nevertheless, despite this current Intelligent Network buy-in, if big business finds that it is better served by the Stupid Network and premises based intelligence, it will not hesitate to switch.

The Telecom Act of 1996 and the World Trade Organization telecom agreement of 1997 can be seen as attempts to preserve oligopolistic hegemony of the telephone companies. The thrust of both is to allow big companies to band together to create a marketplace dominated by a few large players in place of government control. Will there be unintended consequences of these agreements? Count on it! Will they hasten or impede the advent of the Stupid Network? Hmmm.

THE STUPID NETWORK'S NEW VALUE PROPOSITION

The shift from scarcity to plenty is often the harbinger of new value propositions. For example, as computer power got cheaper and cheaper in the 1980s, there was much talk of a shift in value from hardware to software, but it was not easy to see how the shift would unfold. In fact, it appears that only one person (Bill Gates) understood it fully. The changes that now portend the Stupid Network are likely to shift the telecommunications value proposition from "network services" to something else. If I knew what it was, I would not be wasting my time writing these words.

Given that disclaimer, I have three brief observations:

1. It is rare that a market is completely killed by the next generation of technology. Neither TV nor the VCR killed the movies. Neither the minicomputer (alas, remember them?) nor the PC killed the Mainframe. We still have ships and railroads, though their markets are both diminished and changed by the car and airplane. The "paperless office" exists - but mine is cluttered with books, memos and magazines that are printed on paper. So it is likely that the Stupid Network and the Intelligent Network will exist side by side for some time, or even share merged definitions, functions, and value. It is also likely that "deliver the bits" companies will exist in a Stupid Network world, but given much lower profit margins, they will not look much like telephone companies.

2. Telephone companies themselves could cannibalize their own product. Smarter companies often field new products that replace current profitable product.

- Sony does this several times a year - it tries to learn from its own mistakes faster than its competition, fielding new products that improve on its old before such improvements become obvious to their foe.
- Boeing does it - the 757 and 767 cut into the top of its 727 market and the bottom of its 747 market

with fuel efficient, and crew efficient new designs - we can only hope that Boeing does not become complacent now that it is has beat out its strongest competitors.

- Intel does it - having been the first to articulate Moore's Law, it now drives it with a new, more powerful chip every 18 months or so, long before the old chip is obsolete - it realizes that if it stops, there are other chip makers that would be glad to take leadership of that market.

Telephone companies could do it too, but it is unlikely as long as their senior managers prefer to talk with lawyers, regulators, consultants and financiers more than with experts in their own employ.

3. Telephone companies could reinvent a place for themselves as purveyors of new values propositions brought by the Stupid Network. They will have to, because their old value proposition will erode as the Stupid Network grows. In a "deliver the bits" world, so much information, and so many courses of action, will be available, that there will be a great need for known, trusted authorities. Businesses with brand reputation and staying power will be guarantors of transactions, holders of critical information, organizers and filters of information, and even voices of reason, leadership, and "objectivity." (Of course, they will need to HAVE reason, leadership, and objectivity to do this.) There will be other roles for big companies in the world of the Stupid Network, and "forgetting organizations," who are able to abandon old models when new ways no longer support old assumptions, will find them.

> "So has the Internet evolved into the Stupid Network and if so, what does this mean for the traditional Telco service provider?"

THE CHOICE BETWEEN LIVING AND DYING

Former Shell Group Planning Head, Arie deGeus, in his master work, "The Living Company" (Harvard, Boston, 1997), examined thousands of companies to try to discover what it takes to adapt to changing conditions. He found that the life expectancy of the average company was only 40 years - this means that telephone company culture is in advanced old age. De Geus also studied 27 companies that had been able to survive over 100 years. He concluded that managing for longevity - to maximize the chances that a company will adapt to changes in the business climate - is very different than managing for profit. For example, in the former, employees are part of a larger, cohesive whole, a work community. In the latter, employees are "resources" to be deployed or downsized as business dictates. As the Stupid Network arrives, as the business idea shifts from scarce physical infrastructure to something more knowledge based, company culture will need to adapt to the truth that, "Nobody knows as much as all of us."

Whatever we discover to be the new Stupid Network value proposition, my working hypothesis is that it will be based on intelligent end user devices, intelligent customers, employees whose intelligence is valued as a corporate asset, and companies that can learn. isen@isen.com - www.isen.com

'AFTERWORDS' FROM JOHN NOLAN

So did Isenberg get it right, as he did predict the rise of companies such as Skype and also foresaw the beginnings of the net neutrality debate? Skype is interesting and, as an example of its growth, Phil Wolff [3] recently reported that 21.5 million people were logged in on Skype's network on 11 January 2009. Figure 1 above illustrates this. The same source also charts Skype's growth and is reproduced in Figure 2 above.

Apologies for concentrating on Skype, but no mean feat for a non-Telco voice service and maybe this demonstrates how intelligence is moving to the edge?

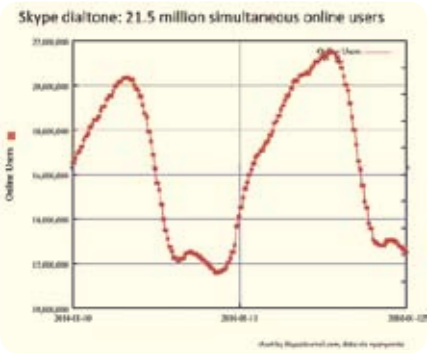

***Figure 1**: Isenberg predicted the rise of companies such as Skype. Was he right? Well, 21.5 million people were logged in on Skype's network on 11 January, 2009. Day snapshot of Skype users (Reproduced with the permission of the Skype Journal)*

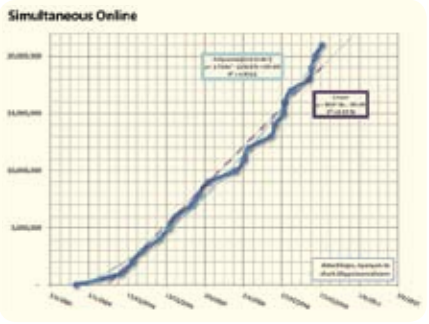

***Figure 2**: The blue line speaks for for itself. Growth of Skype users 2004 to date - £21.5 million. (Reproduced with the permission of the Skype Journal)*

So has the Internet evolved into the Stupid Network and if so, what does this mean for the traditional Telco service provider? Maybe the term Telco is no longer valid, since Service Providers (SPs) such as Sky and TalkTalk would probably consider themselves to be something different and maybe we reserve the term "Telco" for incumbents.

This poses some interesting debates as we have seen all the major incumbents (and the "others") proclaiming their move to Next Generation Networks (NGN) , which as I understand, are all "underpinned" on IP . If we revisit IP [4], its original design purpose was as an interconnecting protocol between the gateways of packet networks. Unfortunately, I am not as close as I once was to the circuit versus packet switching debate, but can I now say that circuit switching i.e. the establishment of an end to end path using central

control (the Intelligent Network) is in decline? Does this also mean that bodies such as the ITU have handed over their standards making responsibilities to the IETF?

So is the Telco dead? I think not as they have, in most cases, risen, phoenix like, from the flames but in a different shape. To illustrate this, here in the UK with DSL and its variants (which I call the "fat pipe" model), a healthy market has grown on the back of the incumbent being "forced" to unbundle its copper access network. But the incumbent is a strong player and maybe the investment in NGN has allowed, not only the incumbent, but others to compete strongly.

But are these next generation SPs (be they incumbents or otherwise) using an intelligent or stupid network to deliver the fat pipe model, as all the end user needs in terms of service delivery is access to the Internet and an end user Ethernet interface? I would argue that this has enabled services such as YouTube and Facebook to gain in popularity. In my experience, such services (could they be called edge) tend not to be created by the Telco or the provider of infrastructure.

Before everyone starts accusing me of being a Nethead, I am a telecoms engineer and the fundamental principles of network engineering and support still apply in the fat pipe model (along with all the other processes defined for the service provider, probably best defined in the Telemanagement Forum's eTOM [5]). The point I am making here is that , if the Internet does replace the traditional "interconnected" Telco model, then some of the engineering disciplines, such as capacity planning, are as relevant today as they were 50 years ago and hence should not be discarded. Maybe this implies even closer collaboration than before – Briscoe [6] has some interesting thoughts here.

Another challenge we also face is future demand on infrastructure and the associated economics given that services such as the iPlayer from the BBC are making increasing demands on the capacity of all SPs. Will we see charging for delivering those stupid bits? McKnight and Bailey [7], although dated, summarise the issues extremely well.

The fat pipe model is made even more interesting with the introduction of fibre in the access network and this leads me to where next?

"The fat pipe model is made even more interesting with the introduction of fibre in the access network"

CONCLUSIONS

Where next? My observations on Isenberg's paper have been kept deliberately concise; primarily to meet the editorial requirements of *the Journal*, but also because I am not sure I know the full answer to the stupid versus intelligent debate. Perhaps we are seeing the blend of the stupid and intelligent model but, in the absence of any defining parameters, it is something we will continue to discuss.

However, I must nail my colours to the mast on a possible future evolution by reference to a Feasibility Study carried out by my company on behalf of the Technology Strategy Board with their Digital Britain initiative, which was entitled "Lightpaths on Demand" and can be best illustrated as shown in Figure 3 below.

In our model, end users select services at the wavelength level and suggest that the intelligence to do so should be available at the end user. Obviously, there are a number of "must haves" for such an approach to work, such as a fibre access network (with long reach WDM-PON?); open access at the end user interface; and services supplied via a photonic core switching network. The obvious question is "will that core network be intelligent, or stupid?" Interestingly, Zukerman [8] describes the concept of Fractional Wavelength Switching and further argues that an evolving Internet could return to circuit switching concepts (albeit at the photonic level).

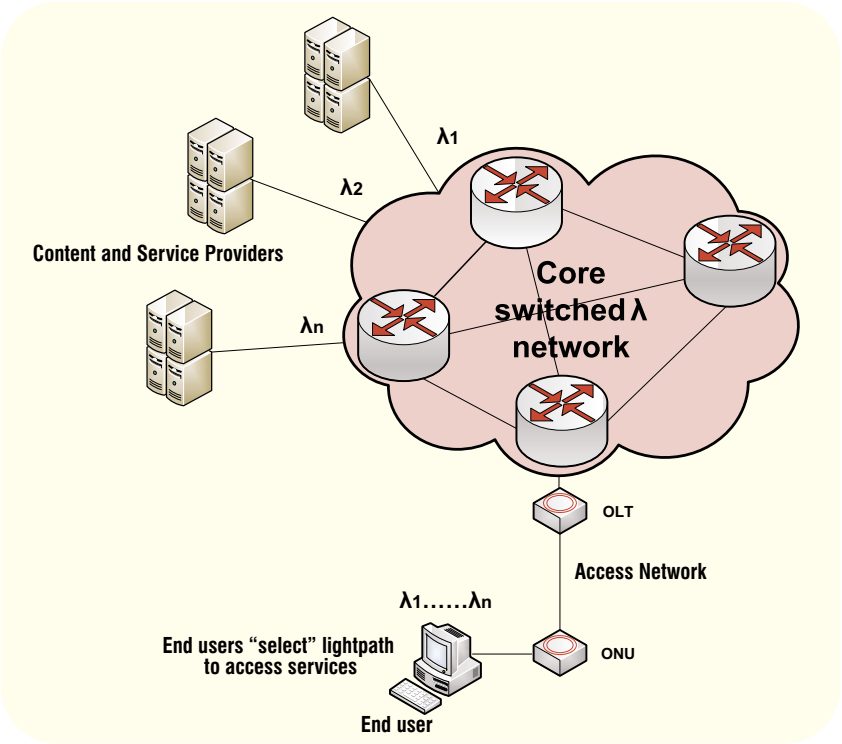

***Figure 3**: Lightpaths on demand*

...AND A CLOSING NOTE

Amazingly, and after almost 40 years since his death, the estate of Jimi Hendrix is to release new material [9] from the original line-up. I will be first in the queue to download (assuming the network doesn't crash on the demand).

ABBREVIATIONS

DSL	Digital Subscriber Loop
IP	Internet Protocol
ITU	International Telecommunications Union
IETF	Internet Engineering Task Force
NGN	Next Generation Network
OLT	Optical Line Termination
ONU	Optical Network Unit
OSI	Open Systems Interconnection
PON	Passive Optical Network
SP	Service Provider
TCP	Transmission Control Protocol
WDM	Wavelength Division Multiplexing

FOOTNOTES

[1] Both Odlyzyko [10] and Cambron [11] offer interesting perspectives on NGN.

[2] On a further historical note, I participated in the protocol wars back in the early 1990s i.e. OSI versus TCP/IP. This was an interesting time for me, particularly in the financial services industry and, needless to say, OSI was swept aside. I was very surprised to see reference to OSI within a recent Ofcom document [12].

ABOUT THE AUTHOR

JOHN NOLAN

John Nolan has been active in the telecoms industry since 1971 and after leaving BT in 1979 he spent extensive periods in senior telecom roles in both Local Government (GLC/ILEA) and subsequently, the Financial Services industry (NatWest Bank Group). In 1995, he joined Intercai Mondiale (www.intercai.co.uk) where he was lucky enough to enjoy many differing cultures during consulting assignments to most of the major Telcos and start ups in Europe, the USA, and the Middle and Far East – rather ironically, not once did he consult to BT. In 2003, John left Intercai Mondiale and co-founded First Mile Networks (www.firstmilenetworks.co.uk) with Dr Chris Lilly, where amongst other projects they competed against Openreach for the Ebbsfleet FTTH project. John has a degree in Electrical and Electronic Engineering from London South Bank University and is proud of the fact that he holds the only true telecoms qualification i.e. a City and Guilds Full Technological Certificate in Telecommunications (Regulation 271). He is a member of the Institution of Engineering and Technology, a member of the Institute of Electrical and Electronics Engineers and is also a member of the Association for Computing Machinery. When he is not watching his beloved Spurs at White Hart Lane, John enjoys the works of F Scott Fitzgerald; E L Doctorow; and William Boyd. (e-mail : john.nolan@firstmilenetworks.co.uk)

References

1. Hendrix, Jimi. Up from the skies. Axis Bold as Love. 1967
2. Uzelac, A. auzelac's blog. Global Crossing. http://blogs.globalcrossing.com/?q=content/8-most-important-publications-internet
3. Wolff, P. Skype Journal. [Online] 11 January 2010. http://skypejournal.com/2010/01/skype-dialtone-215-million-simultaneous.html?utm_source=feedburner&utm_medium=feed&utm_campaign=Feed%3A+SkypeJournal+%28Skype+Journal%29
4. Cerf, V.G., and Kahn, R.E. A protocol for packet network interconnection. IEEE Transactions on Communications. Vol COM-22, No 5. May 1974
5. Business process framework. TM Forum. www.tmforum.org/usinessProcessFramework/1647/home.html
6. Briscoe, R. UCL Home of Bob Briscoe. http://bobbriscoe.net/
7. Bailey, J.P., and McKinight, L.W. Internet economics. The MIT Press, 1997. ISBN 0-262-13336-9.
8. Zukerman, M. Back to the future. IEEE Communications Magazine. Vol 47, No 11. November 2009
9. Hendrix, Jimi. Valleys of Neptune. http://www.jimi-hendrix.com/
10. Odlyzko, A. Next generation network, next generation services, and misleading telecom myths. Digital Technology Center, University of Minnesota. www.dtc.umn.edu/~odlyzko/talks/lightreading-ngns.pdf
11. Cambron, G.K. The next generation network and why we'll never see it. IEEE Communications Magazine. Vol 44, No 10. October 2006
12. Ofcom. Ethernet active line access: Technical requirements (Discussion document). 23 September 2008. www.ofcom.org.uk/telecoms/discussnga/eala/eal/eal_technicalre.pdf