Before the
Federal Communications Commission
Washington, D.C. 20554

AT&T Petition to Launch a Proceeding Concerning the TDM-to-IP Transition	)))))))	GN Docket No. 12-353

COMMENTS OF INTERISLE CONSULTING GROUP

AT&T's Petition is fundamentally flawed. It begins with a false premise, that the use of the Internet Protocol (IP) should determine the regulatory status and business model of a service. Once this premise is rejected, the many conclusions that AT&T bases upon it are also left without a factual basis.

The PSTN is not just "plain old telephone service" (POTS) delivered using analog technology. The PSTN is a legal, social, and business construct that has evolved over the past century and a half, relying over that time on a variety of technologies . It can continue to evolve to incorporate IP technology without any change to its core model. The Internet is something entirely different, a legal, social and business construct based on a very different core model, also relying on a variety of technologies. These differences remain fundamental even if both happen to make use of IP.

ILEC Burdens are not disproportionate to their market power

AT&T asserts that the regulatory requirements placed upon their ILEC operations are somehow impeding their ability to evolve their networks to make use of IP technology. This completely ignores the actual market power that ILECs still have. Only the ILECs have outside plant that reaches the vast majority of urban and rural households and businesses. The ILECs maintain the access tandem switches that remain central to PSTN interconnection, though most of them are now out of date due to underinvestment on the ILECs' part. Only the ILECs provide standalone voice single-line residential and business telephone service in most areas. While cable and other service providers have taken some retail market share, often

only available in bundles, ILECs still have substantial monopoly power over the wireline market as well as the wireless-tower backhaul market.

The core fallacy of AT&T's argument is that somehow the regulatory requirements now incumbent upon them are tightly coupled to their use of TDM technology. This is a deliberately disingenuous argument; the internal technology used to deliver a service to a retail subscriber, or to deliver wholesale interconnection, is always subject to evolution, and no more defines the service than horseflesh defines what it means to be transported from one place to another. Todays switched access tariffs, for example, still use verbiage from the 1980s, when analog voice-frequency trunks were the norm. Those trunks are no longer in widespread use and virtually all access tandem interconnection moved to TDM by the 1990s, and almost all multi-frequency (MF) Feature Group signaling has been replaced by Signaling System 7, but the tariffs and the rules have remained largely unchanged. Likewise, the PSTN can evolve to include some kinds of IP interconnection, without requiring a wholesale change in the legal, business, or regulatory model.

What is POTS anyway? And what is the PSTN?

AT&T's Petition confuses two terms. POTS is a vaguely-defined service concept that is generally understood to include the basic analog telephone line that most households and many businesses still have. One can argue that POTS is thus obsolete, *if* one defines POTS by the use of legacy technology inside the network to provide a legacy analog interface. If the network *implements* POTS using newer technology, but still provides consumers with the familiar line-powered service, is it still POTS? Since there is no legal obligation tied to the term, the debate is purely academic. It appears that the most vociferous supporters of an all-IP telephone network view POTS as the legacy technology rather than the subscriber interface. Thus POTS using 1980s-vintage equipment need not be preserved, just as POTS using 1950s-vintage equipment was not preserved in the transition to TDM.

The PSTN, however, is not just POTS. It is the network itself and all of the services, business arrangements, and rules associated with it. The PSTN evolved technologically from cord switchboards to

electromechanical dial-pulse direct-control switching (Strowger/step-by-step) to electromechanical common control switching (crossbar) to computer-controlled electromechanical (1ESS et al) to computer-controlled time division (current TDM) switching. Each stage of evolution was implemented with backwards compatibility. Retail services evolved separately from the underlying technology, often taking advantage of new capabilities but not always seeing the changes "in the back".

Indeed one major difference between the PSTN and the Internet is that the PSTN maintains a clearer distinction between the line (UNI -- user-network interface, the subscriber demarcation) and the trunk (NNI -- network-network interface, or the point of interconnection between switches or carriers). Thus the two can and do evolve separately. The Internet's current technical model lacks the same distinction. While certain protocols (such as BGP) do not propagate to the retail subscriber, changes in the protocol stack beginning at the IP layer generally do impact the end user. This makes major changes in the Internet harder to implement than in the PSTN![1] While IP may play a growing role in the PSTN's NNI, and may also play a role in some subscribers' UNI, the PSTN still decouples the two. This decoupling should be retained as the PSTN accommodates further use of IP.

Thus AT&T's assertion that maintaining TDM equipment is burdensome is irrelevant. If TDM is to evolve to some kind of IP-based interconnection, then the TDM equipment can be retired. This should not, however, change the fundamental rules of the industry any more than the retirement of cord switchboards did. Thus the concept of a "PSTN sunset" is itself fundamentally flawed. The PSTN is still a vital national service. The TDM legacy switches may well be ready for retirement, and perhaps that can be called the sunset of POTS, but a PSTN should continue to exist even if it uses IP.

ILECs still have monopoly power

The obligations that accrue to AT&T and other ILECs are a result of the huge advantages granted to them as Incumbent Local Exchange Carriers. They control a vast network of conduits, poles, wires, pedestals, manholes, and wire center buildings. These were largely built under conditions of pre-1996 *de jure*

[1] An example of this is the ever-futile effort by many IP vendors to get users to transition to IP version 6. This imposes an enormous cost on end users, with negligible benefits. PSTN transitions have been far smoother.

monopoly and rate of return regulation. They still retain control of these physical facilities, giving them huge competitive advantages. Not only would duplicating all of these be prohibitively expensive for a competitor, but it would be foolish. "Natural monopoly" economics still apply to most outside plant: The incremental cost to the incumbent to serve an additional customer is far lower than the cost for a competitor to do so. If the ILECs were to cede control of their outside plant via functional or structural separation to a wholesale-only utility, then competition would have a chance to provide sufficient market discipline to merit regulatory relief for the ILECs.

Thus most of the actual competition that allegedly faces the ILECs comes from two sources. One is CMRS wireless, which is not a true substitute at all, thought there is some overlap. Indeed the Commission itself noted in its *Phoenix Forbearance Order* that the wireline and wireless markets are fundamentally separate. The other is from cable, which was granted its own monopoly in the past and was for a time protected from ILEC competition by regulations and laws dating back to the 1970s. Thus a duopoly, but not open market entry, exists in most places for some basic residential wireline services, though some customers have only a monopoly supplier.

A small and declining percentage of lines are served via unbundled network elements, a key feature of the Telecommunications Act. But a combination of FCC forbearance, ILEC intransigence and unfavorable rulings over the years have limited their utility to shrinking niche markets. Thus the actual potential value of the monopoly outside plant is limited to what the ILECs themselves choose to do with it, not what competitors could do with it. The AT&T Petition, on the other hand, asks for them to be allowed to discontinue the copper loops and TDM transmission facilities, virtually the only elements still required to be unbundled. This would end such competitive provision of both telephone services and the remaining specialized data services still provided by CLECs using DSL technology.

The net impact of the Petition, then, would be to cement ILEC dominance by removing the market-opening obligations that Congress created in 1996. This has nothing to do with technological transition, which serves as little more than a transparent cover story.

IP itself is not well understood

Internet Protocol itself dates back to 1974, when it was created as part of TCP version 1, and it became "IP" in 1978 when it was split off of TCP version 4 (hence IP itself began life as version 4). Thus it is older than MS-DOS, Signaling System 7, SONET, TDM central office switching, and even practical fiber optics. Thus such a well-seasoned protocol that is assumed to be the future of telecommunications should be better understood than it is. Yet it is not. The Petition, as well as other Commission proceedings discussing transition to IP, reflects either a misunderstanding of how IP works or an attempt by certain parties to sow confusion via inaccurate references to IP.

To begin with, Internet Protocol is not really even a protocol! A protocol generally has defined *elements of procedure* and associated *semantics*, as well as a *syntax* for how it is expressed. IP lacks elements of procedure. In the classical Internet context, it is typically used in a single-queue "best efforts" model of connectionless datagram relaying. But that is not inherent, nor is it the way that voice is properly carried over IP. It is more accurate to describe IP as a multiplexing header. It contains, in lieu of actual addresses, two "address" fields that represent lower-layer points of attachment. (The rest of the multiplexing header is in the next protocol up, UDP or TCP, where the port number identifies the specific connection between the IP endpoints. That these are nominally in different layers is a quirk of the sub-optimal way that IP was separated out from TCP in 1978.) Most elements of procedure are instead left to higher-layer (e.g., TCP) or lower-layer (e.g. MPLS) protocols.

In the classical TCP/IP model used on, for instance, the Internet, IP is used to direct the relaying of packets across some number of routers, with no guarantees that packets will be delivered in-order or even delivered at all. In fact, IP lacks any kind of *flow control* or *congestion management*. Routers silently discard packets when a congestion condition (i.e., full buffer) occurs. TCP detects packet loss as a sign of congestion and reduces its rate of transmission, gradually building it back up until loss occurs again. Packets that are not acknowledged on time are considered lost and are retransmitted by the sender. The resulting sawtooth-like flow pattern is reasonably well tolerated by "data" applications such as web

browsing, email, and file transfer. IP traffic on the public Internet is thus lossy and unpredictable, and that information service is not ideal for telephony. IP is also widely used on private networks; how they are managed and what type of service they provide to their users varies with local management practices.

Voice, on the other hand, is a streaming application. It does not use TCP. It does not use retransmission, nor does it adjust its rate in response to packet loss. It is also sensitive to loss and delay Since TCP is more loss-tolerant than VoIP, networks typically use some kind of traffic prioritization or segregation method to allow VoIP to bypass the often-long, lossy buffers that are used for TCP traffic.

Most IP telephony is not Internet telephony

Only over-the-top voice over IP telephony, which represents a small share of the overall market, mixes VoIP onto the actual best-efforts Internet. Most VoIP, including current carrier-to-carrier interconnection, intra-carrier transmission, enterprise managed services (such as VoIP Centrex), and PacketCable, isolates the traffic *from* both the Internet itself and from other best-efforts traffic. Voice over LTE is not yet in commercial service but will use IP and *will* isolate that traffic from the Internet, albeit differently.

This isolation is typically done at a lower layer. In the case of PacketCable, the necessary prioritization is part of the DOCSIS standard. In the case of enterprise services, it is typically done via MPLS, a connection-oriented shim-layer protocol that creates separate, prioritized streams for different IP flows. Within a corporate network, it may be done via Ethernet prioritization; within a metropolitan Carrier Ethernet network, it may be done via separate VLANs with Committed Information Rate provisioned. None of these VoIP applications are *Internet* voice. They make use of modern IP-enabled equipment to create isolated, managed transport paths for voice.

A further distinction between the Internet and most VoIP is in the way that networks are interconnected to one another. The Internet generally uses IP as the layer at which interconnection occurs. ISPs *peer* with one another, or a small ISP purchases *upstream service* from a larger one, at the IP layer. TCP and higher-layer applications are transparent to the interconnection. Such interconnection is almost always

"best efforts", which is to say the lowest possible quality of service offering. It also requires that the ISPs make use of "public" IP address space – addresses advertised across the Internet backbone using BGP. This makes the peered ISPs' networks function as part of one global internet, and there is usually no isolation of any traffic between the two. Most VoIP, on the other hand, is passed between separately-managed IP networks. These are typically firewalled off from the public Internet, and may make use of private IP address space (which is the norm in both corporate and home networks).

In order to pass the voice traffic without breaching IP-layer security, VoIP traffic is generally relayed between networks (public and private NNI) at the application layer, not at the IP layer. The device that performs this function is a Session Border Controller (SBC). An SBC may also include call (session) route selection based on the dialed number, making it in fact the modern version of a telephone switch, not an IP router. Thus VoIP is exchanged via what are essentially telephone switches, whose interfaces multiplex calls onto a protocol stack that includes IP (as well as including RTP, UDP, SIP, and often MPLS) but which differs substantially from ISP to ISP peering.

The SBC is central to VoIP interconnection, and must play a key role in the transition. Yet it is not mentioned even once in the Connect America Fund Order and Further Notice (FCC 11-161). Nor does AT&T mention it in their Petition. Instead they seek to keep IP-based PSTN interconnection "free of legacy regulation", as if it were the Internet itself exchanging packets at IP. This is simply a disingenuous way of pretending that IP-enabled PSTN interconnection is something that it is not.

Because of this distinction between the Internet and non-Internet managed IP networks, it makes no technical sense to use the Internet as a model for telephone network interconnection. Nor does it make business or regulatory sense, as the Internet is a different entity that has achieved its own success based upon a different business and regulatory model. Of course AT&T and other price cap LECs would prefer to be regulated (or more accurately *un*regulated) like ISPs! But when they are carrying telephone calls, serving telephone subscribers, operating the outside plant network, and interconnecting with other PSTN carriers, they aren't acting as ISPs. They are still the telephone company. There is simply no justification

for major rules changes based on a gradual transition to IP-multiplexed voice. While AT&T notes that "IP-enabled services" are generally treated as information services, *these* specific telephone services, which just happen to also use IP, should not be—they are Title II PSTN telecommunications, pure and simple.

By the same token, the Internet is not the PSTN. Even though a little bit of over-the-top voice is carried across it, the Internet has never promised to be a reliable transmission medium for streaming services. Its success largely stems from its independence from traditional telecom regulation, its status as an enhanced service under the Computer II rules that were in effect until 2006 (and which should be restored) and as an information service under Telecom Act definitions. Attempting to treat the Internet as the backbone of the PSTN would almost certainly place additional burdens on the Internet, reducing its flexibility. The PSTN has responsibilities, such as reliable 911 calling that are simply not well suited to the Internet's operational model. AT&T's petition, then, by confusing all VoIP with the Internet, does nothing to aid Internet freedom, but in fact gravely threatens it.

Business as usual

AT&T's Petition begins with an advertisement for its planned network upgrade, for which it plans to spend $14 billion. Of this, $8 billion is for its wireless operations, a relatively competitive business in which major expenditures are being made to roll out LTE nationwide. The smaller wireline portion is to expand U-verse, the late-life kicker for its copper-based network. When viewed in the context of the company's size and historical levels of expenditure, these numbers aren't huge. They are more like business as usual. They only seem large because capital expenditure on the wireline side has been so low in recent years. Where but in the wireline network is it common to see computers still in use that are two decades old?

Indeed the entire trajectory of ILEC industry investment has been downhill since price caps replaced rate of return regulation. Under the old regime, capital expenditures were encouraged, as a larger rate base meant more allowable profits. It was the regulator's job to decide how much was enough. Under

alternative regulation, price caps, and post-cap full deregulation, the incentive has been to minimize expenditures, both operating and capital, in order to maximize quarterly profits. Hence most investment has been directed at wireless operations, whose margins are higher.

But the Telecom Act's universal service obligation interacts with price cap regulation in an unfortunate way, which would be greatly exacerbated if AT&T's wish to be relieved of its Carrier of Last Resort obligations were granted. The Act essentially divides the country into three zones. One zone comprises the urban and suburban areas with low cost of service. These are profitable, so the ILECs continue to want to serve them. Another zone comprises the rural areas served by the rate-of-return carriers that receive explicit USF money. As these carriers are still on rate of return, they have every incentive to invest, and indeed they do. As the Commission has noted, and began to address in its Connect America Fund Order, there has been no brake on such expenditures. State regulators no longer care, because the federal USF makes up the difference. Hence subscribers in these areas get good service at low prices, but the rest of us pay.

The third zone is the most problematic one here—high-cost (mostly rural) areas served by price cap carriers such as AT&T. The Telecom Act structure simply assumed that the carriers would use internal cross-subsidies to maintain quality service to these areas, while Carrier of Last Resort (COLR) obligations and service quality plans ensured that they did. But these carriers have no incentive to spend above-average amounts to serve anyone that they don't have to serve. If they were relived of COLR and abandoned their own high-cost areas, they could average down their costs and thus raise their profits. This is not what the Telecom Act had in mind, but it is what the petition calls for. Some states have already granted this relief; the Commission should not be pressured into doing so on a federal level.

Hence the proposed expenditures are no reason to strip away over a century's worth of PSTN regulation simply because AT&T suggests, but does not actually propose, that a connection exists between the extreme deregulation that it is calling for and its capital plan. Rather, the capital plan for wireless is based on prudent business judgment, and the capital plan for U-verse still represents a lower level of

expenditure than was promised 20 years ago as part of the various state AFOR plans. The connection is thus not real, and the bright shiny object of proposed spending should not blind the Commission's eyes.

For these and other reasons, the AT&T Petition is without merit and should be summarily rejected. The Commission should continue to investigate modernization of the PSTN, including increased use of VoIP technology. But that should not relieve carriers of their PSTN obligations.

Respectfully Submitted by the Interisle Consulting Group LLC

Fred Goldstein, principal author