per minute cost of longer calls. As explained in the Montgomery Declaration (at 8), there are many types of calls other than ISP-bound calls that are of longer than average duration. Whenever rates are based on costs for a call of average duration, there will be by definition calls of longer than average duration as well as calls of shorter than average duration. One appropriate solution for Verizon's concern might be to apply a two-part reciprocal compensation rate that separates call set-up costs from duration costs. *See* Montgomery Declaration at 8. Sprint proposes such a solution in its comments.

As to dedicated capacity, Verizon assumes that CLECs incur no centum call second ("CCS") costs. Verizon Comments at 26, Taylor Declaration at 15-17. Verizon confuses dedicated, non-switched facilities costs, like private lines, with the high volume ISP-bound traffic which is routed through switches and consumes switch resources at the busy period. Montgomery Declaration at 10. Also, Verizon and SBC overlook the higher costs for additional equipment including the switching fabric and additional line concentration modules. *Id.* at 13.

With regard to call direction, Verizon argues that because ISP-bound calls are all one-way, switching costs for features and functions on the originating end will not apply. Taylor Declaration at 17. Yet, as shown in the Montgomery Declaration (at 14), cost studies have never looked at "one-way" call switching costs separately from other local calls.

Finally, Verizon claims that peak traffic loads for ISP-bound calls are different than for other types of calls and that this results in different – and lower – average costs for CLECs serving ISPs. Taylor Declaration at 17. As shown in the Montgomery Declaration at 11, Verizon's "assertion is wrong on its face because it is not the percentage (or fraction) of busy hour traffic that drives switching costs but rather the absolute load measured in

CCS." The Montgomery Declaration goes on to highlight a Pacific Bell study that "[a]s Internet traffic continues to grow on the PTSN, this traffic will have an every increasing effect on peak-usage with corresponding cost increase." In short, the evidence is that network congestion is most likely to occur at the CLECs' switches that terminate calls to ISPs.

In short, the ILECs have not convincingly demonstrated any meaningful cost differences between ISP-bound calls and other local calls.

B. The ILECs' Policy Arguments Are Misguided

Rather than address the statutory test that *Bell Atlantic* announced as dispositive, the ILECs advance a number of policy arguments in support of their position that the Commission should not require reciprocal compensation. According to the ILECs, reciprocal compensation for ISP-bound traffic (1) leads to a distortion of the market, (2) discourages CLECs from providing residential service and advanced service, and (3) encourages "sham" CLECs whose only reason for existing is to take advantage of the reciprocal compensation "arbitrage opportunity." The ILECs are misguided in raising these arguments. As discussed above, there is no room for policy-making under *Bell Atlantic*. The Commission should thus disregard all of the policy arguments advanced by the ILECs. Even if the Commission were to consider them, however, they are all unavailing.

1. Reciprocal Compensation for ISP-bound Calls Does Not Distort the Market

SBC contends that reciprocal compensation is an "arbitrage opportunity" that distorts the market by shifting investment "towards that opportunity and away from others." SBC Comments at 42. SBC is simply wrong.

ICG serves ISPs to the extent that it does because they represent a technologically-advanced, fast-growing market segment with specialized needs that ICG is much better suited to address than are the ILECs. For their part, ISPs choose CLECs because CLECs offer packages of services tailored to meet ISP needs at competitive prices. The ILECs dispute this, suggesting that the real reason that ISPs choose CLECs is that CLECs pass on the benefits of the "arbitrage opportunities" created by reciprocal compensation. *See, e.g.*, SBC Comments at 44-45. ICG, however is on record that it does *not* pass through any reciprocal compensation revenue to its ISP customers.

In this regard, SBC's citation to the record in ICG's arbitration with Ameritech in Ohio is incredibly galling. SBC claims that there "ICG revealed that it offered service for free to ISPs that agreed to collocate in its switching centers." SBC Comments at 45. This is simply false—and SBC knows it. As ICG witnesses testified during the Ohio arbitration, the tariff that SBC cites in support of its claim was an obsolete tariff that had been replaced because no carriers used the service and which did not, as SBC claims, provide for free service to ISPs as a pass-through of reciprocal compensation revenue.

ISPs are also a natural market for CLECs because they are an ideal bridge to broader market entry. ISPs are high-volume users that create a significant revenue stream. That revenue stream allows CLECs to defray their capital expenditures as soon as possible, and to fund expansion into the traditional residential and business markets.

In other words, competition is functioning precisely the way it should: new entrants are targeting niche customers to gain market share and develop a revenue base that they can then leverage to expand their service into the wider business and residential markets. This is exactly how competition in the long distance and customer premises

equipment markets began. There the claim was that new entrants were "cream-skimming;" today's refrain is that new entrants are on a "gravy train." But the Commission should see these complaints for what they are: once-entrenched monopolists fighting to stave off competition.

2. ICG Is an Active Provider of Advanced Services and Residential Service

The ILECs contend that reciprocal compensation for ISP-bound results in CLECs having less incentive to serve the residential and advanced services markets. SBC Comments at 40-42; Verizon Comments at 12-15. This is simply not the case.

ICG actively provides residential service. While ICG may serve relatively few residential customers, it is not because ICG has focused on serving ISPs in order to collect reciprocal compensation revenue. Rather, one of the chief reasons that ICG has not been able to serve more residential customers is the ILECs' poor track record in opening their market to competition. As a recent Denver Post article stated "providing telephone service against big monopolies such as Bell Atlantic and US West . . . never reached [its] potential, as the monopolies dragged their feet in opening local markets and profit margins for local and long distance calling grew slimmer."[20] If the ILECs had not done everything in their power to make competitive entry into its markets so difficult, ICG and other new entrants would be even further along in serving traditional business and residential customers.

As for advanced services, ICG is increasingly becoming known as a major data services company. ICG has partnered with various DSL companies to provide high-capacity Internet access, which is the quintessential advanced service. ICG is not unique in

---

[20] Andrew Backover, *Massive Makeover: ICG's new team sees a data-services future*, Denver Post, May 15, 2000, at 1E.

this regard. Dozens of CLECs are rushing to fill the nationwide need for advanced services.[21] Those companies had the same reciprocal compensation "arbitration opportunity" available to them as any other CLEC, yet chose not to pursue it. Were reciprocal compensation as irresistible as the ILECs make it out to be, no CLEC would rationally pursue any other strategy. This is, of course, not the case. ICG and other CLECs are well aware that broadband services are a market of the future and that it would be short-sighted to ignore that market.

3. Reciprocal Compensation for ISP-Bound Calls Does Not Lead to "Sham" CLECs

As for the ILECs' argument that reciprocal compensation leads to the creation of CLECs that specialize in ISP traffic and provide little or no genuine local exchange service, *see, e.g.*, SBC Comments at 42-44, several points are in order. First, as discussed above, ICG serves ISPs because they are a fast-growing market niche with specialized needs that ICG as a technologically-advanced new entrant is ideally suited to serve, *not* because reciprocal compensation represents an "arbitrage opportunity."

Second, ICG has over 250,000 non-ISP local exchange lines in service and expects that number to grow to 400,000 by the end of 2000. It is thus simply not the case that ICG specializes in ISP traffic to the exclusion of other local exchange customers.

Finally, the ILECs' reliance on *In the Matter of BellSouth Telecommunications, Inc. v. US LEC of North Carolina, Order Denying Reciprocal Compensation*, Docket No. P-561, Sub 10 (N.C.U.C. March 31, 2000) ("*US LEC*") is completely misleading and is

---

[21] Accordingly to the trade press, the Commission has just adopted a report finding that deployment of advanced services "is reasonable and timely at this time." *TR Daily*, August 3, 2000.

offensive. SBC and Verizon both cite it as proof that reciprocal compensation leads to the creation of 'sham' LECs whose sole purpose is to deliver traffic to ISPs and generate reciprocal compensation payments. SBC Comments at 43; Verizon Comments at 17-18. What SBC and Verizon fail to mention, however, is that there was no Internet traffic at issue in the proceeding; nor was there a terminating end user. Instead, the defendant carrier was simply making calls to a router that did not go past the router. This pure fraud is a far cry from ICG choosing to address the ISP market niche. The ILECs also opt not to mention that the same North Carolina Commission that decided *US LEC* also held in ICG's arbitration that *ISP-bound traffic is properly subject to reciprocal compensation and awarded reciprocal compensation to ICG. See Petition by ICG Telecom Group, Inc., for Arbitration of Interconnection Agreement with BellSouth Telecommunications, Inc.*, Docket No. P-582 Sub 6 (N.C.U.C. Nov. 4, 1999).

As for Verizon's other examples of the "bad behavior" that reciprocal compensation leads to, all they prove is that some companies engage in questionable practices, including fraud, to make money. Such behavior is nothing new in the telecommunications industry—or any industry for that matter—and does not prove that the underlying activity is "socially wasteful" as Verizon suggests. Verizon Comments at 16. Saying that CLECs should be denied recovery for the costs they incur in terminating calls to ISPs because a handful of carriers may have engaged in fraud is like saying that there should be no competition for long distance because a handful of carriers have engaged in slamming. To the extent that the Commission believes fraud is a problem it should combat that problem as it has slamming—through vigorous enforcement actions, *not* by denying the vast majority of law-abiding CLECs the opportunity to recover their costs.

C. CLECs Are Limited in Their Ability To Raise Their ISP Customers' Prices

Notwithstanding that it is the ILECs that have the most proximate relationship with the cost-causing caller, the ILECs insist that the ISP should be responsible for reimbursing a CLEC for the CLEC's costs in delivering an ISP-bound call. SBC Comments at 49; Verizon Comments at 22-23. For the following reasons, that is not a practical suggestion.

The FCC's ESP exemption allows ISPs to take service from the LEC's local business service tariffs. *Amendments of Part 69 of the Commission's Rules Relating to Enhanced Service Providers, Order*, 3 FCC Rcd 2631, 2635 n.8, 2637 n.53. *See Declaratory Ruling*, ¶ 5. The ESP exemption limits CLECs' ability to look to their ISP customers for cost recovery for two reasons. First, if a CLEC sought to raise its rates to ISPs above its standard business rates, the ESPs would assert their right to buy service out of the CLEC's business tariffs.

Second, the ESP exemption effectively caps the rates that CLECs can charge ISPs. Since CLECs must compete with ILECs to win ISP customers, the prices that ILECs charge their ISP customers function as a price ceiling for CLEC. If a CLEC were to attempt to significantly raise ISP rates to recover the costs the CLEC incurs when delivering traffic from ILEC customers, the CLEC would risk losing its ISP customers to the ILECs.[22] Even though the CLEC offers ISP customers packages of services unmatched by the ILECs, it is a reality of the marketplace that CLECs must compete first and foremost on price. Furthermore, since ISPs assert the right to buy out of a CLEC's local exchange

[22] Since ISPs can assert the right to buy out of a CLEC's local exchange business tariff, in order to raise its rates for ISPs, a CLEC would have to raise its rates across the board for *all* business customers. Given ILECs' dominant market power, doing so is not practical.

business tariff, in order to raise its rates for ISPs, the CLEC would have to raise its rates across the board for *all* business customers. Given ILECs' dominant market power, doing so is not a practical business strategy.

IV. CONCLUSION

ISP-bound calls are "telephone exchange service" and therefore subject to reciprocal compensation under Section 251(b)(5) and the Commission's rules. Even if the Commission were to find that ISP-bound calls are not subject to reciprocal compensation, the Commission should prescribe the same inter-carrier compensation arrangements for ISP-bound calls that apply to other types of local calls. The Commission should do so because the characteristics and costs of handling ISP-bound calls are very similar to the characteristics and costs of handling other local calls. To establish different arrangements for ISP-bound and other local calls would result in market distortion that would impair service to ISPs and harm the growth of the Internet, all to the detriment of the consuming public.

Dated: August 4, 2000

Respectfully submitted,

By: /s/ Allan C. Hubbard
Albert H. Kramer, Esq.
Allan C. Hubbard, Esq.
Jacob S. Farber, Esq.
Robert Felgar, Esq.

DICKSTEIN SHAPIRO MORIN
& OSHINSKY LLP
2101 L Street, N.W.
Washington, D.C. 20037-1526
(202) 785-9700
Attorneys for
ICG TELECOM GROUP, INC.

Before the
FEDERAL COMMUNICATIONS COMMISSION
Washington, D.C. 20554

In the Matter of	)	
Implementation of the Local Telecommunications Act of 1996	)	CC Docket No. 96-98
Intercarrier Compensation for ISP-Bound Traffic	)	CC Docket No. 99-68

DECLARATION OF DR. ROBERT A. MERCER

My name is Robert A. Mercer. I am a principal of BroadView Telecommunications, LLC. I provide strategic planning and education related to public and private telecommunications infrastructures, with a particular emphasis on local exchange competition, broadband integrated networks, intelligent networks, and private enterprise networking. Examples of my current work include the analysis of competitive alternatives for the provision of local exchange services, and advising a national association on alternative methods of obtaining and managing its telecommunications services. My previous positions have included President, Hatfield Associates, Inc.; Department Head, AT&T Laboratories; Assistant Vice President, Bellcore; and Director of Network Architecture Planning, Bell Laboratories. I have a Ph.D. in physics from John Hopkins University and a B.S. in physics from Carnegie-Mellon University. A complete statement of my qualifications is attached hereto.

I. Introduction and Summary: What Is a Dial-Up Call to an ISP?

This Declaration discusses the issue of whether a dial-up call to an Internet Service Provider ("ISP") is analogous to exchange access or to local telecommunications traffic.[23] Exchange access is defined by the Act as the offering of access to telephone exchange services or facilities for the origination or termination of telephone toll services." 47 U.S.C. §153(16). Under Section 51.704(b) of the Commission's rules, local telecommunications traffic is interconnected traffic that "originates and terminates within the same local service area established by the state commission." 47 C.F.R. § 51.704 (b). Termination in this context is defined as "the switching of local telecommunications traffic at the terminating carrier's end office switch, or equivalent facility, and delivery of such traffic to the called party's premises." 47 C.F.R. § 51.701 (d). From the technical perspective of this Declaration, the distinction is whether the call terminates when it reaches the ISP premises or continues across the Internet.[24]

This Declaration concludes that in all key respects, a dial-up call to an ISP represents local telecommunications traffic, notwithstanding the fact that the computer connected to the ISP by the call subsequently triggers the flow of information across the Internet. The subsequent Internet events represent a fundamentally different kind of enhanced process-to-process communications related to the information services nature of the Internet. The flow of information that takes place across the Internet during these events cannot reasonably be considered to be part of the same dial-up call – indeed, there is no such thing as a connection or call across the Internet in the first place.

---

[23] Hereafter, I will refer to this as a "local exchange call."

[24] In its comments in this proceeding, ICG Telecom Group, Inc., has dealt with the issue of whether an ISP even offers telephone toll service.

The situation is analogous to what happens when a company's employees dial a local call into the company's computer network from remote locations. Such calls are connected to modems in a pool of modems provided by the company. Over the dial-up connection and through the modems, the employees' computers have access to the company's computing and database resources. Using these resources, the employees may exchange email, store and retrieve data, carry out computing tasks beyond the capabilities of their own computers, and so on. Some traffic on these connections may even be directed to the company's Internet links, and subsequently traverse the Internet. Whatever the nature of the subsequent interaction, there is no sense that the dial-up calls are thereby extended across the computer network. The calls terminate at the modem pool; all of the subsequent information flow is considered to be part of the computer network's traffic, not "calls" in the sense one normally thinks of that word.

Returning to the issue of dial-up calls to an ISP, the flow of information across the Internet subsequent to the establishment of the dial-up connection to the ISP differs from a long distance call over an IXC network in all of the following respects:

The Internet operates in a fundamentally different fashion than does a long distance voice network over which the dial-up connection to the ISP is established, because voice calls transpire over a circuit switched connection[25] which is established by the calling party, maintained for a finite period, then terminated by the calling or called party, whereas there is no notion of a connection over the Internet, and thus no connection to be created or terminated;

---

[25] Circuit switching refers to a form of network protocol in which the network commits switching and transmission resources, collectively called a "circuit," solely to a single instance of communications – two humans talking, or two computers exchanging data – and maintains that commitment until the communicating entities indicate the communication is completed.

Unlike ordinary telephone calls, the information supplied by the user to initiate a dialog with another computer is not a network address; it is an Internet "Domain Name," which must be interpreted and translated into an Internet address before the routers that are responsible for Internet transport can start routing information to the right destination;

Finally, and perhaps most importantly, there is no voice network equivalent of the Internet process in which multiple application-to-application sessions involving multiple sites take place sequentially over a single dial-up connection without subsequent involvement of the user.

In short, any attempt to claim that a dial-up call to an ISP, and the subsequent instances of multiple communications across the Internet using that connection, are in any way analogous to exchange access is strained and artificial.

The rest of this Declaration is organized as follows. First, I will provide a brief tutorial on the nature of the Internet and its protocol suite in order to differentiate between basic transport over the Internet and the application-to-application communications that take place using the Internet. In the tutorial, I will also describe how different is the process of conveying information across the Internet compared to a normal voice call. Next, I will describe the transactions that occur during two common Internet events – the sending of email and access to a website to obtain "pages" of information – and show that these transactions involve computer processing far different than the underlying network transport service. Based on these examples, I will expand on the above summary points in order to show there is a marked technical distinction between the dial-up call to the ISP and what happens subsequently across the Internet. In fact, the Internet situation is much more akin to a remote user making a local exchange call into his/her company's computer network and thereby accessing the company's computer and database

resources than a call through the local exchange network and across an IXC network. I will therefore argue it is only meaningful to consider an Internet access call to be a local exchange call terminated on the ISP. Finally, I will critique the contrary claims laid forth by Charles L. Jackson on behalf of Verizon, demonstrating that his conclusions are based on a mischaracterization of Internet communications, and are therefore erroneous.

II. Brief Tutorial on the Physical and Protocol Structure of the Internet

A high-level view of the Internet is provided in Figure 1a. It consists of a set of interconnected physical networks, shown as clouds in the figure, that are "glued" together by routers between the networks to create an entity that to users of the Internet appears to be one logical network.[26] operates like one large network of worldwide scope. Routers are a key element; they are responsible for examining the destination Internet address contained in the control header part of each packet, or "datagram," that is being transmitted across the Internet and, based on that address, determining the next router to which the packet should be sent in order to get it closer to its ultimate destination.

Figure 1b provides a view of the Internet as it exists in the U.S. today. It shows the role of local and regional Internet Service Providers ("ISPs") in providing links between the end-users and a set of interconnected Internet backbone networks.[27] End users, shown at the bottom of the figure, access the Internet in one of two ways: they are attached to local area networks, which in turn are linked to an ISP through a router-to-router link, or they dial into dial-up Internet hosts belonging to the ISP. A dial-up host is not simply a router;

---

[26] E.g., while computers are attached to specific networks and have addresses on those networks that are only recognized by other computers on the same network, the Internet makes provision for a universal addressing scheme by which a computer is identified to all computers on all networks that are part of the Internet.

[27] There is often a considerable overlap between local, regional, and backbone portions of the Internet.

it is a server that provides protocol conversion, error checking, authentication, and negotiation functions such as the assignment of a temporary IP address to the user's computer for the duration of the dial-up connection. The key conclusion to reach from this figure is that resources of the Internet are available to Internet users through their respective connections – dial-up or dedicated router-to-router links --to the ISPs that serve them.

Figure 2 shows the four-level protocol "stack" employed in the Internet. At the lowest level of the stack, Internet devices—hosts,[28] servers, routers – are attached to a physical network such as a LAN, dedicated circuit, or ATM network, and must communicate with other devices on that network using the protocol(s) of that network. The network protocol (or protocols) utilized is called the network interface protocol.[29] Next up the stack is the Internet Protocol ("IP") layer, which handles the routing function whose purpose is to deliver datagrams to the correct destination IP address. A key aspect of IP that is profoundly different than the voice network is that it is "connectionless." That means each datagram is transmitted across the Internet in a best-effort, but not guaranteed, fashion. Individual datagrams can be lost, received out of order, or corrupted. IP itself will not correct for such failures, although it does contain mechanisms to inform the communicating hosts that failures may have occurred. The advantage of this approach is that it is simple and fast; the disadvantage is that the protocols higher in the stack must address such potential deficiencies or at least be robust in the face of their occurrence.

---

[28] Traditionally, the computers that communicate over the Internet, including everything from PCs to mainframe computers, have been referred to as hosts. Computers serving various specialized functions such as name translation, mail processing and web hosting are often distinguished by referring to them as servers. Routers are not the beginning- or end-point of user traffic; their role is to route user traffic from its source to its destination.

[29] In some texts, this is called the host-to-network or link layer protocol.

The remaining two protocol layers run in the hosts and servers attached to the Internet. The Transmission Control protocol ("TCP") provides reliable end-to-end flow of information from the sending host to the recipient host. For example, if datagrams are lost, so parts of the data being transmitted between the computers are lost, TCP is responsible for detecting the loss and ensuring the missing data is retransmitted. Similarly, because different datagrams may follow different routes from source to destination, data can be received out of order, and TCP is responsible for reordering the data so the communicating applications receive data in the same order it was transmitted. Because TCP involves a considerable amount of overhead processing and transmission, and because not all applications require the degree of reliability guaranteed by TCP, there is an alternative protocol at the same layer of the stack, the User Datagram Protocol ("UDP"), that provides less reliability than does TCP at a considerable savings in complexity. As Figure 2 suggests, the IP layer of the Internet, and therefore routers, are essentially oblivious to the host-host dialog taking place at the higher levels.[30]

Various application programs comprise the highest layer of the Internet protocol stack. These applications communicate email and other messages, provide for the transfer of files between hosts, and allow a user at one host to log onto another host as a remote user. Since the early 1990's a key application has been the World Wide Web ("WWW"), which provides standard methods for accessing linked documents, called pages, spread across machines distributed around the Internet.

---

[30] In recent years, the division between the IP and higher layers has become blurred to an extent. For instance, the IP layer may in some cases examine the content of the higher layer protocol information in order to detect the type of applications that are communicating and assign a higher priority to datagrams serving particular applications. Also, routers are sometimes themselves the origination or termination point of Internet communications, such as when they receive instructions from, or provide monitoring data to, a network management workstation.

The key conclusion to reach from Figure 2 is that the Internet protocol suite, which is often called the TCP/IP suite in honor of what are arguably its two most important members, consists of two distinct parts. One part, comprising the lower two layers, deals with what one normally thinks of as data transport functions across a network. The other part, comprising the higher two layers, has nothing to do with transporting data across the network per se, but with distributed processing functions involving two or more hosts connected to the Internet. While the network transport functions may be somewhat analogous to the processes that take place in sending information across the voice network, the distributed processing functions have no analogy in the voice network.

III. Three Examples of Internet Applications

A. Electronic Mail

Suppose an Internet user wants to send an electronic mail message to two friends, Mary Jones and Don Williams, whose mail addressees are, respectively, mjones@ISPxxx and don.Williams@ISPyyy. The user instructs his/her computer to dial the ISP; the computer is linked to the ISP dial-up host via the dial-up connection.[31] The user then opens an email program, and as prompted by the email program, enters the addressees, types the message, and identifies the file name(s) of any attachments to be appended to the message. When this information is entered, the user clicks on the SEND button on the email application screen. At that point, the email client application on the user's PC carries out a dialog defined by the Simple Mail Transfer Protocol ("SMTP") protocol with the ISP's SMTP server to convey the message's recipient, body, and attachment information to that server.

[31] This describes the case of dial-up connections to the ISP. For users on a LAN with a dedicated connection to the ISP, the connection is already in place, and the user can immediately start using the Internet.

When that dialog is complete, the user can terminate the dial-up connection unless there are messages to receive or additional messages to send.

At this point, The SMTP server knows only the names of the hosts (the characters to the right of the ampersand in the mail address – ISPxxx and ISPyyy in the above examples) that provide mail service to the recipients, not the Internet address of those hosts. It thus invokes another Internet process, the name resolution process described below, to translate the host name into an Internet address. The user does not initiate this process, is not involved in it, and is oblivious to it – in fact, by time it takes place, the user may have terminated the dial-up connection to the ISP! The process involves strictly application-to-application dialogs, and the application in this case – the name translation process – is not even the one the user originally invoked.

When the recipient host names are successfully translated ("resolved") into Internet addresses, the SMTP server enters into an SMTP dialog with the recipient mail servers over TCP sessions established between the SMTP server and the recipient mail servers. The original user's call to the ISP is thus responsible for two different sessions with two different mail servers. Again, the user did not specify the network address of the two recipient mail servers, plays no role in establishing the TCP sessions, does not participate in the structured SMTP dialog itself, and in fact is not even aware the dialog has taken place. Note, too, that if either or both recipients were on the same mail server that serves the sender, there would be no need for any information to flow across the Internet.

At the receiving end, some time later, Mary Jones (and, separately, Don Williams) establishes a connection to her ISP, opens her email program, and invokes the "check mail" function. The email software has been configured to know the so-called Post-Office Protocol ("POP") server that is holding Mary's mail, and will initiate a dialog with that

server using the POP protocol to retrieve mail messages and store them on Mary's machine. At that point, they are available for her to view in a fashion defined by the email software. Note that Mary's POP server may be at the ISP premises she dials into, so she can receive her email with no further information flow across the Internet.

At no time in the process described here is there ever an end-to-end call, or simultaneous communication of any sort, between the sender and either recipient. When all is said and done, there is no simultaneity between the dial-up call to the originating ISP, the transmittal of the message from the originating SMTP server across the Internet to the recipient mail server, or the dial-up call initiated by the recipients to receive their message.

B. World Wide Web

In the WWW application, a user opens a WWW client, or "browser." Either manually, or from a list of "bookmarked" or "favorite" sites, the user specifies the Uniform Resource Locator ("URL") of a site to be accessed. The URL contains both the name of the WWW server, the file name of the document, or "page" to be acquired, and the file access protocol to be utilized (e.g., http). The URL does not contain the Internet address of the WWW server; again, this means that the browser must invoke the name resolution process to learn the address, just as in the e-mail example, and in other Internet applications.[32]

The browser creates a TCP link to the WWW server host designated in the URL. Over that link, the WWW server transmits the page identified in the request. A page may contain text, icons, line drawings, maps, photographs, or even audio tracks and video clips.

[32] A user with sufficient knowledge of the Internet may know the Internet address of a target host machine, and use that address rather than a name, thereby obviating the need for the name resolution process.

It may even contain forms that request the user to enter information; this is prevalent in e-commerce applications of the WWW.

Whatever the contents of a page, the browser interprets the contents, including formatting instructions, and displays the properly-formatted page on the screen using a graphical user interface. The display process itself is often innovative – for instance, an image that takes a long time to download and process may first be displayed with coarse resolution, then successively finer resolution until the final image appears. Or, text may be transferred and displayed first so the reader has something to look at while the more content-rich parts of the page are transferred and displayed.

Various elements of a page may be linked to other pages residing at the same or different servers. Such links are referred to as "hyperlinks." The user can click on these hyperlinks, which typically appear as underlined text (or as instructions that may, for instance, tell the user to click on a map or other image to obtain more detail). Often, an element in a page may be only a hyperlink to another page where the information is contained,[33] and cause the browser to initiate new TCP sessions to obtain and display the linked pages.

No description can adequately capture the richness, dynamism, and flexibility of the WWW application – it has to be experienced to be appreciated. Returning to the subject of the Declaration, exercising the WWW application typically causes many different Internet information flows to occur. While some of these may be under the control of the user – for instance, when a user may optionally click on an image to magnify it or obtain more

[33] This is common, for instance, when one website – say, the website for a news program – displays advertising for another company. The advertising part of the page may contain a hyperlink to the advertisers site where the contents of the advertisement resides, to avoid cluttering the main site with too much data.

information – many occur without any action, involvement, or awareness on the part of the user.

C. The Name Resolution Process

Users typically specify target hosts – WWW servers, email servers, and the like – using their Internet name of the hosts, not their Internet addresses. These names are assigned under the auspices of the Internet Name Domain System ("DNS"), which specifies a mechanism for assigning globally unique names. Because information is routed across the Internet using Internet addresses, not names, there must be a process of translating, or "resolving," names into addresses. The DNS also describes this mechanism, which involves interactions between a hierarchy of name domain servers, each of which has translation responsibilities for some part of the internet name space. At the upper end of this hierarchy, there are "root" name domain servers for each of the major divisions of Internet names – ".edu", ".com", ".net", and so on. At the other end of the hierarchy may be a name domain server responsible for a single company or university, or even individual departments and divisions within a university or corporation.

When an application must translate a name into an address, it invokes a procedure called a "resolver," or DNS client, passing it the name to be translated. The resolver knows the address of a DNS server, to which it sends a request to translate the name. The server either knows the translation information, for instance because it has cached translation information for that name during an earlier transaction, or it sends the request to another server. The process repeats itself in what can be multiple steps that generally proceed up the DNS hierarchy to the root server for the major branch in question, then back down the hierarchy to a server that can ultimately make the translation. DNS servers generally cache

the information they learn in order to respond more quickly to future requests, although cached information is purged frequently since it can quickly become outdated.

The user that invokes the original Internet application is aware of nor involved in the translation process that takes place in the fashion described. Thus, this Internet process and associated information flow takes place entirely over the Internet, and at many steps, does not correspond to any flow of data through the dial-up connection.[34]

IV. Differences Between Calls to ISPs and Long Distance Calls

The previous applications examples demonstrate the following points about the communications that takes place across the Internet in support of various Internet applications:

Over dial-up connections to ISPs, users invoke Internet applications. The subsequent sessions across the Internet involve application-to-application, not user-to-user, dialogs;

The user does not initiate, nor is he/she even aware of, much of the processes taking place across the Internet;

The dialogs that take place often involve an application – the name resolution process – that the user does not invoke;

---

[34] One might claim that a similar process occurs in the voice network when, say, toll-free 800 numbers are translated into regular telephone numbers by elements of the SS7 network. However, like many such analogies, this one is strained and inaccurate. First of all, the Internet process is translating names with multiple parts, not telephone numbers structured like all other telephone numbers. Secondly, there is often a large number of interactions involving multiple hosts required to translate a single Internet name, not a single processor as in the SS7 network. But most profoundly, the voice network translation process is done internally to the network and is invoked by the voice switches themselves, whereas the DNS process involves hosts attached to the Internet, and is invoked by applications, not by the Internet routers, which are completely oblivious to the process that is taking place.

The supposed end-end "connection" between two hosts or humans – say, between an email sender and email recipient – consists of more than one dialog, and the different dialogs usually do not occur at the same time, and in fact may occur after the dial-up call to the ISP has been terminated;

As shown in Figure 3, over a single dial-up connection to an ISP, the user's host is typically involved in multiple sequential internet communications to different hosts and servers; and

In the case of electronic mail via SMTP, there is communication between mail servers belonging to the originating and destination ISPs that does not involve the sender's or recipient's host computers.

These differences lead to the summary points made at the beginning of this Declaration, which are restated below with additional amplification:

The Internet operates in a fundamentally different fashion than does a long distance voice network over which the dial-up connection to the ISP is established, because voice calls transpire over a circuit switched connection[35] which is established by the calling party, maintained for a finite period, then terminated by the calling or called party, whereas there is no notion of a connection over the Internet, and thus no connection to be created or terminated;

The Internet is a packet switching network that processes and transmits units of information on a one-by-one basis as they are sent to the network, interleaving the units from different communicating devices in the routers and over the transmission links of the

[35] Circuit switching refers to a form of network protocol in which the network commits switching and transmission resources, collectively called a "circuit," solely to a single instance of communications – two humans talking, or two computers exchanging data – and maintains that commitment until the communicating entities indicate the communication is completed.

network. Going even further, not only does the Internet use packet switching, it employees a connectionless version of packet switching that means there is no network awareness of any connection. During the times when an application is not sending any data, the same network resources that might otherwise be used to handle data from that application are instead are being used to handle data from other applications. In fact, if an Internet user dials a call to an ISP but then delays any subsequent action to initiate Internet communications, the Internet is oblivious to the dial-up connection – it does not monitor it in any way, and does not commit any resources to it. On the other hand, as Figure 3 shows, there can be multiple sequential (or parallel) Internet sessions that take place over a single dial-up connection.

Thus, to characterize communications across the Internet as being like one "continuous" end to end voice call is clearly a misnomer. Between the time a dial-up connection is established and the TCP or application session is established across the Internet, would the analogy with voice calls mean the call has terminated at the ISP? Worst yet, when the user's host communicates with multiple hosts to perform multiple tasks over the duration of a single dial-up connection, are we to picture this as a new invention of the voice network that allows the far end of an end-to-end call to be reconnected to different places multiple times?

Unlike ordinary telephone calls, the information supplied by the user to initiate a dialog with another computer is not a network address; it is an Internet "Domain Name," which must be interpreted and translated into an Internet address before the routers that are responsible for Internet transport can start routing information to the right destination;

No communication takes place over the Internet until high-level processes having nothing to do with the transport network itself (i.e., the routers) are completed.

Furthermore, the user is not involved in the translation process that takes place. If one were to think of communications across the Internet as being a "call," this means that before the real call a user intends, such as acquiring data from another Internet computer, there must be what is often a long sequence of "calls" between the user's computer and various servers involved in name-to-address translation. This further emphasizes that events subsequent to the establishment of the dial-up connection do not constitute an instantaneous continuation of the original dial-up call, either in the sense of time or in the sense that the network proceeds to establish a call to the intended destination. The closest analogy to this situation in the voice world might be some sort of electronic directory service that a telephone company provided that would allow a caller to indicate electronically[36] he/she wanted to reach "Joe Smith at 560 main street in Milwaukee," a computer would search for that information in a database, determine the telephone number – which was not provided by the caller – and place the call, all without the involvement of a human operator. Such a service, if provided by the telephone company, would manifestly not be part of the basic telephone call, but would be an enhanced information service. Furthermore, it would involve the caller at every step, whereas the Internet user has no involvement in or awareness of the name translation process that is occurring.

Finally, and perhaps most importantly, there is no significant voice network equivalent of the Internet process in which multiple application-to-application sessions involving multiple sites take place sequentially over a single dial-up connection without subsequent involvement of the user.

---

[36] In a fashion that is barely feasible with ordinary telephone devices in the first place, due to their limited repertoire of control instructions they can convey to the telephone network – basically, twelve tone combinations, a switchhook flash, and hanging up.

This means in effect there is a one-to-many relationship between the dial-up connection to the ISP and the sessions taking place in the Internet, and the extra sessions are outside the control of the user.[37]

Based on the foregoing, it is clear that dial-up calls to ISPs are in no meaningful way analogous to long distance calls. Long distance calls are directed to a single destination for the duration of the call based on a single number dialed. There is a one-one coupling between the originating and terminating exchange access part of the connection and the IXC part of the connection. And long distance calls do not invoke multiple applications communicating with multiple end-points during the process of the call, to which the caller is oblivious. But all of these phenomena can and do occur in the Internet.

V. Critique of the White Paper by Charles L. Jackson

As part of its filing, Verizon submitted a paper by Charles L. Jackson that purports to show that in an Internet connection, like in a telephone call, "the information flow – the content, the communication – passes unchanged from the originating computer to the destination computer."[38] One is supposed to conclude from this that the dial-up call to an ISP is merely the initial part of an end-to-end connection, and thus essentially identical in nature to a long distance call.

---

[37] One might claim that the credit card service offered by many IXCs, in which a user can use the # key to end one call and initiate another without breaking the connection from the user to the IXC POP or losing the billing information, is an equivalent service. But it is not – unlike the Internet process we have described, it involves the user, and it amounts to different instances of user-user communication, whereas the multiple Internet processes we have described pertain to the same user-user communication but different applications that must be invoked for that user-user communication to take place.

[38] Jackson White Paper, para. 12.