

Sponsored By

SPECTRUM

Telephony's Next Act

By: Paul V. Mockapetris



ILLUSTRATION: ALEX NABAUM

There are about 1 billion fixed telephone lines and 2 billion cellphones in the world. Most calls still travel across traditional systems based on proprietary software and hardware. Soon, though, they will move to networks based on open protocols—known as voice over Internet Protocol (VoIP).

The same thing is happening all around the world—in some cases very quickly. Last winter, the chief technology officer at BT Group PLC, in London, said that during the next three to five years, BT would like to "turn off the public switched telephone network."

If that happened all over the world overnight, and all those phones connected to the Internet instead, it would double or triple the number of objects joined by that network of networks, wreaking havoc on some parts of it. Done right, though, over the next several years, the move will open doors to new telephone services that will make all our lives easier and better.

Today it is by no means obvious how best to call a friend or colleague. Should I call your home, your office, or your cellphone? Should I send an instant message to your cellphone or a text message to your computer? If you're not home, should I leave you an e-mail or voice-mail message? In the ideal world, I would simply let the network figure out how best to get to you.

To realize this ideal we must master the problem of signaling, that is, keeping track of all the potential communicating parties, their equipment, and their services, and selecting the right combination for each contact. As we shall see, there are two technical issues. First, we must merge the telephone system with the Internet through VoIP; this trend is already under way in many companies. Then we must strengthen the portion of the network that handles signaling across the entire world and the Internet, so it can stand up to the resulting load. That task, in fact, is the brief of my company, Nominum Inc., in Redwood City, Calif., which provides software for the domain name servers that lie at the heart of the Internet.

The traditional telephone networks and the Internet handle signaling very differently. Voice customers use traditional phone numbers, while the Internet is based on domain names, such as my company's "nominum.com," and numerical Internet addresses, such as "81.200.68.193." Voice callers will want to keep their phone numbers, even as the calls move to the Internet. At the same time, Internet users will want to place calls by clicking a browser link or perhaps by selecting a menu item on a PlayStation or other game machine. They'll also want to start phonelike conversations and Internet chat sessions, and to use videoconferencing software on their laptops and PDAs. So VoIP will have to support e-mail and Web addresses as well. In many ways, the problem of merging these two different leviathans comes down to merging their respective databases, which is to say their searchable collections of data records. A few years ago, Internet standards bodies approved a new type of data record—called ENUM, for electronic number—that could accommodate the legacy information of both telephony and the Internet. ENUM will be the bridge between the two global communications networks. Among other things, it provides a simple way to turn a phone number into an Internet address.

*The very flexibility of the ENUM record
makes it a tempting target for abuse,
such as voice-over-IP spam*

As we will see, ENUM offers an opportunity to revisit in a fundamental way the Internet's design, creating an open standards-based approach to the way we identify network devices and the people who use them. Openness is key if we are to ensure that tomorrow's VoIP world is as open and interoperable as the Internet itself has been. There's a danger that regions will instead be cordoned off by some private VoIP providers, such as Luxembourg-based Skype, that encourage their customers to use names specific to their systems.

In order to understand how ENUM can unify these two disparate worlds, though, we need to understand how traditional telephony and the Internet handle signaling today.

The conventional telephone system, jocularly called POTS, for "plain old telephone service," is the product of more than 100 years of evolution. The first commercial office, the Strowger exchange, which began operating in the early 1890s, was the first to replace a human operator with an electromechanical switch. (Strowger, a mortician, invented the system to get around a local operator, the wife of a competing mortician, whom he suspected of redirecting calls for Strowger to her husband.) During the last few decades, parallel work in international, commercial, and national standards has led to Signaling System 7, or SS7, which replaced electromechanical switches with electronic switching nodes.

Today, we expect a phone system to do a lot more than just complete phone calls. It also must provide toll-free numbers; services that locate the caller in an emergency (such as 911 in the United States and Canada and 999 in the United Kingdom); caller ID; automated payment for collect calls; do-not-call lists to ward off telemarketers; the ability to carry your number with you when you move; and so on. To support these services, the system needs database nodes to manage information about who is using the telephones and whom to bill for their use. The databases are owned by carriers or operated for them by specialized companies. Although a customer may influence these databases, say by changing carriers or blocking caller ID, the companies regard the databases as business assets, keep them under wraps, and charge for database lookups.

One of the basic concepts is a fully qualified telephone number, such as +1 650 381 6115 or +41 22 730 5111. This is defined by an International Telecommunication Union standard known as E.164. All ITU-T E.164 numbers are 15 digits or fewer and start with a country code (+1 for countries on the North American numbering plan, +41 for Switzerland, and so on), suballocated by the corresponding national bodies under nation-specific policies. The "+" denotes a complete phone number. Don't, however, be confused by the common U.S. phrase, "call 1-800 [number]"; services such as the toll-free numbers, 911, and directory assistance are processed in a country-, carrier-, and region-specific way and routed to their final destinations.

The Internet, on the other hand, employs several kinds of identifiers. Perhaps the most basic is the domain name, supported by a highly distributed network called the DNS, for Domain Name System. If, for instance, you point your browser at the domain name maps.yahoo.com, your computer will query your local DNS server for all the IP address records—numerical codes signifying physical machines in the Yahoo network—that are associated with that domain name. And if you write me at pvm@nominum.com, your mail server will ask for the corresponding codes (called MX, or mail exchange records) that are associated with mail for the domain name "nominum.com."

Let's look at the latter case in some detail. If your query is a very common one, then perhaps the local network server in your company will have cached the answers to it, obviating the need to relay it beyond the building's walls. If, on the other hand, the answer to the query isn't cached, your server will relay the query through the Internet to Nominum's name server, prompting the server to send you a response. That response, represented in text, would read as follows:

nominum.com. 3600 IN MX 500 mx1 -above.nominum.net.

nominum.com. 3600 IN MX 5 mx2.nominum.com.

nominum.com. 3600 IN MX 10 mx1.nominum.com.

These three records identify three different mail servers; Nominum maintains two to provide redundancy against the possibility of the main mail server's failure. Each MX record contains the name of the mail server and a preference value that tells the order in which the servers should be used, beginning with the lowest value ("5," in this case).

How did your name server find Nominum's? Internet names are interpreted from right to left, beginning with the domain name's suffix—.com, .org, .arpa, .uk, .de, .aero, and so on. A local server looking for information about maps.yahoo.com, and starting with no cached information, would go to one of the 13 root servers spread out across the globe and find the identity and address for name servers having information about .com addresses. The server would then redirect the query to a yahoo.com DNS server, which would send back the correct numerical address for maps.yahoo.com.

The Internet uses domain names to create easily remembered Uniform Resource Locators (URLs) and Uniform Resource Identifiers (URIs) such as <http://www.nominum.com> and <ftp://nominum.com>. The first is the commonly seen hypertext transport protocol for Web pages; the second is the venerable file-transfer protocol. A URI is seen in <sip://paf@swip.net>; SIP (Session Initiation Protocol) can set up audio streams that are functionally identical to phone calls.

One of the more remarkable aspects of DNS is that the ownership, control, and operation of the system are distributed widely. Nominum.com is one of the approximately 40 million domains that are separately owned under .com. Other suffixes carry, in total, about an additional 100 million such domains, all registered to companies, groups, and individuals in return for annual fees.



ILLUSTRATION: ALEX NABAUM

So where does ENUM fit in? Originally, it was conceived as little more than a way to use the DNS to keep track of a few phone numbers that had been moved from one carrier to another. It works, basically, by taking a telephone number, putting periods between the digits, and flipping the resulting string around so it matches the specific-to-global (or right-to-left) format of domain names. The resulting numbers go into a single global directory of telephone numbers in the DNS, located under "e164.arpa." The .arpa top-level domain has been limited to database lookups.

For example, the globally unique telephone number: +1 650 381 6115 becomes the globally unique domain name: 5.1.1.6.1.8.3.0.5.6.1.e164.arpa.

The ENUM standard specifies the use of a resource record called the Naming Authority Pointer (NAPTR) resource record, which can encode far more information than a simple 32-bit value. It does so by means of algorithmic rules which, when sent back to the requester, can be used to compute a customized answer.

ENUM is an international phenomenon. Here's an example for the U.K. phone number +44 1632960083, with data from domain name 3.8.0.0.6.9.2.3.6.1.4.4.e164.arpa:

NAPTR 10 101 "u" "E2U+email:mailto" "!^.*\$!mailto:info@example.com!" .

NAPTR 10 102 "u" "E2U+sms:tel" "!^.*\$!tel:+441632960083!" .

NAPTR 10 100 "u" "E2U+sip" "!^.*\$!sip@example.com!" .

The key variable here lies in the space occupied by "mailto," "tel," and "sip." These elements refer, respectively, to an e-mail application, POTS, and the SIP, the Internet Engineering Task Force's signaling protocol. (Among other things, SIP is what allows an instant-messaging buddy list to say which of your friends is online the same time you are.) So these database statements tell any inquiring network that the customer in question can be reached via SIP, via e-mail to info@example.com, or via conventional telephony. The numbers specify that the caller (or a

calling software program) is asked to use SIP first, and if that doesn't work to try mail or telephony.

Multiple NAPTR records associated with a single number can provide complicated priorities and preferences for contacting users both across devices, including cellphones, PDAs, fax machines, and voice-messaging systems, and within individual ones. A cellphone alone, for example, can be used for voice, voice mail, instant messaging, text messaging, and e-mail, and each of these can have different rules for when the phone is turned on or off. One might even want to send copies of a message in several formats at once.

Of course, the very flexibility of the ENUM record makes it a tempting target for abuse, such as e-mail spam, instant-message spam, and even junk telephone calls, which will surely become common as voice telephony moves to the Internet. Quite apart from such intrusions is the threat that outsiders may try to get their hands on proprietary information, such as the contact numbers of a given company's customers or vendors.

For these reasons, communication providers may choose to put their NAPTR resource records in a secure domain outside the e164.arpa domain, that is, in a private ENUM, accessible only to privileged devices. A level of accessibility in between these two domains may also be created for carriers that wish to exchange ENUM data among themselves, to take advantage of the lowest-cost method of routing telephone calls through the Internet. We call this version Peered Infrastructure ENUM.

Public ENUM trials are under way in a number of countries, including the United Kingdom, Germany, Austria, and Australia; in some cases, these trials have to resolve political and privacy issues as well.

So, WILL the Internet handle the load? The answer is yes, with a large caveat. First of all, the "pipes" themselves—that is, the fiber-optic cables in the Internet backbone—are fat enough, and new technology is always coming along to make them effectively fatter.

But the other part of the system, the DNS, is not quite so robust as many like to think. First, a merged world will require vastly more storage than today's DNS system. NAPTR records will be larger, typically dozens of bytes, compared with today's 4 bytes for IP addresses. In addition, there will be hundreds of millions of names to track in some servers, versus today's entries, which generally run in the thousands. Even the .com database—by far the biggest—now has only 40 million names.

Second, it takes time to answer a DNS query, and such latency, as it is called, has so far been overlooked because it doesn't matter for e-mail and is generally tolerated in Web page displays. Latencies in the tens of seconds are not unusual. Telephone callers, however, expect latencies of less than half a second.

Particularly in public ENUM, subscribers will want to view and change their own contact information and will expect changes to take place immediately. Often these changes will be automated or implicit as a phone comes within range of a Wi-Fi access point, for example. Many of today's DNS servers have complex and lengthy update requirements; they will simply not meet these needs. Further, DNS servers will need to continue answering queries while they are being updated. Indeed, configuration and provisioning often require Internet service providers to take DNS servers off-line, at least briefly. People won't stand for such outages in telephone service. In short, today's DNS servers cannot handle the challenges that ENUM will pose.



ILLUSTRATION: ALEX NABAUM

To measure the effects of ENUM, my company has studied the DNS server implementations that telephone companies commonly use: BIND Version 9.3.0 from the Internet Systems Consortium Inc., Daniel J. Bernstein's DJBDNS, PowerDNS's PowerDNS, and Nominum's ANS, or Authoritative Name Server.

Two of them, PowerDNS and DJBDNS, do not support the latest security protocols for DNS, nor can they be dynamically updated, so they might be unsuitable for ENUM use in any case. In our first tests of the others, we attempted to load 200 million NAPTR records onto 32-bit servers, each with 2 gigabytes of main memory, half the maximum they can use. Only ANS was able to load the

data, so we retreated to 50 million records, and then 10 million, to get results for more than one software program. At this size, the most popular DNS software, BIND, could answer just 143 queries per second. DJBDNS answered 6992 queries per second, and ANS answered 45 135 queries per second. While there is no absolute metric for success here, it's obvious that we need to continually reduce the cost of DNS operations if we are to see the same results in VoIP.

Why did BIND do so poorly? We believe the primary reason is that while it does well for names such as www.ieee.org with three labels, it is significantly challenged by ENUM names, which, by virtue of having dots separating each of the digits in the name, have at least a dozen labels. Similarly, when it comes to dynamic update, BIND achieved 69 updates per second, whereas ANS achieved 467 updates per second, and both need further controls to allow a balance between network control (updates) and network user performance (queries).

Moving to computers with 64-bit processors—which allows for much larger memory spaces—helped, but it wasn't a panacea. In tests on a 64-bit Opteron processor, with test data of 10 million names, organized into groups of 10 000, BIND delivered 3805 queries per second and ANS 57 000 queries per second. Interestingly enough, BIND performance degraded by a factor of 50 when ENUM data were organized into separate groups as they might be for public ENUM.

The tests suggest some necessary improvements for ENUM deployments: they may not be sufficient if ENUM data grow to 10 or 20 NAPTRs for every number instead of the one NAPTR per name used in these tests, or if a new security protocol, DNSSEC, is used and adds its significant overhead to queries. Using existing DNS servers, ENUM is cumbersome to set up and manage. Replacing the DNS protocols with another protocol is not the problem, but re-engineering the way they're implemented in a network is. We at Nominum estimate that a service provider with 200 million records would need to install up to 20 times as many servers if it stayed with an in-memory database and the memory limitations of 32-bit processors. Even with 64-bit machines, there are other challenges, such as the right balance between update speed and query speed.

If these issues aren't faced, Internet users are likely to face higher costs and decreased quality of service, delayed call connections, and dropped calls. Alternatively, the provider could opt for a scalable solution, such as Nominum's. (Of course, we believe that our solution is preferable.) This isn't merely a service provider's problem—it's everyone's problem. We all would like to be able to dial a phone number without worrying about what kind of service the recipient has. The Domain Name System needs to be upgraded with software tested against the high-volume, constantly changing loads of ENUM. Only by strengthening these DNS capabilities will we be ready for the demands of ENUM and other new network technologies. As the demands we face are new, so must our approach be new.

About the Author

PAUL V. MOCKAPETRIS is chairman and chief scientist of Nominum Inc., in Redwood City, Calif. He invented the Domain Name System at the University of Southern California's Information Sciences Institute in the early 1980s and has served as the Internet Engineering Task Force chairman, as program manager for the Advanced Research Projects Agency, and as chief technology officer for several start-ups. He was the recipient of the 2003 IEEE Internet Award for his DNS work.

To Probe Further

A useful primer on ENUM is at <http://www.fcc.gov/realaudio/presentations/2004/110404/McGarryTom.ppt>.

An article in the March 2005 *IEEE Spectrum*, "[Seven Myths About Voice Over IP](#)," offers some additional background information about the protocol.

Readers seeking more information about DNS can start with an article published by the Internet Society, "DNS Root Name Servers Explained for Non-Experts," at <http://www.isoc.org/briefings/019> and follow its links.