

March 1999

By Peter Rysavy, Rysavy Research

Copyright 1999 Peter Rysavy and Data Communications

All rights reserved

Tech Tutorial

Wireless IP: Ready to Lift Off?

Corporate networkers who really want to run IP everywhere need to be sure about what wireless can really do and what's just hot air

Wireless data isn't exactly a front-burner application for most network architects. But substitute IP for data-as in wireless IP-and suddenly things start to simmer a bit. In fact, wireless IP might just be the enabling technology that finally realizes the popular protocol's rallying cry: IP everywhere. Wireless IP isn't just a blue-sky transport (pun intended). It can be accomplished right now.

Increasingly available and affordable wireless technologies like CDPD (cellular digital packet data) and data-over-PCS (personal communications services) do a fine job shuttling IP packets. But real-world applications are more than the sum of their packets. Corporate networkers who want to put wireless IP to work need to clear some high hurdles, starting with limited coverage, low throughput, and high latency. Wireless middleware could be another wrestling match: Products are just starting to emerge that don't force net managers to reengineer their current apps. Wireless service bureaus

also are starting to appear that can offload the hassles and headaches for a fee. And bear in mind that linking wireless end-users—who could literally be anywhere—still means dealing with old-fashioned landline links.

So is wireless IP worth it? That depends. The untethered technology can make a lot of sense in many business situations. Giving mobile workers access to corporate information and the Internet from anywhere can make them far more productive. But wireless IP is not just another type of WAN connection. It has its own unique characteristics, ones that were not necessarily considered by developers of communications applications and networking protocols. Any evaluation of return on investment not only has to take into account the cost of wireless equipment but also must factor in learning curves and special situations.

Coverage Considerations

No wireless IP network can match the coverage of today's analog cellular network, so let's begin there. With the right combination of PC Card modem, cable, and cellphone, any end-user can dial an ISP (Internet service provider) and establish a 9.6-kbit/s Internet session.

This is wireless IP, but the wireless net isn't actually providing any IP functionality. What's more, compatibility problems between modems and cellphones have kept this scheme from catching on, as has the cellular carriers' inconsistent deployment of the gateways needed to translate between wireless and landline modem protocols. But for the right application and situation, this can be the best solution simply because analog cellular offers such broad coverage.

The next best coverage option is with the two oldest packet-data networks: Bellsouth Wireless Data (Woodbridge, N.J.) and Ardis, now operated by American Mobile Satellite Corp. (Reston, Va.). Though both networks reach over 90 percent of the U.S. population, they are not truly IP networks since they use proprietary protocols.

Nevertheless, IP gateways are now available that let these networks operate as wireless extensions of the Internet or corporate intranets. Trouble is, their volume-based pricing, relatively low effective throughput (approximately 4 kbit/s for Bellsouth Wireless and 2.4 or 9.6 kbit/s for Ardis), and round-trip latencies of 3 seconds or higher don't suit them to general-purpose IP networking.

CDPD is the next option, a digital packet overlay to the analog cellular system that transmits IP at an effective throughput of about 10 kbit/s. Each mobile end-station has a fixed IP address and is a true Internet host.

CDPD is offered by a number of the major cellular providers. Coverage has improved tremendously over the past several years, with service now available in most major cities. But big holes remain, including Atlanta and Los Angeles. Providers also emphasize denser population areas, so expect coverage downtown, around airports, and in commercial districts. But as end-users drift through the suburbs and into rural areas, coverage becomes spotty-even though cellphones keep working fine. As with any wireless offering, the easiest way to check coverage, as well as speeds and fees, is to pay a quick visit to the provider's Web site. Most post maps and other information.

The Personal Touch

The technology that could tap wireless IP's true potential, however, is data-over-digital PCS. Today's offerings are quite limited, restricted to circuit-switched connections to GSM (global system for mobile communications) networks. Though barely used in the U.S., GSM is very popular in Europe, exceeding 15 percent of cellular traffic in some Scandinavian countries.

PCS should pick up some popularity as new services are rolled out. This year should see 14.4-kbit/s Internet access for both GSM and CDMA (code-division multiple access) nets. Though still circuit switched, these offerings will feel like packet thanks to the faster connections

enabled by completely eliminating analog modems.

Beginning late next year, CDMA, GSM, and IS-136 (interim standard 136) networks will all start offering high-speed packet data ranging from 64 to 384 kbit/s.

But as PCS data services are rolled out in the U.S., coverage will become a critical issue. Competing digital technologies are available, extensive roaming agreements still need to be hammered out, and far too many so-called digital providers still rely heavily on analog cellular in sparsely populated areas. Things will be better in Japan, Korea, and many European countries, which have standardized on one digital technology and have the dense population needed to justify the buildout of an all-digital network.

Corporate networkers with sites in the San Francisco Bay area, Seattle, Wash., or Washington D.C., also would do well to check out Ricochet, the wireless IP offering from Metricom Inc. (Los Gatos, Calif.). At 28.8 kbit/s, it outperforms any other wireless IP offering out there; its flat-rate pricing also is attractive. The provider says it will expand into other cities this year and indicates it's working on a 128-kbit/s network.

The Need for Speed

Faster wireless IP offerings are on the way. So-called 3G (third-generation) cellular will deliver throughput as high as 2 Mbit/s in the local area and 144 kbit/s for mobile. But the earliest these systems will arrive is next year, and that's on a very limited basis. Most corporate networkers who want wireless IP now are going to have to learn to live with 9.6 to 14.4 kbit/s—a big consideration given that even 56-kbit/s transfers can seem agonizingly slow on today's Internet.

They're also going to have to deal with something else: latency—and its potentially crippling effect on client-server applications.

Consider an SQL (structured query language) database or groupware/e-mail apps like Exchange from Microsoft

Corp. (Redmond, Wash.) and Notes from Lotus Development Corp. (Cambridge, Mass.). These send a large number of messages back and forth in the course of executing transactions. A single query, for instance, could generate dozens of messages. With latencies (and round-trip times) ranging from 0.5 to more than 5 seconds, a new screen of information that takes a few seconds to update over a LAN could take a half a minute or more over a wireless IP link .

Even this wouldn't be a complete disaster. But many end-users, after staring at their screens for what feels like forever, assume the application or machine has hung and terminate it. Or they reboot their PCs.

Then there's the problem of intermittent connections. These are not uncommon with wireless services, particularly if the end-user is mobile. (It's not much different from having a cellphone call dropped in the middle of a conversation.) Some applications can shrug this off; others may become unstable or damage data.

Set Up to Shut Up

So what can net managers do? For starters, find out how to configure chatty applications so they only communicate essential information over the wireless link. For example, by minimizing the amount of upfront synchronization, it's possible to reduce a remote Microsoft Exchange login from over five minutes to about one minute. Similarly, Lotus Notes can be configured to replicate just the subset of the databases that end-users need to work with, rather than all of them.

Another timesaver is to set up e-mail applications so they screen out attachments or don't automatically download messages above a certain size. Database apps also can be programmed to load data dictionaries from a local cache rather than over the airlink. And even Web browsers can be set to operate in text-only mode.

The Middleware Muddle

In some cases, the foregoing fixes may be sufficient. Even

if they're not, they'll get net architects thinking in the right direction. But if applications are still plagued with problems, it's probably time to turn to wireless middleware.

The benefits are real: Wireless middleware reduces both the amount of information communicated over the air and the number of messages needed to complete transactions. In many cases, it can actually queue messages when end-users are outside of a coverage area. Further, it may use transport protocols that are not only more efficient than TCP over airlinks but also more resilient to wireless idiosyncrasies, like variable delay, that can confuse TCP timing algorithms.

However, the downsides to wireless middleware are equally real. It adds cost and complexity.

Wireless middleware is based on an architecture comprising software loaded on the mobile computer (the middleware client) and software that runs on a server (the mobile server), which acts as a proxy to access information from other services on behalf of the mobile client using standard LAN and WAN protocols (see Figure 1). The middleware client and server components use special protocols to communicate with each other over the wireless network.

The challenge is marrying the application to both the client and server portions of the middleware. There are two basic approaches, and the difference between them is huge.

In the traditional scheme, the middleware is supplied by the vendor in a software developer kit (SDK) with APIs (application programming interfaces) for the mobile client and server. It's up to the network manager and IS staff to reengineer existing applications or develop new ones from the ground up. This involves a substantial development effort. It also assumes that corporate networkers already have or will develop the source code for their apps—clearly not an option if they plan on using off-the-shelf packages. (No wonder that wireless

networking has chiefly been used in vertical markets.)

To the Rescue

Fortunately, a new type of wireless middleware has recently emerged. Like its SDK-based counterpart, it comes with client and server components. But it requires almost no programming, other than simple setup; once that's done, it automatically intercepts application and networking calls and optimizes communications.

Unfortunately, there are only a few products out there now. The selection should improve, however, as wireless data starts to enter the mainstream.

Smart IP from Nettech Systems Inc. (Princeton, N.J.) optimizes IP-based communications, potentially making any IP application operate more efficiently over wireless links. It also enables these apps to operate over packet-data networks like Bellsouth Wireless and Ardis.

Laplink Enterprise from Traveling Software Inc. (Bothell, Wash.) is another product to watch. The package includes an accelerator for Microsoft Exchange that not only speeds up mail downloads but also gives end-users far more control over what is downloaded and how. The vendor doesn't have an accelerator available for Lotus Notes, however.

Whither Wireless?

There are a couple of other developments to watch in wireless middleware. For instance, a number of large wireless providers have banded together as the Wireless Application Protocol (WAP) Forum. The group is developing standards that make Web content readily available to mobile wireless devices like smartphones. AT&T's Pocketnet is a precursor to this equipment; it melds a cellphone capability with a CDPD modem and a small-display browser. Users interact with specially prepared content hosted on Web servers, communicating via a gateway operated by AT&T.

The other development is the growing number of wireless service bureaus, like Goamerica Communications

Corp. (Hackensack, N.J.; www.goamerica.com) and Wireless Knowledge LLC (San Diego; www.wirelessknowledge.com) the recent joint venture between Microsoft and Qualcomm Inc. (San Diego). These offer a variety of services, including hosting middleware.

Here's how it works. Rather than linking directly to the corporate intranet, mobile workers establish connections to the service bureau, which has secure, high-speed connection to the company intranet. The bureau accesses the requested information and then presents it in a manner optimized for wireless communication, employing formats like WAP or XML (extensible markup language).

Getting Grounded

Wireless IP may conjure up visions of packets streaming through the air, but how that data reaches its destination on the landline network is a critical concern. For most corporate networkers, this will be familiar turf; the issues are common to remote access in general.

Wireless networks link to the rest of the world in three basic ways: over private connections, over the Internet, and over the PSTN (public switched telephone network). Private connections are the preferred approach for Ardis and Bellsouth Wireless, customers using X.25 over leased lines. Large CDPD customers, in contrast, use IP over frame relay PVCs (permanent virtual circuits) to connect to the carrier.

Increasingly, though, wireless IP providers advocate the Internet as the preferred connection method. With CDPD, as well as all future packet services for PCS, the wireless network has a seamless link to the 'Net. It's much easier for the carrier to point customers to this connection than to deal with individual private connections.

Be aware that going with the Internet means that remote users on a wireless IP network will have the same problems getting through the firewall as dial-up users coming in through an ISP. Fortunately, VPNs (virtual

private networks) can deliver the requisite security, and an increasing number of vendors are selling solutions. But not all the security standards, including IPsec itself, are complete. And while many companies are evaluating VPNs, most do not yet allow Internet-based remote access.

The third avenue of access, the switched public network, is the least problematic for companies, since they already have dial-up systems in place for remote users. If calls come in over a circuit-switched data service on a cellular network (via either an analog or a digital link), they will be switched into the PSTN and appear identical to other modem connections.

Many of today's cellular data services are, in fact, circuit-switched. This will continue to be the case for the next few years, until wireless packet-switched services start rolling out in earnest. Network architects who start getting ready now are going to be in the right place to exploit new services and technologies.

Peter Rysavy is president of Rysavy Research (Hood River, Ore.), a consulting firm that helps companies research, develop, and deploy communications technologies.