Roaming Far And Wide With Mobile VPNs

Lisa Phifer

Users on the go want continuous connectivity. Increasingly, you'll have the tools to make it happen.

If you spend much time away from your desk, you're probably familiar with this drill: Find network access, log in, rush to resync email and files before losing connectivity, then work offline until the next connect opportunity. Repeat *ad infinitum*. This common strategy may be (mostly) sufficient, but it's hardly efficient, much less optimal.

High-speed public Internet access is starting to change the drill by offering opportunities to get and stay connected in more locations. Today, hotels that cater to business travelers frequently offer in-room high-speed Internet access. Wireless LAN hotspots have surfaced in hotel lobbies, airport boarding areas, cafes and conference centers. Wireless WAN packet data services like Single-carrier Radio Transmission Technology (1xRTT) and General Packet Radio Service (GPRS) have become common in metropolitan areas, and trials are now under way for faster services like 1xEvolution-Data Optimized (1xEV-DO), Enhanced Data rates for Global Evolution (EDGE), and Universal Mobile Telecommunications System (UMTS).

As these high-speed access networks ramp up, mobile workers are changing their habits, remaining on line for longer. It may take several years to reach that "always connected" nirvana, and connectivity in less populous areas will lag high-tech corridors. However, workers who *do* enjoy ubiquitous high-speed connectivity are starting to demand devices and applications that can roam seamlessly across heterogeneous networks. Accomplishing this requires many improvements, including security measures that can facilitate, not fight, network roaming.

Going Mobile

Most enterprise surveys rank features like security and performance considerably higher than network roaming on the "wireless wish list." But according to an April 2004 survey of IT managers conducted by iGillot Research, nearly three out of four companies with more than 1,000 employees have deployed a wireless data solution. Of those, two-thirds report using a combination of wireless LAN and wireless WAN access methods. Companies that are just now getting their feet wet with wireless have not yet felt the pain of cross-network roaming—but eventually they will.

Organizations often try to extend their existing remote access VPN solution to devices that connect over wireless. On paper, this sounds like a good idea—why not leverage existing IT infrastructure to secure new access methods? In practice, early adopters are reporting this approach falls short of expectations. Reasons are numerous, including poor performance on wireless WANs, the inconvenience and latency associated with repeated login after loss of connectivity, and appli-

Executive Summary

In their search for anytime/anywhere connectivity, road warriors need to roam among diverse networks, both public and private. That means network managers have to find a way to stitch these networks together to maintain connections.

Options include Layer 2 mechanisms such as VLANs; techniques for roaming between subnets within a private network; and technologies for roaming between networks with different owners. The Mobile IP standard creates a framework for roaming among different networks—though it has both supporters and detractors.

Network managers must also take into account concerns like session persistence—the user's need to roam seamlessly, without losing active sessions. Also, users' devices may be equipped to access multiple types of networks, creating the need to prioritize which is the preferred type of connection.

Ultimately, if users are to have seamless mobility for their remote access connections, the technologies described in this article must be accompanied by business arrangements among different network owners, as is the case today in cellular voice networks□

Lisa Phifer is vice president of Core Competence, Inc., a network security consulting firm based in Chester Springs, PA.

cation session disruption when roaming both within and beyond any given wireless network.

A number of wireless-aware products and features have emerged to address these challenges. Some products focus on seamless roaming within privately-operated networks. For example, wireless gateways and switches from vendors like Bluesocket and Trapeze Networks offer mobility features that help 802.11 stations roam between access points inside a company's network without losing traditional VPN connections.

On the other hand, mobile VPN products focus on roaming across diverse networks and managing differences in performance and cost associated with public networks. Examples include AppGate's Security Server, Columbitech's Wireless VPN, Ecutel's Viatores Mobile VPN, ipUnplugged's Mobile VPN and NetMotion's Wireless Mobility XE. Although these product architectures and security protocols differ considerably, most mobile VPNs are capable of operating as shown in Figure 1.

Here, a wireless laptop, PDA, or smartphone with mobile VPN client software uses a hotel's hospitality LAN to connect to a mobile VPN server over the Internet. The worker logs in, either to a gateway in the company's DMZ, or directly to a VPN server inside the company's firewall. Upon successful login, the client is bound to a virtual address and data is sent through an encrypted tunnel between the mobile VPN client and server. If the client then connects to an application inside the company intranet, that session can survive loss or change in physical connectivity as the worker hops a cab, uses a wireless WAN en route to the airport, then upgrades to a faster hot spot connection at the airport.

Organizations that are launching a mobile data initiative should evaluate their business requirements for secure network roaming. Carefully considering the consequences and available solutions can be worthwhile, even when cross-network roaming isn't an immediate business need. Wireless roaming does occur all the time, even inside standalone networks. And choices made at the start of any deployment have a way of creating a legacy that must be accommodated for years to come.

Network Neutrality

Public Internet access networks incorporate security measures to meet the business needs of network operators. For example, wireless hot spot login portals and smartphone authentication by

Carefully evaluate your enterprise's requirements for secure network roaming

FIGURE 1 Network Roaming In A Mobile VPN

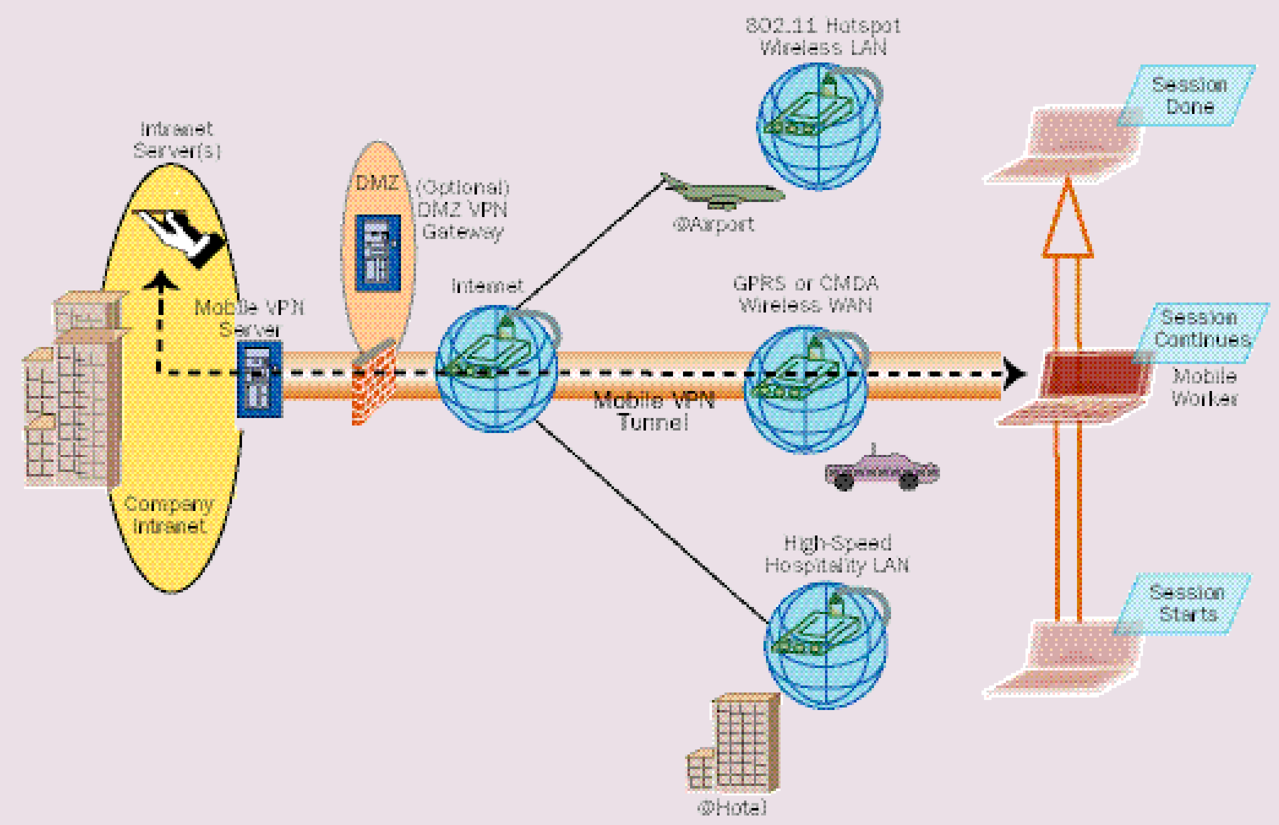

IPSec connections are encrypted end to end

Subscriber Identity Module (SIM) are primarily intended to enable subscriber billing and prevent theft of service.

Wireless hotspots rarely use 802.11 encryption (WEP, Wired Equivalent Privacy or TKIP, Temporal Key Integrity Protocol), but even if they used WEP or TKIP, this only protects the wireless access link; data would still be vulnerable to eavesdropping, modification and replay as it passes unencrypted over the Internet backbone. Carrier networks do encrypt packet data services between mobile devices and base stations, but IP packets usually exit the carrier's network through an Internet gateway as cleartext—that is, unless a higher-layer VPN is used.

VPNs based on secure tunneling protocols like IPSec are mostly (but not completely) network-agnostic. IPSec tunnels create an overlay network independent of the routing, addressing or security (if any) applied anywhere in the tunnel's path. For this reason, IPSec is well-positioned to provide "end to end" security for traffic originating from any kind of public network—wired or wireless. The company that owns the VPN gateway retains responsibility for authenticating users, authorizing access to private resources behind the gateway, and dictating the level of security required for mobile devices that use the VPN.

Unfortunately, IPSec VPN tunnels depend on the tunnel endpoint's IP address. As shown in Figure 2, this dependency commonly results in tunnel disconnection (and application session reset) when a mobile device roams from one subnet to another and the VPN client's IP address changes as a result.

This problem occurs because IPSec provides per-packet authentication based on source IP address. When an IPSec VPN client is successfully authenticated, security parameters for the tunnel are bound to an IP address and index. Whenever a packet is received, an IPSec VPN gateway uses the packet's source IP address and an index carried in the packet to locate the negotiated parameters and take appropriate action (e.g., decrypting and validating the packet with crypto keys agreed for this tunnel).

FIGURE 2 Subnet Roaming In An IPSec VPN

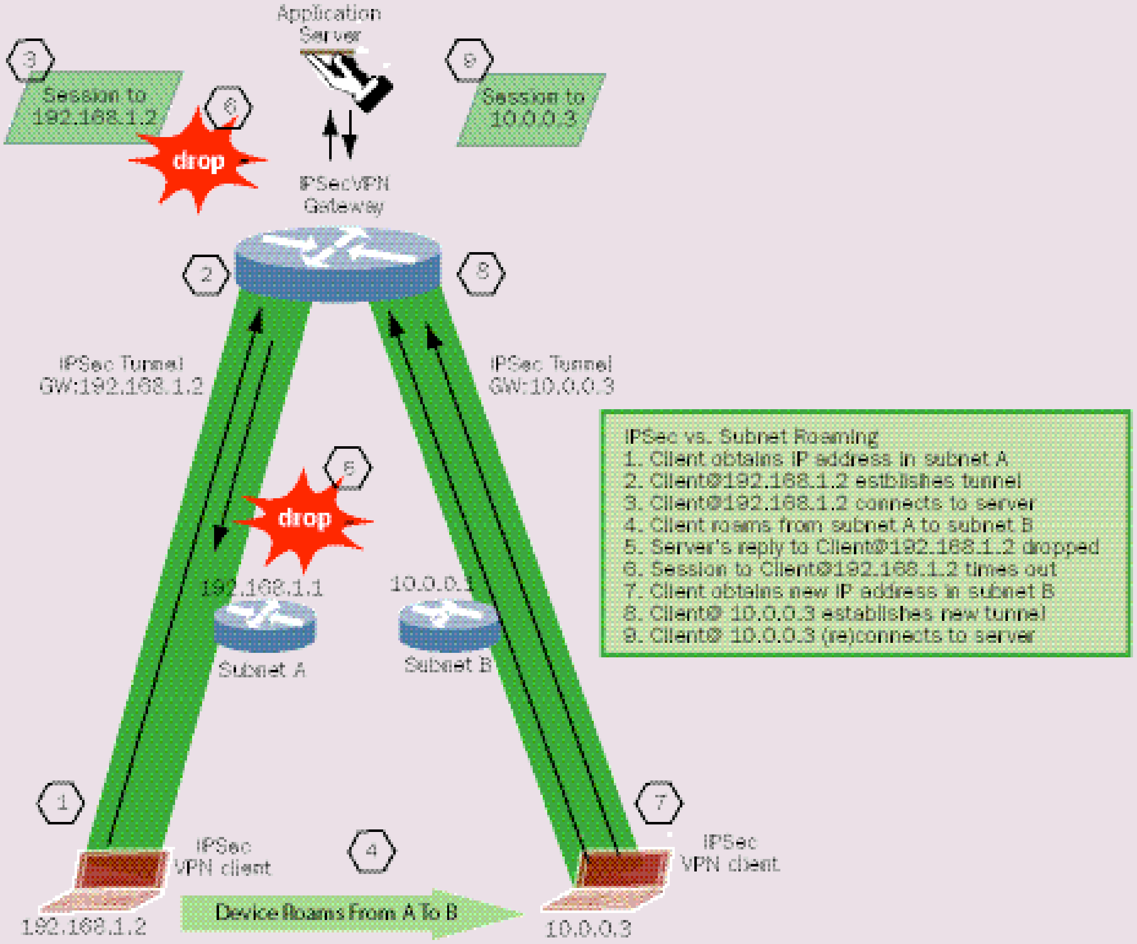

IPSec vs. Subnet Roaming
1. Client obtains IP address in subnet A
2. Client@192.168.1.2 estbilshes tunnel
3. Client@192.168.1.2 connects to server
4. Client roams from subnet A to subnet B
5. Server's reply to Client@192.168.1.2 dropped
6. Session to Client@192.168.1.2 times out
7. Client obtains new IP address in subnet B
8. Client@ 10.0.0.3 establishes new tunnel
9. Client@ 10.0.0.3 (re)connects to server

When a mobile device's point of network attachment changes, it may shift from using one network interface to using another (e.g., moving from Wi-Fi to GPRS adapter). The now-active interface will obtain a different IP address than was previously assigned to the now-inactive interface. After all, these are two entirely different networks, operated by independent providers or carriers, being routed through completely separate access networks.

When an interface becomes disabled, most IPSec VPN clients drop any tunnel using that interface. Any traffic sent by applications to the device's old IP address will be discarded since that address is no longer active. If that continues for a couple of minutes, any TCP sessions open to that destination will time out. When the VPN client does send traffic to Intranet servers once again, a new VPN tunnel will be established, tied to the client's new IP address. Application requests sent by the client over the new tunnel result in new TCP sessions.

Roaming Alternatives

Organizations that deploy wireless network technologies can use various solutions to enable secure roaming.

- **Layer 2 Roaming—**Stations that move between wireless across points (APs) in the same VLAN or connected to the same LAN switch can stay in the same broadcast domain when they roam. If stations roam quickly at Layer 2, a DHCP renew operation may return the same IP address. VLANs can be used to avoid subnet roaming in smaller, privately-operated networks. In networks that are too large for one VLAN, or where VLANs are used for another purpose, another method may be needed.
- **Subnet Mobility—**Stations can roam across subnets without losing VPN connectivity by tunneling through wireless gateways or switches that employ subnet mobility features. For example, Bluesocket's Secure Mobility Matrix uses Generic Routing Encapsulation (GRE) to relay traffic between Bluesocket wireless gateways when a station roams between APs connected to cooperating gateways.

This approach is best used when the entire network is under one administrative domain—for example, roaming between multiple office subnets belonging to the same company.

- **Cross-Network Roaming—**Stations that roam across networks operated by different organizations may require VPN solutions that are IP address-independent. Here again, mobility features are vendor-specific. Some are based on proprietary tunneling protocols, like NetMotion's UDP-based Internet Mobility Protocol. Others combine standard protocols like Mobile IP, Wireless Transport Layer Security (TLS), or Secure Shell with "secret sauce" for competitive differentiation. For example, ipUnplugged and Ecutel both use IPSec for security over Mobile IP for network roaming. Figure 3, (p. 46) illustrates how Mobile IP facilitates roaming between independently-addressed networks.

Mobile IP

With Mobile IP, a Mobile Node (MN) can change its network point of attachment while continuing to use the same IP address for inbound packet delivery. In addition to the Mobile Node, key architecture components are the Home Agent (HA), the Foreign Agent (FA), and Care of Address (CoA). The Home Agent is a router in the Mobile Node's original network, responsible for tracking the node's current location. The Foreign Agent is a router in the visited network, responsible for relaying forwarded traffic to Mobile Nodes. The CoA is a valid destination IP address, inside the visited network, which receives traffic for delivery to each Mobile Node.

A Mobile Node begins on the same subnet as the Home Agent. When the Mobile Node roams to a foreign subnet, it detects the presence of a Mobile IP-capable Foreign Agent by listening for announcements generated by FAs. The Mobile Node then sends a registration request to the Foreign Agent, which relays that request to the Home Agent. These messages between the Mobile Node and Home Agent are authenticated using HMAC-MD5 (a hash algorithm) and a shared secret. If the Home Agent accepts the Mobile Node's registration, the Foreign Agent and Home Agent establish a tunnel between them, using GRE, IP-in-IP, or another tunneling protocol.

Thereafter, IP packets sent to the Mobile Node are received by the Home Agent and forwarded over the tunnel, through the Foreign Agent, to the CoA. IP packets can be sent by the Mobile Node directly (i.e., using triangular routing) or routed through a reverse tunnel. Reverse tunneling increases overhead and latency, but can address certain security and topology issues—for example, when the party with whom the Mobile Node is corresponding has a private IP address.

The CoA may be separate from the Mobile Node (e.g., on a standalone FA) or may be co-located on the MN itself. When the visited network is not Mobile IP-aware, a co-located CoA is required. In that case, the Mobile Node assumes some of the responsibilities of a Foreign Agent and serves as the Mobile IP tunnel endpoint for a single CoA. Because this implementation is more network-independent, it is better suited for public network roaming, but requires a unique CoA for every Mobile Node.

Note that Mobile IP does not itself secure data exchanged between mobile devices and private networks or the application servers within them. However, Mobile IP tunnels can carry IP traffic protected by IPSec tunnels. Because a Mobile Node communicates with all correspondents using a constant home network IP address, that

Mobile IP adds overhead, but it is standardized

Different methods are used for session persistence

address can serve as an IPSec VPN tunnel endpoint without being affected by network roaming.

Mobile IP is just one approach used by some mobile VPN products. Detractors argue that Mobile IP is a high-overhead solution that adds to network infrastructure complexity. Proponents argue that Mobile IP is a standard that has undergone industry scrutiny and is well-positioned to leverage emerging IPv6 infrastructure.

Session Persistence

Whether using Mobile IP or another approach, keeping VPN tunnels active while roaming solves only part of the problem. If a mobile device roams quickly, some applications may not notice the brief interruption that occurs during this transition. But what happens when a mobile device leaves network coverage altogether? In the example shown in Figure 1, it's very likely that the mobile worker will lose network connectivity altogether at some point during his travels, probably more than once.

Depending upon the application, communication may be disrupted immediately or not at all. Store-and-forward datagram applications may be designed to retry automatically at a later time. TCP session-based applications use retransmission to survive some packet loss, but will time out within minutes when there is data to be sent. Latency-sensitive applications like voice over IP (VOIP) or streaming video can degrade noticeably if the outage exceeds tens of milliseconds.

Individual wireless technologies incorporate mechanisms to speed handoffs that occur at Layer 2. For example, 802.11i defines options for key caching and pre-authentication for use by voice over WLAN handsets that roam between 802.11 APs. Network-dependent mechanisms like this can help, but other approaches are required when roaming occurs between heterogeneous networks.

Many mobile VPNs provide some type of persistence. Methodologies differ, and their details are closely-guarded intellectual property. However, objectives are well-known and common: Create the appearance of continuous connectivity so that application sessions stay active when a mobile device temporarily becomes unreachable. For example:

- The Mobile VPN server may retain user credentials and employ single-sign-on techniques so that clients are not required to log in repeatedly when sessions are lost and must be restarted. This isn't session persistence, but it's helpful.
- The Mobile VPN server (and client) may "store and forward" session streams by queuing up packets for later delivery when reachability returns. This method works well only during relatively short outages, or with applications that transfer data mostly in one direction and which also are delay-insensitive.

FIGURE 3 Mobile IP And Inter-Network Roaming

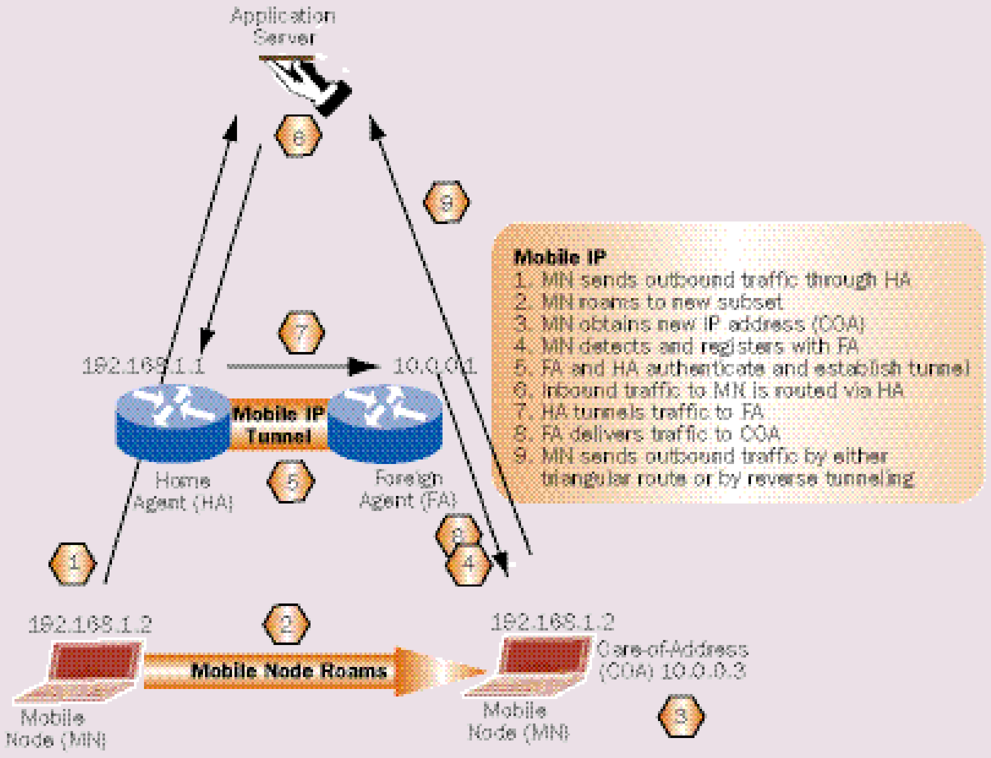

Mobile IP
1. MN sends outbound traffic through HA
2. MN roams to new subnet
3. MN obtains new IP address (COA)
4. MN detects and registers with FA
5. FA and HA authenticate and establish tunnel
6. Inbound traffic to MN is routed via HA
7. HA tunnels traffic to FA
8. FA delivers traffic to COA
9. MN sends outbound traffic by either triangular route or by reverse tunneling

- The Mobile VPN server may proxy applications, terminating TCP connections made to the client's virtual IP address, then forwarding that application data over a separate-but-related connection to the mobile device (and vice versa). Proxies better insulate the application from network congestion or disconnection, but there are still limitations, especially for applications that require continuous user interaction.

Organizations that require cross-network roaming should consider the applications and devices they intend to use, network coverage in locations frequented by mobile workers, the ability of any individual product to support those specific applications, devices, and networks, and determine the VPN server resources required to do so effectively.

Network Selection, Optimization And Policy Enforcement

When a mobile device roams between networks, it's conceivable that one network connection is lost, then another connection is found. But often, a mobile device equipped with more than one type of network adapter must choose between active adapters and associated network routes.

In the absence of other software to make this decision, IP hosts choose between available routes based on cost metrics. For example, the metric for a GRPS adapter's default route could be hard-coded to exceed the metric for a Wi-Fi adapter's default route, biasing traffic towards Wi-Fi whenever that adapter is active.

Mobile VPNs often automate this process, choosing the "best available" network connection based on factors like adapter priority, data rate, recent experience and/or configured policy. For example, connections to the corporate WLAN may be fast and free, but connections to commercial hot spot WLANs may incur hourly charges. Mobile users with unlimited-use 3G services like EDGE or 1xEV-DO may prefer those networks over for-pay Wi-Fi hotspots, even though they are slower. IT organizations may wish to centrally-configure network selection policies used when roaming, to strike the best balance between performance and cost.

Some mobile VPNs can use policies to control the applications available over a given network connection. For example, the Honolulu Police Department recently equipped its cruisers with laptops that run Ecutel Viatores mobile VPN software. The Honolulu department's policies allow very selective application use over GPRS, while permitting further applications like Web browsing over Wi-Fi.

Depending upon the product, mobile VPN policies may also be able to dictate security measures or bandwidth optimizations required over any given network. For example, IPSec tunneling might be disabled when connected via Ethernet to the corporate LAN, but enabled when connected to that same LAN via Wi-Fi. Data compression may be applied before encryption when connected via wireless WAN.

In addition to compression, some mobile VPNs include bandwidth optimizations that improve application performance over low-speed and/or high-latency networks. For example, NetMotion Mobility employs a variety of techniques to reduce TCP "chattiness," including selective and bundled TCP acknowledgements, reduced retransmissions, and fragmentation optimizations.

Next-gen smart phones and VoWLAN will drive production of "multilingual" devices

Beyond Roaming

In this article, we've described how intra- and inter-network wireless roaming works, considered alternative solutions to facilitate roaming and examined just a few of the features found in mobile VPNs that make inter-network roaming more effective.

Initiatives that incorporate roaming between diverse wireless networks require many additional ingredients that we haven't discussed here. For example, laptops with both WLAN and Ethernet adapters have become common, but devices with both WLAN and WWAN (wireless WAN) adapters are far less common. Next-generation smartphones and VoWLAN will drive the production of "multilingual" wireless devices.

Although some carriers offer both WWAN and WLAN data services, they often treat them as separate services. For example, users may be required to log into a WLAN hotspot portal, but authenticate by SIM to the WWAN. Consortiums like iPass do exist, but roaming between Wi-Fi hotspots today is far less seamless than roaming between wireless voice networks. After all, inter-network roaming is just the plumbing; inter-carrier roaming requires business agreements and accounting infrastructure.

Ultimately, wireless users will demand persistent sessions with LAN-quality performance. If wireless coverage is too slow or too spotty, even the best network roaming solution won't be enough to keep workers happy and productive. Network roaming can help by making the best use of available networks, but carriers must still deliver faster, broader, more reliable wireless network coverage□

Companies Mentioned In This Article

AppGate (www.appgate.com)
Bluesocket (www.bluesocket.com)
Columbitech (www.columbitech.com)
Ecutel (www.ecutel.com)
ipUnplugged (www.ipunplugged.com)
NetMotion Wireless (www.netmotionwireless.com)
Trapeze Networks (www.trapezenetworks.com)