



Building Manageable, Scalable, Seamless WLANs:

Bluesocket's Secure Mobility[®] Matrix



Wireless LANs are breaking through the physical boundaries associated with traditional networks, providing convenient non-stop access to business services and data. But, in a large multi-site deployment, enabling mobility in a manner that's both friendly and secure can be a challenge.

This paper introduces Bluesocket's patent-pending Secure Mobility® MatriX, an innovative distributed architecture that offers seamless secure mobility and automatic policy synchronization without depending on a central point or requiring change to existing networks, devices, or user behavior.

Understanding The Challenges

Wireless LANs deliver network access to mobile users, reaching company Intranets and the public Internet from hard-to-wire locations. WLANs can improve business productivity by making data more readily available, streamlining business processes, and speeding network deployment.

To fully realize these benefits, desktops, laptops, PDAs, handsets, and other devices using 802.11 Access Points (APs) must communicate without interruption by transient fluctuations in radio signal. Users that are mobile must also be able to stay connected when moving throughout floors, buildings, and campuses. Both scenarios are referred to as *roaming*.

Roaming within a single subnet with uniform radio coverage can be largely transparent, but roaming between subnets or through radio gaps change the device's network reachability and address. Those changes interrupt applications and VPN tunnels, frustrating users and slashing business productivity. Bluesocket's Secure Mobility MatriX overcomes this by enabling seamless mobility for users as they roam between APs, across subnets, and through short "dead spots."

Mobile users also require consistent network access, no matter where they attach to the corporate WLAN. As network size and complexity increases, maintaining multiple WLAN configurations gets tougher. Re-entering the same security policies at each site becomes both impractical and error-prone. Discrepancies eventually cause security gaps and user frustration when permissions vary

between floors, buildings, or sites. Bluesocket's Secure Mobility MatriX overcomes this by automatically replicating configured policies to every gateway in the network.

Finally, as wireless access becomes pervasive in corporate networks, network utilization and dependence will grow. Businesses can't afford to shift mission-critical activities onto wireless unless those networks can provide robust, non-stop access under heavy loads. This challenge must be addressed at every point, from the APs that provide radio coverage to the gateways that control wireless access. Bluesocket's Secure Mobility MatriX can flexibly distribute or concentrate workload across any combination of Wireless Gateways, including load-balanced clusters.

Introducing The Secure Mobility MatriX

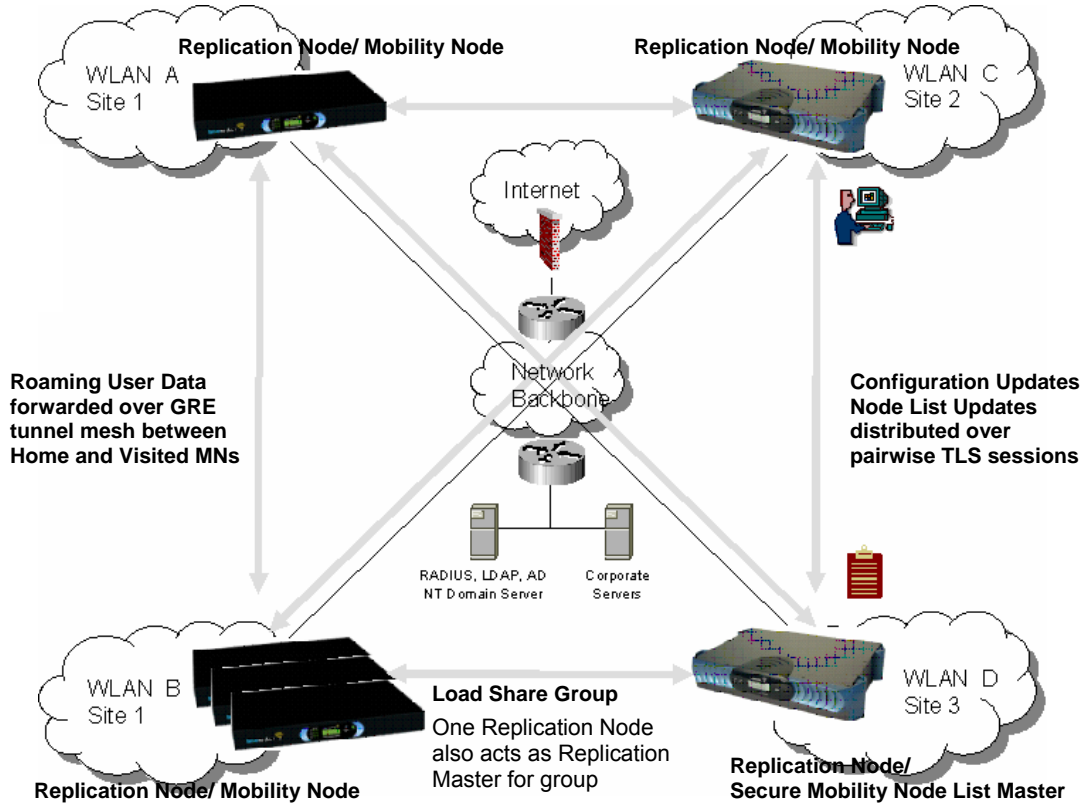
A Secure Mobility MatriX is a highly manageable, scalable, seamless multi-WLAN network, built upon a foundation of Bluesocket Wireless Gateways (WGs) that together provide:

- single-point administration for the entire network, no matter how large or small;
- automated configuration replication and synchronization;
- load sharing across gateways when needed for scalability; and
- seamless secure mobility that is transparent to users, applications, and VPN tunnels.

A Secure Mobility Matrix can include up to 50 Wireless Gateways, communicating across any LAN or WAN backbone in a peer-to-peer fashion (see Figure 1). Matrixed gateways may be distributed across floors, buildings,

campuses or countries, so long as Generic Routing Encapsulation (GRE) and HTTP over Transport Layer Security (TLS) can be switched or routed between all matrixed gateways.

Figure 1:
Secure Mobility MatrixX
Architecture



Configuration Replication

To enable centralized configuration, with automated policy distribution and synchronization, any gateway can be designated as the *Replication Master*. The Replication Master is configured with a list of addresses corresponding to every other gateway and a replication key. Every other gateway (*Replication Node*) is configured with the address of the Master and that same key.

Every Replication Node downloads a configuration snapshot from the Replication Master. Thereafter, except for a few Node-specific parameters, all policy changes are

made in just one place: the Replication Master. Whenever changes are made, the Replication Master relays changed values to every Replication Node, sending XML over an encrypted TLS session. This communication is immediate and brief, ensuring continuous policy synchronization with minimal network overhead. Gateway certificates, the replication key, RC4 encryption, and SHA message authentication protect the privacy and integrity of these messages.

Replication Nodes query the Replication Master hourly or upon reboot for a list of recent updates. If a Node should miss an update due to routine maintenance or

backbone network outage, it automatically retrieves missed updates to resynchronize itself. Replication status is always visible at the Master, making it easy for the administrator to check whether all Nodes are synchronized or to identify an offline gateway or failed update.

In the unlikely event that the Replication Master goes offline, the entire matrix continues to operate normally — only configuration changes must be deferred. If desired, a high availability pair can be designated as the Master, enabling fast fail-over to a hot-standby.

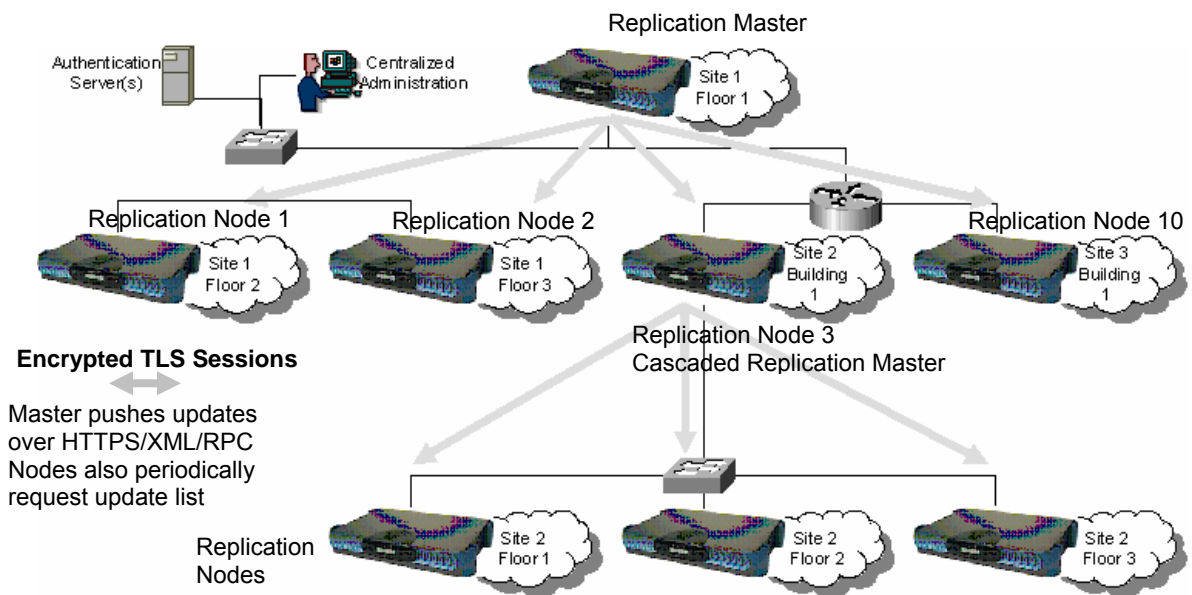
Cascading

In large multi-site networks, high availability and scalability are essential. Bluesocket Wireless Gateways range from 50 Mbps to 1 Gbps of throughput (20 to 400 Mbps encrypted). A MatriX deployment can include any mix of gateways, chosen to meet the capacity demands of each floor, building, or site.

All gateway models are managed in exactly the same manner and can support the same kinds of policies. Configuration Replication makes it possible to administer these policies through a single point, with automated distribution and synchronization to up to 50 gateways. However, to minimize the overhead generated by replication traffic in larger networks, *Cascaded Replication Masters* can be used.

As shown in Figure 2, cascading retains all of the benefits of replication while distributing the associated workload and optimizing network use. All changes are still made at one location: the Replication Master at the root of the matrix. But, in a cascaded deployment, selected gateways serve as both a Replication Node and Replication Master. Those gateways receive policy updates from an upstream Replication Master and redistribute them to downstream Replication Nodes. Cascading can reduce traffic sent over WAN links and prevent any single resource (gateway, firewall, router, link) from being over-burdened with update or status check messages.

Figure 2:
Scalable
Replication
in a
Secure
Mobility
MatriX



Load Sharing

Bluesocket WGs can be added as needed to support network growth, leveraging configuration replication to effortlessly-administer each new gateway. When the workload borne by any single managed subnet exceeds the capacity of one gateway, a Load Sharing Group (LSG) can be deployed.

In a Load Sharing Group, one gateway is designated as the *Load Sharing Master*. Once an LSG is configured, *Load Sharing Nodes* support a common managed subnet by distributing new Active Connections across the group, reflecting weights assigned to each Node. If any Load Sharing Node goes off-line, another Node automatically takes responsibility for the lost Node's managed IP address. This fail-over method is fully automated, fast, and seamless, ensuring high-availability of gateway services in high-density managed subnets.

Secure Mobility

To permit seamless roaming between APs and subnets in a campus environment using a Bluesocket MatriX, any gateway can be designated as the *Secure Mobility Node List Master*. The Secure Mobility Node List Master is configured with a list of all gateways in the matrix and a secure mobility key. Every gateway (*Secure Mobility Node*) is configured with the address of the Secure Mobility Node List Master and key. Replication and Mobility can be used together or separately, with the same or different gateway serving as master; these features are completely independent.

The Secure Mobility Node List Master assigns a unique virtual IP address to each Secure Mobility Node, creating a full mesh of GRE tunnels for logical one-hop routing between every Node pair. These tunnels are used to relay roaming user traffic directly between Secure Mobility Nodes, without depending on the Secure Mobility Node List Master (see Seamless Roaming, below).

The Secure Mobility Node List Master distributes the list to every other Node in the matrix. If a Secure Mobility Node is added or

removed, the Secure Mobility Node List Master immediately propagates that update to all remaining Nodes. Secure Mobility Nodes also query the Master hourly or immediately after reboot to ensure list synchronization. Mobility status, including traffic counts between each pair of Nodes, is tracked by the Secure Mobility Node List Master. This makes it easy for the administrator to eyeball roaming traffic flows and mobility tunnel status, at any time, from any point in the matrix.

Whenever a user joins or leaves a WLAN within the matrix, the gateway on that managed subnet informs every other Secure Mobility Node, passing an Active Connection Status update over an authenticated, encrypted TLS session. This method ensures that all Secure Mobility Nodes remain continuously aware of connection state, letting users roam seamlessly across subnets, anywhere in the matrix, while retaining the same IP address, authenticated role, and network access rights.

Seamless Roaming Within A Secure Mobility MatriX

To fully understand and appreciate the benefits of secure mobility, let's take a closer look at how wireless roaming works, with and without the assistance of Bluesocket Wireless Gateways.

Roaming Without Secure Mobility

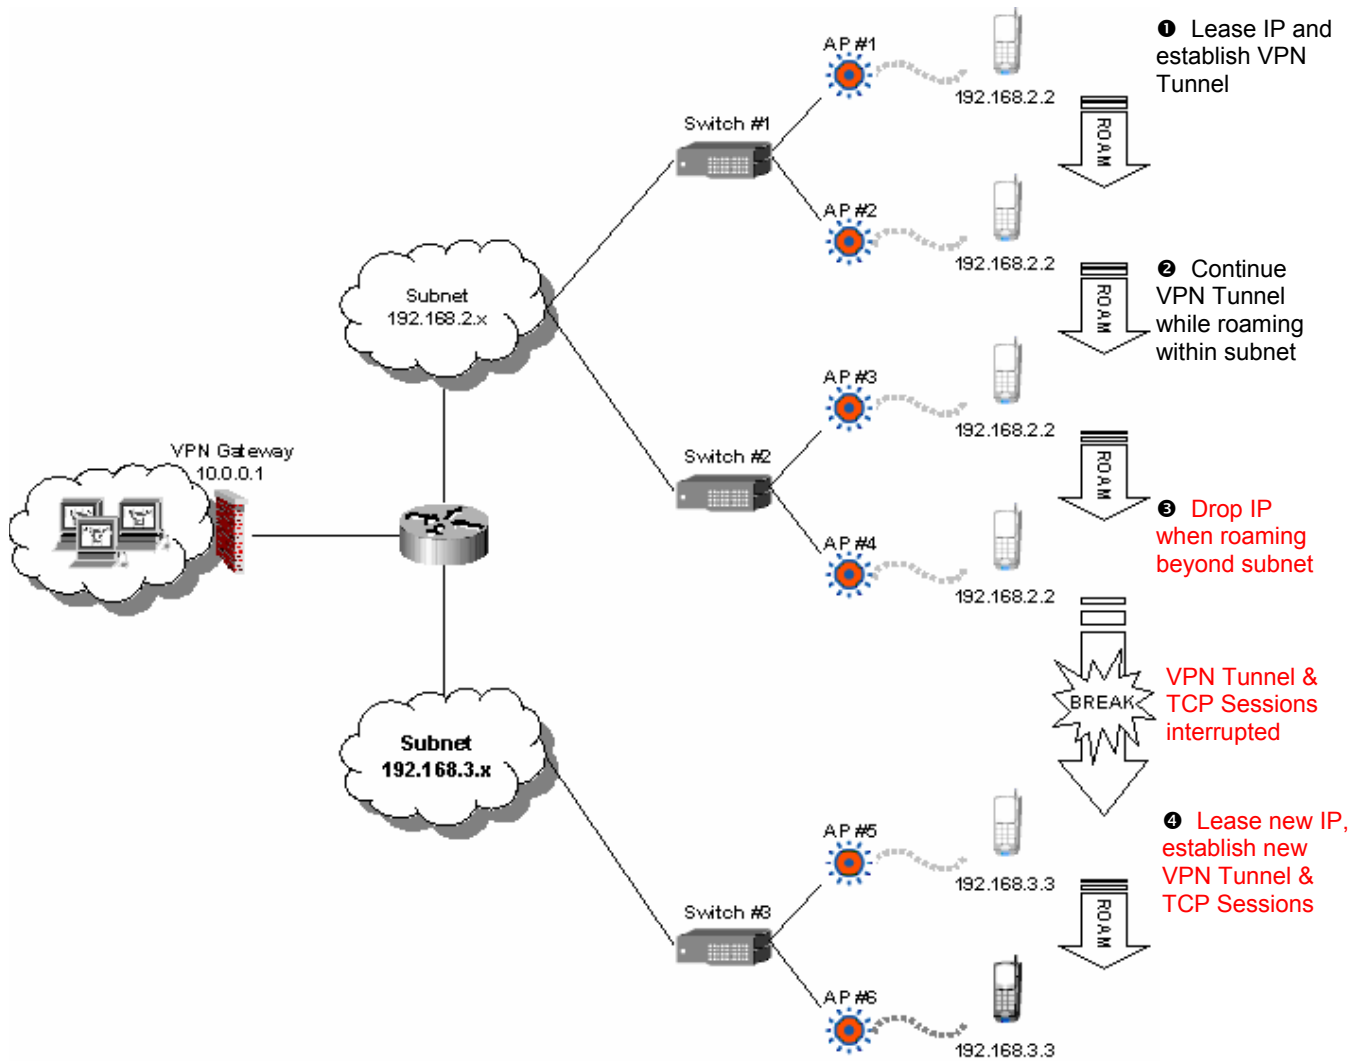
The 802.11 standard uses roaming to adapt to changes in physical location or environmental conditions. Every device with an 802.11 network adapter constantly re-evaluates the best available AP, based on received signal strength, error rate, and other criteria. Whenever a better AP is detected, a device may (re)associate with a new AP. This is known as "roaming."

During a roam, link interface status changes at least briefly. On most operating systems, this status change also triggers a DHCP renew request. If the user roams quickly within one LAN broadcast domain, there may be no adverse impact. For example, a user

roaming between APs connected to the same hub or switch or VLAN, within a ubiquitous radio coverage area, may retain its IP

address, VPN tunnel, and TCP sessions (Figure 3, top).

Figure 3:
Wireless
Roaming
Without
Mobility



However, there are many other scenarios where roaming causes significant handoff delay or change in IP address (Figure 3, bottom). For example:

- There may be gaps in radio coverage or transient loss of signal — for example, when the user enters an elevator or walks between buildings.
- The user may roam from one type of wireless network to another (e.g., 802.11b to 802.11a).
- Handoff delay may be increased by 802.1X re-authentication or by re-keying in WLANs that use dynamic WEP, TKIP, or AES link layer encryption.
- The WLAN may be too large to use a single switch, due to physical limitations or geographic distribution, leading to roaming between switches.
- The WLAN may be too busy to use a single physical or virtual LAN, due to congestion and collisions, leading to roaming between subnets.

- VLANs may already be used for another business purpose, like segregating Intranet traffic to reflect group access rights, preventing VLAN use to avoid subnet roaming.

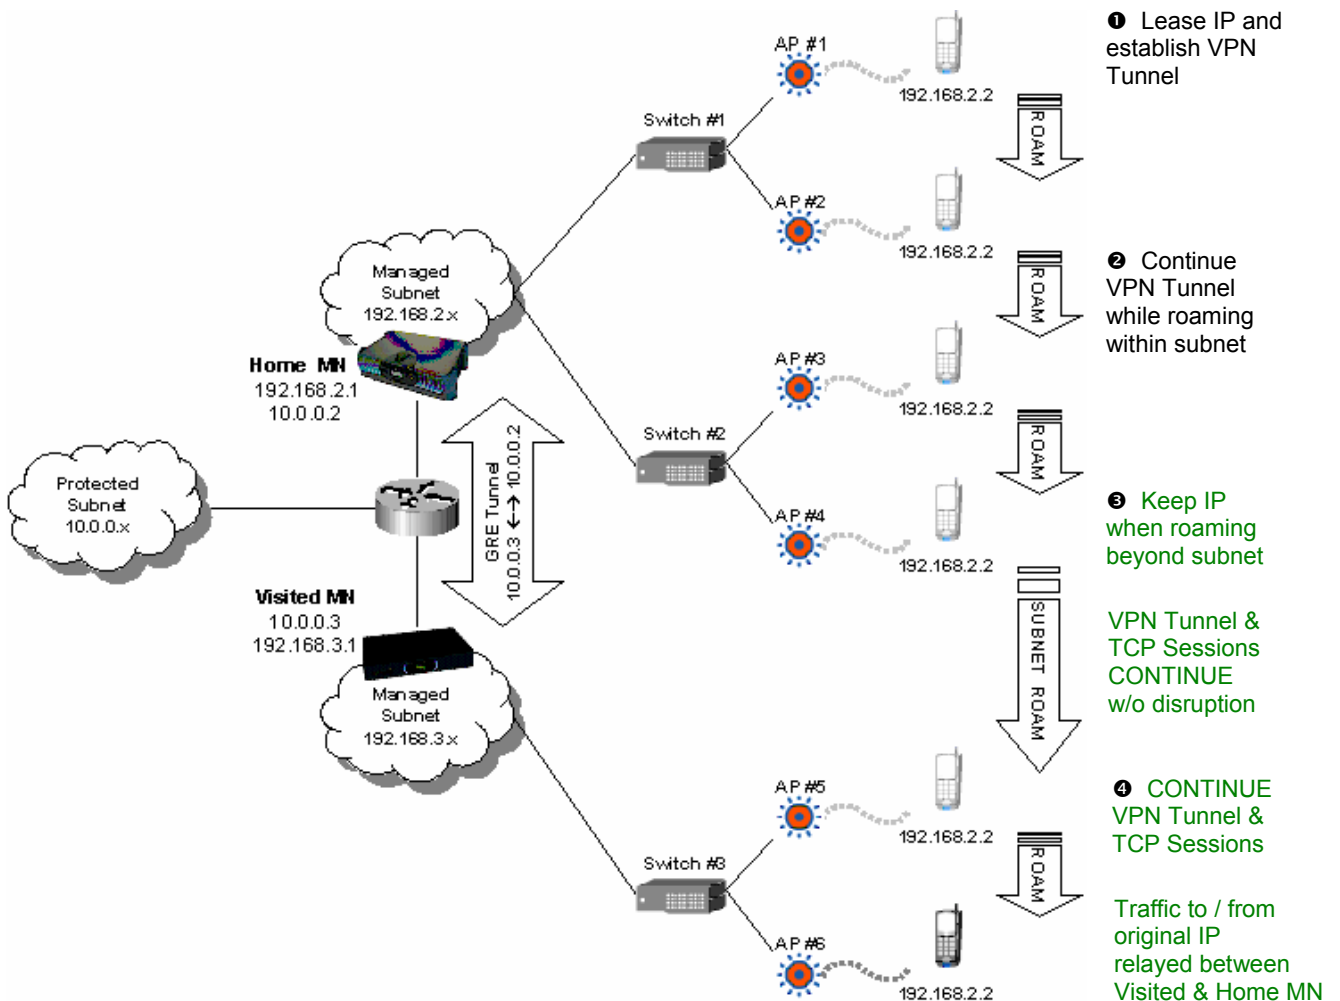
As shown in Figure 3, users that cannot keep the same address when roaming lose all active TCP sessions and VPN tunnels. Inbound traffic routed to the user's former subnet will be discarded, since the user is no longer reachable at its old address. Outbound traffic sent by the user from its new address will (eventually) cause new TCP sessions and VPN tunnels to be established.

As a result, the user may be required to re-authenticate to the domain, VPN, or applications, in some cases responding to

prompts or even rebooting. A roam like this can easily take minutes, causing user frustration and loss of productivity. Transactional applications like email can be bearable, but lengthy file transfers and latency-sensitive voice/video applications can be difficult or impossible.

Certain measures can speed link layer roaming between APs, including proprietary Inter-AP Protocols and emerging 802.11i Key Caching and Pre-Authentication options. Today, these measures are available only in single-vendor WLANs, used with applications like VoIP that require handoff within milliseconds. However, none of these link layer measures can support transparent roaming beyond a single LAN broadcast domain, across subnets.

Figure 4:
Wireless
Roaming
Within A
Secure
Mobility
MatriX



How Secure Mobility Helps

Bluesocket Wireless Gateways use an innovative, patent-pending approach to provide seamless roaming throughout a Secure Mobility Matrix. To maximize interoperability, compatibility, and ease of use, this approach requires no software on roaming devices, and no change to the managed network. As Figure 4 illustrates, a wireless device operating within a matrix keeps the same IP address, whether roaming between APs or across subnets. Existing VPN tunnels and application sessions are no longer disrupted by roaming.


To accomplish this, every wireless device is initially given an IP address in the managed subnet of the nearest Wireless Gateway, referred to as the Home Mobility Node (*Home MN*). It does not matter whether that device is wireless or wired, or whether the address is configured statically or leased through DHCP. That initial IP address will “stick” to the device, wherever the user roams, until the device is turned off, the address is explicitly released/reconfigured, or a connection check time limit is exceeded. MNs can be connected

to the backbone network in any fashion, but must have unique managed subnets.

With Secure Mobility, MNs share connection status with other matrixed MNs. Whenever a new device is detected, the Home MN adds an entry to the Active Connections List, identified by its MAC address, IP address, and authenticated role (if any). All traffic to and from that device is routed through the Home MN's protected and managed interfaces.

If that device roams to a subnet belonging to another gateway (*Visited MN*), traffic to and from the roaming device is automatically relayed between Home and Visited MNs. Virtual interfaces used as GRE endpoints are predefined in the Secure Mobility Node List, so all necessary routes are already in place when roaming occurs. This approach keeps traffic flowing during subnet roaming, adding just milliseconds to the handoff incurred at the link layer. Because proprietary GRE tunnels are only used when there is roaming data to send, there is no overhead in the absence of roaming traffic. For efficiency, a single GRE tunnel is used between each MN pair.

Figure 5:
Outgoing
Traffic
Between
Mobility
Nodes



A Peek Under The Hood

Figure 5 offers a more detailed example of what WLAN administrators can expect to see when roaming occurs within a Secure Mobility Matrix. Here, User B associates with an AP on the 192.168.2.x subnet and obtains IP address 192.168.2.2 from its Home MN, 192.168.2.1. All traffic sent and received by User B is routed through the Home MN's protected interface, 10.0.0.2. When User B roams away to an AP on the 192.168.3.x subnet, the Visited MN detects the presence of this device. By searching the Active Connections List for this device's MAC and IP address, the Visited MN finds that User B belongs to the Home MN at 10.0.0.2.

When the roaming device next sends outbound data, those packets will flow through the Visited MN's managed interface. There, pre-defined routes will automatically relay those packets, through the Visited MN's virtual interface, over the GRE tunnel, to that user's Home MN. For example, as shown in Figure 5:

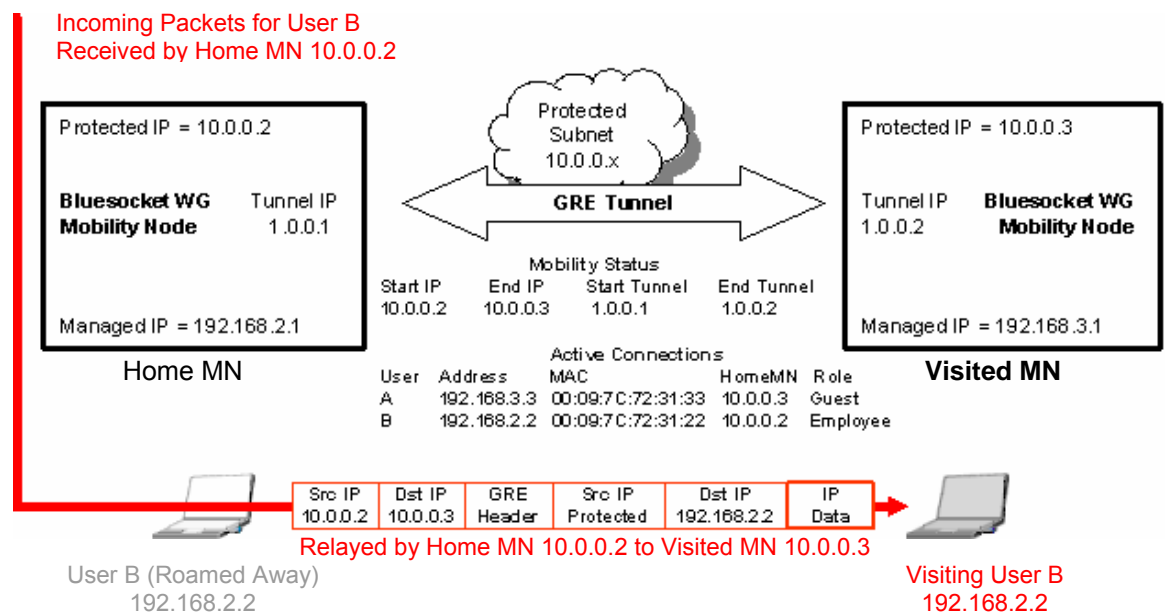
1. The Visited MN detects outbound traffic from User B (192.168.2.2), using pre-

defined routes to direct the packet through a virtual interface leading to the Home MN (10.0.0.2).

2. For each IP packet sent by User B, the Visited MN appends IP/GRE headers to tunnel the packet from the Visited MN (10.0.0.3) to the Home MN (10.0.0.2).
3. The encapsulated packet is sent across any backbone LAN or network in the normal fashion; GRE tunneling is completely transparent to any other router along that path.
4. Upon receipt, the Home MN strips the outer IP/GRE header and routes User B's packet to the inner destination IP address, enforcing policies defined for User B's authenticated role.

Thereafter, the Home MN automatically routes any inbound traffic received for User B to the virtual interface of the Visited MN. The Home MN encapsulates the IP packet for User B as described above and relays it in the opposite direction. Upon receipt, the Visited MN strips the outer IP/GRE header and forwards the packet through its managed interface to User B, associated with an AP located somewhere on this subnet. This inbound flow is illustrated in Figure 6.

Figure 6:
Incoming
Traffic
Between
Mobility
Nodes



No matter where the user roams, role-based policies of the Home MN (not the Visited MN) continue to be enforced, exactly as though roaming had never occurred. For example, the user may send either cleartext or encrypted data over a PPTP, L2TP, or IPsec tunnel. There are no policy limitations imposed by Secure Mobility — any user can roam throughout the matrix. The only controls enforced by the Visiting MN are those defined by the administrator for the user's assigned role. This approach provides the user with a consistent network experience throughout the matrix, and also ensures consistent security policy enforcement.

In addition, because the device retains the same IP address and authentication status, secure subnet roaming is completely transparent to applications, VPN clients, and end users. There is no need to reestablish a TCP session or VPN tunnel or reauthenticate the user, because the device remains continuously reachable through the same externally-visible addresses and routes.

Roaming user connections can even remain active on the WG during loss of radio connectivity, bounded by application timeouts and configurable status check intervals. Specifically, each WG tracks connection status for every device in its managed subnet. A device is considered active if it sent traffic during the last check interval. In the absence of user traffic, the WG sends standard ICMP Ping queries to solicit a response. A device that transmits nothing during N check intervals is deleted from the Active Connections List. If that device should later reappear, it must lease a

new IP address and authenticate to the nearest gateway as a brand new Active Connection.

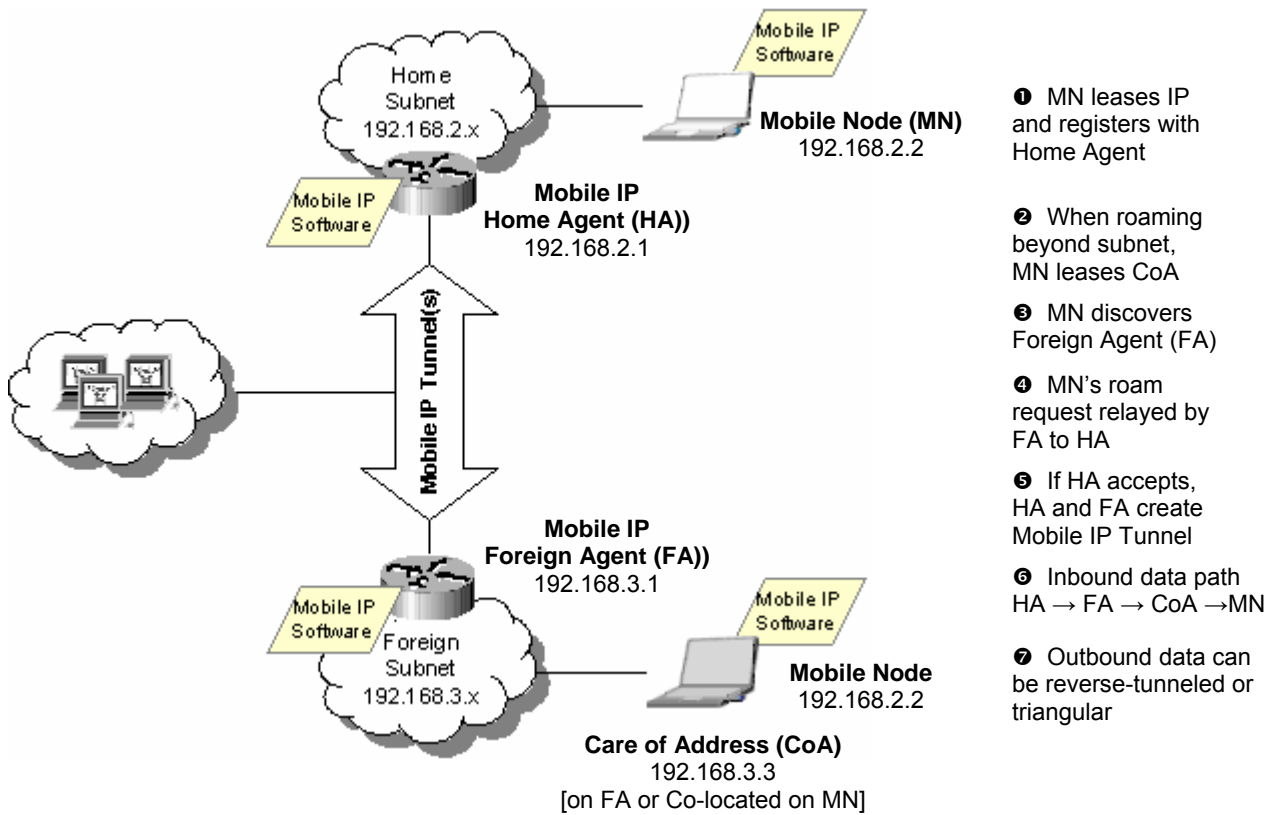
For example, suppose that User B in Figure 6 roams out of WLAN coverage for 100 seconds. If the check interval is 60 seconds and retry limit is 3, User B has up to 3 minutes to associate with a new AP and send data or reply to status checks. Many TCP-based applications and VPN tunnels can survive a transient period of packet loss like this. Latency-sensitive voice/video applications may be degraded during shorter outages, but User B still benefits from avoiding reauthentication, resuming communication faster and easier. This approach requires nothing more than standard TCP/IP on each device, while adding minimal overhead to an active subnet.

The Competitive Edge

Bluesocket's Secure Mobility MatriX architecture was designed from the ground up to avoid short-comings and pitfalls associated with alternative subnet mobility approaches.

For example, RFC 3344 defines an extension to IPv4 to provide transparent routing of IP packets to mobile nodes. Like Bluesocket's approach, a device using Mobile IP can change its point of network attachment while continuing to operate with same IP address. Unlike Bluesocket's approach, Mobile IP requires at least two and sometimes three Mobile IP-aware devices: a Home Agent (HA), a Foreign Agent (FA), and a Mobile Node (MN), as shown in Figure 7.

Figure 7:
Wireless
Roaming
With
Mobile IP



A Mobile IP Home Agent is a router in the Mobile Node's original network, responsible for tracking current location. A Mobile IP Foreign Agent is a router in the visited network, responsible for relaying forwarded traffic to the device. The Care of Address (CoA) is an IP address, inside the visited network, which receives traffic for delivery to the roaming device. The CoA may belong to the FA or be co-located on the roaming device.

When a Mobile IP device roams to a foreign subnet, it detects the presence of a Mobile IP FA by listening for announcements. As shown in Figure 7, the roaming device sends a registration request to the FA, which relays that request to a Mobile IP HA. If the HA accepts that request, the FA and HA establish a tunnel between them, using a tunneling protocol like GRE or IP-in-IP. Thereafter, inbound packets received by the HA are forwarded over the tunnel, through the FA, to the CoA. Outbound packets can be sent by

the device directly (triangular routing) or routed through a reverse tunnel to the HA.

Bluesocket's approach incorporates the benefits associated with Mobile IP, while avoiding Mobile IP's limitations. By predefining routes, the MatriX significantly shortens roam time as compared to Mobile IP. Moreover, by requiring only standard TCP/IP on roaming devices, the MatriX is optimized for interoperability and painless integration. In contrast, Mobile IP requires upgrading network routers and roaming devices to run Mobile IP software. Even when a co-located CoA is used, the roaming device requires client software installation and configuration. This may be labor-intensive but possible for laptops that you own. But it is impractical or impossible for guest devices and specialized wireless devices like VoIP handsets or point of sale terminals.

The MatriX also avoids well-known weaknesses associated with other hierarchical approaches to solving the subnet mobility problem. MatriX

protocols have been carefully designed to avoid creating traffic bottlenecks, unnecessary overhead, or single point of failure. In hierarchical solutions, a central controller is often required to establish new connections, to re-route traffic when devices roam, or even to act as a relay for tunneled data. Controller failure may bring the entire network to a halt, and resources in or near the controller are critical for both performance and availability.

In contrast, the MatriX distributes workload among participating gateways in a highly-distributed peer-to-peer manner. Although Replication, Mobility, and Load Sharing Masters are responsible for certain control tasks, they are never required as part of the data path. Every Node is capable of continuing business-as-usual on its own, establishing new connections, authentication users, and passing traffic.

Finally, the MatriX avoids imposing any requirements on devices or the surrounding network:

- Some alternatives require specialized APs or switches to obtain mobility features; the MatriX provides the same benefits, independent of AP brand or model. Best-of-breed products can be used, both in your WLAN and in your network backbone, without compromising mobility.
- Some alternatives require VLANs to switch roaming traffic across the backbone or to

keep roaming devices in the same logical subnet; the MatriX works seamlessly within your existing VLAN design. Gateways can be configured to pass or assign VLAN tags, but Configuration Replication, Secure Mobility, and Load Sharing do not depend upon use of VLANs.

- Some alternatives require the use of specific VPN clients, or cannot work in combination with industry standard VPN protocols like PPTP, L2TP, and IPsec. The MatriX approach provides fast, seamless roaming for both secure and cleartext traffic, without requiring any specific VPN gateway/appliance or protocol.
- Some alternatives impose constraints on physical or network addressing topology. The MatriX adjusts to your existing topology instead of requiring your topology to adjust to it. Features like Cascaded Replication and Load Sharing bring the benefits of Bluesocket's Secure Mobility MatriX to even very large, very dense, or very distributed networks.

In this paper, we have shown why using voice and data applications over wireless effectively and efficiently requires subnet roaming that is both secure and transparent to end users. Bluesocket's Secure Mobility MatriX can help any organization realize this vision by deploying a solution that's clientless, network topology and device independent, scalable, and easy to manage.

Prepared by
Lisa Phifer
Core Competence Inc.,
for Bluesocket, Inc.



www.bluesocket.com