WhitePaper

Security: Multiple Lines of Defense

Security: Multiple Lines Of Defense

When all connections relied on wires, accessing a company's network required access to its building. Attackers had to find their way through locked doors, around security guards and past crowds of people who may not notice a stranger in their midst. Once inside the building, they had to find an Ethernet port or a PC and hope that no electronic access controls are in place. Because of the difficulty of gaining physical access, attackers often resorted to social engineering: lying to innocent authorized employees in order to gain some access electronically.

However, wireless networks change this paradigm and allow attackers to gain access to the network from outside the physical perimeter. Attackers have no need for brute force or social tricks, and wireless networks can now be threatened from nearly any location that is within radio range of an access point.

To understand the unique vulnerabilities of wireless networks, it is necessary to analyze threats from the enterprise perspective. Wireless security vulnerabilities can be divided into four layers , depending on the attacker's level of access to the network.

Perimeter Security: Attackers without physical access to the facilities

Perimeter security usually involves physical barriers such as walls, ceilings, doors or distance. Traditional Ethernet LANs relied entirely on perimeter security, assuming that malicious users simply had no way to get into a building. Conversely, this kind of security is almost absent in wireless LANs. Except for rare attempts at limiting propagation by playing with radio power level, perimeter defense has been largely and conspicuously unaddressed. This has left many wireless networks open to "parking lot attacks" and "wardriving".

Connection Security: Attacker with physical access but no login credentials

Once an attacker has entered a building, he or she needs to connect to the network itself. With Ethernet, this can sometimes be accomplished simply by plugging in a cable, though many networks use VLAN and switch-based protections to limit exposure. On the other hand, wireless networks have no choice but to expose the connection to anyone and any number of devices in range.

Network Security: Attacker with physical access and login credentials

Wired and wireless networks both go to great lengths to partition the network, firewalling different categories of users and filtering traffic. Additionally, secure wireless networks require 802.1X authorization, a practice so successful that many organizations are now moving towards using it for wired networks too. However, authorization and authentication alone do not protect against authorized users with malicious intent, intense curiosity or a compromised client device. Network security must focus on the entirety of what a user or device is allowed to do even after gaining access.

Remote Security: Piggybacking Attackers

Because employees need to access networks from anywhere, the frontier of defense has expanded to encompass far more than it used to. Employees routinely use corporate laptops outside of the office, exposing each machine and employee to unknown numbers of attackers. Users may not be aware of the problems that lie on networks outside of the corporate domain, but cannot be cut off from these other networks without also eliminating the ability to perform useful work. Thus, wireless security must be pushed to areas outside the campus.

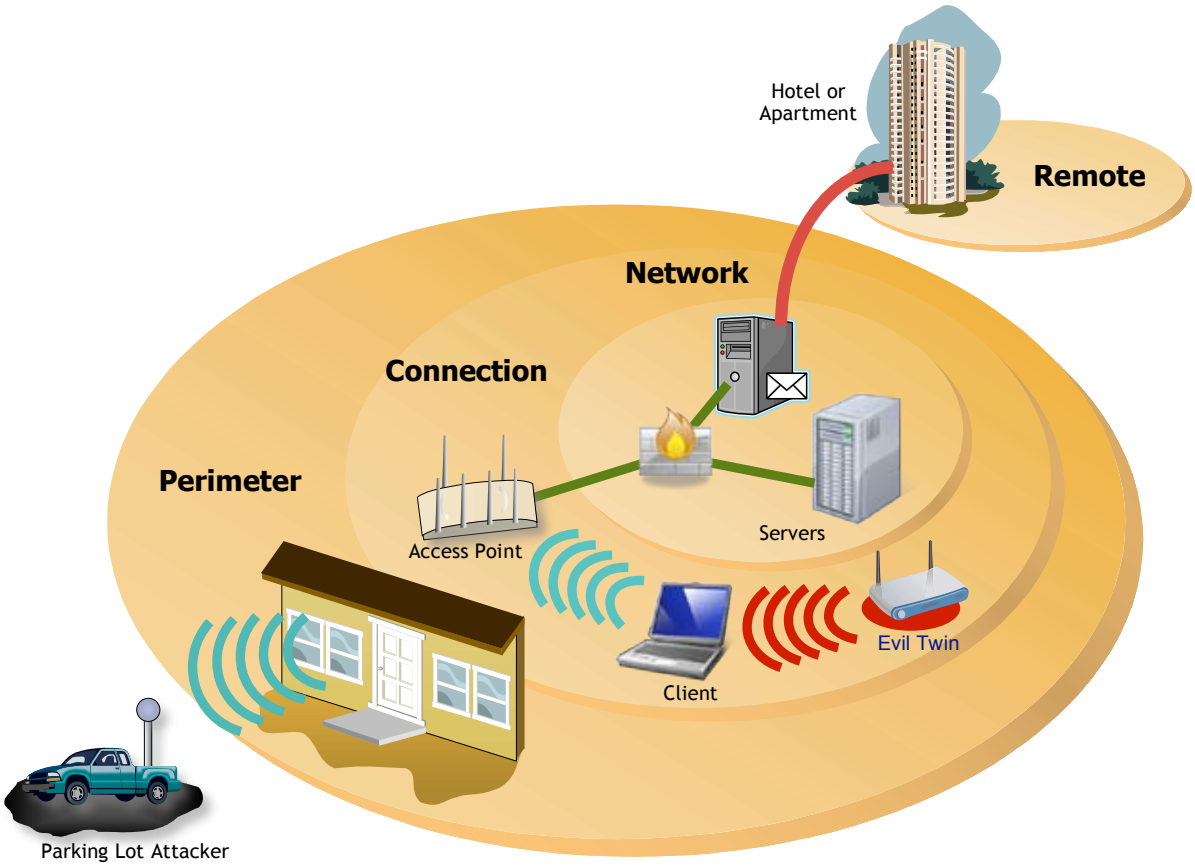

Figure 1: Four lines of defense for wireless networks. This is the same as Joe's figure on the site.

These four layers of active threats are not consistently addressed by most product offerings, but all four must be addressed by security-conscious organizations. Each represents a separate line of defense, all of which an intruder will need to penetrate in order to access a well-protected network.

1: Securing The Physical Perimeter

Perimeter security protects a network against attackers who don't have physical access to a building. It is inherent in the design of Ethernet LANs, which rely on the same physical security systems that protect other assets: walls, keys and guards. However, it is usually absent from wireless networks. Radio waves don't respect borders or boundaries.

Until now, there been no cost-effective solution to this problem. However, with the advent of perimeter wireless security, it is possible to limit wireless networks' exposure proactively. Effective perimeter security prevents attackers who are outside a network from obtaining data held within it.

Outsider Threats: Parking Lot Attacks and War Driving

The most common wireless security breaches involve attackers who never enter a building, simply sitting within radio range. Attacks may involve actively trying to connect to Wi-Fi networks, but the most dangerous are those in which an attacker passively eavesdrops on traffic as they are completely undetectable.

Most Wi-Fi networks now use cryptography to prevent eavesdroppers from understanding the traffic they sniff, with all recent hardware including secure AES encryption. However, many legacy devices do not yet support this as the effective life of some handhelds is much greater then that of older WLAN infrastructure equipment. To work with these devices, networks must allow unencrypted access or support older and flawed technologies such as WEP. Many other networks that could theoretically be made secure still have holes due to user error or unpatched vulnerabilities in hardware or software.

One way to prevent such attacks is to use a physical barrier such as RF-blocking paint or wallpaper to create what is essentially a Faraday cage around the building. Cost and complexity limit such protection to military or government installations—and even then, only to a few selected cases. Worse, this sort of solution is not selective: it blocks all wireless LANs and almost all cellular signals as well, preventing signals from both going in and going out. This is impractical for most users. Meru Networks' new RF Barrier product uses RF technologies to contain the wireless network within a physical boundary, providing the first cost-effective solution.

Perimeter security should not be confused with location-based wireless security (geofencing), which can limit a client's ability to connect to the network from a specific region. This is clearly a useful feature, but it only prevents active intrusion, not passive eavesdropping.

Meru Networks RF Barrier

The RF Barrier acts like an invisible shield around a building, actively scrambling selected Wi-Fi radio transmissions that originate from the organization's own wireless network. Unlike the radio jammers used in some high security military installations, RF Barrier blocks only selected 802.11 signals that originate from one particular wireless LAN, allowing traffic from cell phones and even other wireless LANs to pass unhindered.

When RF Barrier is activated, outsiders won't even be able to detect that a Meru wireless network exists, let alone be able to sniff traffic. Even if an attacker knows about the network through other means and has managed to obtain a username and password, attempts to connect will be futile as any responses from the network will be unable to pass beyond the perimeter.

Unique Technology

RF Barrier uses directional antennas mounted on each side of a wall, inside and outside, selectively voiding signals as they exit a building. It operates on a per-channel basis, inspecting the metadata in the header of every Wi-Fi frame. If it detects anything that matches the internal network, it transmits a counter-signal, a process that happens in real time. Frames from the internal network simply disappear when viewed from outside a building, as shown in the figure below.

Because RF Barrier uses directional antennas, it has no impact on signals within a building or on signals from other networks. Internal clients continue to connect as normal and internal APs continue to serve them at full speed. The only effect on the network is a potential improvement in performance for users within a building, as APs no longer have to waste airtime denying connections to users outside the perimeter.

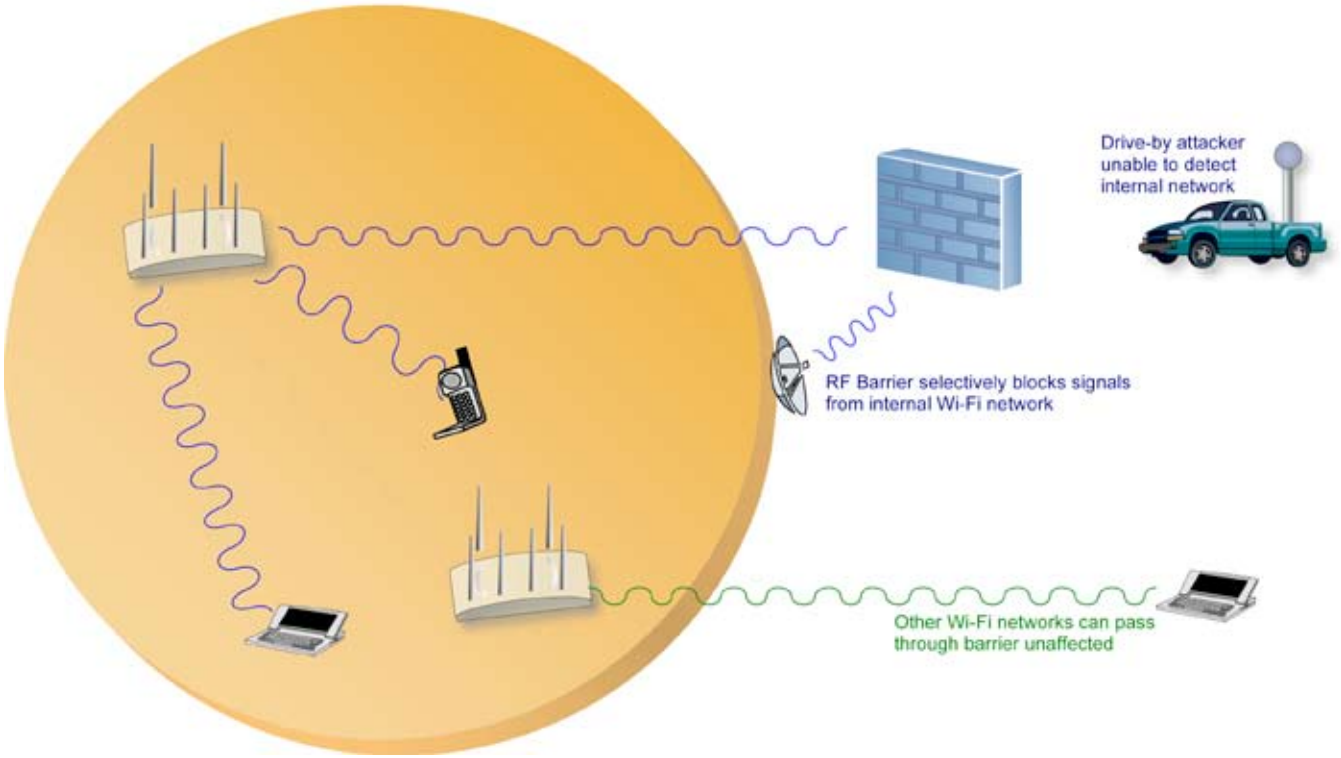

Figure 2: RF Barrier selectively blocks one Wi-Fi network. Others are unaffected.

Application Flexibility

Unlike Faraday cages, the RF Barrier can be turned on and off through the same management system that governs the rest of the network, perhaps allowing outdoor access at some times but not at others. Because it determines whether to block each frame individually, it is extremely adaptable and can be applied to many different use cases. For example:

- A warehouse that must support legacy WEP-based devices can restrict them to a network that only works inside.
- An office can allow Wi-Fi access in outdoor areas during the daytime so that people can work outside, but block it at night when intruders are harder to see and most employees have gone home.
- An enterprise sharing a multi-tenant unit with others can block only the signals that emanate from its own network. Other building tenants' networks are unaffected.

From a security perspective, the greatest benefit of wireless can also be its biggest problem: the signal goes everywhere. With RF Barrier, wireless signals can be made available only to users within a certain physical boundary. Wi-Fi gains physical perimeter protection equivalent to that of Ethernet, a critical consideration as users increasingly rely on wireless as their primary network access method.

Geofencing

CBecause modern wireless LAN systems are able to determine a user's position, location information can also be used for access control. This enables granular control of access policies with fine-tuning down to the level of individual rooms. It also allows specific users or types of traffic to be given access in areas and at times where others are denied it.

Geofencing is only as good as the positioning technology that underlies it. Meru's location-tracking is particularly accurate due to RF Fingerprinting, a technology that goes beyond triangulation to calculate distance and direction based on the known radio propagation characteristics of different materials. Because a signal will be noticeably different depending on whether or not is has passed through a wall, RF Fingerprinting can almost always determine which room a user is within. The same applies to other physical dividers such as cubicle partitions in large open-plan offices or to rows of shelving in a warehouse.

But despite its benefits, geofencing is not physical perimeter security of the kind that RF Barrier achieves. This is because it only prevents intruders from actually accessing a network, functioning at the connection layer. In contrast, RF Barrier forms an impenetrable physical perimeter beyond which the network's radio signal cannot propagate. Geofencing is ideally suited for prohibiting different types of access within a network's boundaries, RF Barrier for the external border.

The choice of which technology to use will depend on the individual application. For example, a network protected by geofencing could allow people to use their Wi-Fi phones from the parking lot but prevent laptop-based drive-by attacks. In contrast, RF Barrier prevents any traffic from a designated network getting out at all. Many networks may also need both: RF Barrier as a hardened external perimeter, geofencing for fine-tuned internal access controls.

2: Securing The Connection

Connection security protects a network against intruders who have physical access to a network but no legitimate login credentials. In a network that protects its physical perimeter using Meru's RF Barrier, connection security acts a second line of defense that intruders must cross if they manage to gain access to a building. In most wireless networks, it is the first security system that an attacker will encounter. The first sign of an active attack is either the network failing or the IPS reporting a problem.

Most enterprise 802.11 networks use some form of encryption, authentication and intrusion prevention to enforce access control, but a traditional IPS only works once an attackers' packets have reached an access point. Meru's is different, combining more effective detection with two separate firewalls. One can drop unwanted packets while they are still in flight, while another targets insider attacks that originate from otherwise legitimate users.

Intrusion Prevention

Most wireless networks treat the airwaves as a medium similar to an Ethernet hub, only much more extensive: Every wireless device within radio range can associate to an access point, which must then decide whether to allow the client to connect to the network. Meru Networks has developed a unique technology that can prevent unwanted devices from connecting even if they do have physical access to the network.

Real-Time Scanning

The first step in intrusion prevention is detection of potential threats. This requires over-the-air scanning – listening for transmissions that don't match the legitimate wireless network so that action can be taken against them. There are two main ways to do this, both of which traditionally require some compromises:

- Overlay networks use dedicated sensors that are separate from access points. These provide very high radio sensitivity but at added cost. They also lack integration with the main wireless network.
- Integrated sensors use additional radios built in to APs. These reduce costs, but due to interference they are often unable to scan continuously: A transmission from an AP will temporarily overwhelm the sensor, drowning out signals from potential rogues. The only way around this for most systems is to stop transmitting when a scan is taking place, something that both interrupts traffic and reduces total network capacity. Some vendors actually recommend that security scanning be disabled when voice or other real-time applications are running.

Meru offers integrated sensors within APs, meaning that intrusion detection and prevention features are able to leverage the Meru management system's full knowledge of the wireless network. But unlike other integrated sensors, Meru's can operate in the presence of real-time traffic, ensuring that there is no compromise between performance and security.

Meru AirFirewall

The final step in intrusion prevention is to respond appropriately to attempted attacks. With Meru's IPS, the response takes place at multiple levels: The AirFirewall blocks illegitimate attempts to connect to the network, then deeper network-layer firewalls target unwanted traffic from clients that manage to connect.

Meru offers the only true wireless firewall. An Ethernet firewall is a necessary feature in any network, but it does not protect against wireless denial-of-service attacks or against evil twins – attackers who impersonate an AP to lure clients away from the legitimate network, bypassing its wired security features.

To avoid these and similar problems, Meru offers the AirFirewall, a feature within its IPS that actually intercepts rogue wireless traffic over the air. Like a regular firewall, it works by dropping packets. Unlike a regular firewall, it does not require that the packets touch the wired network. Traffic from an unwanted device is automatically blocked over the air before it even reaches legitimate clients or APs, as shown in Figure 3 below.

Like the RF Barrier, the AirFirewall acts as a physical line of defense, preventing attackers from reaching the network-layer at all. In addition to deliberate attacks, it also offers protection from accidental connections that can leak information or degrade performance:

- Unauthorized networks such as those in neighboring buildings. The AirFirewall can make these disappear from the enterprise client's perspective. Because AirFirewall targets individual packets rather than committing a DoS attack, this is accomplished without interference to the neighboring networks themselves.

- Unauthorized clients that try to connect to any network in range. Many of these may not be malicious; they may belong to people who would prefer to connect to a neighbors' network, or just to passers-by whose Wi-Fi device is automatically scanning for available networks. However, all consume bandwidth when they send packets to APs. Such connections can be prevented through a combination of RF Barrier and AirFirewall.

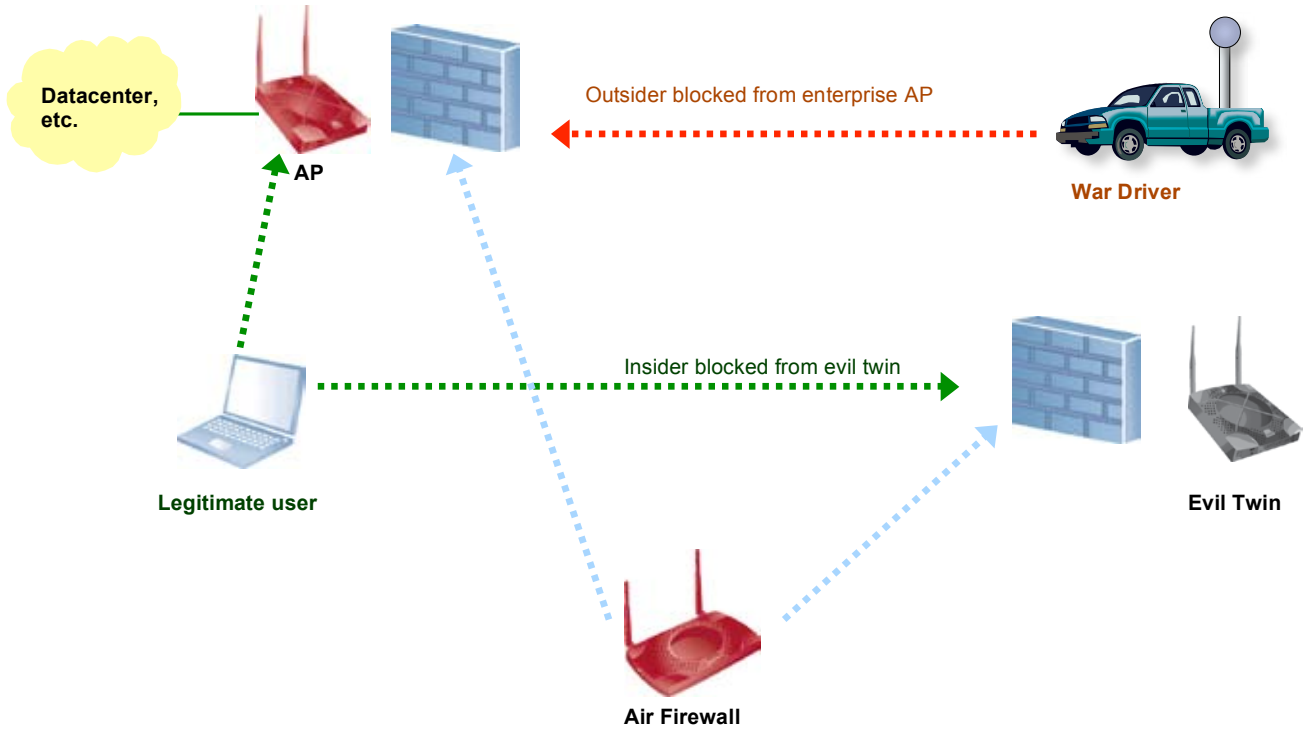

Figure 3: Two use cases for AirFirewall, blocking traffic coming in from a war driver and going out to an evil twin

Rogue Access Point Prevention

When Wi-Fi networks first became popular, rogue access points were among the most feared security threats. Because they didn't involve the IT department's official wireless network they threatened everyone, even companies with a strict no wireless policy. And most weren't deliberately malicious: Employees would buy a cheap Wi-Fi router and plug into the corporate network to help them access data while away from their desks, not realizing that it also helped everyone within radio range access their data too.

The best way to prevent rogue APs is to build a secure wireless network. When employees already have Wi-Fi access that combines mobility with Ethernet-grade performance, there's no reason to plug in a cheap consumer product. But even then, rogues remain a problem. Some are deliberate attempts to violate a security policy, while others may be set up by people wanting to try out new technologies.

Wired and Wireless Correlation

Meru's rogue AP detection uses the same continuous, real-time scanning as its IPS. However, it also monitors the wired and the legitimate wireless network constantly and correlates data from them with its scans. This enables faster detection of rogues and eliminates the risk of false positives, which are APs belonging to other organizations that mistakenly show up as rogues.

When a rogue is detected, Meru uses the AirFirewall to prevent clients from connecting to it. Any packets sent to or from the rogue are dropped, causing it to disappear from clients' lists of available Wi-Fi networks. At the same time, Meru's location technology pinpoints the device so that IT or security personnel can be alerted to disconnect it.

In addition, network administrators can also set whitelists of allowed APs to support organizations who use multiple wireless networks or environments where there are known neighboring APs.

Protection from 802.11n Rogues

Many of the technologies used in 802.11n first appeared in consumer routers, leading to a new category of threats that can't be detected by sensors based on 802.11a/b/g. This means that even enterprises who aren't yet ready for a wholesale migration still need 802.11n-aware security infrastructure to prevent rogue 802.11n APs from plugging into the network. With most newer clients featuring 802.11n, there is also a risk of 802.11n evil twins.

One solution is to use a dedicated overlay network of 802.11n sensors, but this has the same weaknesses as other overlay networks: high cost and low integration with the main Wi-Fi network. Another is to upgrade the entire network, but that too can be expensive in the short term – especially if most clients on a network are legacy devices that can't support the new technology.

For enterprises that aren't yet ready for 802.11n but still need full protection, Meru offers the AP302: a low-cost, dual-radio 802.11a/b/g access point that also supports 802.11n security scanning. For investment protection, the AP302 is also software-upgradeable to full 802.11n access capability using one or both radios.

3: Securing The Network

Network security protects IT resources from attackers who have already gained some access. Most traditional IT security products fall into this category. It includes technologies such as encryption, authentication and firewalls.

In addition to supporting IEEE 802.11i for authentication and encryption, Meru Networks uses both signature-based and behavior-based techniques to block attacks as soon as they are initiated. Meru also offers the industry's most accurate location-tracking system for pinpointing attackers.

Encryption

Most intrusion begins with passive eavesdropping: Attackers capture packets form a network to learn about potential security holes. In some cases, attacks may remain passive and undetectable, with information recorded for later cracking. For example, password-based encryption can often be broken

through dictionary attacks, while the WEP security in older 802.11b devices contained a flaw that allowed attackers to determine any encryption key once enough packets had been captured.

Fortunately, all recent Wi-Fi hardware is now certified to support WPA2, a subset of the IEEE's 802.11i security specification that includes wire-speed AES encryption. This has no known weaknesses, so when deployed properly it can ensure that wireless traffic is unintelligible to eavesdroppers.

However, WPA2 alone is not a panacea. While it is secure, it does not by itself meet the US governments' FIPS standard for provable security. Legacy devices that don't support 802.11i can mean that holes exist in a network. Encryption keys need to be generated and exchanged during authentication, a procedure left out of the standard. Meru offers solutions to all of these.

FIPS 140-2 Compliant Secure Gateway

FIPS 140-2 is a U.S. government computer security standard used to accredit cryptographic systems. It covers both hardware and software and officially applies to all US Federal Agencies that use cryptography to protect sensitive data, though similar requirements are in place in many state agencies and other countries. For example, the government of Canada and Israel both accept FIPS 140-2 certification as meeting their requirements.

Because of the rigorous testing required to achieve FIPS certification, most computer security products are not FIPS compliant. However, this does not mean that they cannot be used in an environment subject to FIPS rules, merely that cryptographic functions must be offloaded to a separate device that is FIPS 140-2 compliant. This is how Meru implements the standard: through a dedicated, physically-hardened appliance.

Unique among controller-based gateways, Meru's FIPS product meets the demanding standards of FIPS 140-2 Level 3 certification in the critical areas of roles, services and authentication. Unlike competitors' products that only meet the requirements of Level 2, the Meru gateway features a strengthened casing to prevent physical attacks, role-based access control for network administrators and two-factor key input.

Some vendors try to implement FIPS within APs or controllers themselves. However, this is less flexible, as it means that the APs or controllers need to be recertified with each major software release. This recertification process is expensive and time consuming, and vendors tend not to recertify every release, meaning that FIPS users often suffer from outdated equipment. It also means that the entire network must be dedicated to FIPS, refusing to connect to client devices that are not FIPS compliant.

The Meru FIPS Secure Gateway allows both FIPS and non-FIPS traffic to run over the same infrastructure, as shown in the Figure below. Traffic from FIPS clients is sent to the FIPS Gateway for cryptographic functions, while unencrypted traffic from non-FIPS clients goes directly to controllers as normal. A single FIPS Gateway can serve multiple controllers, all of which can use the latest version of their software as only the Gateway needs to be FIPS certified.

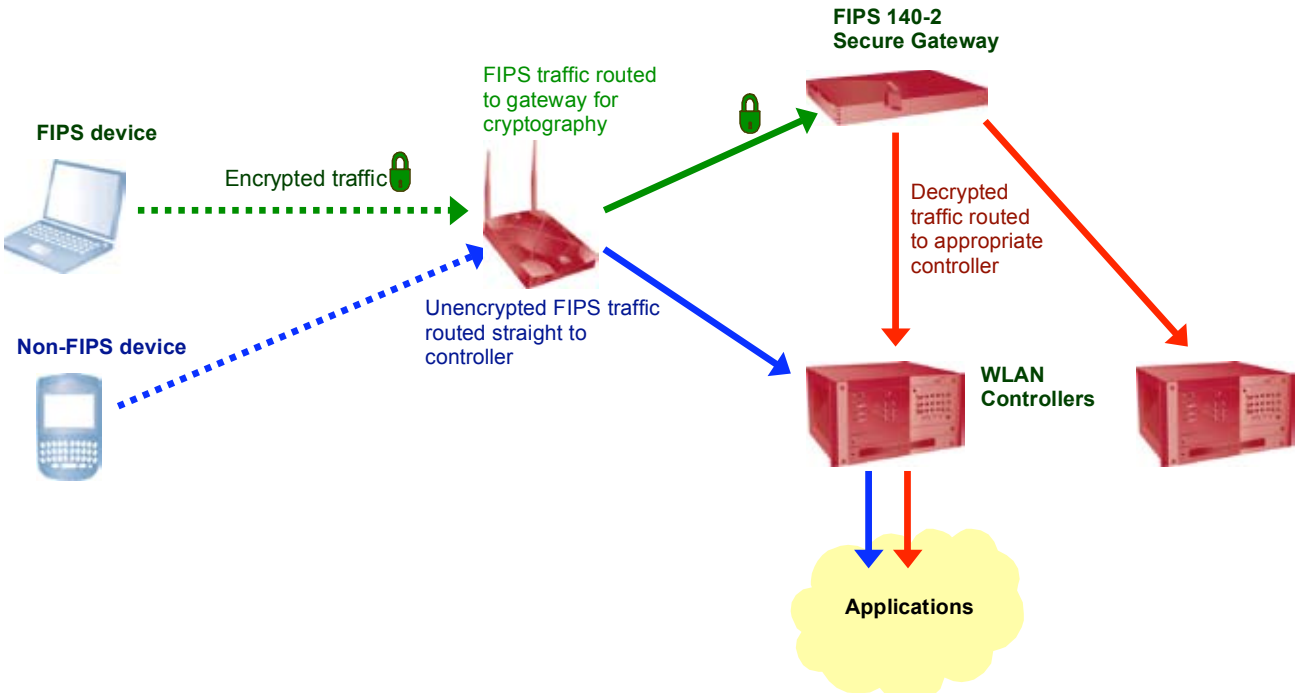

Figure 4: FIPS and non-FIPS traffic on the same physical APs.

Authentication and Authorization

WPA2 is available in two versions: Home and Enterprise. Though both use the same AES cryptography, they differ in how the encryption keys are generated and managed and how users are authenticated.

WPA and WPA2 Home use static keys: Simple knowledge of the key is enough for connection to the network, so there is no authentication of the user or the device. This is a security risk, as the same key must be shared amongst all devices on a network. If one device is stolen, the key needs to be changed for every one. Because there is no way to exchange keys securely, many users choose a key derived from a simple password.

WPA Enterprise uses randomly-generated keys that are changed regularly, with a new and unique key chosen for each packet. Keys are exchanged during user authentication, which takes place via the 802.1x protocol. All Meru products support WPA and WPA2 Enterprise, as well as multiple different authentication technologies Meru's authentication technology ties into RADIUS servers from many vendors, supporting secure, standardized algorithms including EAP-TLS, PEAP and GTC.

Network Access Control

NAC enables authorization decisions to be made based on a client device's hardware or software configuration. For example, PCs that are not running the most up-to-date version of an OS may only be authorized to connect to a software update service until they are patched.

Meru supports multiple vendors' NAC offerings, including those that verify a machine's state at the time it connects to the network and those that use client-side agents to monitor each device continuously. Two are particularly important:

- Microsoft Network Access Protection. Meru has partnered with Microsoft to support NAP, a NAC technology that verifies a Windows PC's software configuration as part of the authentication process.
- Trusted Network Connect. Meru is a member of the TCG (Trusted Computing Group), which enables cryptographic keys to be stored on hardened security chips that are now included as standard on most enterprise laptops.

Policy Enforcement

Meru's firewall functionality begins before users even connect to the wired network, with the AirFirewall dropping packets in-flight. But it doesn't end there. Meru also offers a per-user firewall within its wireless infrastructure, able to prevent insider attacks as well as enforce security and network policies for legitimate users. The two firewalls form two separate lines of defense, as shown in the figure below.

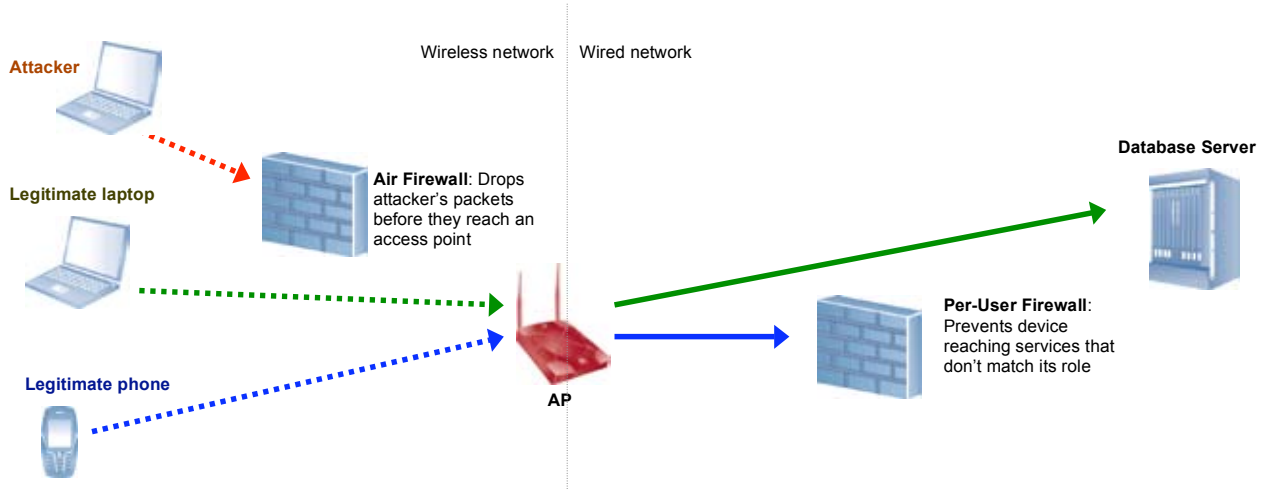

Figure 5: AirFirewall and Per-User Firewall

Role-Based Access Control

Instead of a one-size fits all firewall, Meru's can be configured on a per-user basis: Each device gets a customized firewall that only allows access to the specific services authorized for that particular device and user. The access privileges associated to roles can all be derived seamlessly from a user's RADIUS profile.

For example, the firewall protecting a network printer can be set so that it can only receive print jobs. This means that if attackers manage to connect to the network by impersonating the printer, they will be unable to access the sales database. Similarly, the HR servers are only accessible to laptops that have been issued to HR personnel.

The per-user firewall enables role-based access control, in which only specific categories of users can access each service. In addition to outside attackers, this protects against two potentially more important classes of threats:

- Malicious insiders. Whether employees or contractors, even legitimate users may have interests that aren't aligned with those of the organization as a whole. Role-based access control ensures that they cannot step outside the boundaries of their job.
- Malicious software. Machines belonging to innocent people often become infected by viruses or spyware, especially if they have been carried out of enterprise boundaries. A per-user firewall prevents the infection from spreading to other machines or probing the network for data while the user is unaware that the machine is infected.

Application-Aware QoS and Firewalling

Meru's network firewall does not simply block traffic based on origin or destination. It can also classify and categorize applications based on two different technologies:

- Deep-packet inspection (stateful firewall) at Layers 4-7
- Flow signatures that detect applications even when traffic is encrypted

The signature-based firewall goes beyond deep packet inspection, allowing network administrators to block or restrict applications such as Skype. This is particularly useful to enterprises that face strict regulatory compliance requirements. For example, the firewall that can block unauthorized peer-to-peer voice calls while enabling and prioritizing SIP-based calls to authorized voice services.

The same rules used in the firewall can also enforce quality of service settings. Once detected, applications can be blocked, throttled, logged or prioritized depending on network-wide policies or individual user roles. For example, the bandwidth available to users of Web sites such as YouTube can be restricted without blocking access altogether.

Guest Access

Allowing guests access to a network poses challenges to any security system based on credentials. By definition, guests are both insiders and outsiders: They will often lack login credentials, yet need to be given access anyway. And while most networks used for guest access must be open and unencrypted, they must also be secured.

Once connected, guest access needs to be constrained. The most important consideration here is usually ensuring that guests can only access the Internet and not resources on the LAN itself, but organizations may also need to limit the kinds of applications that can run and the bandwidth available to each user. Meru offers several technologies that can address these issues, all of which can be used alone or together:

RF Barrier. The physical security offered by the RF Barrier product is a powerful way of controlling guest access, as it restricts the geographic area in which network connectivity is available. For example, a conference center or restaurant could provide unrestricted Internet access to customers inside while ensuring that the network is not available to freeloaders outside.

Application Firewall. The role-based access control enabled by Meru's per-user firewall gives enterprises a way to govern exactly how guests use the network. It can classify, isolate and block unapproved applications, whether or not they are encrypted. This can be used to limit guests to Web access only, or to prevent specific bandwidth-hungry applications such as peer-to-peer file sharing. Similarly, guest access can be limited by data rate, location or time of day.

Secure GRE Tunneling. Guest traffic is tunneled directly to the Internet, ensuring that guests do not have access to the internal network. Employees can connect to the internal network as normal, bypassing the restrictions on guests by going through secured SSIDs that don't offer guest access.

Captive Portal

Most guest access sessions begin with a captive portal: a Web site through which all traffic from a user is initially directed. Depending on the network's requirements, the portal may ask users to register using an email address, authenticate using a pass code that has been provided offline, or simply agree to an acceptable use policy. Once the user has authenticated, the session proceeds as normal. Meru's Captive Portal includes several important features:

- RADIUS integration. The Meru captive portal can use RADIUS authentication, authorization and accounting services, meaning that the same credentials and policies used for accessing other resources can be used to pass through the portal.
- Third-party interoperability. If necessary, the portal can use other security products such as RSA two-factor authentication.
- Fully customizable. Organizations can adapt the portal to match their needs, both for branding and for what users must do in order to connect to the network.
- Multiple SSIDs per network. This allows unencrypted guest access and secured employee access networks to coexist on the same infrastructure.

Legacy Devices

Though properly deployed WPA2 has no known weaknesses and is present in all new Wi-Fi hardware, many legacy clients don't support it. Organizations with these may require that portions of the network use WEP or even no encryption at all, opening a security hole that can be an attractive target to attackers. This is why Meru's approach is particularly critical for networks that must support non-WPA2 clients, giving three separate lines of defense beyond encryption:

- RF Barrier prevents attackers from even finding the network unless they are inside a building.
- AirFirewall prevents attackers from connecting to the network at all.
- Role-based Access Control limits the services available if an attacker manages to connect to the network by stealing or impersonating a legitimate client.

4: REMOTE SECURITY

Remote security is intended to protect against attackers who try to access a network through the Internet. However, the ubiquity of wireless clients means that it is not the same as Internet security, posing new threats above and beyond the usual attacks that come through an Internet connection.

The killer app for wireless technology is mobility, which often extends beyond the enterprise premises. No matter how strongly protected the wireless network or the datacenter, devices that leave it may have to connect over other Wi-Fi networks. Meru's solution is to extend the enterprise wireless network into telecommuters' homes or business travelers' hotel rooms, making enterprise-grade security and management as mobile as employees themselves.

Attacks on the Extended Perimeter

Mobility means that enterprise boundaries now extend way beyond the corporate firewall, encompassing remote offices, telecommuters and traveling employees. However, perimeter security rarely moves with them.

Once outside the enterprise firewall, users are unprotected. They may deliberately visit Web sites that expose them to Trojans, or accidentally connect through insecure wireless networks that expose their data. When they require secure access to the enterprise, most rely on software-based VPNs that don't work with all devices. Meru's Telecommuter Access Point means that these users now have a simpler and more secure option.

VPN Clients

Software VPNs are the traditional technology of choice for remote access. Most are secure, but they have disadvantages that become more important as a greater number of employees spend an increasing amount of time working remotely:

- Limited device support. Most VPN clients are designed only for laptops, preventing secure access through phones and other devices.
- Employee workarounds. Because devices connected via VPN are subject to enterprise policies even while out of the office, many employees try to avoid VPNs wherever possible. Some enterprises may try to improve the user experience by enabling split tunneling, a potential security risk.

When the VPN client is not used, sensitive enterprise data must be sent through the employee's home Wi-Fi network with no additional security. This is a potential attack vector, as many home networks still rely on WEP or have security turned off entirely. Although most home routers now offer WPA, they do not support WPA Enterprise. The static, pre-shared keys of WPA Home are more easily compromised than the rotating keys of enterprise-class wireless networks.

Insecure home networks pose two threats to enterprise security. Data transmitted over the network from an enterprise laptop may be compromised through passive packet-sniffing, or active attacks may install malware on corporate laptops.

Meru Telecommuter Access Point

The Meru Telecommuter Access Point extends the enterprise LAN right into an employee's home or hotel room wherever there is an Ethernet connection available with internet access. This gives the remote employee seamless access to all the same features available within the enterprise headquarters, along with the same security policies that are applied to the enterprise LAN.

Like other Meru APs, the Telecommuter AP is centrally managed by Meru's Controller appliances. For users, it is plug and play, with none of the configuration hassles that affect traditional software-based VPNs. It can support all the same Wi-Fi client devices as the main office network, allowing users maximum flexibility to use technologies such as VoFi.

Because the Telecommuter AP is essentially an extension of the enterprise LAN, security policies should be no different from those in the office. For example, access can be denied unless devices authenticate using 802.1x along with whatever encryption method the enterprise has chosen, while the client's abilities to access network resources will be constrained according to the same policies that affect local devices. Instead of a disconnected island, the remote user is truly part of the enterprise LAN.

For maximum flexibility, the Telecommuter AP uses the same hardware as Meru's popular enterprise access points, with the only difference being software. This allows enterprises to reuse existing infrastructure, moving APs between locations as necessary

CONCLUSION

As organizations move from wired to wireless as a primary mode of network access, wireless security becomes paramount. Security technologies from Meru Networks mean that for the first time, Wi-Fi can truly be as secure as Ethernet.

Whereas most vendors offer point products that address only certain categories of threat, Meru effectively addresses all four. Its market-leading network security technologies protect networks from insiders and others who have managed to gain access, while its unique RF Barrier and AirFirewall bring true physical security to wireless networks.

When deployed together, these form three separate lines of defense around a network, all of which an attacker must penetrate in order to steal information or disrupt services. The Telecommuter AP adds a fourth layer while outside of the office, offering remote users the same combination of freedom, mobility and security that Meru wireless networks have already brought to local users.

Security: Multiple Lines of Defense

Threat Level Architecture	Meru Security Network	Traditional Wi-Fi Ethernet	Traditional Wired
Attacker without physical access to facilities	RF Barrier	No protection	Walls, doors
Attacker with physical access but no login credentials	AirFirewall WIPS Rogue detection / mitigations	WIPS Rogue detection / mitigations	Limited protection
Attacker with physical access and login credentials	Policy Enforcement Module with Per-user Firewall FIPS gateway with military grade encryption	3rd party firewalls	No protection
Attacker connecting through the Internet	Telecommuter Access Point	WEP or WPA Home	Software

Glossary

Air Firewall
A physical security technology from Meru that drops attackers' packets while they are still travelling through the air. This advancement in wireless security exceeds the capabilities of existing wireless intrusion prevention systems (WIPS) by being able to eliminate the threat, rather than just contain or mitigate it.

Airtime Fairness
Method of governing access to the airwaves so that all clients are able to transmit for the same amount of time, meaning that performance is higher for 802.11n users than for legacy clients. Without airtime fairness, slower clients can hog the airwaves as they take longer to transmit each packet.

Air Traffic Control
Meru technology that exercises a high degree of control over all transmissions within a wireless network. Unlike superficially similar technologies from other vendors, Air Traffic Control technology coordinates uplink and downlink transmissions on a single 802.11 channel in such a manner that the effects of co-channel and adjacent channel interference are eliminated and all access points on a network can share a single radio channel.

Co-channel Interference
Radio interference that occurs when two transmitters use the same frequency without being closely synchronized. Legacy wireless systems cannot achieve this kind of synchronization, so access points or cell towers that transmit on one channel must be spaced far apart. The result is coverage gaps that must be filled in with radios tuned to another channel, resulting in an inefficient and complex microcell architecture. Air Traffic Control technology avoids co-channel interference by tightly synchronizing access point transmissions, enabling that adjacent APs to use the same channel.

Channel Bonding
The combination of two non-overlapping 20 MHz. channels into a single 40 MHz. channel, doubling the amount of data that can be transmitted in a given time but halving the number of available channels. Along with MIMO, it is a key innovation in the 802.11n standard.

Channel Layering
Wireless LAN architecture in which several Virtual Cells are located in the same physical space but on non-overlapping channels, multiplying the available capacity. This additional capacity can be used for redundancy or to support higher data rates or user density. Channel Layering can also segregate different applications or client types, for example keeping 802.11n clients on separate channel from legacy 802.11b/g so that the .11n network can operate at maximum capacity. Channel Layering can be enabled through multiple radios on one AP or by using multiple AP close together, so the total capacity is limited only be the number of non-overlapping channels available.

Channel Reuse
A pattern in which different APs can use the same channel. In microcell networks, such APs need to be placed far apart to avoid co-channel interference, meaning that contiguous coverage requires multiple channels. In networks using Air Traffic Control technology, the same channel can be reused throughout the network, meaning that only one channel is required and others are left free for other purposes.

Connection Security
Measures that prevent an attacker from connecting to a network. In wired LANs, this is accomplished using 802.1X, or requiring ports to be enabled on an as-needed basis. In wireless LANs, connection security includes authentication, encryption—such as WPA2/AES—and methods to protect an established connection from attach, such as Air Firewall.

Controller
Appliance that manages a wireless network and (usually) aggregates traffic. Controllers were introduced with third-generation wireless LAN systems as a way to manage access points. In a fourth-generation system, the controller also governs client transmissions, deciding which AP each client is connected to. Controllers are sometimes referred to as switches because early versions took the place of Ethernet switches and had to be connected directly to access points, though this is now rare as most can now be placed anywhere in a network.

Denial of Service
An attack on the network that is intended to consume all available resources. Typical DoS attacks take the form of flooding the air with signals meant to cause disruption, such as Deauthentication frames or long, high-priority data frames. DoS attacks can attack at any layer in the OSI stack. Meru's Air Firewall physical security method can prevent A DoS can take place at all levels of the OSI Stack. Meru's WIPS is able to report the presence of DoS attacks. Furthermore, channel layering is able to circumvent many DoS attacks by offering multiple channels of spectrum, each one available in every square foot of the network's coverage..

Evil Twin
An attacker who impersonates the enterprise network, essentially pretending to be an access point, often so that it can intercept authentication credentials and other traffic. Clients can be prevented from connecting to evil twins by Air Firewall.

FIPS 140-2 Certified Encryption
Hardware and software that has passed government-specified tests for compliance with FIPS 140-2, a U.S. computer security standard used to accredit cryptographic systems. The AES encryption algorithm used in 802.11i-based FIPS products is fully interoperable with Wi-Fi Certified devices and transparent to applications. NIST provides FIPS certification to offer assurance against physical tampering and attacks and is required for many government applications.

Fourth Generation
Term coined by analyst firm Gartner to describe a wireless LAN system in which the controller governs handoffs, such as one utilizing Virtual Cells. This is contrasted with third generation (micro-cell architecture) systems, in which the controller is only responsible for managing access points and clients must decide for themselves when to initiate a handoff. Second generation systems lacked a controller altogether and were designed for standalone operation, whereas the first generation lacked any enterprise management features.

Handoff
The transfer of a link from one access point to another as a client moves through a network. In legacy microcell networks, Wi-Fi clients are responsible for handoff, meaning that the quality of the link and the overall network performance is dependent on each client's implementation of 802.11 roaming algorithms. In Virtual Cell networks, the network itself governs handoffs as clients remain connected to a single virtual AP.

Intrusion Prevention System (IPS)
A system that recognizes attacks and automatically institutes countermeasures to remove the attack from the network. The term is slightly misleading as it usually refers to defenses that are initiated after an attacker has already gained some access to a network, not a device that literally prevents intrusion. This is a result of the IPS's evolution from the Intrusion Detection System (IDS), an earlier product category that lacked automated countermeasures. The System works through pattern recognition, recognizing either known attack signatures or behavior. The unique characteristics of wireless attacks have led to the wireless IPS (WIPS)

Line of Defense
A security product or technology, each of which represents an additional barrier to an attacker. A well-designed security architecture will include multiple lines of defense to ensure protection even if one barrier falls. However, most wireless security vendors still focus on only one line: defenses that are not activated until an attacker has already connected to a network. Meru Networks also offers physical security products that provide two additional lines of defense, proactively preventing attackers from connecting to a network at all.

Geofencing
Access control based on a user's physical location within a building. Meru's geofencing is particularly accurate thanks to its use of RF Fingerprint and its single-channel architecture, which means that every AP within radio range of a client will be able to detect a signal for triangulation purposes. Note that geofencing is not the same as the physical perimeter security enabled by RF Barrier. Compared to RF Geofencing allows very fine-grained access control such as limiting certain users or devices to access from particular areas while allowing others, whereas RF Barrier can block a signal completely to prevent even passive eavesdropping.

Microcell
Wireless architecture in which adjacent APs must be tuned to different, non-overlapping channels in an attempt to mitigate co-channel interference. This requires complex channel planning both before the network is built and whenever a change is made, and uses spectrum so inefficiently that some co-channel interference still occurs, especially at 2,4 GHz.. Microcell architectures were common in 2G cell phone systems and legacy wireless LAN systems. They are not used in 3G cellular networks or in wireless LAN systems that use Air Traffic Control, as these allow all access points to share a single channel.

Overlay Network
A dedicated network of radio sensors that are similar to access points but do not serve clients, scanning the airwaves full time for security or management issues. Overlay networks lack the flexibility of AP-based scanning, as radios cannot be redeployed between scanning and client access. They also lack deep integration with the main wireless network, necessary for real-time management and intrusion prevention.

Parking Lot Attack
An attack on an enterprise LAN launched from outside a building, taking advantage of radio waves' ability to pass through barriers such as walls. Many serious wireless security threats fall into this category. Parking lot attacks are often undetectable, because, when properly executed, they generate no traffic. Proactive security means are required to prevent them, such as RF Barrier.

Perimeter Security
Protection of the physical perimeter around a network. For Ethernet LANs, it is accomplished surveillance cameras, security guards, key-locked doors and other measures that prevent physical intrusion. In wireless LANs, it is usually absent due to the difficulty of blocking radio waves. However, Meru's RF Barrier offers an alternative that can selectively prevent Wi-Fi transmissions from leaving a building, restricting the geographic availability of enterprise networks without affecting other wireless networks.

Per-Station SSID
A unique SSID (network name) for each device that accesses the network. This means that each client effectively gets its own customized network, with performance and access privileges adjusted to match.

Per-User Firewall
A firewall that has access privileges customized to fit each user or device. This limits the damage that can be done by attackers who steal hardware ,malicious insiders who try to access resources that they don't need to do their jobs, and innocent but compromised insiders that are infected with viruses or other malware.

RADIUS Authentication
Remote Authentication Dial In User Service, a networking protocol that enterprise-class 802.11i hardware uses for AAA (authentication, authorization and accounting) functions.

RADIUS itself does not specify how users verify their identity, so Meru supports several standardized authentication techniques that run on top of RADIUS. These include EAP-TLS, PEAP, GTC and secure tokens.

RF Barrier
A security technology from Meru that protects the physical perimeter of a building. It uses directional antennas to block Wi-Fi signals on one channel or network, preventing transmissions from leaving a facility. The result is that people outside will not even know that a network exists, let alone receive data from it. Unlike expensive techniques such as RF paint and Faraday caging, RF Barrier is both cost-effective and selective, allowing signals from other wireless networks to pass unimpeded.

Role-Based Access Control
Technology that gives each user or device access to a defined set of resources. For example, guests may only permitted to access the Internet or the public Web, while phones may only be permitted to access a SIP server or PBX. Meru's management system also supports multiple levels of access for administrators that can be assigned to different roles to ensure that network managers can maintain a high degree of control over their staff's access.

Rogue Access Point
An AP that is not part of an official enterprise wireless network, making it unlikely to comply with security policies. Rogues often provide attackers an easy route into a network. They remain a threat, whether set up due to ignorance or malice.

Rogue Detection and Prevention
The first and most effective line of defense against rogue APs is to give employees a high-performance enterprise Wi-Fi network so that they have no incentive to plug in other access points. The next is an affective scanning system. In networks that carry real-time traffic such as voice and video, it is important that rogue scanning does not interrupt client connectivity. With the rise of 802.11n, its is also important that scanners are able to detect 802.11n rogues even if the network itself is still using 802.11a/b/g.

Rogue Mitigation
A technique used by many vendors that tries to disconnect rogue APs by performing what amounts to a DoS attack. While this is effective against rogues, it also disrupts legitimate traffic on the enterprise wireless network. Meru's Rogue Prevention technology is based on continuous scanning and Air Firewall rather than DoS, so the network can continue to function normally even as rogue packets are dropped.

Scanning
The process of checking the airwaves for rogue access points or attackers. Scanning APs are typically implemented as an Overlay Network, as most APs can not scan and serve traffic at the same time. Meru's APs are able to scan the airwaves and serve clients simultaneously. Meru's single-channel architecture improves accuracy when scanning for intruders, as all APs are able to detect signals from all clients.

Spectral Efficiency
The ratio of data rate to radio spectrum usage. Of the new technologies introduced in 802.11n, MIMO and new modulation techniques represent higher spectral efficiency, but channel bonding does not (as it uses twice as much spectrum.) With any Wi-Fi variant, a Virtual Cell is much more spectrally efficient than a microcell architecture, as the microcells consume at least three non-overlapping channels to provide the coverage that a Virtual Cell offers with just one.

Single Channel
Term sometimes used to describe a network in which all access points operate on the same channel, such as one using Air Traffic Control technology. Single channel operation is more spectrally efficient than a microcell architecture and necessary for the use of Virtual Cells and network-controlled handoff. Single Channel improves security by making intrusion detection easier and location tracking more accurate, as every AP automatically receives transmissions from every client within range. It also enables the RF Barrier to function with as little as one radio, because only one channel needs to be blocked from outside access.

Virtual Cell
Wireless LAN architecture in which a client sees multiple access points as just one, all sharing a single MAC address, BSSID and radio channel . Air Traffic Control. Because clients remain connected to the same virtual AP as they move through a network, no client-initiated handoffs are necessary. Instead, the network itself load balances traffic across APs, maximizes bandwidth, simplifying network management and conserving radio spectrum for redundancy.

VoFI (Voice over Wi-Fi) or VoWLAN
(Voice over Wireless LAN)
Voice over IP links that run over a wireless network. VoIP does not usually require high data rates, but it stresses wireless networks in other ways by demanding low latencies and smooth handoffs.

VoFI Security
Defense against intruders or rogue access points that can function in the presence of real-time traffic. Many vendors' radio management systems cannot, as interference between scanners and transmitters means that access points cannot scan and transmit at the same time. Meru's scanners do not suffer from the weakness and are able to scan even while transmitting, fully protecting networks that carry voice and video.

War Driving
The practice of driving around with a laptop or other Wi-Fi device in search of wireless networks. Most war drivers are only interested in free Internet and networks that transmit in the clear, but some run WEP cracking tools or use more sophisticated hacking techniques that can lead to parking lot attacks. In dense urban environments, often replaced by war walking.

Wi-Fi
Brand name for wireless LANs based on various 802.11 specifications. All products bearing the Wi-Fi logo have been tested for interoperability by the Wi-Fi Alliance, an industry group composing every major 802.11 client and infrastructure vendors. The Wi-Fi Alliance's tests now require support for much of the 802.11i security standard which includes strong AES encryption.