WhitePaper

From Centralization to Virtualization

MERU NETWORKS

The Evolution of Wireless Infrastructure

The introduction of 802.11n means that wireless networks can now exceed the throughput of switched Ethernet. This has led many organizations to consider making wireless the preferred network access method, relied on for constant connectivity rather than just occasional portability.

The case for wireless becoming the standard method of network connection is powerful. Compared to Ethernet, building a wireless network requires far less cabling and network hardware. The benefits aren't confined to cost savings. Wireless also means greater flexibility and responsiveness to changing organizational needs. And it's what end users want.

Until now, there have also been good reasons not to use wireless. Given the vagaries of radio frequencies and the limitations of the "microcell" WLAN architecture, wireless is rarely as predictable or manageable as its wired counterpart. Until Meru Networks' introduction of WLAN Virtualization, choosing wireless over wires involved large extra costs and operational uncertainty.

WLAN Virtualization brings to edge networks the same benefits that virtualization has brought to server farms and storage: resources are pooled to achieve economies of scale and efficiency, then partitioned among users to ease management and make the network link exactly match the application that it serves. The result is a wireless network that, for the first time, can offer the predictability, reliability, and certainty of switched Ethernet.

Five Steps in the Evolutionary Path

Wireless network vendors have always made bold claims about the mobility and performance enabled by their products, so WLAN Virtualization's advantages are clearest when contrasted to those of previous WLAN technologies. According to industry analysts, wireless LANs have evolved through four distinct generations, starting as proprietary standalone access points and progressing to Meru's single-channel Virtual Cell architecture. With the introduction of full WLAN virtualization, Meru moves on to the next generation.

As Figure 1 shows, each generation of Wi-Fi technology has built and improved upon the features of its predecessors. While each stage was innovative, it increased user expectations and drove further innovation. The evolution of wireless LANs is a progression from isolated islands of connectivity to a technology that can challenge fast Ethernet as the preferred technology for connecting to the network.

Access	Standardization	Centralization	Coordination	Virtualization
Pre- 802.11	Fat AP (Microcell)	Thin AP (Microcell)	Virtual Cell	Virtual Port
				SLA
			Predictability	Predictability
			Reliability	Reliability
			Deployment	Deployment
			Seamless Mobility	Seamless Mobility
		Coverage	Coverage	Coverage
		Management	Management	Management
	Interoperability	Interoperability	Interoperability	Interoperability
	Security	Security	Security	Security
Connectivity	Connectivity	Connectivity	Connectivity	Connectivity

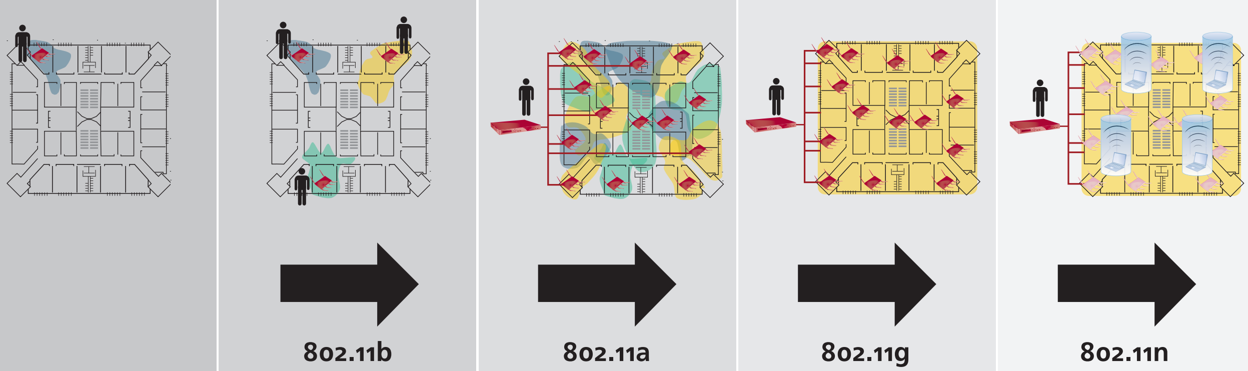

First Generation: Basic Access (pre-1999)

Description: Pre-802.11

Example: Proxim RangeLAN

Main Innovation: No wires necessary to connect to the network

New Demands: Interoperability, standards, security

Current Status: Obsolete

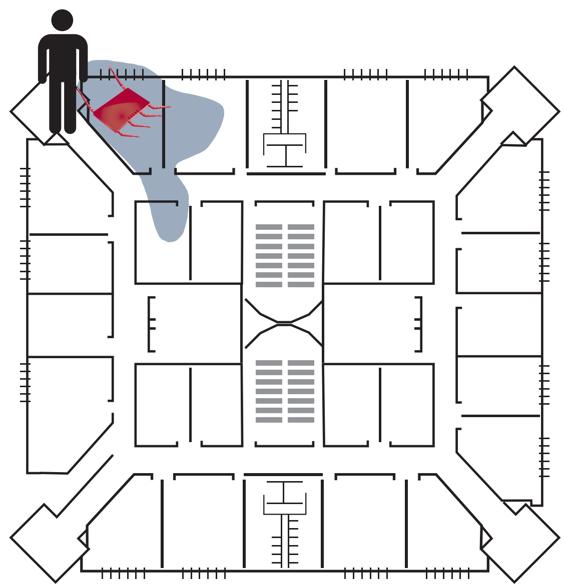

Early wireless access points were designed to provide basic connectivity. Because there were no widely-adopted standards, most vendors used their own proprietary systems, with some even relying on infra-red rather than radio bands. Client devices needed proprietary NICs, and more advanced features like security, management and roaming were generally absent.

Early WLANs were largely confined to environments where installing cable was difficult or impossible. Most users had a single AP, so roaming was impossible: Users who wanted access had to make sure that they were within the coverage area of that AP.

Second Generation: Standardization

Description: Fat APs

Examples: Cisco Aironet / IOS, Orinoco

Main Innovation: Standards-based, standalone AP

New Demand: Central management of multiple APs

Status: Still in use at many enterprises

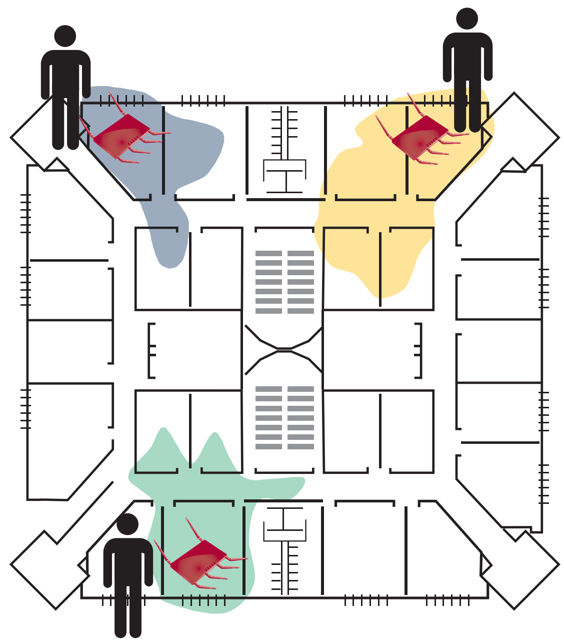

With the issuance of the 802.11b standard and the formation of the Wi-Fi Alliance in 1999, wireless LANs became mainstream enterprise products. Laptops began to include built-in wireless capability, while many networking vendors began to offer APs with some management features. Because security was recognized as important, several companies offered VPN encryption or proprietary systems while the industry worked on a standard.

Demands for wireless coverage grew, and customers began to build out networks that resembled a mosaic of overlapping "micro-cells", so-called because they were based on a scaled-down version of the cell phone networks of the time. Most users still saw wireless access as something to be tapped only when Ethernet was unavailable and not as a ubiquitous means of network access. APs were standalone devices, designed to be managed and operated independently. Although many enterprises deployed quantities of them to form microcell networks, these APs were not designed to work together and had to be configured one at a time.

Client-side issues became equally important as networks scaled out. Each wireless device tries to ensure that it achieves and maintains the strongest connection possible, a strength in an isolated area with just one AP but a weakness in a crowded environment with multiple APs and even more clients contending for access to the airwaves. When two nearby APs had roughly equal signal strength, a laptop could become confused, switching back and forth between them.

Third Generation: Centralization (2002)

Description: Thin APs, Centralized controllers

Examples: Cisco Airespace, Aruba, Nortel, HP ProCurve, Colubris, Symbol, Motorola

Main Innovation: Central control of APs for large networks

New Demands: Better reliability and connectivity, efficient RF management, wider coverage

Status: Currently the dominant architecture, though its limitations are becoming clear

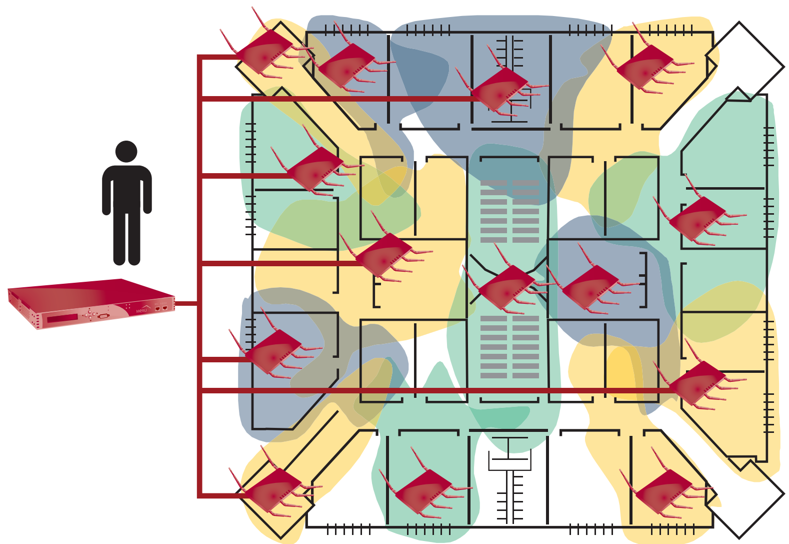

As networks of APs became denser, managing each one separately became increasingly difficult. This led several vendors to impose a centralized management scheme on top of the microcell architecture. Abandoning their role as independent networking devices, APs were slimmed down to be more like simple radio transmitters (often called "thin APs"), with their intelligence moved back up the line to be housed in and managed by a central controller. The technology was sometimes described incorrectly as switching because early controllers were integrated into Ethernet switches and needed a direct link to each AP. Most vendors soon saw the limitations of integrating wireless appliances into wireline devices, and moved beyond this to make the controller a separate appliance that could be placed anywhere – usually in the data center.

Because there was no longer any need to configure each AP separately, centralized controllers enabled very large wireless LAN deployments. Theoretically, coverage holes could be filled by adding more APs. However, the reality was not so simple, because each new AP added to fill a coverage hole created a new microcell, which had to avoid interfering with its neighbors. Because cells must overlap to ensure continuous coverage, each AP must be tuned to a different radio channel in order to avoid channel conflict that both reduces the available bandwidth and leads to dropped packets and a poor user experience.

Planning this channel mosaic pattern in a large deployment is very complex, requiring software that tries to predict the area covered by each AP – something impossible with absolute accuracy. This problem is magnified with 802.11n, as the new standard relies on multipath effects that are inherently unpredictable and lead to spiky, non-contiguous coverage zones when deployed using microcell architecture.

Organizations trying to increase WLAN capacity run into further limits. Especially at 2.4 GHz, the limited radio spectrum available means that there are simply not enough channels to accommodate the architecture's requirement of non-overlapping channels to avoid radio interference. As a result, microcell APs must have their power turned down, forcing customers to buy and deploy additional APs. Because so many channels are consumed mitigating interference, there is little room for expansion to denser networks. As with planning a network, this issue is most severe in 802.11n networks, whose highest data rates require wider channels.

This becomes most critical when the traffic running over a wireless network expands to include telephony, high-bandwidth video, or mission critical, time-sensitive applications such as those needed in hospitals. In these and areas with high user density, microcells' inefficiency and lower quality of service has led to demands for improved service and an end to dropped sessions.

Fourth Generation: Coordination (2003)

Description: Virtual Cell delivered through thin APs and controllers

Examples: Meru Networks AP150, AP200 and AP300 series

Main Innovations: Control of RF resource use, single-channel operation, network-initiated handoffs

New demands: Switch-like reliability for 802.11n

Status: Used by thousands of Meru customers worldwide

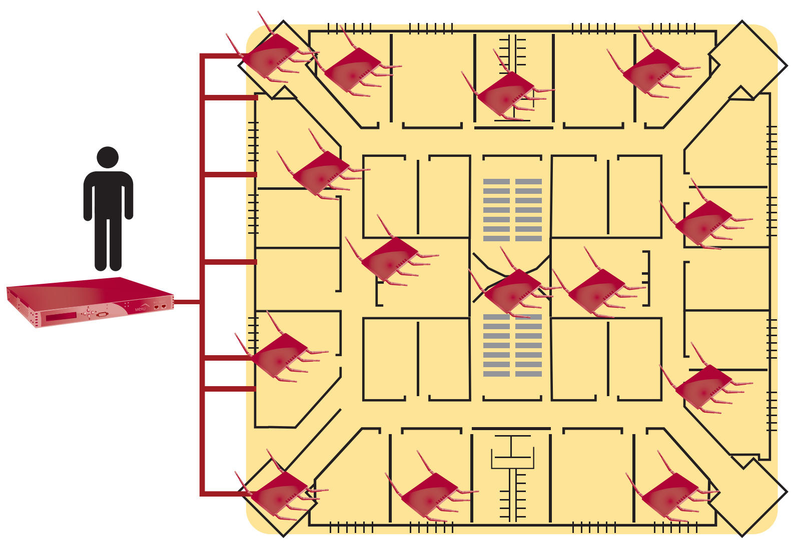

The Virtual Cell architecture achieves a much greater level of control over the wireless network and the client experience than a microcell network can. This is because, instead of seeing a network of separate APs, a wireless client sees only a single, large "virtual" AP that represents all the physical APs deployed in the network. Because the client device does not perceive multiple APs beckoning to it, it does not try to initiate handoffs as it does in a micro-cell deployment. The client's ability to disrupt its connection is neutralized, letting the network effectively take control of decisions about which AP each client will connect through.

How is this achieved? Like the microcell solution, the Virtual Cell architecture uses many thin APs under the command of a single controller. There the similarity ends. Unlike microcells, the Virtual Cell solution enables seamless mobility and optimal use of radio spectrum thanks to two unique innovations that allow multiple physical access points to be pooled and treated as a single virtual access point: single-channel architecture and network-controlled handoff.

Single-channel architecture permits adjacent access points to transmit on the same channel, eliminating the need for the channel planning, hand-off management, and other software required to mitigate the problems which arise at the interface of any two microcells in a network. Because all the Virtual Cell APs can broadcast on the same radio channel, channel overlap problems do not exist. This simplifies installation dramatically, both before the network is built and whenever a change needs to be made. Instead of engaging in the time-consuming and tedious mapping of a building and predicting the coverage areas of each AP, virtual cell deployment becomes a breeze: it is simply a matter of placing the APs in a manner to achieve full coverage at their maximum broadcast capability, then connecting the APs to the network, and finally letting them download all their settings from the central controller.

If coverage must be extended or enhanced, a new AP can be added without any cascading effect on the rest of the network. Meru's unique Air Traffic Control technology coordinates AP transmissions to ensure that they enhance rather than interfere with each other, allowing each AP to transmit at its full power. Thus, other radio channels that a microcell network could consume to provide just one layer of coverage are free to be used by other Virtual Cells layered in the same physical space, increasing bandwidth and giving each client multiple access choices. Many Virtual Cell deployments needs about 30% fewer APs than a microcell network to cover a given area, although some Meru customers have reported needing up to 70% fewer APs. The benefit: cost savings in all related infrastructure including hardware, cabling, planning, controller capacity, and deployment.

Network-controlled handoff means that the Meru controller uses its network-wide awareness of traffic load and the radio environment to route every client's packets through the access point which provides an optimal experience for the client. This is possible only because the client sees a single AP so it never initiates a handoff which can disrupt the smooth functioning of the network: instead, the client remains connected to the same virtual AP wherever it goes in the network, which allows the controller to manage the accumulated bandwidth of all APs. The result is smoother roaming and a more reliable connection for the end-user. Overall network performance is improved, as time wasted when clients that drop to a lower data rate due to a poor signal is eliminated.

Finally, a Virtual Cell uses only one radio channel to provide coverage to a floor, a building, or a campus. In contrast, a microcell architecture requires a minimum of three channels, plus the channels left unused to provide buffer space to absorb channel bleed. Thus, other radio channels which would be consumed by deployment of a single layer of coverage for a microcell network are free to be used to create other Virtual Cells layered in the same physical space as the first to provide and bandwidth and access choices.. Capacity grows linearly with the number of radios available. Clients are automatically load-balanced across channels to ensure optimum use of the RF resources.

Next Generation: Virtualization (2008)

Description: Virtual Cell and Virtual Port

Examples: Meru networks running System Director 3.6

Main Innovations: Switch-like reliability and management, predictability and privacy

New demands: All-wireless edge networks

Status: Now available to all Meru Networks customers

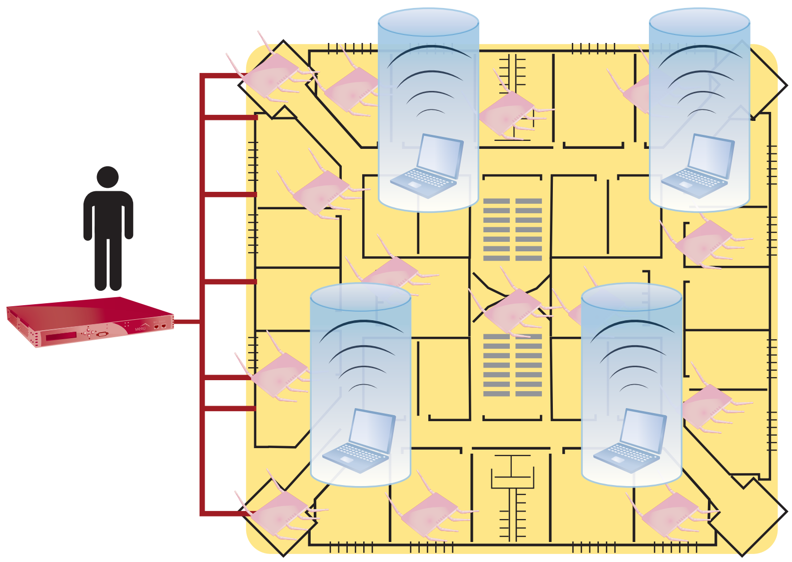

The introduction of Virtual Port completes the realization of Meru's WLAN Virtualization vision. Just as the single-channel architecture allows Meru controllers to manage and allocate their entire resource of channel bandwidth to optimize the client experience, Virtual Port gives organizations the ability to treat each end-user the same, whether they connect through an Ethernet port or a Virtual Port.

Virtual Port builds on Virtual Cell by giving each client its own virtual access point with all the characteristics of an Ethernet link. The key enhancement from the Virtual Cell-only environment is that whereas the Virtual Cell is shared between all clients on a network, just like other wireless APs, the Virtual Port is dedicated to a single device and provides that device with the same kind of uninterrupted and unconflicted experience that applications expect from a switched Ethernet port. Thus a Meru controller and its APs behave more like a wired network switch than a traditional wireless AP, overcoming the last barriers standing in the way of deploying the all-wireless edge network.

Because the Virtual Port is unique to each client device, the network can tailor the Virtual Port to match exactly the kind of network that the client needs. Different employees can be given different amounts of bandwidth depending on the applications they need to run. A voice client gets limited bandwidth but high quality of service. A guest is given lower priority and restricted access.

Like clients within a Virtual Cell, a client connected to a Virtual Port sees just one AP no matter how large the network. The Virtual port inherits and builds upon all the Virtual Cell's benefits, including the single-channel architecture and the smooth roaming. Clients never try to initiate a handoff. From their perspective, the Virtual Port travels with them as they move through the network.

Summary

With 802.11n exceeding Ethernet in raw bandwidth terms, wireless connectivity is ready to replace wired edge networks for mainstream enterprise applications, if the wireless network is designed and built to take advantage of all the capabilities of 802.11n. To best understand how this can be achieved, it is useful to review how Wi-Fi networks have evolved, and where Meru's recognized reputation for innovation and design leadership has taken this powerful technology.

Such a review makes clear the differences between different choices in WLAN architecture. It explains how the microcell architecture used in most WLAN deployments contains design limitations which mean that these wireless networks still lack the reliability and predictability of wired Ethernet. Although microcells proved sufficient for the occasional portability that early wireless networks supported, they lack support for seamless mobility and are complex to manage. They are also unable to scale to the high data rates required by new applications or the high user densities which result when wireless is used as a primary network that replaces wires.

The Meru Virtual Cell and Virtual Port eliminate the problems, costs, deficiencies, and performance limitations inherent in the microcell architecture, and which now will be exacerbated by the move to 802.11n. The Virtual Cell eliminates co-channel interference and handoffs, offering smooth roaming as clients move through a network's coverage area. It also makes scalability simple, as new Virtual Cells can be activated by adding more radios. The Virtual Port gives the network fine-grained control over each client, while clients get their own private connection to the network. The result is a network that matches switched Ethernet in every way, combining the connectivity that people expect from wires with the mobility of wireless and the agility of virtualization.

Meru Networks
Corporate Headquarters
894 Ross Drive
Sunnyvale, CA 94089 USA
P 408.215.5300
F 408.215.5301
www.merunetworks.com
info@merunetworks.com