Meru Networks
Powering the All-Wireless Enterprise™

WHITE PAPER

The Evolution of Wireless LANs: From Centralization to Virtualization

Date: May 2010

TABLE OF CONTENTS

INTRODUCTION

The introduction of 802.11n means that wireless networks can now match the high performance of switched Ethernet. The result is the all-wireless enterprise – an organization that uses wireless as its preferred network access method, reliable enough to offer constant connectivity rather than just occasional portability.

The case for wireless is powerful. Compared to Ethernet, building a wireless network requires less cabling, fewer components and less time. 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. Thanks to the vagaries of radio interference and the limitations of the "microcell" WLAN architecture, wireless LANs were rarely as predictable or manageable as their wired counterparts. Before Meru Networks' introduction of the virtualized wireless LAN, choosing wireless over wires involved operational uncertainty over both the service level available to each application and the amount of staff time consumed by troubleshooting.

Virtualized wireless LANs change all this, bringing 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 a connection as predictable, reliable and trustworthy as 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 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 fully virtualized wireless LANs, 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 high-performance technology that is as reliable, predictable and trustworthy as Ethernet, able to challenge wires as the preferred technology for connecting to the network.

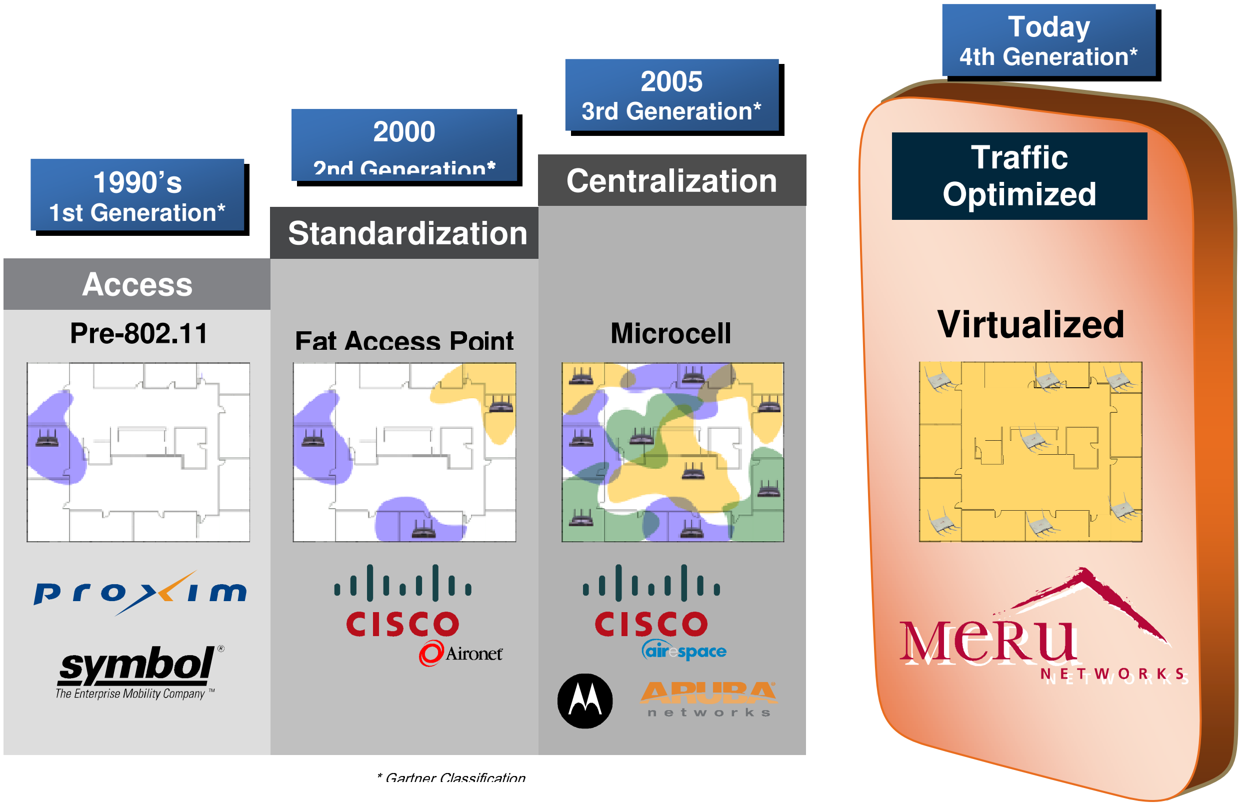

FIRST GENERATION: BASIC ACCESS (PRE-1999)

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, some even relying on infra-red rather than radio bands. Client devices needed proprietary NICs. 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 an 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 acceptance 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 management into wire line devices so the next step was 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, but reality is not so simple. Each new AP added to fill a coverage hole creates 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 and 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 higher performance, assured service and an end to dropped sessions.

FOURTH GENERATION: COORDINATION (2003)

Description: Virtual Cell delivered through thin APs and controllers
Examples: Meru Networks AP 150, 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 would in a micro-cell deployment. The client's ability to disrupt a connection is neutralized, letting the network take control of decisions about which AP each client will connect through.

How is this achieved? Like the microcell system, 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 design and network-controlled handoff.

Single-channel design permits adjacent access points to transmit on the same channel, eliminating the need for channel planning and fine tuning of power settings before the network is built and after changes are made. This simplifies installation dramatically. Instead of engaging in the time-consuming, and tedious mapping of a building and then trying to predict the coverage areas of each AP, virtual cell deployment is plug and play. All APs are equivalent, automatically downloading their settings from the central controller after they are connected. Filling coverage holes is simply a matter of adding a new AP.

When coverage needs to be extended, 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 each other, automatically eliminating the problems associated with co-channel interference. Each AP to transmit at its full power.

Because microcell coverage plans require most of the APs to be powered down so that their broadcast range fits the requirements of the mosaic plan, the microcell architecture requires a significantly higher number of Access Points to cover the same area than the Virtual Cell architecture. 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 than with a microcell network. 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 the entire network as a single AP and so it never initiates a handoff. Instead, it remains connected to the same virtual AP wherever it goes, allowing the controller to manage the accumulated bandwidth of all APs. The result is smoother roaming and a more reliable connection for the end user. Network performance and reliability are improved, as time wasted when clients drop to a lower data rate due to a poor signal is eliminated.

The Virtual Cell is also much more scalable than its predecessor, requiring only one radio channel to provide coverage to a floor, a building, or a campus. In contrast, the microcell architecture requires a minimum of three non-overlapping channels, plus those 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 microcell coverage are free to be used to create other Virtual Cells layered in the same physical space as the first, providing 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 and predictable performance.

NEXT GENERATION: VIRTUALIZATION (2008)

Description: Virtual Cell and Virtual Port
Examples: Meru Networks with a controller running System Director 3.6 or above
Main Innovations: Switch-like reliability and management, predictability and privacy
New demands: All-wireless edge networks
Status: Used by thousands of Meru Networks customers worldwide

The introduction of the Virtual Port technology fully realizes Meru's virtualization vision. Just as the single-channel design 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.

The Virtual Port builds on the Virtual Cell architecture by giving each client its own virtual access point with all the characteristics of an Ethernet link. The key enhancement is that whereas the original Virtual Cell was 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 predictable and contention-free service 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 the all-wireless edge network.

Because the Virtual Port is unique to each client device, the network can tailor each device's Virtual Port to match exactly the client's exact requirements. 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 design and the smooth roaming experience. Clients never try to initiate a handoff. From their perspective, the Virtual Port travels with them as they move through the network for a reliable connection no matter where they are.

SUMMARY

With 802.11n surpassing Ethernet in raw bandwidth terms, wireless is ready to replace wires at the network edge for mainstream enterprise applications. For this potential to be realized, the wireless network must be designed and built to take advantage of all the capabilities of 802.11n. To 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 distinctions 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 o, they lack 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 architecture and Virtual Port technology eliminate the problems, costs and performance limitations inherent in the microcell architecture and 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 predictability, reliability and high-performance that people expect from wires with the mobility of wireless and the agility of virtualization.