UNIFIED MOBILITY SERVICES IN THE NEW CAMPUS NETWORK

Maximizing Mobility and Service Quality on Any Mobile Device

Table of Contents

Table of Figures

Executive Summary

As 802.11n technology finds its way into enterprises large and small, the campus network is ripe for profound change. The new campus network will have mobility at its heart, which will simplify access by minimizing wiring. Indeed, the wired access layer will ultimately disappear, as enterprise network access becomes completely untethered. Such full mobility is as inevitable as it is desirable. As full mobility is realized, worker productivity will increase and operating costs will drop. Employee satisfaction will also increase as full mobility enables workers to achieve an optimal work/life balance with communications anywhere, anytime, and with any device.

Achieving this vision of the fully mobile enterprise is not without its serious challenges, however. Part of the problem is the different and sometimes conflicting needs of different stakeholders. Users want unrestricted flexibility that recognizes them, supports their personal devices, allows total mobility, and works every time. IT management wants a network that is scalable, reliable, and completely secure with centralized control. The CEO naturally wants to satisfy both of these stakeholders, but only with a solution that is cost-effective and enduring, enables business processes that enhance productivity, and can withstand a rigorous audit.

Unified Mobility Services

As enterprises migrate to full wireless coverage and become dependent on wireless LANs for primary network access, network scalability and application reliability are paramount concerns.

This white paper examines the requirements and solutions for enabling service-level guarantees for the growing number of mission critical mobile applications.

Many existing wireless LANs (WLANs) are simply unable to accommodate so many different needs that surface when users come to expect uninterrupted voice, video, and campus-wide mobility—all at the same time. The underlying reason is that most of the business-enabling mobility services (voice, video, guest access, location, etc.) continue to be implemented in isolation. As a result, performance suffers, troubleshooting becomes extraordinarily difficult, and total cost of ownership (TCO) goes up. As more services are enabled, the problem is magnified. Users complain, and for good reason. Organizations waste time and money on worthy mobility initiatives that are ultimately doomed to fail.

It need not be this way. This white paper presents a Unified Mobility Services Architecture capable of realizing the full potential of what many call the “Unwired Enterprise” or perhaps more accurately the ”less-wired” network. The content is organized into three remaining sections followed by a brief conclusion. Background information is provided in *The New Network: Opportunities and Challenges*. This is followed by a detailed description of *The Unified Mobility Services Architecture. Putting It All Together* provides examples demonstrating the need for unifying mobility services in a holistic fashion.

The New Network: Opportunities and Challenges

The performance afforded by IEEE 802.11n now makes it possible for the enterprise to achieve a long desired goal—eliminating wired access in the campus to give all users full mobility. Of course, this does not happen over night. It is usually the culmination of a progression of deployment phases that begin with wireless access as a convenience in only a few locations and for only the most mobile of users.

The productivity and other benefits of mobility are motivating organizations to now provide access in more places and for more users than just those who have long had mission critical mobility needs, or specialty mobile devices, or both. Today it is both feasible and beneficial to provide enterprise-wide wireless access for all users and all mobile devices, including those issued by the organization as well as those owned by individual users.

The potential cost savings and increased productivity of a fully mobile network are highly desirable, and some organizations have already committed to “cutting the cord” entirely. Hospitals and universities are at the leading edge of this trend (see sidebar on *The Unwired Campus*), and a growing number of organizations in other sectors are not far behind. Many new buildings are now being “wired for wireless” in anticipation of this growing trend.

The New Campus Network Saves Power for the People

These days, management of IT assets—from the data center to the user devices at the edge—also needs to factor in one other important consideration—electrical power consumption. Because one of the fastest ways to reduce power consumption is to accelerate the migration from desktop PCs to laptops and hand-held devices, this initiative should go hand in hand with a move toward total mobility sooner rather than later. Users want to be mobile anyway, and laptops consume only a fraction of the power needed by desktop PCs. Access points also consume far less energy than their equivalent 10/100/1000 Mbps Ethernet access switches.

The Unwired Campus

Founded in 1855, Kean University has grown to become one of New Jersey's largest institutions of higher learning with a student population of some 15,000. The University's president issued a mandate to deploy campus-wide connectivity for all students, faculty, and staff—both indoors in all 40 buildings and outdoors throughout the 150 acre campus. The network would also need to support all devices students and staff now use, including laptops, tablets, e-book readers, MP3 players, and VoIP phones. Kean's IT department decided to go all wireless, and chose a solution from Juniper Networks for its scalability, dependability, fast seamless roaming, high performance, and centralized management. Integrated support for a variety of mobility services will enable the university to support voice, video, and data applications, and implement location-based services.

There are many other benefits to be gained from unwiring the enterprise for both the users and the organization. Users gain increased freedom to access any application from any location with any device, provided permission has been granted. This dramatically improves productivity, and it allows for the use of personal mobile devices to enhance work/life balance, although this adds to the total number of devices being managed. The organization also benefits from the enhanced productivity afforded by more effective and efficient business processes, and the removal of barriers to adoption of new mobility applications. Another benefit to the organization is the significant reduction in operating expenses based on eliminating the previous problem of costly moves/adds/changes as an inherent part of mobility. These and other benefits result in a solid ROI in the new network.

As the enterprise goes all laptop, tablet, and notebook, and aging access switches become due for retirement, these changes will provide the final impetus to go all wireless, even for voice communications. Some "power users" may still need an Ethernet connection at their desks, and some may still prefer tethered telephones. But the vast majority of users will find 802.11n access to be quite sufficient for their data, voice, and video communications needs, and will relish the new flexibility afforded by total mobility.

Of course, even in buildings and campuses wired for wireless, there is still a critical role for wired Ethernet beyond connecting and powering access points. Switched Ethernet will not be going away either in the network backbone or in the data center. But the evolution of the access layer does present an opportunity to simplify the wired network by collapsing it to just two layers. Access points will no longer be wired to wiring closet access switches, but instead will be aggregated on switches with 10GbE uplinks to the core. The few remaining devices that will forever remain stationary, such as power workstations, fixed kiosks, teleconferencing stations, and some video surveillance cameras, may also use these switches, making it possible to eliminate an entire layer of switches.

There are three main differences between tethered and mobile connectivity that must be considered in the new campus network—constantly changing locations, shared bandwidth challenges, and higher quality expectations for mobility.

Constantly Changing Locations

The main difference between wired and wireless access is that users can readily change location—and constantly will. This is, of course, rather obvious. But such mobility also has profound implications for the applications and services involved. Consider E911 as an example. With tethered telephones, the location of the phone is programmed into the PBX or IP PBX. Hence the PBX can readily provide the exact location of the caller to direct first responders or other appropriate personnel. But with VoIP over a wireless LAN (VoWLAN), the location of the caller is unknown and must be computed in the moment. Unfortunately, the best some WLAN solutions can manage is to identify the access point being used; the user could be anywhere within a 3,000 square foot radius, including on another floor. The ability to pinpoint a user's location precisely is, therefore, an important requirement for the new network.

Shared Bandwidth Challenges

There are also other important differences between wired and wireless LANs that must be considered for a campus-wide mobility initiative to succeed, particularly related to performance. Unlike switched Ethernet, wireless bandwidth is shared on a common medium, which has the potential to create contention among users and uses. To make matters worse, different users access the shared bandwidth at different data rates depending on their devices' capabilities and their current distance from the access point (the latter changing constantly). Some of the radio frequency (RF) bands may be crowded or noisy from interference, while others remain relatively open and clear. In such an environment, it is vitally important to be able to manage the WLAN infrastructure and the various mobility services in a coordinated and cohesive fashion, in order to utilize available capacity efficiently and deliver acceptable levels of performance to everyone.

What use is it to allow a client to roam to an access point that is experiencing chronic performance degrading interference? The client is obviously at risk of having a substandard experience. But if the management system cannot match demand to capacity, this scenario cannot be prevented. There are many techniques such as band steering, load balancing, spectrum analysis, per user quality of service (QoS), and bandwidth management that can all serve to regulate the loading of the network. But if the management system is not able to use these techniques intelligently, in light of prevailing demands, the WLAN system will always fall short on long-term scalability and optimal performance.

Higher Quality Expectations for Mobility

Performance issues raise another important difference in the new network. Unlike during the early days of cellular phones, users are no longer willing to sacrifice quality for mobility. This new expectation applies equally today in both the WAN and the LAN. Consider, for example, the frustration voiced by iPhone4 users experiencing dropped calls because of antenna issues, or how much the cellular carriers now emphasize call quality. In the wireless LAN, these same expectations translate into seamless roaming with no voids, along with voice call quality comparable to tethered phones. These new expectations have even given rise to a new term—Quality of Experience or QoE. Satisfying the new QoE expectation can be particularly challenging for solutions that are incapable of scaling the required coverage and capacity as the WLAN expands to support all users.

Performance can also degrade dramatically with various network failures. "Failure" can take many forms from a user's perspective, of course, ranging from a VoWLAN call being dropped to a widespread outage. The former can be tolerated occasionally; the latter simply cannot in new networks supporting a suite of mission critical applications.

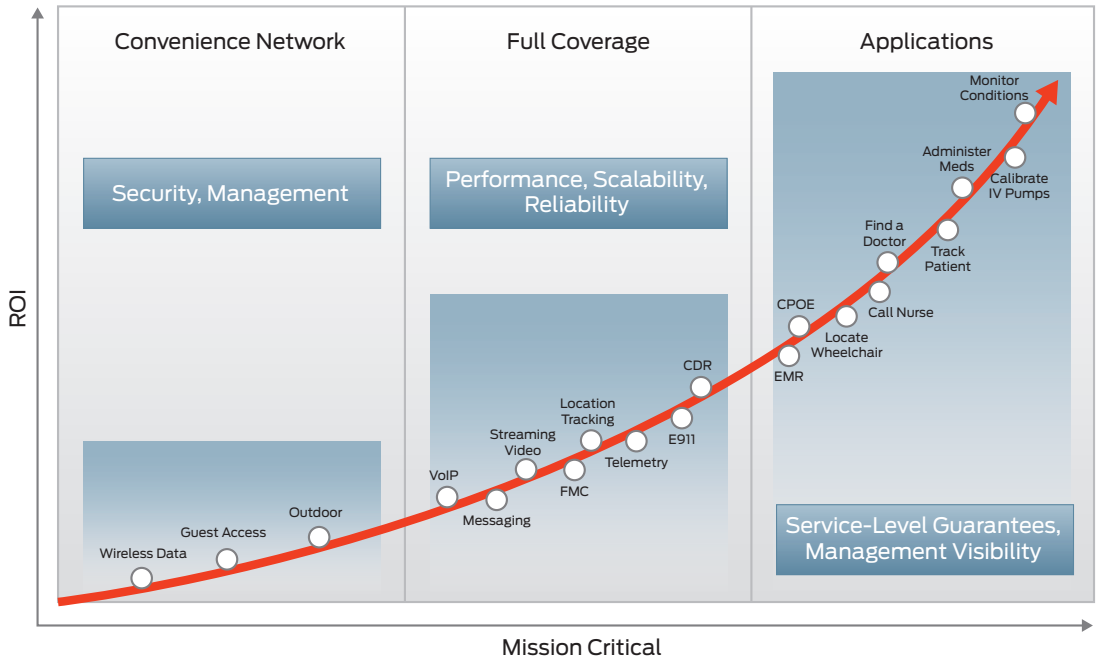

Figure 1: The evolution of wireless LANs in hospitals began with basic connectivity and now must support a wide range of productivity enhancing and mission critical applications.

The Evolution of Mobility Services

Part of the natural evolution of wireless LANs noted above is the recent emergence of a suite of mobility services such as voice, video, guest access, advanced security, and more. Figure 1 depicts the evolving use and corresponding value in the form of return on investment (ROI) over time for hospitals as early adopters of the new network. The same diagram could also be used to show the evolution of the sophistication of the WLAN over time, from basic connectivity in the early days to support of mission critical business processes and applications today.

Enterprise migration to full coverage of the WLAN is accompanied by the need for more advanced features in the form of mobility services like real-time location services. But as more users and more mission critical applications are added, these services begin competing with one another for resources.

The need for mobility with both good performance and quality creates other needs and challenges in the wireless LAN as well. These needs are addressed by the mobility services that are built on the wireless network infrastructure. Examples include advanced security, RF spectrum management, user/device identity management, and application context, along with more traditional wireless intrusion and prevention systems (WIDS/WIPS), voice, video, and guest access services. These and other mobility services must be capable of supporting the full range of applications, from basic Internet access to mission critical business processes, and some of these are rather ambitious applications that have only recently worked on switched LANs.

Different vendors and analysts characterize mobility services and infrastructure capabilities differently. And such differences are to be expected during evolutionary periods. But most vendor offerings today have one characteristic in common—the implementation of each mobility service in its own silo. Implementing mobility services as standalone applications was done to accelerate the time to market for new services from WLAN vendors, and to enable services from third parties to be added to another vendor's wireless LAN.

"Most Wired" Move to Become "Most Unwired"

Concord Hospital, the second busiest acute care hospital in New Hampshire, was recognized as one of America's 100 "Most Wired" by *Hospitals and Health Networks* magazine. Now the hospital's goal is total mobility for every mission critical application.

"Every new application at Concord Hospital needs to be a reliable mobile application," says Mark Starry, director of enterprise architecture and security. "The wired network is a critical utility, and our CIO directed that the wireless network be upgraded to the same level, as it is now equally critical to our mission."

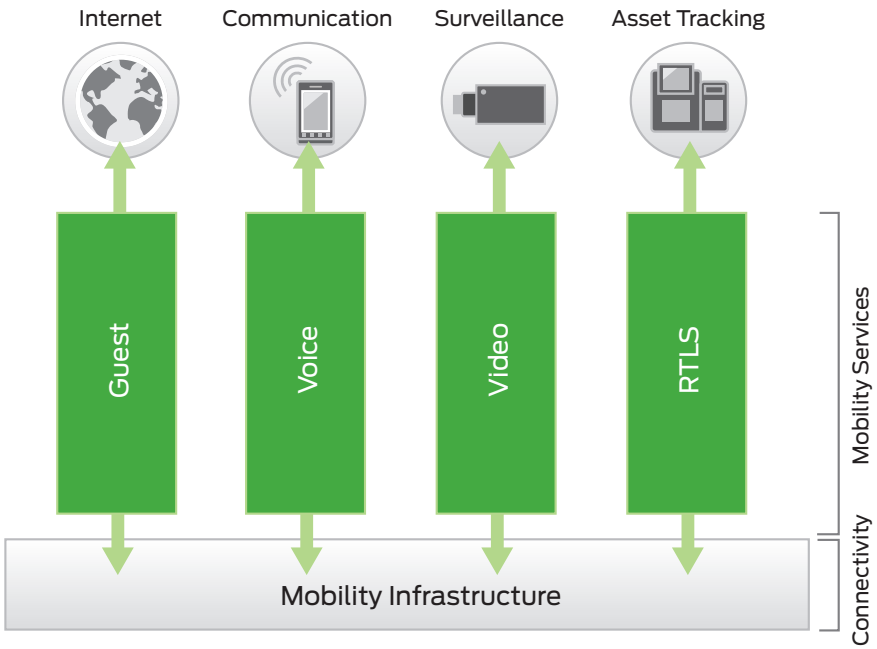

Figure 2: Existing mobility services have arisen out of necessity and speed to market as isolated services, each configured and operating independently.

Mobility Services in Silos

There is a serious problem, however, with implementing mobility services in separate and isolated silos, as shown in Figure 2. Because each is installed, configured, and managed separately, the individual services remain completely unaware of one another as they all compete constantly for the same underlying resources.

Here is an example that illustrates the problem. A WIDS/WIPS capability is configured to automatically turn on simultaneous scanning on an access point that is near a suspected rogue to help track down its precise location and launch a counterattack. Now imagine a user on a VoWLAN call who is about to roam to this access point. The network should be able to detect that the access point is no longer available, and preemptively steer the client seamlessly to another one that is known to have adequate capacity to service the call. With isolated silos, such cross-functional management is impossible.

Sometimes this isolation may be quite intentional, at least initially, and would permit a future enhancement that enables integration with other services. But sometimes the isolation is inherent in the fundamental architecture, thereby precluding future integration. And in both cases, it may be difficult or impossible to integrate the mobility services fully with certain capabilities provided by the resources in the underlying infrastructure. Making matters worse is the fact that user mobility, literally, makes resource management a moving target.

The severity of this problem depends on where the organization is on the evolutionary path depicted in Figure 1. For organizations just beginning to implement basic mobility services on a relatively small scale, the use of siloed services is likely to be adequate. But for the organization implementing a full suite of mobility services—many being business-specific—the separate silos will create an untenable situation that ultimately costs more in manpower to operate, and likely runs out of steam prematurely. As it is, those organizations are typically seeing a two-fold growth in traffic and an almost three-fold growth in devices every 12 months. Wireless resource management can simply no longer be ignored.

Consider hospitals as an example. Hospitals have been on the leading edge of Wi-Fi adoption and a growing number are now going all wireless. Many of the applications are obviously mission critical, even life and death, and some are quite demanding. More and more doctors and nurses are equipped with a VoWLAN handset or voice badge, and many also have specialized tablets for accessing patient records while on the move or at the bedside. Critical portable equipment is tagged to permit locating it precisely and quickly when needed, and increasingly that equipment is becoming Wi-Fi-enabled out of the box. Patients and visitors are also being granted guest access services. Imagine attempting to manage this complex network with all of the mobility services implemented in silos. How, for example, would the doctor in the emergency room needing access to the patient's CT scan be given priority over all other, less critical traffic? Could the doctor rely on location data to see which specialists are on duty, and which of those are currently available to render assistance? When implemented in silos, these services either ignore each other completely, or compete with one another for available resources while being oblivious to resource commitments already made on behalf of other services. What is needed, instead, is unified situational awareness with real-time coordination among all services.

Even in situations with only a relatively small set of mobility services, network administrators continue to struggle with unique management tools for each service, and regularly find it difficult to troubleshoot problems owing to the lack of coordination among services and resources. A recent wireless LAN survey by Aberdeen Group revealed that over half of those organizations classified as best in class (ostensibly the leaders in the use of information technology) continue to lack the ability to allocate bandwidth for specific applications, users, or groups. As a result, users suffer from poor performance and quality by receiving, at best, a best effort service that lacks service-level differentiation.

The bottom line challenge is this. To optimize performance and Quality of Experience, the all wireless access network must be intelligent enough to dynamically provision all network resources for all users based on who they are, where they are, and what they are doing—all in relationship to what others around them are doing. And that requires a unified mobility services architecture.

Unified Mobility Services Architecture

Overcoming the many problems caused by having separate silos of mobility services requires a more unified architecture, one that enables multiple services to be integrated and managed together, not in isolation, so that they can share collective network intelligence. The architecture must be built on a solid mobility infrastructure foundation. It must be both comprehensive and integrated—top to bottom and end to end—to ensure support for the full gamut of potential enterprise applications. It must be open, scalable, and extensible to ensure that the new network can take advantage of new technologies and services, potentially from third parties. And network administrators must be able to manage everything efficiently, effectively, and centrally. Such a robust Unified Mobility Services Architecture is depicted in Figure 3.

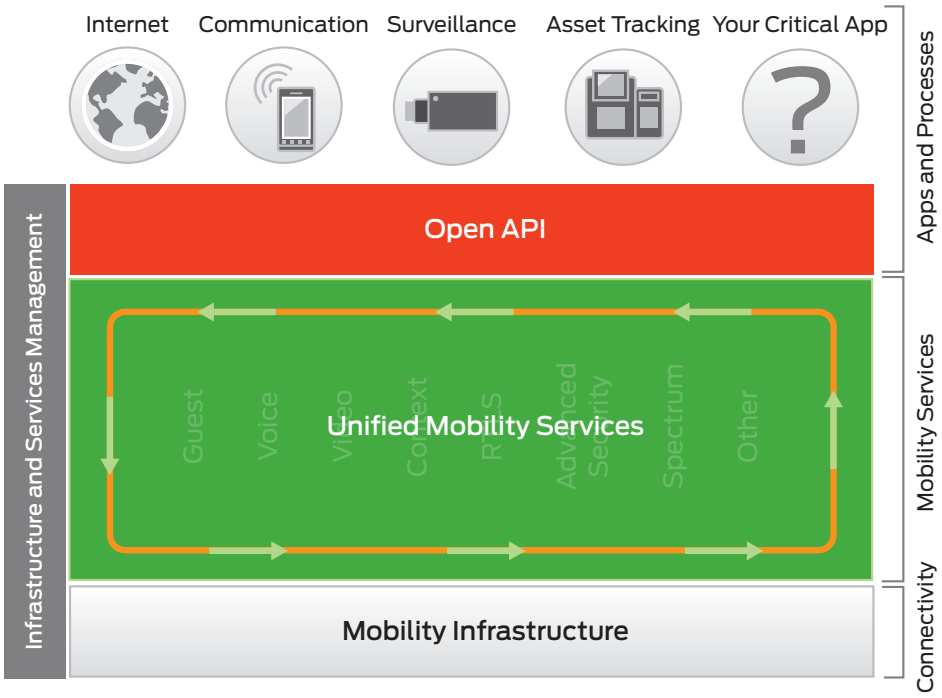

Figure 3: The Unified Mobility Services Architecture depicted here is capable of supporting even the most demanding needs in any organization—today and for the foreseeable future.

The Right Foundation—A Dependable Mobility Infrastructure

Any viable architecture must have a solid foundation. In the case of mobility services, that foundation is the wireless LAN infrastructure itself. Because this mobility infrastructure has been the primary focus of the industry to date, most WLAN vendors have fairly solid solutions. Although a full description of this foundation is beyond the scope of this document, the five essential elements or pillars required to support a robust set of mobility services are:

- Mission critical reliability
- Scalable performance
- Location awareness
- Comprehensive security
- Effective management

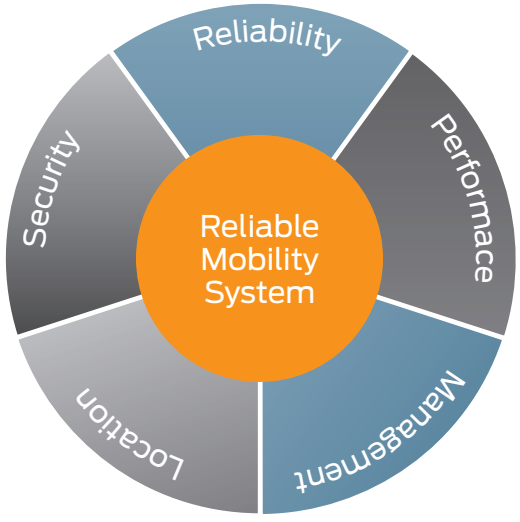

Figure 4: The five essential elements needed to build a solid foundation for mobility services.

Mission Critical Reliability

For any foundation to be solid, it must be reliable, particularly when all users are connected via the wireless LAN. Maximum system-level reliability can only be achieved with load-balanced controller virtualization that supports either N:1 or N:N automatic failover. N:N is a form of many-to-many redundancy where all controllers act as backups for one another, enabling the network to provide hitless failover with 100% application continuity. In addition, the system must be capable of self-calibration, particularly with regard to RF coverage, in the event of catastrophic failure of access points or unexpected RF obstructions or interferers. This is needed to ensure predictable application behavior, even under suboptimal conditions.

Scalable Performance

Achieving peak performance with good QoS is a perennial objective in all WLAN deployments. The advent of 802.11n with its generous throughput, and the various QoS provisions designed for LANs (wired and wireless), help significantly. But there is a wide range of other proven techniques available for squeezing every last byte of available capacity out of the shared medium. The more of these a vendor supports, the better—distributed switching, spectrum-aware dynamic load balancing, band steering, dynamic call admission control (CAC), air time fairness, and Dynamic Frequency Selection (DFS3) support for maximizing available channels.

Together these techniques serve to eliminate controller bottlenecks and other causes of load imbalance and congestion, such as client devices defaulting to the 2.4 GHz band even when the 5 GHz band is idle. When all are used in a coordinated fashion, the result is a much more effective use of all of the available RF spectrum, along with more efficient traffic flows, lower end-to-end latency, superior QoE, and faster roaming.

Location Awareness

According to IDC, "The information generated by Wi-Fi LBS [location-based services] solutions can be used to fundamentally change and enhance business processes, increase operational efficiency, reduce costs, and optimize asset utilization." Indeed, the ability to precisely and quickly pinpoint the location of a device or its user (or both) is vitally important for many mobility services such as E911 for VoWLAN calls, and real-time tracking of critical portable equipment or assets. Although the degree of precision can vary among the triangulation, trilateration, or other techniques employed, all are far more precise than determining location based merely on which access point is being used.

Comprehensive Security

There is no shortage of security standards, and most WLAN vendors support all of the critical ones. The real security challenge in the new network involves accommodating the many different users, many of whom have multiple devices running different applications as they roam to different locations. Some of these devices will be secured with up-to-date antivirus software and other endpoint security provisions, but many won't. It will be necessary, therefore, to treat different devices differently, even for the same user, and to implement mobile device management in a way that integrates with security policies. Further, as users move about, there is often a need to adjust access privileges accordingly. For example, if users step outside the building, their access rights should be reduced, as these users now present a much higher risk than when indoors. The network infrastructure should, therefore, be capable of continuously and dynamically adjusting the appropriate attributes to the user session each time the system observes a significant state change for any user.

Effective Management

Management here refers to tools used for managing the mobility infrastructure itself, and most vendors have a good solution for managing Faults, Configuration, Accounting, Performance, Security (FCAPS). Some vendors also offer one other important capability, a sophisticated 3D modeling tool for planning and optimizing the deployment and configuration of access points. After all, a prudently planned and properly configured network is far easier to manage than one that is neither. Of course, management also has a significant and even broader role to play in unifying the mobility services.

Unifying the Mobility Services

Unified mobility services, operating in a coordinated fashion, will allow the creation and enforcement of application service-level agreements (SLAs). In the unified architecture depicted in Figure 3, the mobility services comprise a separate layer between the mobility infrastructure and the organization's applications and business processes. While the seven categories of mobility services shown represent the set of most commonly deployed generic services available today, more are likely to evolve over time, especially ones that are vertically aligned (such as in healthcare that may include telemetry, and in schools for E911), and these are represented by the use of "Other" as a placeholder. Here is a brief description of these mobility service categories and how they relate to or interact both with other services and the underlying mobility infrastructure. Emphasis is given to the policy considerations involved in the integration or unification of multiple services.

Guest

Guest access has long been a common service in wireless LANs. Some organizations place considerable restrictions on guests, while others are more lenient. In either case, the issues are the same. Should guests be permitted access to voice and video services, and if so, at what levels of performance and QoS/QoE compared to employees in the vicinity? Once inside the premises, should they be allowed to roam around the building, or should their access be limited only to certain areas of the building? Should guests be permitted access while outdoors, and if so, to which set of applications? How will their devices be secured to fully protect the enterprise network?

Voice

VoWLAN is another common service. However, recognizing and classifying voice call flows, and then moving the packets around with the appropriate QoS for the desired QoE, remains a challenge with some WLAN solutions. For example, CAC is a common feature to limit the number of active voice sessions in order to avoid a situation where the network becomes so congested that quality deteriorates for everyone. But some CAC solutions include every VoWLAN-enabled device in their counts, even when these devices are not currently on a call. The use of dynamic CAC solves this problem, but others remain, such as being able to precisely locate an E911 caller or provide support for the wide range of devices being used.

Video

There are three different types of video traffic, and the mobility infrastructure must be able to accommodate the specific demands of each. The first is broadcast video, which normally employs multicast to transmit live or recorded webcasts to large audiences. Video surveillance feeds may also be multicasted to multiple locations. Second is video on demand (VOD) that normally employs unicast for applications like training, webtorials, and other instructional applications. Third is two-way video, which is becoming increasingly common for teleconferencing, webcam calls, and other forms of site-to-site (unicast) or multisite (multicast) collaboration. Because video traffic can consume a significant amount of bandwidth, an awareness of the needs of other mobility services is especially important.

Context

At a minimum, context involves who (the user), what (the user's device), where (the user's location), and why (the application). For troubleshooting, when an event occurred can also be an important factor. The permutations and combinations of different users, devices, and applications create some perplexing policy challenges regarding access capabilities. Guests with vulnerable devices are normally granted quite limited access, whereas a trusted user with a secure, corporate-issued device may be granted unrestricted access. Context also has a role to play in identity/role-based performance policies. The "who" may be a chronic and acute bandwidth abuser, for example, who now deserves a lower priority (at least temporarily). But the "who" could also be the CEO or CIO, who are obviously more "equal" than others.

RTLS

Real-time location tracking systems add the "where" to complete the context, and a user's location is an increasingly important factor in other mobility services. For example, it is vitally important to be able to precisely and quickly pinpoint the location of any user making an E911 voice call, and then identify all others in the vicinity to either request assistance or notify them to evacuate the area, depending on the nature of the emergency. Location-based access control (LBAC) can be used to restrict the capabilities of guests while outdoors (but still connected to an indoor access point) to enhance security. Knowing the precise location of a rogue access point—either an RF interferer or a security problem—eliminates the need (and time required) to track it down with a portable spectrum analyzer. Troubleshooting efforts are similarly assisted by knowing precisely where a problem occurs.

Advanced Security

Advanced forms or layers of security are either beneficial or necessary in the new network. For example, unsecured, user owned devices may need to be subjected to a mobile data management system (see discussion on Open APIs below) or protected by a wireless intrusion detection/prevention system (WIDS/WIPS). Rogue access points can cause a serious threat when being backhauled by a 3G/4G cellular service, and therefore need to be recognized, pinpointed, and removed immediately. And an RF firewall can be used to create a secure perimeter (geo-fencing) around the mobility infrastructure to prevent neighboring users or skilled hackers from gaining access.

Spectrum

There is no shortage of other wireless devices using the unlicensed spectrum today. Many of these exist in personal area networks (PANs) for keyboards, mice, peripherals, and phone "ear buds" using protocols like Bluetooth or ZigBee. Although these PANs operate at lower power settings than 802.11, they can still cause RF interference with Wi-Fi radios. Non-Wi-Fi devices such as movement sensors, wireless video cameras, cordless phones, and microwaves can be even more disruptive interferers. And then there are rogue users who create ad hoc or peer-to-peer Wi-Fi networks in the same spectrum that need to be detected and stopped.

Fortunately, wireless LAN solutions have long had effective techniques for detecting and mitigating against a range of RF interference sources—from powerful planning tools that employ 3D modeling during the installation, to features like automatic channel and power tuning, band steering, and rogue access point detection. As more and more devices enter the spectrum, the industry can expect more attention to this issue in the future as vendors tap into the spectral analysis capabilities in new RF chipsets now being used in 802.11n access points. But the real payoff comes when spectral management goes hand in hand with other services in a coordinated fashion to enforce application SLAs.

Other

The seven mobility service categories just described afford a fairly robust set of capabilities. But as the WLAN continues to evolve, more, and more sophisticated mobility services will certainly be needed. Initially, these may be offered in standalone silos, but the architecture outlined here allows for their full unification with other mobility services and the mobility infrastructure. Consider just one inevitable example—the convergence of wireless LANs with the wireless WAN. With the migration to 4G networks, mobile carriers will begin to utilize VoIP, and this will facilitate a mobile/mobile convergence that will enable users to roam to/from an 802.11n network from/to a 4G service seamlessly.

Open APIs

Having published application programming interfaces (APIs) is key to making the Unified Mobility Services Architecture extensible with an "ecosystem" of value-added services and applications. The architecture depicted in Figure 3 has APIs both above and below the mobility services layer. The API between the mobility services and the mobility infrastructure layers is designed to enable third parties to add new mobility services or extend existing ones. These are, in effect, the "horizontal" enhancements that can benefit any organization. The API between the mobility services and the applications and business processes is intended primarily for "vertical" third-party applications targeted at the specific needs found in different organizations.

An example of horizontal integration is mobile device management (MDM), which many organizations now need to support and secure mobile phones. As smartphones and personal devices continue to proliferate, the need to create MDM as a mobility service will grow. Examples of vertical applications and business processes can readily be found in hospitals with Vocera voice badges, specialized medical tablets, and asset tags affixed to portable equipment.

Infrastructure and Services Management

As the name and positioning of infrastructure and services management implies, this unifying capability applies to both the mobility infrastructure and the mobility services. The critical importance of unifying this aspect of the architecture is highlighted in the following Aberdeen Group best practices for best-in-class enterprises:

- A holistic (versus piecemeal) approach to network architecture, deployment, and ongoing support and management
- A unified network management infrastructure to provide visibility and control over the entire network simultaneously, which reduces redundant control layers and enables support staff to take a proactive approach to maintenance, upgrades and support
- A consistent policy regarding performance optimization, system upgrades, and maintenance across the entire organization
- Effective enforcement of security and compliance mandates organization-wide.

Naturally, this aspect of the architecture must leverage the capabilities of the WLAN management system, as well as any available for individual mobility services. And in doing so, it must also unite these separate provisions into a cohesive and holistic set of capabilities for managing QoE for mobile users in the new network. For only with such a holistic approach will the all wireless access network be intelligent enough to dynamically provision network resources for all users based on who they are, where they are, what device they are using, and what they are doing—all in relationship to what others around them are doing—to optimize performance and overall QoE.

Putting It All Together

Here are two real-world examples that illustrate the need to unify mobility services in a holistic fashion.

CEO on the Move

Imagine this scenario. The CEO arrives at his usual time and receives an important VoWLAN call while walking through the lobby. As he approaches the second floor on his way to his office on the third, his phone will need to roam to another access point near the Training Department. But on this particular day, the training staff is using an unusually high amount of bandwidth for streaming video and telemarketing calls. In fact, the AP is so oversubscribed that it rejects the roam, and the CEO's call is terminated unexpectedly. After getting to his office and completing his rudely interrupted call, you can bet his next call will be to you or your boss.

With robust and holistic mobility services, this scenario would unfold quite differently—much to the relief of the CIO, and to the financial benefit of the company. Based on historical tracking, the management system recognizes that it is the CEO on the phone, and anticipates his usual route through the Training Department. Before he roams to the access point on the second floor, the system throttles back the video streams slightly and reserves a session in the dynamic CAC table, enabling the CEO's call to continue uninterrupted with sufficient quality. As he roams to the access point near his office on the third floor, he continues his conversation, and normal access is restored to the Training Department in a way that is seamles, automatic, and holistic.

"Oh, No!" vs. "I Know"

Troubleshooting problems in a wireless LAN is inherently more complex than in a wired LAN. People are moving about. Conditions change. Resource utilization shifts. When a problem occurs, an administrator's usual reaction is, "Oh, no," signifying the recognition that finding and fixing the cause(s) are not going to be easy. Take an unreliable VoWLAN service or a dropped call in the lobby, for example. Was the user a guest or an employee? Was the problem something wrong with the user's phone? Which access point was the client associated with? Where exactly was the user when the problem occurred? Was the user attempting to roam at the time? Was this a roam that crossed controllers? Which RF band and channels were involved? Was any access point overwhelmed with other traffic, possibly from an active video surveillance camera? Was the CAC limit reached? Was there some interference or a network intrusion? Were there any changes recently to the WLAN's configuration or to the lobby itself? Without unification among the many mobility services involved, the administrator is faced with a daunting task.

The Unified Mobility Services Architecture makes it possible to turn the dreaded “Oh, no” reaction into, “No problem, I know I can find exactly what happened.” By correlating the user’s precise location at the time with everything else occurring in all of the other mobility services that might have been involved, the administrator can quickly and easily find and fix the cause(s) of any problems. There is no need to check the individual management systems for each mobility service or the underlying mobility infrastructure, and no need to manually compare these separate findings in search of some suspicious interaction. Only with the holistic intelligence afforded by a unified set of mobility services can organizations achieve predictable operation of the WLAN, support service-level guarantees and application assurances, and deliver a better QoE for all users.

Conclusion

The new network built for mobility and stripped of wires at the access layer is now as inevitable as it is viable. The advent of 802.11n and full suites of unified mobility services make it easier and more cost-effective than ever before to “cut the cord” and give all employees the power of full mobility. But an attempt to implement mobility services in independent silos is destined—at some point—to fail.

How could it be otherwise? Is it reasonable to expect services to compete in complete isolation with one another for common resources? Is it reasonable to expect an IT department to struggle on a daily basis with no means of holistic management or effective troubleshooting? Is it reasonable to expect users to sacrifice quality for mobility? The only way to achieve the desired result—the ability to access any application from any location using any device, securely and reliably with satisfactory quality—is to manage the network as the inextricably intertwined system that it is. Attempting anything less is a recipe for grief and ultimately failure.

Unwiring the enterprise with unified mobility services promises to deliver peak performance and quality for all users, while simultaneously enabling new productivity enhancing applications, and driving down the implementation and management costs. Users will appreciate having the improved work/life balance made possible with full mobility; the CIO will appreciate being able to provide holistic and centralized control for the IT staff; and the CEO will appreciate that all stakeholders have been satisfied without busting the budget.

To learn more about how your campus network can benefit from a unified approach to mobility services, visit Juniper Networks on the Web at **www.juniper.net/us/en/products-services/wireless**.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at **www.juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King’s Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

2000425-001-EN Aug 2011 Printed on recycled paper