TRAPEZE NETWORKS

A BELDEN BRAND

How Will the "Tablet Generation" Affect Wireless LANs in Primary Education?

WHITE PAPER

The proliferation of wireless-ready notebooks, tablets and smart phones continues unabated. This white paper examines the security, performance, management and reliability implications of supporting the tablet generation in schools.

Introduction

Maybe Apple, Amazon, Google and others are onto something. The evidence for the popularity of new categories of wireless-enabled, Internet-ready, always-on, hand-held devices is mounting. Apple's iPad and iPod, Amazon's Kindle and Google's Android are all soaring in popularity. So too is the enthusiasm for all the many other tablets, eBook readers, netbooks and notebooks, smartphones and plethora of similar portable platforms now on the market—all with wireless networking interfaces built-in.

Sales of laptops overtook desktops in 2007, according to Frost & Sullivan. Of course, laptops long ago dominated the landscape at universities and for mobile workers. Frost & Sullivan, along with other analyst firms like iSuppli and IDC, all now concur that laptops (not including tablets) will dominate the PC market going forward. IDC expects portable PCs will account for 70% of the market by 2012—a number that again excludes tablets, as IDC does not consider a tablet to be a PC. Gartner also predicts that mobile PCs will have a 70% share of the market by 2012.

As for smartphones, including the popular iPhone, nearly half are now "dual-mode" with Wi-Fi built-in, according to ABI Research. ABI predicts that by 2014 a full 90% will incorporate Wi-Fi, even with the advent of 3G and 4G replacing GSM cellular services. In terms of units, ABI is forecasting an increase from 144 million Wi-Fi-equipped smartphones in 2009 to 300 million in 2011 and 520 million by 2014. This trend includes students, naturally, and many children are now getting their first mobile phone before the age of ten.

The availability of so many dual-mode mobile phones and dedicated Voice over IP over wireless LAN (VoWLAN) handsets now provides a practical and cost-effective way to get a much needed phone in every classroom. Integrating mobility with fixed-mobile convergence (FMC) and fully Unified Communications (UC) into the school's IP-PBX holds tremendous potential for improving productivity for teachers and staff, while also improving public safety for all. Users will be able to place and receive calls on the device of their choosing and/or based on their current location. Unified messaging will enable voicemail messages to be converted to emails, or vice versa. These and other advanced calling features will empower users to determine when and how they prefer to communicate.

Enter the tablet. PC World magazine has already declared 2010 as "The Year of the Tablet Computer"—in the year of the iPad's debut! And then there are the eBook readers. ZDNet calculated that Amazon had sold some 1.5 million Kindles by Christmas 2009, and on that very same Christmas day, electronic books outsold their paper counterparts. Not all eBook readers have Wi-Fi, of course, but the trend is clear: People now expect to be both mobile and connected—24x7. With the introduction of the

There is no "one size fits all" approach for deploying 802.11n.

iPad, Apple has helped to define a huge market segment for tablet computing, including as an eBook reader, and has likely kicked off a tablet war that will accelerate innovation and market adoption. And according to Gerry Purdy, the chief analyst of mobile and wireless for Frost & Sullivan, "The way the stars are aligned, it won't be long before someone adapts eBooks out of the consumer space and makes textbooks available on these portable devices."

In today's wireless world, some people prefer never to be tethered at all. And those numbers are growing. Indeed, many of today's youth have never experience the Internet with wires. Fortunately, the high data rates supported by IEEE 802.11n now make an all-wireless school a real possibility. But can the wireless LAN really scale well enough to cut the cord completely? This question is critical because unlike Ethernet, which is now deployed mostly in a switched, point-to-point topology, Wi-Fi remains a shared medium access network. More devices will cause more contention, and that can create new challenges without the right approach. How will the network, including the security provisions, be managed effectively? What impact will a pervasive wireless LAN access have on network reliability? This white paper endeavors to answer these and other questions about the effects of wireless device proliferation on the wireless LAN. The discussion is organized into four sections on performance and scalability, security, reliability and management, followed by a brief conclusion.

Performance & Scalability

The latest, high-performance 802.11n WLAN technology promises to revolutionize wireless networks with substantial gains in throughput and range compared to legacy 802.11a/b/g systems. Supporting data rates of up to 300 Mbps, 802.11n is six times faster than current 802.11a/g technology, and 802.11n's greater efficiency delivers about 10 times the throughput of these legacy protocols.

But as with any new technology, there are new challenges. School systems adopting this new standard have learned that there is no "one-size-fits-all" 802.11n solution because real-world deployments require trade-offs when deploying access points for either maximum coverage or maximum capacity. With the proliferation of wireless devices, the deployment of WLAN access points should clearly favor higher capacity over increased coverage.

Why is the capacity vs. coverage tradeoff such a fundamental consideration when designing a wireless LAN for a high density of users? The reason is price/performance, with both price (i.e. total cost of ownership) and performance (i.e. throughput and quality of service) also being separate considerations. This may seem counter-intuitive at first, but the rationale is worth understanding in some detail.

The problem with configuring access points for maximum coverage is that the wireless LAN ultimately fails to scale without a major redesign—and at a considerable cost. In some situations this may not be a problem, such as in cafeterias, sports fields and other outdoor gathering places that need only support relatively few active users with low throughput demands. The nature of the problem, which is rooted in the fact that wireless networking utilizes a shared medium, is revealed in the following diagram.

When maximizing for coverage, costs are kept low by deploying fewer access points and turning up the radio frequency (RF) signal power as high as possible. But this seemingly cost-effective approach, which may be adequate for the initial wireless rollout, soon becomes insufficient once user adoption reaches a certain threshold, making it inappropriate for supporting the numerous devices and applications now requiring high levels of throughput and quality of service.

The maximum-capacity approach is, therefore, far more appropriate for accommodating the proliferation of high-performance wireless devices. With maximum capacity as the primary goal, the WLAN infrastructure is better able to support a large number of high-bandwidth devices and real-time applications, such as VoWLAN and location services. Fortunately, the flexibility built into 802.11n Multiple-Input Multiple-Output (MIMO) technology enables the WLAN to be optimized for scalable capacity cost-effectively without sacrificing performance.

Deploy access points for maximum capacity and not maximum coverage.

Designing for maximum capacity requires deploying more access points at lower RF power settings and with the lowest data rates disabled in the outlying bands. More access points operating at lower power settings also has the advantage reducing RF interference, especially in the 5 GHz band with its greater supply of non-overlapping channels. (How to best utilize the unlicensed spectrum is covered in greater detail below.) The result of using smaller, high-performance cells is a dramatic improvement in both throughput and quality of service for all WLAN users.

The use of lower power settings has cost benefits, as well, because access points designed for maximum capacity do not require certain expensive features needed for maximizing coverage, and are therefore less expensive. To achieve the desired capacity most cost-effectively, it is best to utilize either 2x2 or 2x3 802.11n MIMO access points in most locations, while restricting the use of more expensive 3x3 MIMO access points only to those areas where a low concentration of users is expected. Designers should also be careful not to waste money on features that are not truly necessary, which is often the case with high-end 3x3 MIMO access points.

Astute readers will quickly realize that deploying multiple access points designed for maximum capacity will cost more than deploying them for maximum coverage. But this is only true when considering cost on a per-square-foot basis. But if the metric is performance or available bandwidth per square-foot, as it should be with the proliferation of wireless devices, then cost should be evaluated on a per-Mbps basis. And by this measure, 2x2 and 2x3 MIMO access points deliver far superior price/performance. And as new applications are added to the wireless LAN, such as multimedia learning, online testing, interactive whiteboards and tailored curricula, designing the network for peak performance from the outset will prevent the need for a costly upgrade later.

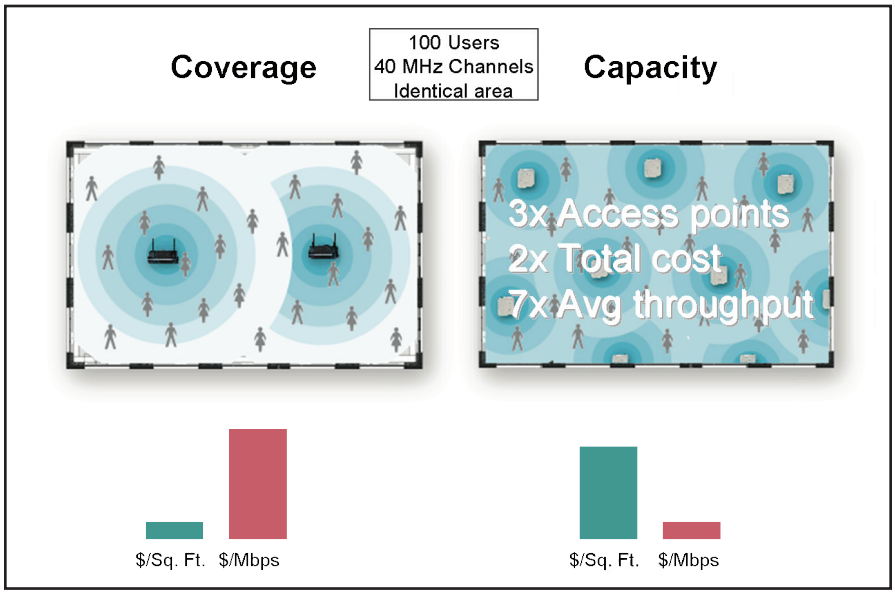

The cost of deploying access points is actually lower on a per-Mbps basis (as well as a per-Mbps-per square foot basis) when configured to favor maximum capacity.

TRAPEZE NETWORKS
A BELDEN BRAND

Band steering is critical to optimizing performance in the wireless LAN.

Achieving Peak Performance and High Quality

Deploying the latest 802.11n MIMO technology is necessary to accommodate the proliferation of wireless devices, but it may not be sufficient in many situations. The reason is that scaling a network to support more users often leads to degradation in throughput or quality of service (QoS) because Wi-Fi is a shared medium. Indeed, is a network really scalable if adding users causes performance and quality to degrade? True scalability requires some additional resource management and provisioning capabilities, such as band steering, distributed forwarding, dynamic load-balancing and special QoS provisions, including Call Admission Control (CAC) for VoWLAN sessions.

Band steering is a feature of some access points for directing clients to either the 2.4 GHz or the 5 GHz band. This can be important with the proliferation of tablets and smartphones, as most of these devices employ the 2.4 GHz band to conserve power. Leaving it to the device to choose bands is not recommended because a substantial majority defaults to 2.4 GHz making this spectrum too crowded, while the higher capacity of the 5 GHz spectrum is left mostly idle. To maximize capacity for all users, as many laptops as possible should be steered to the 5 GHz band. With their larger batteries, laptops can readily provide the additional power needed to operate in this higher frequency range.

The most effective way of achieving maximum performance and scalability is to use a combination of distributed forwarding and dynamic load-balancing. Distributed forwarding involves the ability of some decentralized WLAN architectures to leverage the processing power in access points to make forwarding decisions. By switching as much traffic as possible at the network's edge, distributed forwarding has the advantage of minimizing the central controller's load (while still preserving centralized management and control). And by avoiding the central controller, local switching also provides higher resiliency; if a controller fails, it may not be possible to initiate new sessions, but existing ones can continue to operate.

Dynamic load-balancing among radios, access points and controllers leverages distributed forwarding to help keep the traffic flows throughout the network optimized for peak performance. Balancing the load avoids the congestion that occurs based on the tendency of users to congregate in common areas, where they often aggregate on the same radio and/or access point. This problem is exacerbated by always-on mobile devices, such as dual-mode smartphones. When students enter a building, for example, their devices normally associate with the access point offering the strongest signal, which often results in unbalanced access point and controller loads.

Centralized Forwarding vs. Distributed Forwarding for Voice Traffic

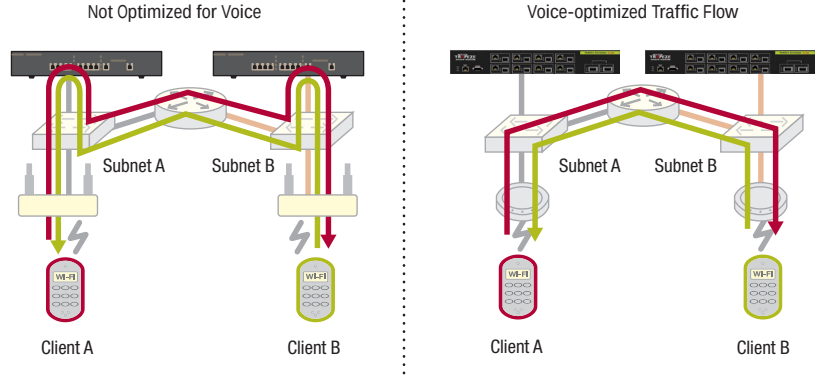

For multimedia and other latency-sensitive traffic, distributed forwarding has the additional advantage of utilizing the shortest possible path to minimize the round-trip latency.

Load-balancing at both the client and access point levels dramatically improves performance.

In addition to band steering techniques, which can reclaim 30-40% extra capacity in a crowded network, two other load-balancing approaches are also available for improving aggregate throughput; a combination of both achieves the best results. The first, Client Load Balancing is at the client level, where clients are transparently forced to associate with alternate access points within their range, instead of all congregating on the one nearest the common point of entry. Note that this is very different from balancing the load among VLANs, which yields little or no improvement in performance for users who continue to aggregate on the same access point. The second, Access Point Load Balancing is at the access point level where the management and control of access points is spread among available controllers, which also ensures that any traffic requiring centralized forwarding is handled more evenly among those same controllers.

Maintaining sufficient quality of service with more users requiring more access to more sophisticated applications can be particularly challenging, especially for VoWLAN and other real-time traffic. One effective means for accommodating high volumes of VoWLAN traffic is the use of call admission control (CAC). CAC's role is to limit the number of active Voice over IP (VoIP) sessions (normally only for teachers and staff) to avoid a situation where the network becomes so congested that quality deteriorates for everyone. The Public Switched Telephone Network (PSTN) also has a form of CAC to limit the number of calls, recognizable by the familiar "all circuits busy" message or "fast" busy signal.

As with anything else, some CAC implementations are better than others, however. The most basic "static" implementations block new sessions based solely on device capabilities and fail to take into account active VoIP calls. This approach unnecessarily keeps certain users from accessing the network even when there is adequate capacity available, thereby undermining the objective of ensuring good performance for all users. The more robust dynamic CAC implementations are fully aware of actual access point resource utilization in real-time. Dynamic CAC recognizes the always-on nature of the many VoWLAN-enabled devices, and only considers those with current VoIP sessions in its call count.

For VoWLAN calls it is also important to take into account the need for more air time for devices operating a lower data rates. For this reason, the QoS provisions should be capable of differentiating between 802.11b and 802.11g devices, and permit sufficiently greater air time for the slower 802.11b sessions.

Finally, the QoS provisions should be able to enforce some degree of "fairness" among the many users on the wireless LAN. This is particularly important in schools where some students become almost addicted to being online. One effective way to accomplish this is to utilize the accounting feature in RADIUS or other directory system to track utilization by all users over time. Rules for "fairness" can then be established and applied by temporarily throttling back the usage of both chronic and acute bandwidth abusers as necessary.

Inside Out

Students and teachers now desire wireless access everywhere: in classrooms, in hallways, while eating lunch, and yes, even outdoors. Unfortunately, some wireless LAN solutions are not seamless between the indoors and the great outdoors. Some have completely different product families built on different architectures for outdoor versus indoor deployment. And some have no or only limited RF planning capabilities for outdoor environments. The reason for the lack of seamlessness is simple: indoor and outdoor networking requirements are quite different. Outdoor wireless LANs require features like bridging, meshing, filtering, special antenna systems and more to overcome issues like the lack of access to wired Ethernet and other technical obstacles related to limited bandwidth and higher range requirements.

Users shouldn't need to care about any of this, of course; they simply expect seamless mobility as they move about, in and out. The IT staff does need to care, though, which makes the need for a common indoor/outdoor architecture compelling, particularly as the number of always-on wireless devices continues to proliferate. The common architecture should also have common RF planning tools and a common management system to minimize the burden on the staff. And the service and security profiles should extend end-to-end across the entire

The proliferation of different devices requires a more flexible and granular approach to security.

network. Then and only then will users enjoy the seamless mobility they now expect—and deserve.

Security

Security is often a primary concern in school systems, and the proliferation of mobile devices can compromise security without the right tools and techniques. The challenge involves accommodating the many different users, many of whom may have multiple devices running a variety of different applications as they roam to different locations. Some of these devices will be secured with up-to-date anti-virus software and other endpoint security provisions. But many won't be secure—ever. It will be necessary, therefore, to treat different devices differently, even for the same user, whether a student or staff member. Then there are the anonymous guest users who come and go, and also expect to be granted at least some access privileges.

For these reasons, the wireless LAN security provisions need to be flexible enough to deliver different levels of service to different device types depending on how secure they are, yet remain manageable to ensure that policies are applied constantly and consistently. These provisions will need to leverage the existing security infrastructure, of course, such as directory services for authentication and access control, and they will need to conform to industry standards for authentication and encryption. Furthermore, it must be possible to provision guest access, without tying up precious IT resources.

The best approach to securing a multitude of devices roaming about the wireless LAN infrastructure is central control with distributed enforcement. The need for centralized control is rather obvious and normally easy to achieve with any wireless LAN solution. The real challenge, therefore, becomes distributed enforcement, which must be as efficient as it is effective. Distributed enforcement is achieved by propagating security credentials among controllers, thereby enabling each controller's access points to instantly recognize existing users that roam into range. Among distributed enforcement's many advantages is its ability to deliver the fast roaming needed to avoid disrupting VoWLAN calls, even when crossing major network boundaries, such as roaming across controllers, or from inside to outside.

Another important element of the required flexibility is granularity. With the growing number of permutations and combinations of different users, devices and situations, greater granularity becomes critical when determining access control rules. Simply put: identity-based networking, while necessary, is no longer sufficient. Access control must now also take into account the user's device, its endpoint security, and potentially the user's location and other criteria.

Rather than simply try to preserve existing user security profiles that give all users consistent access wherever they roam, a more granular approach, based on real-time information gathered by the WLAN infrastructure, enables access privileges to be adjusted dynamically depending on what individual users are doing on which devices, when they are doing it and where they are. More granular and dynamic authentication and access control also allows network managers to, for example, lock-down bandwidth abusers, restrict guest access to controlled areas, such as meeting rooms, and even prevent certain network access based on the time of the day and/or the day of the week.

Here are some more examples of how dynamic authentication and access control provisions can be made more granular to accommodate different users, devices and applications:

- Trusted Network Connect can be used to validate endpoint integrity before granting certain access privileges.
- A restrictive Guest Access account can be used for some or all devices that lack end point security.
- Optionally, wireless intrusion detection/prevention systems (WIDS/WIPS) can be employed to provide a layer of security for users with unsecure endpoints.
- A combination of user, device and location can be used to grant anything from full access to no access—or something in-between depending on the time of day.

Network failure is simply no longer an option in today's wireless world.

- Users with certain devices can be assigned to a specific SSID or VLAN to segment traffic on the network.
- Different firewall rules and filters can be applied to different combinations of user groups and/or devices.

Special situations may call for even more granular security measures. For example, the school may want to allow access by tablets and laptops in an auditorium, but disable VoWLAN devices and dual-mode smartphones—or at least their VoIP calling capabilities. And during a test it may be prudent to prevent all forms of access. Schools normally want to monitor for E911 calls and issue an alert to everyone else nearby as a warning or a plea for assistance. Or in a widespread emergency situation, it may be necessary to shut down all non-essential access after issuing a notification suitable for the various device types. As shown by these few examples, the possible scenarios can grow exponentially with the proliferation of mobiles devices.

With different types of users and devices connecting via shared wireless access points, it may be desirable to segment certain groups in ways that prevent eavesdropping or tampering. This is normally done by utilizing different SSIDs (service set identifiers), and then segmenting the traffic for these groups based on user type/identity onto different Virtual LANs (VLANs). Teachers and staff should, for example, have their own separate VLANs to ensure secure communications.

Real-time location services (RTLS) that employ triangulation to precisely pinpoint a user's location have many useful applications, including for security. The degree of precision can vary, of course, but some users (especially guests) should not be allowed to transmit/receive data from sensitive locations. Or it may be desirable to prevent users who roam to a public area, potentially outdoors, from accessing certain applications or services, while still granting access to the Internet and VoWLAN calling.

The ultimate in flexible and granular access control requires one additional capability: deep packet inspection. DPI is the most accurate and often the only way to determine the specific applications and protocols being used in any session. Normally such awareness is used only to apply an appropriate QoS profile during the session. But this same deep awareness can and should have a more fundamental role in the mix of criteria used to determine the appropriate level of access.

Reliability

With the growing number of students, teachers and staff becoming increasingly dependent upon being mobile with constant (authorized) access, network failure is simply no longer an option. "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.

The best way to make the wireless LAN immune to widespread outages is to implement redundancy in critical systems, especially at the centralized controllers. The traditional way to deploy redundant controllers is to install a "hot standby" or secondary controller that takes over when the primary fails. For this to work seamlessly, the standby must constantly track the state of the primary, including all active sessions and their security associations. The secondary controller must also have some means for detecting a failure in the primary. This may all work well enough, but there is a major problem with the approach: it becomes very expensive as the network scales to accommodate more users.

A far superior way to implement redundancy is by virtualizing the controllers. Virtualization has become increasingly popular for both servers and network-attached storage based on both its cost-saving and reliability-enhancing benefits. The same approach and benefits now make sense for the mission-critical wireless LAN.

Maximum reliability is achieved with controller virtualization that supports either N:1 or N:N automatic failover for controllers, as well as for access points to controllers. While the N:N configuration is more difficult to engineer because it requires a distributed architecture at its foundation, those solutions that do offer it can achieve far superior redundancy for active sessions. N:N is a form of many-to-many redundancy where all controllers act as backups for one another, which requires each to carry a

Controller virtualization affords the redundancy needed for an "always on" wireless LAN.

copy of the configuration of all other controllers at all times. With this holistic, self-healing approach, all controllers are in constant operation, and whenever one fails, one or more others instantaneously takes over control of its access points and current users. The result is exactly what users want: an always-on wireless LAN.

In addition to the hitless failover provided, with no VoWLAN calls being dropped or any data sessions being disturbed, controller virtualization affords another major advantage: the ability to make in-service moves, adds and changes throughout the wireless infrastructure without interrupting service. For example, if another controller is added to scale overall capacity, the new controller assumes its fair share of the work as the total load becomes automatically rebalanced among all available controllers.

Resiliency for access points can be achieved by properly planning the network to ensure that access points are spaced closely enough to enable power levels to be adjusted automatically should a neighboring access point fail. This can be particularly challenging in schools with adjacent classrooms on multiple floors. Of course, computing the optimum placement of access points and reconciling all the variables associated with power settings and channel selection will require some pretty sophisticated 3D modeling capabilities, such as those depicted in the diagram below.

Traditional Approach: Hot Standby Controller

N:1 redundancy using a hot standby has at least one controller sitting idle, while the N:N redundancy afforded by virtualization utilizes all available controllers.

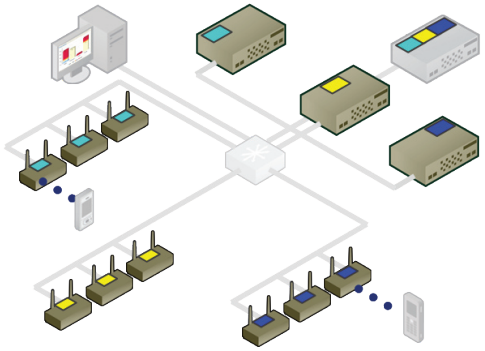

Smart Mobile Approach: Virtual Controller Cluster™

Smart Mobile clusters all MX controllers in the network so they act as a single Virtual Controller Cluster, providing NonStop availability.

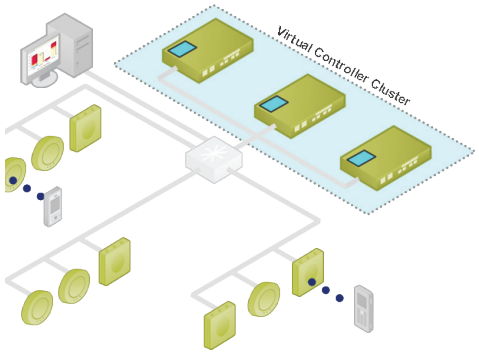

The cost of deploying access points is actually lower on a per-Mbps basis (as well as a per-Mbps-per square foot basis) when configured to favor maximum capacity.

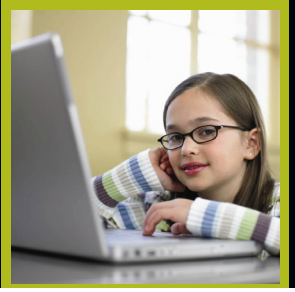

The management system must also scale along with the wireless LAN.

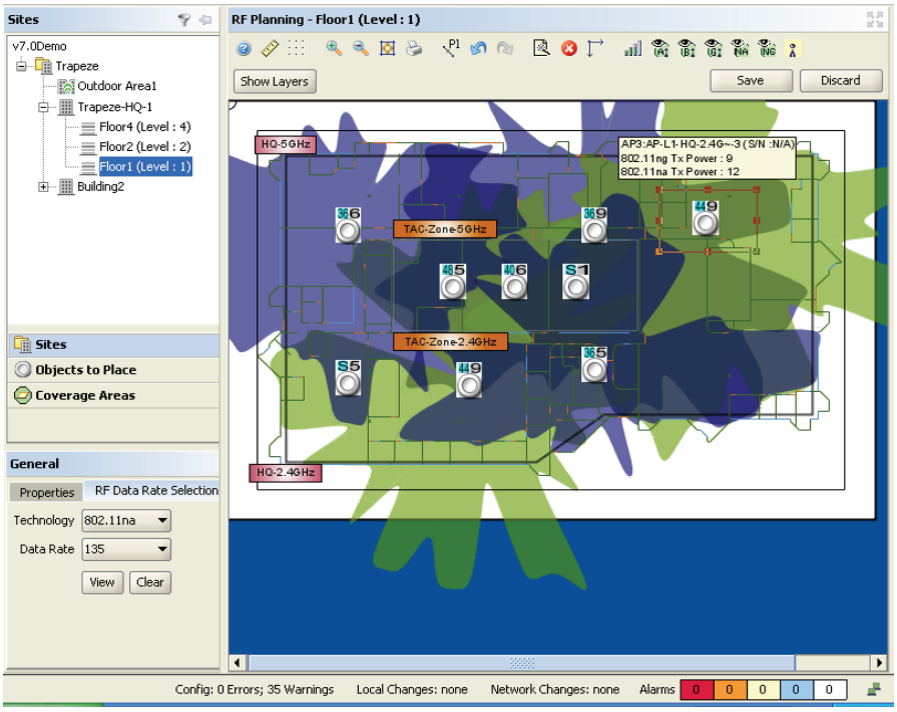

Properly planning for complete coverage with high performance and high resiliency requires sophisticated three dimensional modeling tools as shown here.

There are these additional steps that should be taken to minimize or eliminate other "soft" or hard failures throughout the network:

- Utilizing redundancies at the systems level, such as installing dual, hot-swappable power supplies in controllers and switches supplying Power-over-Ethernet (PoE).
- Using access points with dual Ethernet ports and PoE to provide both a dual-homed net work connection and a backup power source.
- Creating a mesh topology wherever wireless backhaul is needed, which is normally the case outdoors where the access points are untethered.
- Use of access points capable of detecting changes that affect attenuation (the movement of furniture, for example) and automatically recalibrating themselves and/or altering their receive/transmit power levels to adjust for these changes.
- With awareness of user locations, sessions can be switched to another access point within range during routine maintenance.
- Utilize load-balancing and redundancy for the servers used to implement authentication, authorization and accounting services.
- Employ distributed security enforcement to ensure that roaming occurs fast enough to maintain continuity of sessions and security associations as users transition from one access point to another, especially for VoWLAN calls when even a momentary disruption may seem like a network failure to the user.

Properly configuring a large-scale WLAN is virtually impossible without the right tools.

Management

Scaling a wireless LAN to support more users, more devices and more applications will obviously place greater demands on the management system. The proliferation of different devices will require the ability to manage access, QoS and security provisions at a far more granular level with finely-tuned profiles. Databases will grow. Authentication and access control rules will need to take into account more factors. Quality of service profiles will need to accommodate different devices being configured and used in different ways. IT managers will need to drill down more deeply into the details of these and other provisions.

Naturally, this increased degree of complexity makes everything from network design to troubleshooting more difficult without the right tools. The challenge, therefore, is to make the management system scale efficiently and effectively as the network itself scales without placing a proportionally larger burden on the management staff. The management server hardware may need to be upgraded, of course, to keep pace with new demands being placed on it. But the existing staff should be empowered with the tools needed to keep pace with all the planning, installation, monitoring and maintenance, troubleshooting and other tasks required on a daily basis. This is especially important in school systems that lack the budget for large staffs with high levels of expertise.

The Scalable Network Management Toolbox

A prudently planned and properly configured network is far easier to manage than one that is neither. So, good management begins with good planning tools. Planning tools that let network managers model the school building air-space in 3D, and take into account the building materials used, are essential for avoiding many of the common pitfalls when deploying wireless networks. Ornamental mirrors and steel lockers, for example, can play havoc with wireless signals, leaving unexpected coverage holes, if the network is not properly configured. Getting it just right is too complex for rules of thumb, and full site surveys can be prohibitively expensive. Hence advanced planning tools are an essential part of the network manager's toolkit. Predictive planning with automatic, orchestrated deployment simplifies configuring large-scale networks. It should be possible, for example, to ensure that all nodes are properly configured with the same software release, both during the installation and any subsequent upgrades. This is particularly important in large-scale networks because it minimizes and can even eliminate hard-to-diagnose version incompatibility problems.

An even more powerful configuration tool involves a link to the infrastructure vendor's support Web page that checks for firmware upgrades, and downloads and installs them automatically. This is a great way to ensure that each device is running the most current software/firmware image available. Alternatively, the school's IT staff may prefer the "perched" upgrade approach where all the firmware upgrades are downloaded and ready for installation, but await the IT manager's approval before actually being deploying. In either case, in the event an upgrade has unintended and undesirable effects, a good configuration tool will permit a full rollback to the previous stable state.

A good configuration tool will also guide staff through the service profile configuration process to create the many different services needed, such as highly secured data service for full-time staff, guest access, VoWLAN and real-time location-based asset tracking, among many others. The better configuration tools include special VoIP services support features so that technically confusing decisions for VoWLAN, such as whether WMM or WMM-PS should be turned on or not, need not be pondered by the network manager. To further simplify deployment some tools now even include built-in support for the most popular VoWLAN handsets from vendors like SpectraLink, Vocera, Avaya, Ascom, Nortel, etc.

A good security management tool will enforce common policies across the network so that users authenticating at a certain security level with associated authorization attributes will find that same level of authorized access regardless of where they roam. Alternatively, policies that permit different access privileges to different services in different locations should be just as easy to define and apply.

802.11n is necessary but not sufficient when scaling a WLAN for today's tablet generation.

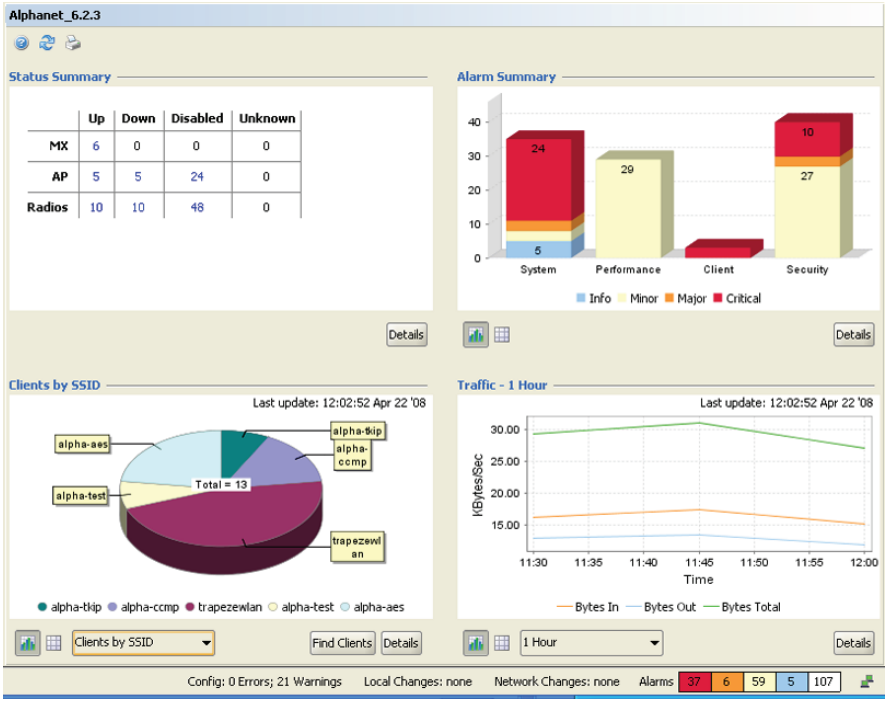

An at-a-glance overview of network status, combined with the ability to drill down into specific details, are absolute requirements for managing wireless LANs that need to support the proliferation of new devices.

The monitoring and reporting tools should provide both historical and real-time information in a variety of useful views and formats. The former is essential for spotting trends; the latter can be invaluable during trouble-shooting. The better management applications include a wide range of predefined reports, including equipment inventory, client session summary, rogue device summary, controller configurations, and more. In addition, the staff should have some means of producing customized reports that enhances their productivity.

The best management applications provide at-a-glance assessments of network status while also enabling the IT staff to drill down deep into the detail. Naturally, the ability to drill deep requires capturing the detailed data. This is why the best monitoring tools monitor network activity both in the aggregate and at a detailed device-by-device and/or client-by-client level. The network traffic patterns revealed can be useful for a variety of tasks, including identifying beneficial network configuration changes, refining access control rules, and adjusting QoS policies.

The better monitoring applications periodically conduct audits to check for serious conditions like missing or incorrectly-configured equipment and services, and automatically issue an alarm or other urgent notification whenever such problems are detected. The best tools further prioritize alarms into useful categories, such as informational only, medium, high and severe, and may even give the IT manager the ability to customize these categories to prioritize alarms based on the specific needs of the organization.

Finally, the toolset should support a variety of different interfaces, including an intuitive management application GUI, a Web-based portal and a command line interface (CLI), to enable staff to utilize whichever tool is most convenient or effective. Naturally, changes made via any interface should be synchronized for use with any others.

TRAPEZE NETWORKS
A BELDEN BRAND

Scaling the wireless network, while extracting 5-7 years of service from it, will be any school's biggest challenge

Conclusion

The proliferation of wireless devices can be expected to continue unabated and even accelerate in schools as the tablet war heats up. Network managers can expect this onslaught to have a profound impact on the wireless network. More users with more devices running more applications, often concurrently (such as a teacher using a VoWLAN phone for voice and a laptop for data applications), will become the norm.

Scaling the wireless network to deliver consistently high performance and quality of service, while extracting 5-7 years of service from the network, will be any school's biggest challenge. The use of 802.11n technology will be necessary but not sufficient. And not all 802.11n solutions will have the advanced capabilities needed to scale the wireless LAN in a way that keeps it secure, ensures reliable 24x7 operation and makes it easy to manage with existing staff.

This white paper has hopefully helped you prepare for the wireless proliferation, whether already upon you or coming soon. Additional information about how you can be fully prepared is available from Trapeze Networks at www.trapezenetworks.com, by sending an email to info@trapezenetworks.com or by calling 925.474.2200.