LEADING EDUCATION INNOVATION

A COSN
BEST PRACTICES
WHITE PAPER

Wireless Best Practices For Schools

Guidelines for School System Leaders

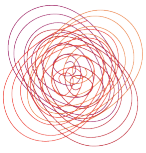

Brought to you by
Smart Education Networks by Design
a CoSN leadership initiative

April 2015

Table of Contents

Executive Overview

The requirements to implement a robust wireless network have vastly changed since wireless was first being deployed. Originally networks were deployed with 802.11b WiFi. The maximum speed of 11Mbps was an order of magnitude less than wired connectivity. Most end users connected with FastEthernet at a speed of 100Mbps. Most devices would primarily connect via a wired connection, with wireless often considered a luxury. During this period of time, most of the traffic traversing the network was text-based email with the occasional document or picture attachment. The typical implementation strategy for deploying wireless was to try to provide as much RF coverage as possible per access point, with limited channel selections. Often the desire was to utilize the fewest number of access points to cover the largest amount of real estate. The wireless network was either an open connection or secured with a pre-shared static key with simple (and what later found to be easily breakable) encryption.

Much has changed since these initial deployments. Maximum wireless connection speeds now equal or exceed those of the wired network and perform at ten times the speed of FastEthernet. The number of wireless devices connecting to the network has greatly increased, often surpassing those that connect via a wired port. In fact, many devices now have wireless connectivity as their only method to connect to the network..

Today's wireless networks are no longer being using for a single or limited number of primarily non-real time, text based applications. The network transports a variety of data including real time voice and video communications, multicast and streaming video, real time data, as well as non-real time communications. The same network that transports email and recreational web surfing is also supporting on line testing application, the delivery of course material to students, learning management systems, VoIP/Video Conferencing, security video as well as other mission critical applications. The same traffic that traversed the wired network is now traversing the wireless network. The wireless network now needs to have the same ability to identify and priorities applications, security, and resiliency one expects from the wired network.

In summary, a robust wireless infrastructure has become a critical requirement. Business critical data now traverses the wireless network such that users need to be authenticated with the network protected at all times to ensure it is as reliable and secure as any wired connection.

In this document we will discuss several considerations one should take into account when deploying a robust wireless infrastructure, so that it will not only address current needs, but also provide flexibility to adjust to future requirements and emerging technologies.

Coverage and RF Considerations

1. AP placement should be designed considering bandwidth and user density as opposed to simply providing RF coverage. The bandwidth requirements of the applications utilized by the student and the number of students in a given area should be used as components in the determining factors.

   For example let's say a school plans to deploy an application that leverages high quality streaming video as part of their curriculum. The average student device is going to require approximately 5Mbps per student to run this application. If there are between 25 to 30 students in a classroom, then one should plan for a shared aggregate bandwidth of about 125-150Mbps per access point for this application alone. Based on this, one can determine the required transmit power as well as cell size that will be needed in determining AP placement. Using a methodology similar to the previous example shifts the focus from solely providing RF signal to creating an infrastructure that can deliver an excellent end user experience when utilizing their required applications.

   This estimate was based off of a single planned application. One should factor for growth (i.e. proliferation of wireless devices), change in applications bandwidth requirements, and potential future applications as part of their planning process.

   Though data rates are important they are not the sole factor in determining AP placement and cell size. Cell size should overlap to a degree to facilitate seamless roaming for applications such as real time voice communication. Cells that overlap also provide the ability for adjacent APs to extend coverage to a hole in coverage due to an AP going off line. However, this overlap needs to be balanced so APs do not become a source of interference to each other. Automated RF management as part of the solution will help prevent interference between APs, but one should avoid the over deployment of APs since these automated RF management have limits based on the non-overlapping channels and power settings available.

2. A general rule with respect to AP cell coverage is that higher-performance 802.11ac has a smaller coverage area than previous technologies, partly due to higher frequencies used. In order to ensure coverage for proper roaming, base the site survey upon a coverage area similar to other 5GHz frequency protocols. APs serving both 2.4GHz and 5GHz frequency protocols should adjust the transmit power of the 2.4GHz to so that the overlap of AP coverage allows for roaming, but does not provide an unintentional source of interference to other APs.

3. The system should incorporate a method for radio frequency management that can dynamically assign AP channels, adjust AP transmit power, and provide coverage lapse

mitigation for the wireless infrastructure. In the past 802.11n only had 20MHz and 40MHz as mandatory channel widths. In the 802.11ac protocol, 80MHz has been added as a mandatory channel width. Due to this new requirement, 802.11ac RF management should be done at 20, 40, and 80MHz channels widths. Different user devices will support different channel widths within the given 802.11 protocols. Devices that support the wider channel widths will be able to support higher bandwidth within the particular protocol. In order to support the 1Gbps+ physical speeds provide by 802.11ac Wave 1, devices (such as laptops) will use 80Mhz channel width, thus this channel width should not be ignored with regards to radio and frequency management.

4. If possible predictive Wi-Fi planning and design should not be used in lieu of actual site surveys especially in challenging physical environments. Passive surveys show the channel plan and interference. Active surveys can show expected client experience. Performing an actual survey of the physical will provide more information about the actual RF characteristics of the geography than solely relying upon simulated or predicted modeling.

5. The system should be able to automatically adjust AP frequencies to avoid sources of interference in the 2.4 and 5 GHz bands. The system should be able to detect, identify, and locate source of RF interference as well as devices using 802.11based protocols. The system should be able to recognize other things that could operate in this unlicensed spectrum (i.e. Zigbee, Bluetooth, 2.4/5 GHz phones, etc.) and adjust accordingly. APs with built in dedicated hardware resources for RF spectrum analysis provide remote visibility into interference without negatively affecting client data performance. In addition to this standard mode of performing spectrum analysis while servicing clients, it is beneficial to have options to configure an AP to only monitor the spectrum for greater fidelity or be configured to act as a remote RF spectrum analysis sensor to facilitate troubleshooting with requiring the administrator to on site. Even if used only temporary to facilitate troubleshooting, these are valuable configurable options to have.

6. One should remember that 802.11ac only uses the 5GHz frequency range. 802.11ac Wave 1 requires 20, 40, and 80MHz channel widths, while Wave 2 increases the available data rates by adding the optional 160MHz channel width as well. These increases in data rates by increasing the channel width come at the cost of reducing the number of non-overlapping channels. Using 40 MHz channel bonding you have (9) non-overlapping channels and using 80 MHz wide channels you get (4) non-overlapping channels. This needs to be taken into account when doing the RF planning, channel planning, and site survey verification.

7. In order to maximize performance the wireless space and simplify deployment, try to minimize the number of SSIDs being broadcast into the environment. A target of three (3) SSIDs provides for a simple yet flexible deployment model. For example, a SSID

using a captive web portal would be used for guest access and 802.1x client provisioning. A second SSID would be used for 802.1x authenticated users and devices. The third SSID would be used for special uses cases or specialized, wireless devices (e.g. Wi-Fi-enabled VoIP phones, non-802.1x capable devices, or specialized network devices). Other special use case SSIDs may be needed depending on your specific needs, but one should strive to reach this three SSID target.

8. Beam forming is a process by which APs use the constructive interference of multiple radios to increase the fidelity of signals being sent to a wireless client. In order for beam forming to occur, there must be at least one extra radio per band than the number of spatial streams they wish to support for a client. For example, in order to support beam forming for a device that has 2 spatial streams, the AP will need 3 radios. The maximum number of special streams supported by high-end clients today is 3 spatial streams. In order to utilize beam forming for devices that support 3 spatial streams, 4 radios would be necessary.

9. In cases where APs can support both 2.4 and 5GHz frequencies, the system should have band steering mechanisms to move devices that are capable of running in the 5GHz spectrum to connect at those frequencies.

10. Plan for support of 802.11n and 802.11ac. Purchase equipment with 802.11ac.

Speed and Bandwidth

1. In order to maximize speed and facilitate roaming, one should disable lower data rates in support of legacy wireless protocols.

2. In addition to using 802.11k/v/r/e as methods to reduce client roaming times, one should look for wireless solutions that incorporates the use of real time data and advanced algorithms that use Radio Receiver Sensitivity levels and Received Strength Signal Indication (RSSI) thresholds to improve roaming for all client devices, including those that are not capable of supporting the aforementioned protocols.

3. Wireless implementations often use designs where user traffic is tunneled to centralized controllers, then forwarded to end destinations from there. Though this may facilitate some positive aspects of network design, in cases were maximum bandwidth and performance are desired, one should consider using a wireless architecture where the user traffic is distributed rather than centralized. Distributed deployment implementations include: wireless controllers that are imbedded in the

switching infrastructure; public cloud based wireless controllers; on-premise controllers that centralize AP configuration but have user traffic switched locally at the wired switch to which the AP is connected; as well as controllerless (i.e. stand alone APs). The first three methods listed still use "controllers" for management, but decentralize the user traffic. The last option may have a higher management overhead since each AP would have to be managed individually. These methods prevent the centralized controller from being a potential bottleneck. This decentralization of user traffic also allows for the full use of the wired switching infrastructure for both data transport as well as the use of potentially more advance QoS and network management features in the wired network.

4. The wired and wireless networks should have the ability to implement QoS (Quality of Service). In addition to a Good, Better, Best traffic categorization, administrators should consider systems that use newer more sophisticated QoS implementations, such as within category fair sharing, so that no single user can consume all of the defined bandwidth within a QoS classification. As additional users connect to the network, those utilizing the most bandwidth dynamically will receive more restrictions, so that everyone gets their "fair share".

5. In addition to traditional traffic marking methods, one should consider the use of application visibility and control to identify and give precedence to specific applications. Application visibility and control uses Deep Packet Inspection (DPI) to recognize and identify applications regardless of port number. This allow for better classification of traffic type other than relying solely on legacy IP port and protocol.

Security and Authentication

1. The primary method of authenticating school users should be WPA2 Enterprise using 802.1x for the authentication and 802.11i for the encryption. One should avoid pre-shared key (PSK) based authentication and security wherever possible. One should implement 802.1x utilizing an appropriate extensible authentication protocol (EAP) type (e.g. PEAP, EAP-TLS, etc.). For additional security one can authenticate access points to the network via 802.1x as well end user client devices. Though it may not be widely deployed currently in K-12 and may increase complexity of deployment, one should also considering enabling 802.1x on the wired infrastructure as well. Doing this mitigates the open wired network from being an attack vector for unauthenticated users as well as the connecting of rogue devices.

2. The RADIUS authentication server should leverage a centralized user directory store such as Active Directory or LDAP. A common user database used for authenticating people connecting to the network as well as access to network applications and servers greatly simplifies account management. It also allow one to use the robust user group categories found in these data stores as the basis of policies used in the RADIUS authentication required for 802.1x.

3. Network security access should be based upon a combination of extensible authentication and authorization policies. Factors in identifying access to the network should include, but not be limited to, the following criteria: user; group/department; membership; device type; time of day; device ownership (school owned vs. personal asset); etc. The policies should be ubiquitous and able to be enforced equally on both the wired and wireless infrastructure.

4. In addition to user authentication, device profiling can be used to determine the type of device and whether it is acting in a predicted manner for that type of device. For example, a network printer is not expected to be originating a lot of HTTP requests. A profiling system that supports multiple device probing strategies helps insure the fidelity of the device classification. Examples of probing methods to look for: network flow information (i.e. Netflow/IPFIX), DHCP/DHCP span, HTTP, RADIUS, DNS, SNMP Query, and SNMP Trap.

5. User access to network resources should be granted based on the aforementioned security policies, and limited through access control mechanisms. It is highly desirable to use new methods based on authenticated group membership. Upon being authenticated, user's data is tagged as being part of a specific group, based upon the identity of the user. This tagging should be done as close as possible to the point of connection to the network. Some technologies allow this to be done in the Ethernet frame itself. Access control policies can then be enforced based group membership irrespective of IP address, MAC address or other network abstract. This allows for policies to be written in an easier to understand way, such that a teacher can access "teacher" network assets while student cannot. In lieu having an identity based security system, security policy can be enforced using traditional network based methods such as Layer 3 ACLs (Access Control Lists), VLAN separation, and traditional firewalls. These legacy methods still require the administrator to be mindful of the specific network, IP, or other network construct when configuring the filter lists.

6. Guest traffic should be completely segregated from student, faculty, and staff data. Since most guest users should only utilizing the public Internet, it is highly desirable to implement a guest networking strategy where guest traffic is tunneled from their point of access directly to an Internet facing DMZ area, bypassing the school network. If desired, traffic can be inspected to determine if it meets the acceptable use policy (AUP).

7. Automated methods of 802.1x supplicant provisioning should be used in order to reduce the workload of the network staff. It is also highly desirable to have a system that supports a user's ability to self-register their devices. The system should also be flexible in the ability to support guest access through a variety of methods. These can include delegating guess account management to one or more users, guest sponsorship by a registered user, or self-registration. If self-registration is a desired method for guest access one should look for a system that incorporates some method of user identity verification, such as emailing/texting guest credentials to a mobile phone or using a social media login as credentials for guest access.

8. Rogue AP detection should be implemented. In addition to looking for rogue APs through wireless broadcasts, an effective rogue detection implementation should also correlate this information with location information from the wired environment as well. The system should optionally support methods of reducing the impact of rogue devices. Though recently considered controversial, techniques such as those involving the use of disassociation messages to prevent clients from associating to rogue access points can be used to contain rogue devices.

9. In addition preventing rogue devices, one may need to prevent users from using their devices to do direct peer-to-peer sharing of information on the wireless network. An effective wireless security implementation should also incorporate methods to prevent these activities as well.

10. The system should incorporate a Wireless Intrusion Detection System (WIDS). If possible, the WIDS should impose an "always on" methodology, allowing for WIDS even while servicing guest access. If an AP is servicing clients while performing WIDS functions, it is beneficial for the AP to have built in dedicated hardware resources for WIDS so not to impede the servicing of wireless client data requests. In addition, there should be optional setting such as WIDS-only and rogue containment modes. The system should also look for "off channel" rouge APs that may be servicing clients on non-standard 802.11 bands in the various frequency spectrums.

Wired and Wireless Infrastructure

1. Currently wireless systems using 802.11n and 802.11ac wave 1 are the most viable technologies to deliver the capabilities defined in the executive summary. The recommendation for any future purchases would be to purchase APs that support 802.11n/ac wave 1. Having the ability to upgrade the infrastructure to 802.11ac wave 2 could also prove beneficial.

2. In support of these current wireless protocols (i.e. 802.11n and 802.11ac Wave 1) one should strive to have access switches with Gigabit Ethernet connections for end user devices and with 10GE uplinks ports. In order to support newer wireless protocols with speeds in excess of 1Gbps (i.e. 802.11ac Wave 1, Wave 2, 802.11ad, 802.11ax, etc.), the access-switching infrastructure should have a migration strategy to support future speeds greater than 1Gbps on existing copper cabling as well as POE to end devices. This new technology is expected to increase speeds across existing Category 5E/6 to 5 Gbps up to 100 meters while still supporting POE.

3. The wired infrastructure should be able to provide power to the APs via Power over Ethernet (PoE). The access switches should at a bare minimum support 802.3af (up to 15.4W) of PoE. It is recommended that access switches support 802.3at (up to 25.5W) to provide power for newer/upcoming technologies such as 802.11ac Wave 2.

4. The wired access switches should have multiple, redundant paths to the distribution/core switches. This provides an increased amount of available bandwidth as well as provides additional reliability of the upstream data path to the core network.

5. The wired infrastructure should be able to support 802.1x as well as policies similar to those being deployed in the wireless network. Users will have a similar experience connecting to both wired and wireless infrastructure. Administrators will have a consistent, common policy that can be applied throughout the network. Additionally, implementing 802.1x on the wired side reduces the likelihood of rogue devices transmitting data into the network.

6. The wired network should support security services to mitigate network threats and man-in-the-middle attacks. Such features include dynamic ARP inspection, DHCP spoofing, IP source guard, control plane policing, denial of service protection, unknown unicast and multicast flood protection, as well as traffic storm control which can be used to protect the network against malicious activities.

7. In order to increase the reliability of the wireless network, one should implement wireless components that insure redundancy throughout the infrastructure. Controllers should be able to be deployed in a redundant fashion. One should strive to implement wireless controllers that can support clustered N+1 redundancy, active/standby, and active/active high availability. If possible the controllers should also be able to support client high availability as well to enhance resiliency of the wireless network.

Management and Location Services

1. A common management platform should be implemented for both wired and wireless device management. Doing this provides a common operating picture when managing the network and simplifies management of policies and configurations across both environments.

2. A network time protocol (NTP) server should be implemented, and NTP should be enabled on all infrastructure devices, both wired and wireless, to insure the synchronicity of timestamps of logging information. This will facilitate troubleshooting across the infrastructure.

3. Location services should be deployed to help facilitate location of assets. Device tracking over time should be an option. Also, the ability to do robust queries should be available (e.g. where are all my iPads? provide the location of all rogue devices.).

4. The location management system should allow for the import of floor plans into graphical representation on the screen. The resulting floor plan maps should be able to be "marked up" to indicate construction materials to increase the fidelity of RF heat map, location services, and assisted site survey features.

5. Advanced users may want to incorporate location-based services beyond asset tracking. Content can be delivered to the end user based on their location. For example, when near a cafeteria, devices could have today's menu available. Visitors could receive tour of the facility based on their location. The same system could be used to provide dwell time and other advance analytics to the administrators.

MDM (Mobile Device Management)

1. The system should be able to work with any of the major commercial mobile device management (MDM) solutions available. The system should be able to utilize the MDM compliance information as part of the access profile.

2. In addition to the ability to integrate with commercial systems, the solution should be able provide basic MDM functionality or a low cost "light MDM" solution.

Preparation for the Future

1. In order to protect the investment of the wireless system, it beneficial to not only evaluate current needs, but have a plan to support future requirements as needed. When implementing 802.11ac, consider the various options to upgrade from 802.11ac Wave 1 to Wave 2. Choosing equipment that is field upgradable or wholesale replacement of components at a later time are alternative upgrade strategies. A modular design that allows for future upgrades such as 802.11ac Wave 2 may be beneficial versus replacing the entire access point. One should perform an analysis of the potential for future upgrades to determine the best strategy for their particular needs.

2. Access points supporting upcoming wireless protocols (i.e. 802.11ac Wave 2, 802.11ad, 802.11ax, etc.) will likely have higher electrical power requirements than existing access points. In order to be prepared for these new technologies the recommendation is to implement 802.11at POE.

3. Newer wireless protocols (i.e. 802.11ac wav 2, 802.11ad, 802.11ax, etc.) have the potential to drive wireless data rates well over 1Gbps. Some of these protocols hold the promise to provide speeds up to 10Gbps.

4. As wireless continues to evolve, being able to provide location-based services in addition to data connectivity is becoming more important. Bluetooth Low Energy (Bluetooth LE or BLE) allows for proximity-based services. Currently many Wi-Fi implementations use triangulation based on received strength signal indication (RSSI) to determine location. These solutions provide location within 5-7m. There are other emerging technologies using hardware base phased arrays and algorithms such as angle of arrival (AoA) increase location accuracy to approximately 1m. In order to support future location base applications, consider solutions that have a forward migration strategy to support these new location services while minimizing the need to replace existing APs.

5. The network is now supporting many other types of traffic than just unicast IPv4 packets. A dual stack network interface with IPv4 and IPv6 both being active is often the default configuration of end user devices. Some devices will even prefer establishing IPv6 connections to traditional IPv4. Some applications use multicast since it allows multiple recipients to receive the same data without flooding to all users. Several consumer devices now use mDNS (multicast DNS) for services announcements. The network, both wireless and wired, should have configurable options to facilitate the deployments of these protocols. Though some of these protocols aren't currently widely deployed, it is beneficial to have the option to enable these features in the network if new applications require these services.

CoSN (Consortium for School Networking)
1025 Vermont Avenue, NW, Suite 1010
Washington, DC 20005
202.861.2672
info@CoSN.org

www.cosn.org