White Paper |

Enterprise

The All-Wireless Workplace Is Now Open for Business:

Using 802.11n As Your Primary Network LAN Deployments

Peter Thornycroft

ARUBA® networks

Introduction

Wi-Fi technology has been steadily improving for some years, to the extent that many workers now rely on wireless as their primary data connection to the corporate network. Wi-Fi infrastructure for manufacturing and retail organizations, hotels, universities and schools is already a $1 billion market, with annual growth in double digits. Adoption to date in enterprise offices, also known as 'carpeted space' has been slower, as many CIOs and users still regard a Wi-Fi connection as inferior to a wired Ethernet connection. 802.11n is a game changer because when properly deployed it has the potential to displace wired networks to enable a completely all-wireless workplace.

It is already accepted that a well-designed Wi-Fi network is more secure than a wired LAN connection. Likewise, millions of Wi-Fi phones are in use worldwide, demonstrating the maturity of multimedia over Wi-Fi technology. The latest Wi-Fi advance, 802.11n has now proven that Wi-Fi can offer higher performance than most wired Ethernet connections: 802.11n access points available before the end of 2007 will support data rates to 300 Mbps, superior to common 100 Mbps Ethernet connections. This 5x increase in speed over older Wi-Fi equipment removes the last serious objection to adoption of the all-wireless workplace concept, where no cables need be run to individual desks and workstations. As a result of 802.11n, the edge of the corporate network will finally become wireless.

The primary benefit of 802.11n is its superior radio performance, allowing connection at much higher data rates with saturated coverage that reduces the 'dark spots' with poor coverage that are sometimes experienced in legacy Wi-Fi networks.

Knowing these significant benefits, should an organization – whether a conference center, a university or an enterprise – immediately move to 802.11n? As with all new technologies in their infancy, there are potential issues that may not be resolved for some time. In the case of 802.11n, these include the risk of changes or legal challenges to the standard, the availability of clients, and the premium charged relative to well-established Wi-Fi options such as 802.11a and 802.11g.

This paper describes the technical advances in 802.11n, and predicts the future path of standards and certifications. It investigates the state of available silicon – the fundamental building block for a Wi-Fi radio – and assesses the risk that future developments make today's products incompatible or obsolete. Beyond its RF capabilities, 802.11n raises potential difficulties in powering access points, and some vendors have used the step increase in performance to justify 'new' architectural approaches. These claims are tested, along with a discussion of options when upgrading existing Wi-Fi networks to 802.11n.

The reader should use the information here to decide when and how to adopt 802.11n: some organizations will want to move quickly, and many should do so; others may adopt the technology at a more gradual rate that is paced, for example, by the availability of suitable clients. Either way, an 802.11n-based network has tremendous disruptive potential relative to legacy, port-based wired networks, offering organizations the opportunity to build a high-performance, robust wireless network with a long practical service life.

802.11n technology

The 802.11n standard is composed of several key technologies and features including:

- A high throughput physical layer (PHY) with new modulation and coding. The new PHY supports OFDM modulation with additional coding methods, preambles, multiple streams and beam-forming. These can support higher data rates, and a much larger range of data rates than earlier 802.11 standards. The MIMO technique that is synonymous with 802.11n belongs in this section;
- High throughput PHY with 40 MHz channels. Two adjacent 20 MHz channels are combined to create a single 40 MHz channel. This relatively simple technique, already used in some point-to-point bridges and consumer equipment, more than doubles the effective data rate for a given set of RF conditions;
- Efficient MAC using MAC aggregation. Two MAC aggregation methods are supported to efficiently pack smaller packets into a larger frame. This reduces the number of frames on the air, and reduces the time lost to contention for the medium, improving overall throughput;
- Efficient MAC via block acknowledgement. Particularly for streaming traffic such as video, this performance optimization technique enables one acknowledgement to cover many transmitted frames: an ack is no longer required for every frame. This technique was first introduced in the 802.11e standard;
- Power saving through power save multi-poll. This is an extension of the U-APSD and S-APSD concepts introduced in 802.11e to extend the battery life of mobile clients.

In developing 802.11n Draft 2.0, the basis for the Wi-Fi Alliance 'Draft-n' certification, the IEEE omitted, or made optional, a number of important features of the original 802.11n specification, notably explicit feedback messages used in advanced MIMO and beamforming (although MIMO is still a significant part of the standard). While the original standard included options for up to 4 transmitting and receiving antennas, as well as support for 4 spatial streams, initial 'Draft-n' products are limited to 3 antenna chains and 2 spatial streams. Thus, whereas the performance of 802.11n is often quoted as 600 Mbps, initial products for 'Draft-n' will only reach 300 Mbps under the best conditions.

Also, although the 'headline' data rates of 802.11n are often used, several effects serve to reduce the effective capacity of a cell. While 802.11n-Draft2.0 advertises rates to 300 Mbps, the expected performance of an actual 802.11n cell is between 100 and 200 Mbps, and it could certainly be less if clients connect over long distances, transmit short frames, or in the presence of legacy 802.11a/b/g clients. However, this still represents an increase of 5x over 802.11a/g technology.

Chip technology

Wi-Fi infrastructure vendors build access point units around RF and baseband chips supplied by various silicon vendors, and the characteristics and performance of the overall system – particularly the RF segment – are heavily dependent on the chips' performance and features. Infrastructure vendors add value by building systems around the RF chips, adding management, control and coordination of access points. Since many chips are programmable, system vendors frequently dig into the driver software to tweak particular features. That said, access points are typically limited by the silicon from which they are composed, and features that are missing in silicon cannot readily be added through workarounds.

A small number of silicon vendors build chips suitable for enterprise-class access points, and the design cycles of the chips mean that they can be classified by 'generations'. Within a generation there is some variety between silicon chip vendors, but capabilities and features are broadly comparable within any single generation.

From the perspective of enterprise network infrastructure, the first important generation of silicon was that built to the IEEE 802.11n-Draft2.0 specification, adopted by the Wi-Fi Alliance as Wi-Fi Draft-n. These chips became available as samples early in 2007, and access points incorporating these chips became the 'test bench' for the Wi-Fi Alliance certification that was announced in late June 2007. Assuming that these are the 'first generation' of enterprise 802.11n chips, their important characteristics can be summarized as follows:

- Feature set to Wi-Fi Alliance 'Draft-n' mandatory requirements (see below):
- Support for 2 or 3 antenna chains;
- Support for 2 spatial streams;
- 20 MHz and 40 MHz RF channels;
- Power consumption of 15W for a 3-antenna design 'single-radio' AP, or 20W for a 'dual-radio' AP;
- Board footprint of 2800 mm^2 for a 3-antenna design;
- Several shortcomings, such as incomplete DFS support in parts of the 5 GHz band, the effect of which restricts the number of available channels under national and international regulations, as well as the use of 802.11n greenfield mode.

These chips were used in - and to an extent defined – the Wi-Fi Alliance certification testbed, so they are all certified and any AP they power is inherently certifiable. However, as described above, these chips include compromises and have not been optimized in all areas. Aruba chose to leapfrog this first generation of 802.11n silicon, by waiting for second generation chips in which issues such as full DFS support have been corrected.

The second generation of chips does not extend the feature set beyond that required for Wi-Fi Alliance Draft-n certification, but it does offer several improvements:

- Power consumption is reduced through silicon integration and optimization, so a 3-antenna 'single-radio' AP now draws 12W, and a 'dual-radio' AP draws 17W – approaching the specified limit for 802.11af standard PoE;
- Silicon-level integration also allows the board area required for the 3-antenna reference design to be reduced to 1800mm^2. This means that a dual-radio AP including one 3x3 802.11n radio can fit in the same form-factor, and use the same mounting bracket as a dual-radio 802.11a/b/g unit;
- Performance is increased over the first generation of chips which ran out of horsepower when faced with certain permutations of traffic;
- DFS, greenfield mode and some other restrictions of earlier silicon are removed;
- Future generations of 802.11n chip will match this form-factor, allowing ongoing silicon and feature upgrades with minimal board-level and enclosure (plastics) redesign.

With the passing of time since early 2007, it is now reasonable to predict future developments in 802.11n chip capabilities:

- A third generation chip in early 2008 will leave the feature set largely unchanged, but continue the trend towards lower power consumption and lower cost. Some basic beamforming and space-time block coding (STBC) features may be introduced with this generation;
- A fourth generation, probably in mid-2009, will extend the feature set based on progress within the IEEE in defining the features omitted from Draft 2.0. Such features may include support for 4 antennas and up to 4 spatial streams, explicit feedback messages to assist in MIMO and enable beamforming functionality, and improved MAC aggregation, block ack, and direct link features.

As we look this far out, it is certainly possible that pressures on and among the IEEE, silicon designers, consumer electronics and enterprise infrastructure vendors will prevent any meaningful improvement to 802.11n-Draft2.0, in which case the cost-reduction, power-reduction cycles will continue until the next significant IEEE or other standard advance, on or about 2011.

Technology risks

All infrastructure vendors want the products they ship to have an adequate life in the field, despite some customers' observations about the financial benefits to vendors of short product lifecycles! With 802.11n, there was widespread concern in early 2007 that the 802.11n Draft 2.0 standard (then a future standard) was not stable, and that changes or corrections to the standard might make early 802.11n-compliant products quickly obsolete.

Events of the first three quarters of 2007 have alleviated these concerns. During this time, the IEEE worked intensively on the 802.11n specification, identifying some features as optional even as they deferred others, such as beamforming, for future definition. The result is that Draft 2.0, while still 500 pages, provides a relatively solid basis for designing interoperable devices with significantly higher performance than earlier versions of Wi-Fi.

The Wi-Fi Alliance took the frozen 802.11n-Draft2.0 and offered a second forum to separate the core requirements from the options, organizing interoperability events that resulted in a set of certification tests. Although the exercise proceeded extremely fast, when the certification was announced at the end of June, there were few loose ends, and no significant subsequent discrepancies have emerged.

There remains the possibility that as networks are rolled out, new interworking issues will be discovered, particularly between products using different vendors' silicon. However, given the amount of testing to date, such issues are likely to be minor and vendors may be able to provide solutions in software updates. The probability of one vendor's 'Draft-n' certified NIC failing to interoperate with another vendor's certified access point now appears to be well-contained.

What then of the second technology risk, that future developments in the standard or Wi-Fi Alliance certification will render current devices incompatible? Whereas this was a significant concern early in 2007, the chances of such an outcome are now also very small. Since all major silicon, client, and infrastructure vendors have announced 'Draft-n' products, these players all have a vested interest in assuring that the products they are shipping will enjoy a reasonable life: they cannot afford the damage to their brand that

would result from premature obsolescence. And as these players' representatives comprise majorities in the IEEE 802.11 and Wi-Fi Alliance groups, it is highly likely that whatever new standards emerge over that period will be backwards-compatible with Draft-2.0.

As discussed in the section on silicon, new features and functions will undoubtedly become available over time, but today's clients can be expected to maintain their current performance against future 802.11n infrastructure, and vice versa. This means that from a technology viewpoint, 'Draft-n' certified infrastructure purchased today can be expected to work with current and future 802.11n clients over at least the next 2-3 years.

Avenues of protection from technology risks: modular access points

The uncertainty over the backwards-compatibility of future 802.11n standards has engendered a refreshing variety of responses from enterprise infrastructure vendors. One major vendor has announced a 'future-proof' modular access point, where the radio units can be removed and replaced in the field. Aruba considered such an approach, but the following reasoning led to an alternate design:

- The risk of discovering a fault in current extensively tested hardware is very small;
- The risk of changes to the 802.11n standard, resulting in incompatible products (within a 2-3 year period), is very low;
- The utility of a modular software architecture for modifications, updates and/or upgrades is much higher, and much more cost-effective than a modular hardware design.

Then one must analyze the financial and other consequences of such a modular hardware strategy:

- The modular access point will be initially more expensive to build (whether or not the vendor sells it with a premium on the price) because modular is by definition more expensive than an integrated unit;
- The size of a modular access point will be larger than an integrated design;
- The cost of an upgrade extends beyond the cost of a new hardware module: the labor expense of reaching each access point (possibly in an above-ceiling location), disassembling/reassembling the unit, adding up to 4 antennas for future modules (3 today) and then retesting the updated unit will be significant. This labor cost, in addition to the capital expense of the new hardware, represents a significant barrier to justification for future upgrade programs;
- The Achilles heel of a modular design is the base unit. The more active components that are included in the base unit, the more likely the base itself will become obsolete. Dual-radio access points include two radio modules and one processor section: in making a modular version, the designer must choose whether to mount the processor in the base unit, risking premature obsolescence, or in the radio modules, further increasing costs.

Aruba weighed the drawbacks and benefits and determined that the higher initial cost and complexity far outweighed the benefits of potential future upgrades. By the time significantly better hardware becomes available, the cost and technology of new access points will make a wholesale replacement more attractive than a modular upgrade anyway. One need only remember that the widely touted strength of the original 802.11g access points was a slot for a future 802.11a radio. Although many thousands of modular access

points were shipped, an astonishingly small number – estimated at less than 3% - were ever revisited for upgrades. The price premium was enormous considering the utility of the modular design.

A much more likely scenario is that in three years time, advances in the standards and in silicon will allow higher performance when a new access point is paired with a new client. At that point, leading-edge technology companies will be ready to refresh either a part or the whole of their access point installed base to take advantage of this capability. And whether the original access point is swapped out for a new one, using the original cable and mounting bracket as Aruba envisages, or whether a new radio module is swapped into the older, modular access point base unit in the other scenario, the costs will likely be very comparable.

Powering access points

The current Draft-2.0 access points are power-hungry. This is a consequence of each 'radio' having up to three RF transmit-receive chains and antennas, chip layouts that were not optimized for power consumption, and the higher-speed packet-handling and encryption needs of 300 Mbps radio units (600 Mbps peak for a dual-radio access point).

A first-generation dual-radio 802.11n Draft-2.0 access point built to a silicon vendor's reference design can draw in excess of 20W. Aruba's second-generation access point draws 17W for an equivalent configuration, while a dual-radio 802.11a/b/g access point draws about 11W. The implications for network design depend on how the access point is powered:

- Most access points accept local DC power and can be powered by a plug-in power supply, but many enterprises do not like to use such supplies because they lack an AC outlet near the access point mounting location, fire regulations prohibit such materials in ceiling plenum spaces, or because of the extra labor required to install the power supply;
- Most LAN edge switches now provide power-over-Ethernet (PoE) to the IEEE 802.3af standard. Power is injected at the switch, and carried over Ethernet cabling to the access point. This was hitherto a satisfactory solution, but the nominal limit for 802.3af (at the device) is 12.95W;
- It is possible to use PoE with a 'mid-span' power injector, usually installed in the wiring closet next to the LAN edge switch. This is often used to avoid a LAN switch upgrade. Injectors with 802.3at-like specifications (see below) are now available, but older PoE units will be to 802.3af specifications.

The increased power draw of a Draft-2.0 access point is not a problem with a plug-in power supply, but it does present a problem for PoE installations. Vendors have again responded with a refreshing variety of solutions:

- Change the PoE specifications. A new standard, tentatively 802.3at, is under development and promises to deliver at least 30W to the client device. However, the standard will not be complete until 2008, and of course implementation will involve upgrades to LAN edge switches and mid-span injectors. An alternative approach is for a LAN edge switch vendor to invent a new, proprietary PoE protocol to achieve the same result, but this has obvious drawbacks. Aruba will support 802.3at on access points when it is ratified;
- Use dual-terminations. It is possible to run two cables to the access point (arguably at less than twice the cost of a single cable) and draw power over both of them. However, for existing installations at least, this would add considerable cost to an 802.11n installation, including the cost of the PoE and possibly an

additional port. Aruba investigated this, but found a superior solution without requiring a second cable run;

- Work within available power constraints. This is Aruba's approach, and it the end result is that a single Ethernet cable, using existing 802.3af PoE, can be used to power the new Draft-2.0 access point in full 3x3 MIMO mode under most circumstances.

To explain Aruba's solution, we must examine the 802.3af standard and its implementation in current LAN edge switches. The standard specifies that the powering device should supply a current of 350mA minimum with a voltage in the range from 44 – 57V. However, in practice the minimum voltage supplied from nearly all PoE switches and injectors is 50 – 52V (802.3at will specify a range of 51-57V). Taking 'typical' values for this voltage and for cable losses (most cables have lower resistance than the specification and shorter runs than 100 metres), the expected power delivered to an access point will be greater than 17W. This allows an Aruba dual-radio Draft-2.0 access point (one radio for 2.4 GHz and the other for 5 GHz) with 3 antennas and supporting two spatial streams, to be powered on most existing 802.3af PoE terminations. Such a design would not have been possible with earlier-generation silicon, which ran even 'hotter'.

However, a combination of in-specification but marginal components might fail to deliver the 17W requirement, so contingencies have been engineered. The access point senses the voltage available, and if it is insufficient to power the access point, it sequences through a series of actions, shutting down parts of the access point starting with one of the antenna chains on each radio. At this point the access point is within 802.3af specifications, so there will be no need to disable further functions. The Wi-Fi network is still functioning better than an 802.11a/g WLAN, typically in 2x2 mode, but the administrator will be alerted so as to deliver more power to the access point if desired. Note that future developments in access point components will continue the trend of reducing the power drawn, although moving from 3 to 4 antenna chains would somewhat increase power consumption.

The result is that Aruba's access point will deliver full 802.11n 3x3 MIMO performance in the overwhelming majority of installations, but where the power supplied is insufficient it will gracefully back-off as many functional blocks as necessary to maintain stability. In most cases this precludes the need to upgrade PoE infrastructure just to support an 802.11n deployment.

Flexible deployment strategies

By 2006, five years into enterprise Wi-Fi deployments, it was established that the 'thin AP' architecture with a centralized controller was the optimum architecture. This architecture solved the major problems of management and configuration, RF coordination and handover, and a host of minor issues such as protection against frequent access point hardware upgrade cycles. It also made possible a user-centric architecture using identity-based security, as centralized controllers are increasingly linked to WIDS and firewall systems.

However, the advent of 802.11n has opened the door to a re-appraisal of the centralized, thin-AP architecture. Some vendors, having built on controller hardware that cannot easily scale to the expected increase in data rates, are forced into a major design change. Others are lagging in the market and hence seek a 'new' and differentiating marketing message. In this search for meaningful messages, some of the engineering facts have been forgotten. The issues that elevated 'thin AP' designs to supremacy in the market remain paramount drivers for this architecture; while 802.11n changes some parameters of enterprise

infrastructure, it does not justify a new architectural approach if the underlying network architecture is already sound. The following section is an assessment of implications of 802.11n for the enterprise LAN and WAN, and the optimum architectural solution for each scenario.

An assessment of traffic generated in 802.11n networks

It is true that the highest data rate of an 802.11n network deployed at the end of 2007 (3x3, 2 spatial streams) is approximately 300 Mbps, compared with 54 Mbps for older 802.11a or 802.11g. Many observers have extrapolated from this to predict a six-fold increase in traffic throughout the network with, for example, a controller for 10 dependent access points having to process 6 Gbps of traffic. A moment's introspection will reveal the logical inconsistency of this argument. As long as the network covers the same number of users, working with the same devices and the same applications as before, this is a completely unreasonable calculation. While users may get faster responses (assuming the WLAN was previously a networking bottleneck), and hence complete more transactions in a given time, the average traffic levels on the wired side of the access point will be very little changed after an 802.11a/g to 802.11n upgrade. Peak traffic will almost certainly increase, but peaks across different access points will be uncorrelated.

Of course, an 802.11n upgrade may be accompanied by other changes. If previously wire-connected users become wireless, or if new applications are enabled on the new wireless network, traffic levels will increase. But they would not be expected to be greater than the traffic generated by the same population of wired users on the LAN edge switch. Since there are accepted design rules for sizing wired Ethernet LANs, we can follow these rules to predict the traffic carried by a wireless network.

The analysis above provides us with two methods of predicting the load generated by an 802.11n network. Where there is an existing 802.11a/b/g network to be upgraded, traffic levels can be expected to increase somewhat, perhaps doubling at the most. Where a wired workplace is to move to an all-wireless configuration, the traffic generated will be equivalent to that on the original wired network. The following section suggests some rules of thumb for calculating traffic.

Dual-radio-set cell with 50% legacy and 50% 802.11n clients, maximum expected traffic load		
AP radio	Client traffic mix	Total client traffic
2.4 GHz AP, 20 MHz channel, mixed-mode operation	50% 802.11g clients @ 36 Mbps	70 Mbps
	50% 802.11n 2x2 clients @ 104 Mbps	
5 GHz AP, mixed-mode operation	50% 802.11a clients @ 36 Mbps	136 Mbps
	50% 802.11n 2x2 clients @ 216 Mbps	
Total traffic (half-duplex)		206 Mbps

Note that the figures above are half-duplex, and include all the 802.11n PHY overhead. When this traffic is carried on Ethernet backhaul from the AP, the Ethernet PHY overhead will be substantially less than 802.11n, and the connection will be full-duplex. One can conclude that a dual-radio, dual-band 802.11n access point with a mix of legacy and 802.11n clients is unlikely to fill a 10/100 Ethernet connection with traffic.

Dual-radio-set cell with 100% 802.11n clients, maximum expected traffic load		
AP radio	Client traffic mix	Total client traffic
2.4 GHz AP, 20 MHz channel, mixed-mode operation	802.11n 2x2 clients @ 130 Mbps	130 Mbps
5 GHz AP, mixed-mode operation	802.11n 2x2 clients @ 270 Mbps	270 Mbps
Total traffic (half-duplex)		400 Mbps

The table above shows that as the number of 802.11n clients increases towards 100%, it is very likely that a 10/100 Ethernet connection will become a traffic bottleneck. The implication for WLAN design is that it will not be necessary to upgrade 802.11n access point backhaul to gigabit Ethernet (GE) until the client population is predominantly 802.11n based.

LAN traffic prediction based on Ethernet LAN design rules

Consider a section of an enterprise – whether a headquarters building or a multi-building campus – where most buildings house offices or cubicles, there are a number of data centers, and the LAN within and between buildings utilizes high capacity links. The standard LAN design for such buildings is edge switches in closets on each floor, distribution switches for each building, and core switches serving the data centers. Most traffic on the network is directed to (or from) users on the edge switches to servers in the data centers, and to the firewall for WAN and Internet connections, also located in the data centers.

With the introduction of 802.11n technology, the data flow above will remain unchanged. The 'standard' design for a centralized WLAN network connects access points to edge switches, and places mobility (WLAN) controllers in the data centers in front of the core network: all traffic arriving over the air at access points is directed via secure tunnels to the Mobility Controller and then on to the data center.

Two segments of this data-plane path deserve our attention. First, the access point connection to the LAN edge switch. While wired users may have been satisfied with individual 100 Mbps Ethernet connections, a number of users now share an access point, so the traffic carried on the wired side of the access point will be the sum of all the served users - perhaps as many as 5 -10 simultaneously. Thus it is important that the connection to the 802.11n access point should be GE wherever possible, as a 10/100 Base-T connection may become a bottleneck in the path.

As we follow the data path from the access point, the LAN edge switch, in-building cabling and distribution layer of the LAN are carrying very close to the same traffic load as before, so we would not expect to need an upgrade. On entering the data center, the switches and cabling again should not see any increased load. However, the Mobility Controller is likely to be subject to considerably more traffic than formerly, depending on the number of new users and applications on the Wi-Fi service.

In the analysis above, traffic flows follow the path from edge switches, through the distribution layer to the core network. In this case, placing WLAN Mobility Controllers at the core does not change traffic paths,

hence loading on network links should be unaffected. However, as some analysts have suggested, there are cases where departmental servers are housed in the same building as the people who use them, and there may be significant traffic within the building, or even across a floor of a building.

If a centralized WLAN is overlaid on this network, with the WLAN Mobility Controller in a distant data center, traffic flows will be changed. Traffic flowing within a building may now be backhauled to a data center Mobility Controller, then returned to the originating building. While the delays introduced by such 'tromboning' are trivial, there may be increased load on network switches and links and this may be undesirable. In such cases, flexible WLAN deployment architectures allow for a fully-managed slave Mobility Controller to be co-located with a LAN edge switch or a LAN distribution switch, providing a point where traffic can take the direct path to its destination.

The flexible deployment model for campus scenarios

Aruba has supported such a hierarchy of centralized and distributed forwarding architectures using Mobility Controllers for many years, with excellent results. As there are always exceptions to the 'standard' architectural model, we offer the following analysis:

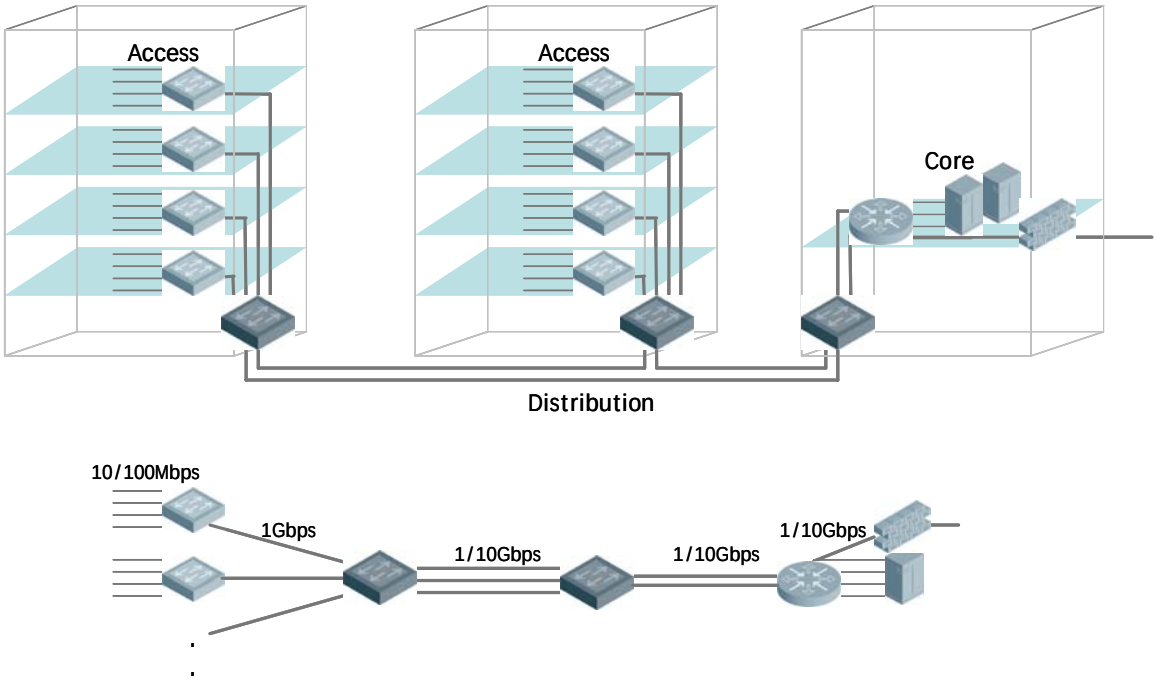

- The management plane, used for configuration, monitoring and central alarm generation, is anchored by the mobility management server – a network manager – that connects over LAN or WAN to all Mobility Controllers (several thousand controllers in Aruba's current architecture) in the network. Traffic rates on this plane are relatively low, the single point of management is the most important aspect. The management plane effectively terminates at Mobility Controllers, as in the 'thin AP' architecture, dependent access points are wholly managed by their parent controllers (but see exceptions below);
- The control plane, used for RF coordination, WIDS, handover between Mobility Controllers and access points and other near-real-time functions, is set up as a secure network of connections between Mobility Controllers. This allows, for example, fast handover between access points homed to different

controllers. The control plane is extended from Mobility Controllers to access points via secure tunnels set up on activation of the access point;

- The data plane has been discussed above. On a campus with high-bandwidth LAN links and high-capacity switches, the 'standard' model is to position Mobility Controllers in data centers. However, if traffic patterns differ from these expectations, Mobility Controllers may be placed in edge closets or co-located with distribution-layer switches, allowing traffic to 'short-circuit' the path to the data center and back to the edge. Aruba's flexible deployment architecture accommodates all of these options.

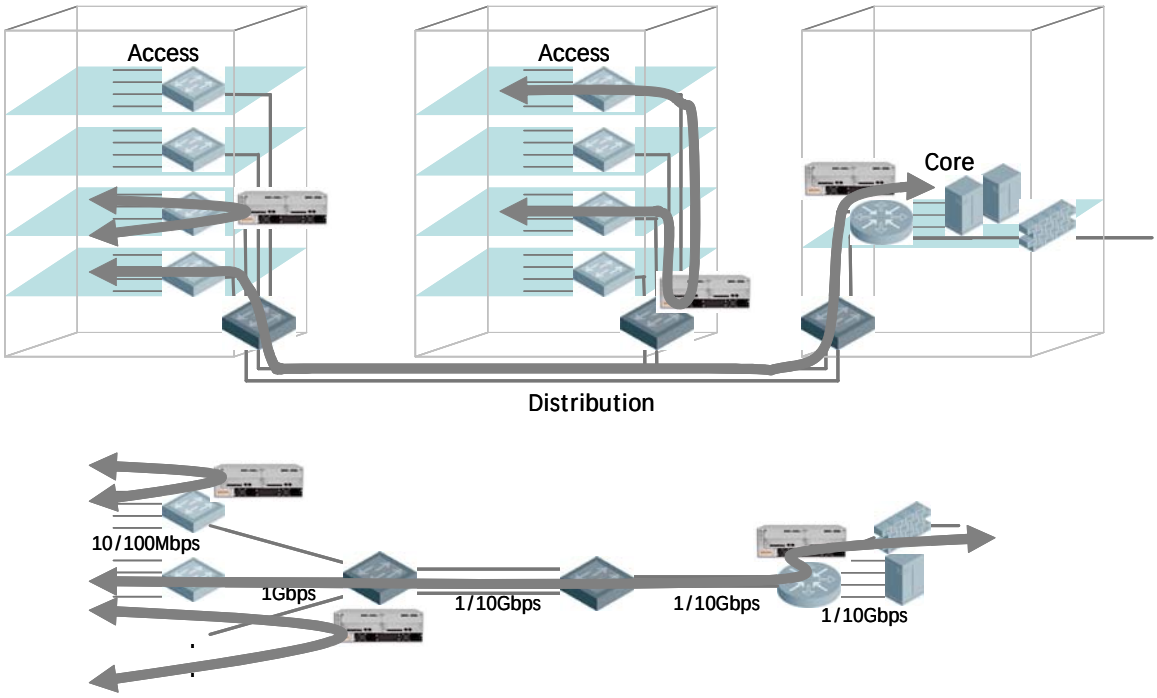

Estimating Mobility Controller load

An arms race has broken out as Wi-Fi infrastructure vendors tout ever more powerful WLAN controllers. But there has been little serious analysis of how much traffic will be processed by such leviathans in real enterprise networks. The discussion above suggested two approaches: when an existing Wi-Fi network is to be upgraded, based on the same users and applications, one can suggest a modest increase in traffic - maybe 2x for a rule of thumb – but there is no widely-accepted figure for the Mbps generated per user on such a network. For a new deployment of 802.11n, where wired users become wireless, one should take the existing wired network as a baseline, and here there is a precedent.

Wired LAN designs generally use a concentration factor of 8x at the edge switch, and 2x at the distribution layer. This means that for every 100 Mbps connection to an edge switch, the distribution side of that switch needs 12.8 Mbps. An edge switch with 48 users each at 100 Mbps would normally be provisioned for 600 Mbps on the riser side.

At the distribution layer the situation is more fluid, but let us assume, for example, that if the sum of the riser bandwidth is 12 Gbps, the connection between distribution switches might be 6 Gbps, for a factor of 2. If we

follow this rule of thumb for an 802.11n network, we can calculate that the Mobility Controller at the core would see approximately 300 Mbps of traffic for every 48 end users, or an average of 6 Mbps of traffic for every provisioned user on the network. Given that at any point in time many users are not connected, or are idle, this allows plenty of bandwidth for active users, to the point (the goal) that the network is not a bottleneck for user-observed application performance. Using this rule, a Mobility Controller with a capacity of 6 Gbps of traffic would be expected to service 1000 users.

WAN traffic considerations and the flexible deployment model

The section above discussed the requirements of a campus, characterized by high-bandwidth, short-distance links. Many enterprises include such a campus, but often they need to connect locations over a WAN, or over the Internet, where bandwidth, transmission delay and reliability may be concerns. The flexible deployment architecture provides solutions for these challenges, too.

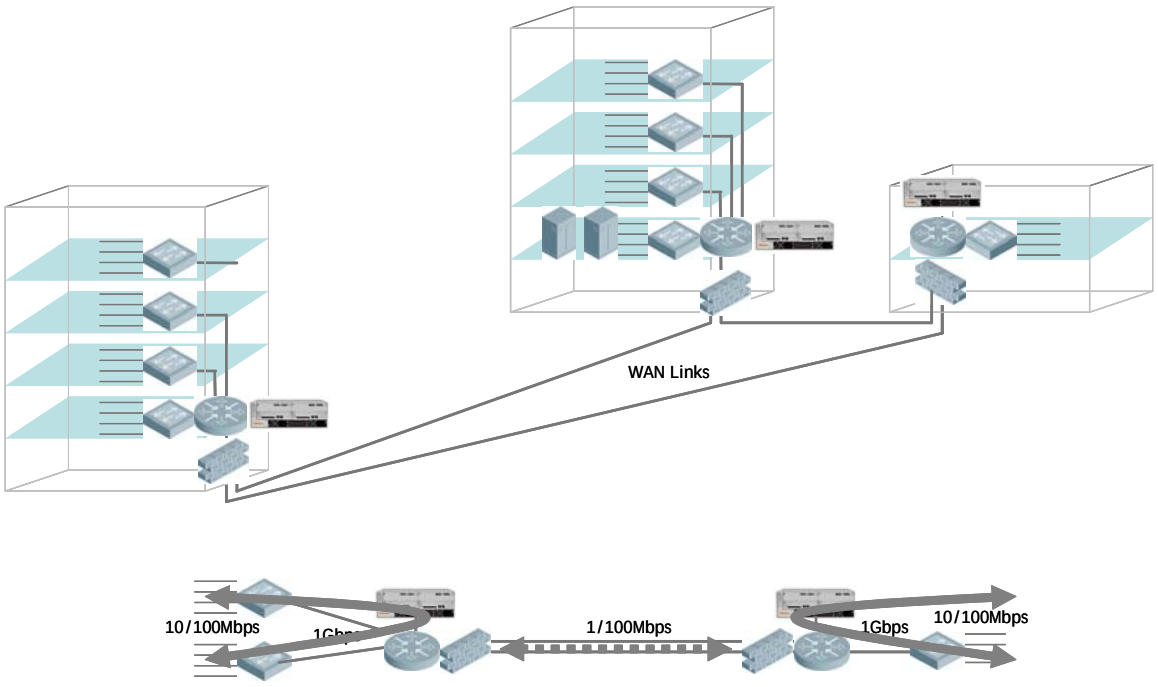

Where locations are large enough to require an on-site Mobility Controller, the model is as suggested above. Each site is provisioned with access points and a local Mobility Controller (or a network of controllers), terminating the control plane. Management is still provided centrally from the mobility management system. Traffic patterns on the data path are unchanged from a wired network.

Internet-connected sites and the flexible deployment model

Small sites, with only a handful of access points or a single access point, are a special case but an important one. For these sites, typified by the branch or home office served by a single access point, Aruba has built a solution known as the 'remote AP' using a split-tunnel forwarding architecture. A remote AP on installation sets up a secure L3 tunnel over the Internet to its parent controller, and can operate as a normal 'thin AP'. However, many branch offices have local servers and printers and much traffic may be destined for the Internet and not corporate servers: since the WAN bandwidth is limited, it makes sense not to trombone all

traffic to the distant Mobility Controller, but instead to pass it to a local LAN or directly to the ISP. For this application, Aruba developed some time ago a 'split tunnel' solution that allows such traffic to be split off directly from the access point, either for a local LAN connection or a direct Internet link, while corporate traffic is directed via the secure tunnel to the central location in the usual way.

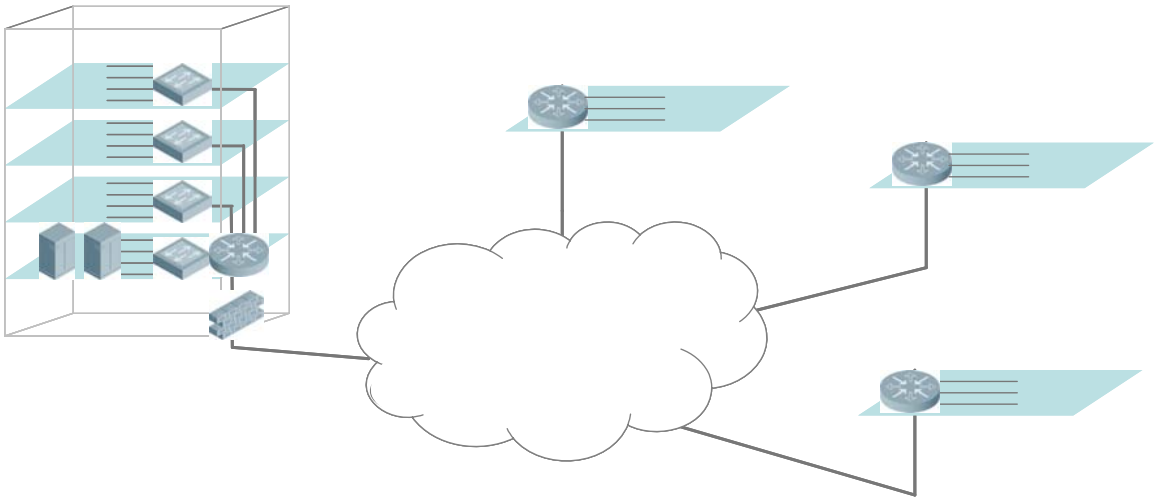

This remote AP feature allows configurations similar to those recently promoted by competing WLAN infrastructure vendors: data-plane traffic from the AP can be directed straight to the LAN connection without requiring backhaul to the Mobility Controller. Aruba believes such a topology is useful for remote offices, as described above, but offers no advantage for a larger office LAN installation. Indeed, it can be seen as a return to the 'fat AP' architecture where, for instance, VLANs must be configured on switches behind every access point rather than only at the Mobility Controller's location.

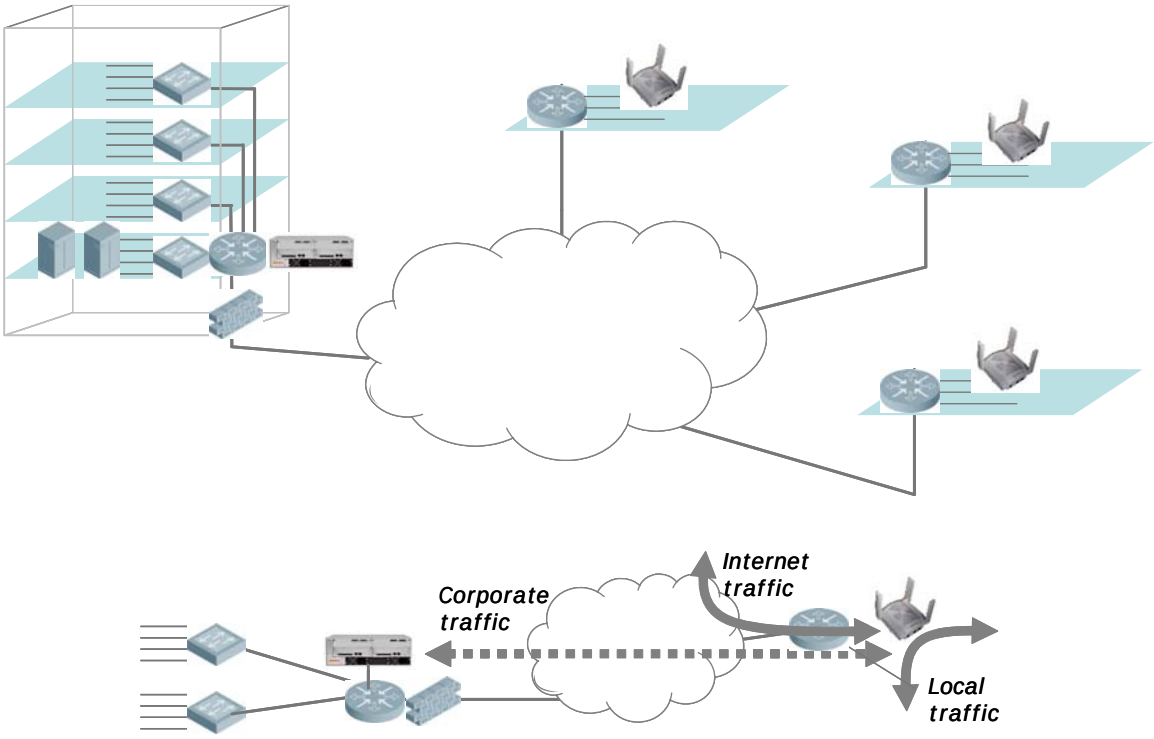

The overriding utility of Aruba's solution is its flexibility: the data-plane can be split-off upstream of the access point, or behind any Mobility Controller in the network as traffic patterns dictate. Yet the underlying network management and security policy engine never changes regardless of the forwarding model selected.

Implementation & migration strategies

Although 802.11n for enterprises is an infant technology, recent experience allows Aruba to make more detailed recommendations on upgrade strategies. While some of the considerations below are general and would apply to any Wi-Fi infrastructure vendor, others assume unique Aruba features and functions.

Wi-Fi clients

First a note about clients, as upgrading infrastructure without 802.11n-capable clients will offer little improvement in performance or reliability. The most common Wi-Fi client is the laptop PC. Most business PCs today ship with integral Wi-Fi cards, and the major PC vendors have already added 802.11n to many consumer and some more capable, business-oriented PCs. When purchasing new PCs, enterprises should ensure they are Wi-Fi certified to Draft-n and ensure that they sport at least two antennas and support at least two spatial streams. Most companies have implemented refresh programs, so over 2 to 3 years the majority of PCs will be 802.11n-capable. However, don't forget that even a small population of pre-n clients will affect the performance of the network, forcing it into 'legacy' mode so long as older 802.11a/b/g clients are active.

An alternative is to purchase 802.11n NIC cards for existing PCs. These will deliver satisfactory performance, avoiding 'legacy mode', but may fall short of embedded 802.11n as it is more difficult to mount multiple antennas on a small card, so MIMO performance may not be optimal.

Many organizations support a variety of Wi-Fi clients on their networks. In retail and manufacturing organizations, Wi-Fi bar code scanners are commonplace, and the use of voice over Wi-Fi with both single-mode Wi-Fi and dual-mode cellular/Wi-Fi phones is quite widespread. It is unlikely that the designers of these products will adopt 802.11n soon. They will face the same challenges of physical size, power consumption, and cost as access point vendors, but unlike infrastructure vendors, are unlikely to command a price premium for the functionality. Indeed, Wi-Fi phone vendors have only moved from 802.11b to 802.11g over the past year or so, and there are still very few 802.11a-capable embedded devices on the market.

We must conclude that in the great majority of enterprises, there will be a need to cater to 802.11a/b/g clients for many years hence. This is not an insurmountable problem, as 802.11n has several features enabling backwards-compatibility. However, it does mean that performance will be somewhat less than promised, and that access point spacing cannot be increased, as an 802.11n access point only has greater range (under Draft2.0 and without beamforming) when paired with an 802.11n client. In many ways this simplifies planning and migration programs, as one should assume using the same RF planning rules and access point locations as for 802.11a/b/g networks.

Migrating from 802.11a/b/g to 802.11n

Organizations with existing Wi-Fi service can choose from a variety of migration strategies when moving to 802.11n:

- *Network-wide substitution*. If the objective is to provide 802.11n coverage everywhere, in one upgrade sweep, it is possible to replace existing access points with 802.11n on a one-for-one basis. Assuming the previous network offered both 802.11b/g and 802.11a coverage, dual-radio 802.11n access points should be used. New cabling should not be necessary: each access point location will be served by Ethernet, and if 802.3af PoE is used, Aruba's access points will be able to take advantage of it. If possible, the LAN termination should be upgraded to GE to accommodate the higher peak data rates available with 802.11n. Access points will almost certainly be set to 'mixed' mode, as non-802.11n clients will be present;
- *Point substitution*. This is a variation on the scheme above, where only certain access points are upgraded for reasons of local traffic patterns or budget. The same considerations apply: removing the old access point, using the same cabling and PoE and, when replacing an Aruba AP 65, the same mounting bracket. The new access point will run in 'mixed' mode, and newer clients will switch to 802.11n mode when associating with it;
- *802.11n Overlay*. Assuming they are from Aruba, the older access points (assuming they are from Aruba) can be managed as part of a mixed network, so it is possible to leave them in place, and to deploy new 802.11n access points either across all locations or just in parts of the network where there is a need for better performance. This is likely to require cabling to the new access points, and new LAN switch ports. The most likely scenario would be to run the 802.11n access points (which could be single-radio) in the 5 GHz band, and in 'greenfield' mode, so they only accept connections from 802.11n clients – older clients would be serviced by existing access points.

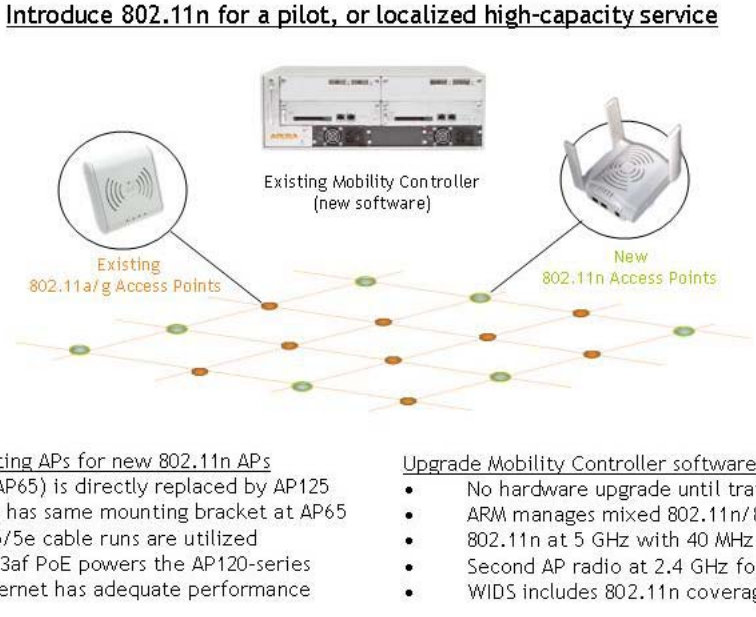

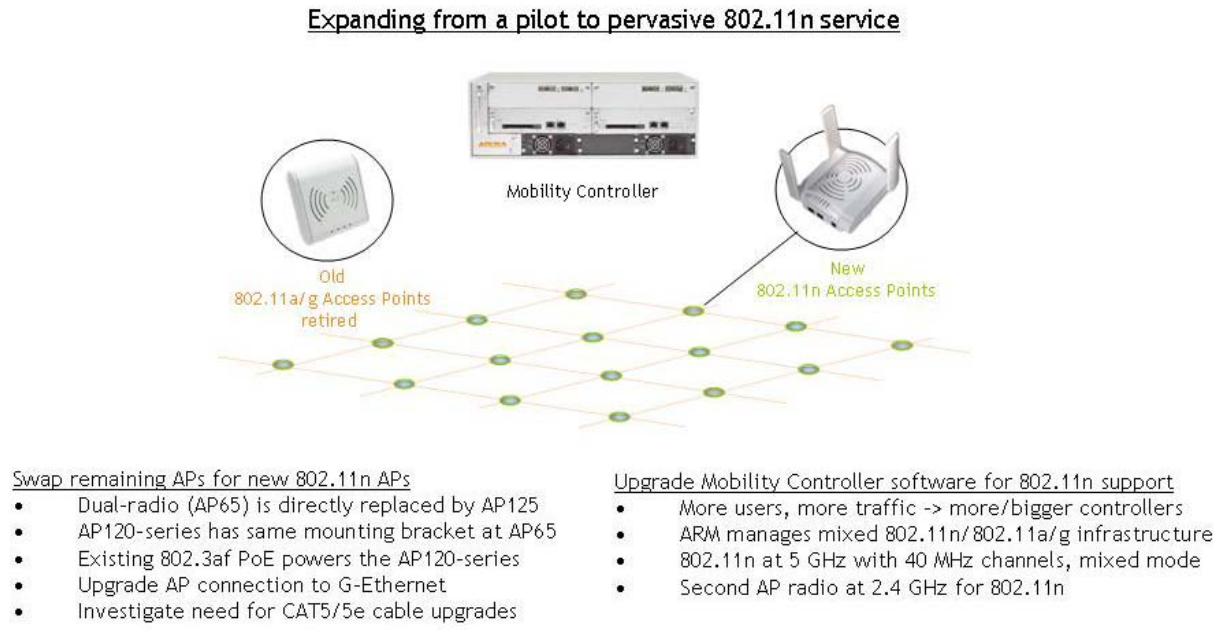

There are other considerations when planning for an 802.11n migration:

- *Strategy in 5 GHz spectrum*. Since the 5 GHz band supports on the order of 20 channels of 40 MHz each, Aruba recommends using 40 MHz channels for 802.11n. This allows operation when interspersed with older 802.11a access points, as many channels are available. Network managers should ensure that RF planning tools incorporate 40 MHz and 20 MHz channel support;
- *Strategy in 2.4 GHz spectrum*. There is some debate about the utility of 802.11n in the 2.4 GHz band, as there are only 3 non-overlapping channels, and the channels commonly used are non-contiguous, making them incompatible with the 20/40 MHz 'mixed' mode mechanism of 802.11n. Aruba recommends using 802.11n at 2.4 GHz, but with 20 MHz channels because this arrangement provides all the benefits of 802.11n, including MIMO, higher data rates and the various MAC enhancements, but allows the new access points to work in the existing 3- or 4-channel RF plan. Mixed-mode operation will be used in the usual case where older clients are present;
- *Gigabit Ethernet edge upgrades*. As mentioned above, an upgrade to 802.11n exposes the LAN edge switch connection as the potential bottleneck in the data plane. If top performance is the goal, for instance because of the need for pervasive streaming video, and especially if there are likely to be a number of simultaneous, heavy users in a cell and dual-radio access points are used, it is important to develop a plan for this upgrade;
- *Traffic patterns and LAN capacity*. While the wired network design should be considered, the discussion above shows that it is unlikely that upgrades will be required. However, an analysis of traffic patterns will identify the optimum locations for Mobility Controllers and provide an indication for the required traffic capacity. Aruba Mobility Controllers are rated with independent limits for the number of dependent access points, number of associated users and traffic throughput in the firewall and encrypted mode: only simple calculations are required to test designs against these limits;
- *Resilience.* This accomplished through redundancy at different levels. Each access point can be connected on dual Ethernet cables, but most networks rely on automatic RF management algorithms (Aruba's is called Adaptive Radio Management, or ARM) to expand the coverage of neighboring cells in

the event of a single AP failure. Mobility Controllers may be configured for redundancy either with modules within a chassis, or by using VRRP between controllers. 802.11n introduces no new dimensions to network resilience.

The following table provides a guide to the various elements that can be upgraded as part of a project to add 802.11n to an existing Aruba installation.

Operation	Phase	Triggering Event
Install 11n APs	Initial	As many as required for 802.11n coverage: replace existing APs, or intersperse
Upgrade Mobility Controller software	Initial	802.11n-aware software on existing Mobility Controller
Upgrade PoE for APs	Defer	Aruba 802.11n APs operate on 802.3af PoE: may require local power brick if PoE is marginal.
Replace AP cabling (for GE)	Defer	Avoid upgrade unless existing cables are long and low-quality
Upgrade edge switch (to GE)	Defer	Avoid upgrade until traffic increases and AP backhaul is the bottleneck: even then, solve by deploying more APs for smaller cells
New/larger Mobility Controllers	Not until traffic requires	Monitor actual traffic whenever traffic patterns change: not necessary unless wireless usage increases considerably.

True greenfield 802.11n sites

When a new site is to be provisioned with 802.11n access points, and only 802.11n capable clients will be used, inter-AP spacing may be increased, allowing a smaller number of access points to serve a given building and realizing some capital and installation savings. A good RF planning tool should include options for this, and in practice one would expect to cover the same area with 50% to 75% of the number of access points in an 802.11a/g design. However, this will be dependent on all clients having at least two antenna chains and supporting two spatial streams, and capacity considerations may dictate smaller cells, even for 802.11n.

In this scenario, it will be possible to use single-radio access points set for 40 MHz channels in the 5 GHz band, providing a very 'clean' RF design in a band with little interference. Of course, other considerations such as WIDS coverage may drive the use of dual-radio access points, even though only the 5 GHz band is used for traffic.

The software component

While much of this paper has focused on data transmission in an 802.11n Draft-2.0 network, there are many more considerations that are no less important, and where infrastructure vendors' software provides considerable differentiation.

- *RF planning*. Aruba believes the majority of 802.11n deployments will be upgrades to existing centralized 802.11a/b/g networks. Therefore it is important for the vendor to offer planning tools

catering for the different permutations of 2.4 and 5 GHz deployment, and optimum channel choice for 20 and 40 MHz 802.11n in 'greenfield' and 'mixed' mode, particularly in the 2.4 GHz band;

- *RF control and coordination.* Most vendors already offer functionality where real-time decisions are made to optimize channel choice and transmit power for each access point. This must be extended for 802.11n so it can handle both uniform and mixed deployments;
- *Wireless Intrusion Detection (WIDS).* Most Aruba customers use WIDS functionality, in which the network scans all channels in the background, identifying and neutralizing rogue APs and other intruders. Since rogues are now likely to be 802.11n access points, it is important that the network identify them even if it is not offering 802.11n service: and attack signatures, while they may differ little from 802.11a/b/g, must also be applied to 802.11n. For instance, a client blacklisted for behavior on 802.11a should also be locked out of 802.11n access. Over time, 802.11n-specific attacks will be developed and must be countered;
- *Location services*. Location is an important dimension of Wi-Fi networks, being used to identify coverage issues, the position of rogue APs and tracking Wi-Fi asset tags, among others. Most location algorithms to date use signal strength received at the access point as an indication of distance from that access point, and triangulate several measurements to place an object on a floorplan. MIMO complicates this approach as there are now several signal strength measurements at each access point, leaving infrastructure vendors with an opportunity to develop better algorithms for 802.11n;
- *MAC and PHY control*. Modern Wi-Fi chips are substantially programmable, and the drivers provided by the chip vendors can be customized and improved. For instance, Draft-2.0 offers a plethora of data rates, coding and modulation options. Designers may find it useful to limit or influence the data rate chosen, perhaps preventing an up-change when it appears that better conditions will be short-lived or to maintain a lower error-rate for certain traffic types. Other areas for improvement include optimizing the MAC aggregation or block acknowledgement options for certain applications, making video smoother or more error-free, and tweaking QoS parameters (802.11n incorporates 802.11e/WMM-PS);
- *High bandwidth mesh applications*. One of the difficulties of designing Wi-Fi mesh networks is that bandwidth requirements on the 'thicker' branch connections, near the center of the mesh, are much higher than those at the leaf connections. The high bandwidth available with 802.11n is a good solution for this topology puzzle, even if the client devices using the mesh are not 802.11n-capable.

Conclusion

802.11n technology and products are developing at lightning speed, and whereas the picture was unclear early in 2007, it is now possible to identify the characteristics of early 802.11n equipment, and to predict the developments of the next few years with more assurance. The IEEE 802.11n-Draft2.0 standard and the Wi-Fi Alliance Draft-n certification, appear to be relatively stable, and the benefits of 802.11n in terms of high performance and reliable coverage are well-established in testing of early certified products. The chances of future changes to the standard that are not backwards-compatible have significantly diminished, in part because all the players in standards groups would suffer commercially from such a development. Therefore the risk of early Draft-n equipment requiring upgrades is very small. We believe that the performance envelope of Draft-2.0 will stand till at least the end of 2008 however, it is probable that subsequent products will incorporate more features and offer higher performance than the Draft-n generation.

Organizations will implement 802.11n earlier if they see the initial performance as adequate, and they have a budget to match: an 802.11n access point currently commands a price approximately double its 802.11a/b/g equivalent. Also, benefits will only be seen when both clients and access points are 802.11n-capable, so

many of the early 802.11n networks will be running in legacy mode until the client base has been refreshed with 802.11n-capable devices.

802.11n affords organizations justification for a comprehensive network audit, particularly in identifying traffic volumes and patterns. A flexible deployment architecture allows wireless traffic to follow 'natural' paths over the LAN, WAN, and the Internet with minimal re-direction or disruption.

Providing power and high-speed connectivity to access points may represent a significant challenge when upgrading to 802.11n. Aruba allows the new access point to be powered from 802.3af sources, but for the ultimate in performance it will be necessary to provide Gigabit Ethernet ports from the edge switch.

Reasons to delay an upgrade to 802.11n include performance, convenience and cost. As noted above, we do not expect significant additional performance improvements until at least the end of 2008, but some advances will certainly be achieved. The next generations of silicon are likely to focus on lower power consumption (alleviating the issue noted above), with higher levels of integration and production volumes leading to potentially lower cost. It is not unreasonable to imagine prices approaching today's 802.11a/b/g access points over time.

Even with early 802.11n products, it is already clear that performance has indeed overtaken the 10/100 wired Ethernet connection, opening the way for all-wireless workplaces in which there are no cables to the desk. One early Aruba 802.11n adopter envisions saving millions of dollars by deploying an all-wireless 80-2.11n network in lieu of upgrading the cable plant. Thus we mark the beginning of the end for traditional wired office networks and the port-centric products that serve them.

About Aruba Networks, Inc.

Aruba securely delivers the enterprise network to users, wherever they work or roam, with user-centric networks that significantly expand the reach of traditional port-centric networks. User-centric networks integrate adaptive WLANs, identity-based security, and application continuity services into a cohesive, high-performance system that can be easily deployed as an overlay on top of existing network infrastructure. Adaptive WLANs deliver high-performance, follow-me connectivity so users are always within reach of mission-critical information. Identity-based security associates access policies with users, not ports, to enable follow-me security that is enforced regardless of access method or location. Application continuity services enable follow-me applications that can be seamlessly accessed across WLAN and cellular networks. The cost, convenience, and security benefits of user-centric networks are fundamentally changing how and where we work. Listed on the NASDAQ and Russell 2000® Index, Aruba is based in Sunnyvale, California, and has operations throughout the Americas, Europe, Middle East, and Asia Pacific regions. To learn more, visit http://www.arubanetworks.com.

WP_AWW_US_071217

1322 Crossman Ave. Sunnyvale, CA 94089-1113
Tel. +1.408.227.4500 | Fax. +1.408.227.4550 | info@arubanetworks.com
http://www.arubanetworks.com