

Technical Brief

Advanced RF Management for Wireless Grids

Enables Wireless LANs that Match the Performance of Wired Networks and Reduce Deployment Costs by Taking APs Out of the Ceiling and Into Structured Cabling Infrastructure

HEADQUARTERS

180 Great Oaks Blvd.

San Jose, CA 95119

USA

Tel +1 408-227-4500

Fax +1 408-227-4550

www.arubanetworks.com



INTRODUCTION

When corporate network administrators start to plan for a wireless network deployment, they often face the complexities of RF engineering for the first time in their careers.

Depending upon the approach taken, comprehensive in-building RF planning can require expensive site-surveys. In addition, there are often concerns that such a network will generate many time-consuming and expensive trouble-calls from users, needing dedicated and expensive wireless experts for resolution.

It is true that there are many aspects of wireless LAN (WLAN) networks that are different from, and less predictable than, wired LANs. The solution lies in recognizing the aspects where wireless is different and difficult, and in building a structured network that can shield the administrator from such aspects, while providing the level of visualization and control necessary to have confidence in managing the network.

Aruba Wireless Networks (Aruba) wireless grid architecture solves many of the problems that have in the past required expert RF expertise, specifically problems in providing coverage, assuring capacity and reacting to interference.

MANAGING CORPORATE WI-FI NETWORKS

Early wireless network architectures used autonomous, individually-managed access points (APs), installed above the ceiling tiles to provide wireless signals. It quickly became clear that automated tools were necessary to coordinate the RF configuration of these APs, as altering the settings in one cell would affect neighboring cells, with a ripple effect until the whole network was reconfigured as a result of a single AP change. (In this paper we use the widely-understood term 'cell' to refer to an 802.11 BSS and its coverage area).

Vendors have included such automated RF modeling and configuration tools - of varying sophistication - in their product portfolios for some time. However, the number and placement of APs through the building limited these tools, and since the installation of an AP was an expensive procedure, there were never enough for the management system to configure an optimum RF plan.

The Aruba Wireless Grid architecture uses inexpensive APs, called grid points (GPs), installed in the user-space (i.e. the cubicle) to bring down the cost of procuring and installing a WLAN. This allows for a structured model for the WLAN, similar to the wired LAN, where GPs are installed to a floor plan rather than an RF survey, and form a dense wireless grid. This grid of radio resources provides the density necessary for the RF management system to compute and configure the optimum RF plan for the building, unconstrained by the usual shortage of AP resources.

In turn, this allows the management system to deal with the complexity of the RF network, while giving the administrator an accurate, high-level view of the network, and the tools necessary to tailor it.

MANAGING COVERAGE

In RF terms 'coverage' is the ability of a network to provide a good RF signal level wherever it is required: in this case, through the corporate building. There are several aspects of RF which are important here:

- **Unpredictable propagation.** While the range of a radio wave in free space ('line of sight') is known and repeatable, it is nearly impossible to accurately predict range in a building. This is because walls and other obstructions attenuate the signal, while reflections can locally amplify or nullify it. Many vendors have built models that import floor plans, require administrators to estimate wall and cubicle construction and then derive RF models, but these have proven cumbersome to use and only partially accurate.
- **3-D.** Because wireless penetrates floors and ceilings, the radio plan for a building must reflect its 3-D nature: the best signal for a client may be from an AP on the next floor. This is another reason why floor-plan-based tools have been difficult to use.
- **Connection speed varies.** As the signal drops off, the link automatically down-speeds in order to maintain a usable connection. However, this variation in link speed can result in trouble-calls because of inconsistent connection speed. This has been one of the key problems with past wireless networks.
- **Dark spots in coverage.** While it is easy to program an AP or a client to 'listen' to its surroundings and report on weak or interfering signals, it is difficult to establish the

absence of a signal at various points through the building - except from site surveys or trouble-calls.

- **Limiting excessive coverage.** To enhance security, some administrators would like to prevent the radio signal from spilling out of the target area - the building, or lab, for instance - by limiting its range.
- **Symmetrical / asymmetrical cells.** Today most standard enterprise-class APs are installed with omni-directional antennas. In a 'perfect' world, this would result in regular, circular coverage cells. However, obstructions to the radio signal are seldom symmetrical. One of the results is that, to properly cover an area behind a wall on one side of the cell for instance, an AP's transmit power may need to be turned to maximum. However, this can push out the coverage on the other side of the cell, creating difficult conditions there such as interference.

MANAGING CAPACITY

In contrast to 'coverage', which is the ability to deliver a good-quality signal to all desired locations, 'capacity' concerns the number of users that can be simultaneously supported.

- **Limited resources.** Since there is a single radio channel available in a Wi-Fi cell, all users in that cell contend for that shared channel. This is analogous to early Ethernet, where all clients on a coax cable contended for the medium (using CSMA/CD). The MAC algorithm used by Wi-Fi, CSMA/CA, is very similar to Ethernet. One effect of this algorithm is that as the load in a cell increases, performance decreases almost exponentially. The solution adopted in Ethernet switches is to dedicate an Ethernet link to a single user. A similar approach in wireless would reduce the cell size, towards the limit where each cell supports a single user.
- **Uneven load.** A traditional corporate wireless network is set up with static cell sizes and RF channels. This reflects the set of requirements used in original network planning, but the demands on a network vary over time. Cubicle movements, conferences and seminars all change the characteristics of the load, often overloading certain cells while others are under-used. Reacting to such changes - for instance by expanding adjacent cells to share some of the load - is very difficult without automated support.
- **Interference between cells.** This is sometimes called 'co-channel interference.' Since there are a limited number of non-overlapping RF channels (three for 802.11b and 802.11g) there may be many APs in a building using the same channel. For a station associated with a particular AP, the signals from the other APs are unwanted, but in the same band: they are interference. Historically this interference has been managed by reducing the AP's transmit power. However, this may not be an option if it reduces coverage in another part of the AP's cell.

- **Interference from clients.** It is not just other APs in the building whose signals can interfere: clients transmitting at high power can radiate signals into other cells, causing interference. However, this is not as great a problem as AP interference, since 80% of traffic in corporate wireless networks is from the AP to the client.

MANAGING INTERFERENCE

Even if the radio signal is strong enough for high-quality reception, there are times when interference can increase errors significantly, or make connections unusable.

- **Sources of interference.** Common interferers in the 2.4GHz band (802.11b and g) are cordless phones and microwave ovens. These can be difficult for an administrator to identify, since they are intermittent, and if inactive at the time they can easily be missed in a site survey.
- **Localizing.** One common problem with an interferer, once it has been identified, is localizing it. Many second-generation wireless architectures have a function for this, based on triangulating the signals received at different APs, but such measurements are often inaccurate, as they rely on the signal being 'audible' at a number of APs, and accuracy is proportional to the distance between APs.
- **Disabling.** Once a source of interference has been identified and localized, it should be disabled or avoided. If it only radiates on a number of RF channels, it can be avoided by changing the working channels of the APs in the vicinity, which may necessitate re-engineering the plan for the entire building. Alternatively, a site visit may be required to identify and disable the device. This is one case where accurate location and characterization from the management system can save considerable on-site engineer time.
- **Rogue APs.** A 'rogue AP' is an unauthorized access point, often an AP brought into the building - with no ill intent - by employees, to make their work easier. Such APs are serious security risks, but they can also cause significant interference for the rest of the network.

ARUBA'S APPROACH TO MANAGING RF

The difficulties listed above have been significant since the very early corporate wireless networks. There have been many advances in RF management, but it has not yet reached the standard of wired LANs on a number of levels. The objective of Aruba's wireless grid architecture is to use the same architectural model that supported the evolution of the wired LAN. It works as a non-disruptive overlay to the existing wired LAN and supports the three core elements needed to deliver an enterprise-class wireless LAN architecture:

1. A predictable high performance solution.
2. A cost-effective deployment model, and
3. A future-proof architecture that supports evolution to newer and faster technologies.

The four key components of a wireless grid are:

1. Wireless APs called grid points (GPs) that are deployed next to existing wired network jacks or even integrated with them,
2. Power over Ethernet (PoE) delivery infrastructure built into regular Ethernet switches or integrated with patch panels into the structured cabling infrastructure itself,
3. A 3-dimensional grid addressing hierarchy that is built for identifying geographical locations in a building, and
4. High-performance and resilient WLAN switches that deliver automated grid management capabilities.

This architectural approach borrows heavily from the structured cabling architecture where the components involved are wired network jacks, cabling patch panels, a hierarchical addressing and management model for the structured cabling infrastructure and a LAN switch for delivering high-performance LAN services. From the viewpoint of RF management, there are several significant aspects to this architecture:

- Since procurement and installation of the new GPs is inexpensive, it is now practical to install them 'on plan,' eliminating the need for an RF site survey. This results in a dense wireless grid of GPs.
- Since there are, by design, more GPs installed than would be required for minimum coverage, GPs can be configured to operate at reduced power, reducing cell sizes, or switched to become passive 'air monitors' (AMs) where they monitor their RF environment and report measurements to the RF management system.
- The central RF management system has a view of all the GPs in the building, which allows it to calculate and configure the optimum RF model, taking account of cell-to-cell interference.
- The network administrator is given a comprehensive, high-level visualization of RF coverage and capacity in the building, without having to understand all the technical details, or get involved in the management of individual GPs.

GRID POINT FEATURES AND OPTIONS

The wireless grid can be built from a range of GP products. These include:

- Flexible choices for radios, including single-radio GP or dual-radio GPs
- GPs can be mounted on cubicle walls, under desks, inside cubicle walls

- Physical security provided by security mounting screws and cable locks
- Powering by power over Ethernet (PoE) or AC adapters
- Snap-on antennas, with a choice of 14 types for special needs, if required

MANAGEMENT FEATURES

The wireless grid management system has a number of important features:

- **Dynamic reconfiguration of GPs.** GPs can be configured in GP mode, with dynamically-assigned transmit power and RF channel, or monitor mode, where they can 'listen' across all channels, gathering data for use in measuring RF coverage, detecting interference and security risks such as rogue APs or reporting client signal strength for location services.
- **Collecting statistics.** The network elements report a host of RF parameters on a per-client, per-GP or per-channel basis. These include fragmentation, retransmissions, low-speed frames, broadcast/multicast frames, received signal strength and more.
- **Managing statistics.** Because the network can diagnose and automatically correct problems in real-time, the network manager is offered high-level visualizations, and is only alerted to specific problems if thresholds are exceeded.
- **Estimating coverage.** The management system detects coverage holes by analyzing clients associating at low rates, or unable to associate, although this should not be a problem in a wireless grid network.
- **Dynamic load balancing** detects overloaded cells and directs clients to move from heavily loaded cells to adjacent GPs.

DEALING WITH TRADITIONAL RF PROBLEMS

All of the problems previously discussed are solved using Aruba's wireless grid architecture consisting of a dense array of GPs and a WLAN switch supporting a centralized management system.

1. Coverage

Since the wireless grid includes GPs much closer together than in traditional networks, there are few difficulties in providing coverage to all desired parts of the building. The presence of many AMs allows for continual monitoring of signal strengths at many locations, giving the network administrator a comprehensive view of coverage. Also, since cells are small, it's feasible for the IT staff to specify minimum connection speeds, making network performance more predictable.

2. Capacity

Since cells are small, each supports a smaller number of users and the degree of sharing between users is reduced: contributing again to more consistent, predictable performance. Co-channel interference would be a problem in such a network, were it not for the centralized RF management system that uses live calibration data to derive an accurate and complete RF plan for the building, including the RF channel and transmit power for every GP. The centralized RF management system can also detect short-term load peaks, and dynamically reconfigure the network, for instance by turning AMs into GPs, to react - by reconfiguring a monitor to GP mode, for example, or balancing the load across several GPs. Thus the wireless grid of GPs is a flexible resource which is used by the centralized RF management system to deliver wireless capacity when and where needed.

3. Interference

The centralized RF management system uses measurements from the many GPs and AMs to identify any sources of interference, even if intermittent. The dense wireless grid allows accurate localization of such sources, and in most cases automatic reconfiguration can allow the network to avoid such interference without site visits. In cases where dispatch is required, the location of the source will be accurately known, reducing on-site time and expense.

CONCLUSION

The wireless grid network solves the RF management problems posed by conventional corporate Wi-Fi network deployments. The low cost of sourcing and installing the new GP means it is feasible to use many more GPs in a network than was previously the case.

This means that the placement of GPs is no longer critical, so the traditional site survey can be eliminated. The network can be deployed and managed in a 'structured' manner, similar to a wired LAN network. The dense array of GPs in a wireless grid network allows better coverage of the building, eliminating dark spots without coverage. The smaller cells made possible by the dense array of GPs also support more traffic per user, giving more consistent high-data-rate service, while enabling re-configuration on the fly for load balancing.

Once the wireless grid is deployed, many of the APs will be automatically switched to monitor mode, where they can detect and localize interferers and security risks in the network, as well as detecting and automatically correcting problems with coverage and capacity.

This enables the network manager to avoid day-to-day involvement in RF engineering, as the wireless grid network handles routine problems automatically. The network manager is presented with a high-level visualization of the RF network, and alarms only where thresholds or exceptions are detected.