

Centralized/Infrastructure Spectral Assurance: *Evaluating Total Cost of Ownership*

A Farpoint Group White Paper

Document FPG 2010-166.1
April 2010



While many forms of wireless-LAN assurance and verification tools have become available over the past decade or so of the history of the WLAN and Wi-Fi, one of the last to be addressed is the ability to examine activity at the physical layer (PHY) – in other words, to look at what’s happening in the air itself, what we call *spectral assurance (SA)*. SA is particularly important not just because the frequencies used by WLANs are unlicensed, and thus subject to arbitrary radio signals and interference, but rather because WLANs are now mission-critical in so many applications, serving as primary and even default access, and the need to detect and remediate spectrum-related challenges has become equally vital. The device used to address this challenge is called a *spectrum analyzer*, and, with the initial availability of Cisco’s notebook-PC-based *Spectrum Expert* in 2005, a low-cost but highly-effective tool was finally available to network operations staffs everywhere. But a new challenge has emerged: as wireless LANs have grown in importance and coverage, the *ad hoc* approach to spectrum analysis, what we call the “walking around” (“WAM”) model of operation enabled by mobile spectrum analyzers, while still valuable, is less than optimal. Given the mission-critical nature of contemporary WLAN installations, there is an obvious need for full-time, *pervasive* spectral analysis, and thus a clear requirement to move this functionality *into the WLAN infrastructure*, thus creating what we call the “centralized/infrastructure” (“CIM”) model of spectral assurance. With recent introduction of Cisco’s CleanAir product line, the CIM strategy is now a reality, leading to some very interesting questions about the total cost of ownership of such an approach – the subject of this White Paper.

Spectral Assurance: The Benefits

Why is spectral assurance so important? Let’s look at some of the key benefits:

- *Interference detection* – The unlicensed bands can be filled with interfering signals from a broad range of devices, from cordless phones to Bluetooth products, wireless video cameras, microwave ovens, and more. Most interference is unintentional, weak, and harmful to Wi-Fi signals only intermittently. But, as our own experiments have shown, interference from common, commercial wireless devices can be very damaging to Wi-Fi transmissions, and, without spectral analysis, there is no way to verify the nature of a specific problem.
- *Optimal operations* – While most enterprise-class WLAN systems can today automatically assign radio channels and set transmit power levels, having knowledge of the radio environment can avoid channels with significant interference (Wi-Fi or other), thereby optimizing (potentially dynamically) a given installation for throughput, reliability, and overall capacity.
- *Security and system integrity* – While full-blown PHY-layer denial-of-service (DoS) attacks are thankfully rare, such remain a possibility. Being able to detect

and localize illicit transmitters of any form on a firm's premises is today an essential WLAN assurance capability.

- *Policy enforcement* – Firms should have policies with respect to what wireless devices (including cordless phones, wireless video cameras, non-Wi-Fi wireless LANs, and more) can be used on site. Being able to detect prohibited devices is vital to this end.
- *Forensic troubleshooting* – Getting to the root cause of disruptions (intermittent or not) to wireless operations quickly, efficiently, and at minimal cost is a key capability of spectral assurance. Ultimately, as overall productivity depends upon the LAN and network services, it thus depends upon optimal wireless services as well.

As we noted above, spectral analysis in Wi-Fi applications has traditionally been implemented via an engineer or similar technical professional walking around with a notebook-based analyzer, looking for the source of radio interference or otherwise prohibited transmitters. While this approach can and does work, it's far from optimal. The biggest issue is that this technique is usually applied only when a problem is suspected. By then, of course, the problem might be gone, or, if intermittent or bursty in nature, undetectable. And, by definition, the WAM strategy covers only a portion of a given infrastructure at any moment in time, and it is always labor-intensive, involving a trained, experienced engineer or RF-qualified technician, and often with long-distance travel and significant time-on-site required. If one desires to localize the source of an interferer, such is often difficult with only one sensor. And, finally, the WA technique is completely independent of WLAN management and other assurance tools, meaning that decisions resulting from spectral assurance exercises must be made manually – with additional complexity and indeterminism often inherent in the process.

While we do in fact recommend that a PC-based spectral analysis tool be available (for ad-hoc analysis and pre-installation RF surveys, for example), incorporating SA functionality into the infrastructure of an enterprise-class WLAN installation introduces a broad range of benefits. These include, most importantly, integration with the management console and management system of the wireless LAN, thus merging all key functions onto a single screen with a corresponding improvement in both staff productivity and overall operations. Radio resource management can be similarly optimized, with channel assignment and transmit power logic now able to take advantage of another key variable, and with automation in the bargain. Reconfiguration, as required, is rapid and automatic. Event recording and trend analysis become simple and consistent. Locating an interferer or unauthorized device is automated, accurate, and rapid. And, most importantly, large areas – even multiple buildings, campuses, and widely-dispersed facilities can all be monitored and managed from a single location on a continual (24/7/365) basis. The idea is very similar to that applied to intrusion detection and prevention (IDS/IPS) services and other elements of a WLAN assurance solution, which also began with the walking-around model and evolved into enterprise-class, infrastructure-based solutions.

Farpoint Group thus believes that the network optimization and productivity enhancements of the centralized/infrastructure model are sufficient justification for implementing a CIM spectral-assurance solution. To introduce the impact on total cost of ownership, think of the CIM strategy as converting expensive, personnel-centered operating expense into greater but more efficient capital expense. But it's also important to consider that spectral assurance is also a bit like an insurance policy, with an opportunity cost assigned to not having it at all. So we assume that an SA capability is essential, and that the fundamental cost analysis is between the WAM and CIM approaches, not with respect to whether SA is valuable or not. It is, and very.

Building a TCO Model of Centralized Spectral Assurance

Any model of total cost of ownership has two components: *capital expense (CapEx)*, and *operating expense (OpEx)*. Capital expense includes all required equipment, planning, installation, functional verification, and related non-recurring engineering charges, which are most often only a small percentage of the cost of the equipment required. Operating expense includes all expenditures related to ongoing operations, including network management, troubleshooting, remediation, user support, maintenance, and many other functions.

Over time, the OpEx is usually much larger than CapEx, and for a simple reason. While CapEx largely depends upon the cost of manufactured products, which almost always declines over time and which benefits from improved price/performance regardless (often noted in the press as the “faster/better/cheaper” typical of high-tech products), OpEx is essentially *labor-intensive*, and costs here demonstrate the opposite property – they almost always *increase* over time. The best way, of course, to address rising labor expense is to improve productivity. The centralized model of spectral assurance can most certainly do this, as it converts staff-oriented walking around into centralized, continuous, infrastructure-based monitoring and analysis, along with appropriate automation in response to detected problems and issues. And, of course, the requirement for direct labor to do all that walking around is dramatically lessened and even eliminated in many cases.

In a centralized spectral-assurance solution, assuming that the required sensors are resident in devices that can also function as access points, the bulk of CapEx involved in greenfield (pervasive) deployments is the differential between the cost of a spectrally-enabled AP and one that is not. Depending upon model, this might amount to a few hundred dollars per AP. In augmenting existing deployments, Farpoint Group generally assumes a 1:4 to 1:6 relationship between spectrally-enabled APs and others. While some compromise in capability is inherent in a partial deployment, the cost can be quite low. In addition, it may be necessary to perform functional verification on a client-type basis against two different APs in the partial case, but these costs should be quite low and are regardless a function of local policies and procedures. The only other cost is that required for an appliance to implement location functionality; again, this is quite minimal given the overall cost of the installation. The incremental cost of adding centralized spectral

assurance is thus only a very small percentage of the overall capital cost of any given installation, and this percentage declines with overall scope. But the CapEx of CIM will often be much higher than that required in the WAM case, which involves only a notebook and an inexpensive spectrum analyzer tool.

On the operating expense side, recall how the walking around model is fundamentally labor-intensive, and often of indeterminate scope and thus of indeterminate cost. It can be quite frustrating to hunt down intermittent sources of interference (we have significant direct experience with this kind of work!), and being in the right place at the right time can never be guaranteed. The CIM strategy eliminates almost all of these costs, and thus we expect to see overall TCO benefit from centralization.

As we noted above, equally important, but much more difficult to quantify, are a broad range of potential opportunity costs associated with failures in the network. These can include lost productivity in the event of user-visible problems or wholesale downtime, as well as potential consequential costs, including legal and other regulatory violations and their potentially enormous impacts, the fallout from security breaches and related failures, and lost prestige, customer confidence, and business that can result from customer-facing service problems arising from network failures. Even degraded throughput can factor in here, if such affects the responsiveness of the enterprise. As most businesses are today heavily dependent upon the corporate LAN (wired and wireless) for essential operations, it makes sense to implement improvements in reliability – like centralized spectrum assurance – that contribute to the reliability and performance required to avoid these potential consequential opportunity costs. The network is, after all, a tool to optimize the productivity of the organization. Assuring its efficient operation no matter what the challenge is essential, again why we believe that SA functionality is essential in *all* enterprise-class WLAN installations. The TCO differential between any form of spectral assurance and no spectral assurance capability whatsoever are the large opportunity or insurance costs associated with all of the exposures noted above. Quantifying these will depend upon the specifics of a given installation and industry, but they can clearly be very large indeed. The analysis we will thus focus on below will be limited to a comparison of the costs of an infrastructure-based spectral assurance vs. the walking-around model.

TCO of Spectral Assurance Solutions: An Example

While the specifics of any given installation can vary, Farpoint Group, in consultation with Cisco Systems, has developed a general TCO model (See Table 1) for the deployment of spectral assurance capabilities in the enterprise. In this model we consider three specific scenarios: the walking-around model (WAM), as discussed above, and two variants on the centralized/infrastructure model (CIM). The first of these is a *partial* (sometimes called an “overlay”) strategy, which might be used in the case of an existing deployment, and which is implemented via the *addition* of some number of spectrally-enabled APs, in this case the Cisco 3500 series. This type of deployment can provide continuous monitoring and the location of interferers, the latter enabled via Cisco’s

	Walking Around Model	Centralized/ Infrastructure Model - Partial	Centralized/ Infrastructure Model - Pervasive
CapEx			
Cost of notebook PC	\$1,000		
Cost of spectrum analyzer	\$5,000		
Cost of additional spectrally-enabled AP		\$1,300	
Differential cost of AP			\$300
Total number of APs deployed		1000	1000
Ratio of sensor APs (1:n)		5	1
Number of spectrally-enabled APs		200	1000
Cost of spectrally-enabled APs		\$260,000	\$300,000
Cost of MSE for context-aware location		\$7,000	\$7,000
Total CapEx	\$6,000	\$267,000	\$307,000
OpEx			
Hourly cost of engineer	\$80	\$80	\$80
Number of incidents/month	10	10	10
Number of hours/incident	24	1	0.25
Total OpEx (annual)	\$230,400	\$9,600	\$2,400
Three Year TCO	\$697,200	\$295,800	\$314,200

Table 1 – Sample model of the total cost of ownership (TCO) for three different deployment scenarios. Yellow fields are variables. *Source:* Farpoint Group.

Mobility Services Engine (MSE) appliance. The second case is a *pervasive* deployment of 3500-series APs, which includes all of the benefits of the partial deployment plus automated remediation (also called “self-healing”) of interference-related problems.

As can be seen in Table 1, the WAM has a very low capital expense component, but a very high operational expense because it is fundamentally labor-intensive. Based on conversations with key staff at several large organizations now deploying spectral-assurance solutions, as well as our own experience, we believe that a budgetary number of 24 hours from incident to resolution is appropriate here. Note that additional personnel latency and travel expense may be involved in any particular incident. Three-year TCO (CapEx plus annual OpEx times three) is thus very high, and it should be again noted here that continuous monitoring is not possible in this case.

The CIM – Partial case assumes an existing deployment and the addition of new 3500-series APs in a 1:5 ratio (one 3500 to five existing APs). As was noted above, Farpoint Group has found that this ratio will usually be between 1:4 and 1:6. Note, however, that 24/7 monitoring is now enabled, and personnel time required is vastly reduced.

The CIM - Pervasive example need consider only the differential cost of the 3500 AP, \$300, since no APs are already deployed in this case. Note, though, that the personnel time required in this case is again significantly reduced due to the self-healing capabilities enabled. Thus the three-year TCO of the Pervasive strategy is only inconsequentially higher than that of the Partial case, making the choice and easy one for greenfield deployments.

Both CIM examples assume 1000 total APs deployed, and one spectrally-related incident per 100 APs per month. All pricing is based on Cisco's list prices.

Note that this model does not quantify any maintenance costs, and also does not consider the opportunity costs noted elsewhere in this document. We believe the former is of little consequence, and the latter so potentially significant as to demand the installation of a spectral-assurance solution regardless.

Conclusions

There is no doubt in our mind, based on years of experience in planning, installing, and troubleshooting wireless LAN installations, that spectral assurance is a required element in *any* enterprise-class installation. The potential costs associated with a lack of visibility into Layer 1 and resulting from a failure here can be enormous. And, similarly, there is little doubt, based on our cost analysis above as well as the other benefits we noted (particularly global, continuous monitoring), that the centralized/infrastructure model is the preferred strategy. The conversion of OpEx into CapEx and the improved reliability and performance should easily result in a positive return on investment in a relatively short timeframe. Again, it's ultimately best to think of spectral assurance as an insurance policy covering both productivity and liability. The benefits in terms of WLAN optimization, security, policy enforcement, and troubleshooting are undeniable. The availability of this capability in WLAN infrastructure is clearly a huge leap forward for the industry and its customers, and one that believe will become standard and indeed common over the next few years.



Ashland MA USA

508-881-6467

www.farpointgroup.com

info@farpointgroup.com

The information and analysis contained in this document are based upon actual testing and publicly-available information sources believed to be correct as of the date of publication. Farpoint Group assumes no liability for any inaccuracies that may be present herein. Revisions to this document may be issued, without notice, from time to time.

Copyright 2010 – All rights reserved

Permission to reproduce and distribute this document is granted provided this copyright notice is included and that no modifications are made to the original.