Avaya WLAN 2300 Series Design and Implementation Guide

Avaya WLAN 2300

7.1

Document Status: **Standard**

Document Number: **NN47250-200**

Document Version: **04.01**

Notices

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of this documentation unless such modifications, additions, or deletions were performed by Avaya. End User agree to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked Web sites referenced within this site or documentation(s) provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on this product. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language, as well as information regarding support for this product, while under warranty, is available to Avaya customers and other parties through the Avaya Support Web site: http://www.avaya.com/support

Please note that if you acquired the product from an authorized reseller, the warranty is provided to you by said reseller and not by Avaya.

Licenses

Copyright

Third Party Components

Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements ("Third Party Components"), which may contain terms that expand or limit rights to use certain portions of the Product ("Third Party Terms"). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: http://support.avaya.com/Copyright.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the documentation(s) and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the documentation(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party. Avaya is a registered trademark of Avaya Inc. All non-Avaya trademarks are the property of their respective owners.

Downloading documents

For the most current versions of documentation, see the Avaya Support. Web site: http://www.avaya.com/support

Contact Avaya Support

Avaya provides a telephone number for you to use to report problems or to ask questions about your product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: http://www.avaya.com/support

Contents

Customer service

Visit the Avaya Web site to access the complete range of services and support that Avaya provides. Go to www.avaya.com or go to one of the pages listed in the following sections.

Navigation

- "Getting technical documentation" on page 6
- "Getting product training" on page 6
- "Getting help through an Avaya distributor or reseller" on page 6
- "Getting technical support from the Avaya Web site" on page 6

Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to www.avaya.com/support.

Getting product training

Ongoing product training is available. For more information or to register, you can access the Web site at www.avaya.com/support. From this Web site, you can locate the Training contacts link on the left-hand navigation pane.

Getting help through an Avaya distributor or reseller

If you purchased a service contract for your Avaya product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

Getting technical support from the Avaya Web site

The easiest and most effective way to get technical support for Avaya products is from the Avaya Technical Support Web site at www.avaya.com/support.

Acknowledgements

This Engineering Guide was adapted from the document titled “Wireless LAN Design, and Implementation Guidelines“.

Introduction

This document is intended to provide best practice guidance on Avaya WLAN 2300 Release 7.1 product design and implementation, based-upon engineering design limits along with real-world practices.

The three main topics discussed in this document includes:

- System scaling and limitation guidelines
- System design
- Implementation

System scaling and limitation guidelines

General

Table 1: Mobility Domains

Attribute	Limitation Guideline
Maximum number of switches in a Mobility Domain	32

➡ **Note.** Several factors effect scalability in a Mobility Domain (RF Detect table size, countermeasures configuration, bandwidth between WSS, etc.). This number reflects expected scalability for WSS 2350 capacity switches.

Wireless LAN security switch

➡ **Note.** The number of WAPs supported on WSS 2380 depends on the configuration and maximum is 120 WAPs.
The WSS 2382 supports 120 WAPs regardless of the configuration.

Table 2: WLAN security switches

Model	Ports	Maximum Active sessions	Maximum FDB entries	Maximum Active APs	Maximum Configured APs
2350	2 x 10/100 1 x Uplink 1 x PoE 1 x Console	75	8,192	3	8
2360 2361	8 x 10/100 2 x Uplink 6 x PoE 1 x Console	300	8,192	12	30

Table 2: WLAN security switches

2380	4 x GigE (GBIC or RJ45) 1 x Console	2,500	16,383	40 80 120	300
2382	2 x GigE SFP 1 x 10/100 Mgmt. 1 x Console	3,200	8,192	32 64 96 128	320

Wireless LAN 2330/2330A/2330B access points

Table 3: WLAN APs

Model	Ports	Number of Radios	Radio 1 Bands	Radio 1 Antennas	Radio 2 Bands	Radio 2 Antennas
2330	1	2	802.11b/g	Internal	802.11a	Internal
2330A	2	2	802.11b/g	Internal	802.11a	Internal
2330B	2	2	802.11b/g	Internal	802.11a	Internal
2332	2	2	802.11b/g	Internal	802.11a	Internal

Wireless LAN management system

WMS general

Table 4: WMS general

Multi-threading	Both the client & server are multi-threaded
Throughput Requirements	WMS clients accessing Server: • LAN Throughput = 100Mbit • LAN Latency = 100ms or less • WAN = Not supported WMS Server communication to switches: • WAN throughput = T1 or better required for single switch WAN link. • WAN Latency = 100ms or less.
Location services. How many feet can the WMS software locate rogue AP/users.	10-30 feet, if the plan/obstacle data is accurate.

WMS limits

Table 5: WMS limits

Number of Simultaneous WMS Client Sessions	8
Total Number of Managed Switches	500 Switches – Configuration Only 64 Switches – Full Monitoring & Configuration
Total Number of Managed APs	1,000

WLAN 2350 sized devices, assumption is that monitoring features are not enabled other than basic device status.

WMS databases

Table 6: WMS database

Client Session Database	1,000 entries/switch
Radio Stats Time-Series Database	For all radio stats: -1 Hour (12 samples) -24 Hour (24 samples) -7 day (28 samples) -30 day (30 samples)
Device Log Database	1,000 – 5,000 entries/switch
Device State Database	Per MX snapshot of -System Status -AP Announce Status -Radio Status -RF Neighbor table -RF Attributes table
Rogue Database	100 activity events / rogue Per rogue time-series (up to 6 listeners per entry): -1 Hour (12 samples) -24 Hour (24 samples) -30 day (30 samples)
Radio Activity Database	Per Radio: -20 Power/Channel change events -20 Counter measure events

Note: The database automatically rolls up sampled data at specific intervals. The table above indicates how many samples are displayed for each specific time period. For example, in a 1 hour time period there will be 12 sample points displayed, each indicating the average value over 1/12th of an hour

WMS processing : polling & traps

Table 7: WMS processing

Area	Polling Interval (minutes)	Updates from SNMP Traps
Status Monitor	1-60 (default 1)	Yes
Log Monitor	1-60 (default 5)	Yes
Client Monitor	5	Yes
RF Monitor	5	Yes

Table 7: WMS processing

Rogue Detection	5, 10, 15, 30 (default 5)	Yes
Configuration Poll	1-Max In (default 15)	No

WMS hardware requirements

Table 8: WMS monitoring service hardware requirements

Hardware Requirements for WMS Monitoring Service	**Minimum**	**Recommended**
Processor	Intel Pentium 3.6 GHz or equivalent	Intel Core 2 Extreme, 3.2 GHz or equivalent
RAM	2 GB	4 GB
Hard drive space available	50 GB	150 GB
Monitor resolution	1024x768 pixels, 24-bit color	1600x1200 pixels, 32-bit color
CD-ROM drive	CD-ROM or equivalent	CD-ROM
Hardware Requirements for WMS Monitoring Service on the Sun Solaris platform	**Minimum**	**Recommended**
Processor	Sun UltraSPARC 10 or equivalent	Sun UltraSPARC III or equivalent
RAM	1 GB	2 GB
Hard drive space available	1 GB	2 GB
Monitor resolution	1024x768 pixels, 24-bit color	1600x1200 pixels, 32-bit color
CD-ROM drive	CD-ROM or equivalent	CD-ROM

Note: Running the client on the same hardware as the server requires adding the minimum memory requirements for both and makes the recommended processor for the server, the minimum processor.

Recommended guidelines for hardware requirements and memory allocation based on the number of radios and switches

1-25 Switches, up to 200 APs	3.6GHz P4, 2GB RAM, 50 GB HD
25-64 Switches, up to 1000 APs	Intel Core 2 Extreme, 3.2 GHz, 4 GB RAM, 150 GB HD

Scaling WMSfor Large Deployments

Current limits

In Release 6.0 a single WMS server supports a maximum of 64 switches or 1000 APs, whichever limit is reached first. To manage such large deployments, you need to tune and/or disable polling options. If all log, client session, rogue, and traffic statistics polling is disabled, a single WMS server has been tested to support up to 500 WSS 2350 sized devices for all configuration management and basic monitoring features.

To manage even larger networks, you need to deploy additional WMS servers. The typical model is to have a server manage a "region" or logical group of switches.

Hardware and Operating System

For large deployments, the server must be installed on an adequate hardware platform. Here is a typical recommended configuration:

Dell PowerEdge 2950:

Primary Processor: Dual Core Intel® Xeon® 5130

Additional Processor: Dual Core Intel® Xeon® 5130

Memory: 4GB 667MHz (4x1GB)

Operating System: Windows Server® 2003 R2, Standard Edition

Server Settings

The following settings must be adjusted after WMS is installed. The installer may overwrite some settings with default values; hence these must be reset after each installation. Once these changes are made, the server must be restarted.

- Memory - During installation, set the server memory to 1536MB, and set the client memory to 512 MB.
- Thread Count - By default, the server allocates 20 threads for device communication. For scaled deployments, the server thread count should be calculated as 1.5 x (number of MXs) or a maximum of 200, and the lower value used for the thread count.

➡ **Note.** A higher thread count increases CPU and memory usage and for lower performance computers, this can cause degradation in performance.

To configure the thread count, locate the *services-conf.xml* file in the *<installdir>/conf/* directory on the WMS server. Then edit this file and replace the string threads="20" with threads="N", where N is the calculated thread count.

- Polling Settings - Locate the polling settings under Setup->Monitoring Settings on the WMS server's web interface:

  1. Change the Status Monitor Polling Interval to 60 minutes
  2. Disable Log Monitor
  3. Disable RF & Traffic Trending
  4. Disable Rogue Detection Polling

On every applicable WMS client, launch the application and navigate to the "Devices" panel. Click on the "Options" icon on the toolbar (looks like a notepad with a pencil):

a Change Device Configuration Polling interval to 60 minutes.

System design

System design – RF coverage

Warning! Network plans generated by the WMS planning tool use a predictive RF model. This information is an **estimate** of the amount, configuration, and placement of WLAN 2300 equipment and is intended to facilitate project cost projections. A Wireless LAN site survey service captures actual RF data to accurately model RF environments and is the best method to determine correct location of APs, and therefore installation locations of cabling. A WMS network plan is an **approximation only** and does not substitute for a Wireless LAN site survey.

You can use the following six techniques for RF coverage:

- Auto-tune Snapshot (with or without Modeling)
- RF Planning with Ekahau Site Survey
- Manual Site Survey
- Auto-tune
- RF Planning
- Auto-Tuning with Modeling

Auto-tune Snapshot (with or without Modeling)

- All the same attributes of RF auto-tuning (with or without model)
- Use the auto-tuning function for a limited period of time (enough to provide a baseline) then use WMS to convert the dynamic data into a static configuration.

RF planning with Ekahau site survey

- All the same attributes of RF planning
- Ekahau site survey import can be used to used for a new model or to calibrate/optimize an existing RF model ability to combine "real-world" data with planning algorithms. Obstacle attenuations are adjusted based on the measured data. A site-survey report can be used to aid the process.

Manual site survey

- Uses the placement of live APs into the environment to measure real-world RF characteristics, to be able to plan permanent AP locations, based upon measured coverage.
- Does not use WMS for planning location, power or channel, but instead uses it for monitoring and management.

Auto-tune

- Lets you use the default auto tuning feature to select power and channel settings for RF signals in your RF coverage area. You upload the WSSs into WMS, configure the APs, enable Auto-RF, and deploy.
- This is a quick way to install a MX and some APs, and observe how the network operates.

RF planning

- Is a technique used to create a detailed network plan that provides powerful monitoring and visualization benefits. It fully models the geographic location with detailed information about the floors, and specifies the RF coverage areas and existing RF obstacles. Each of these methods is described in the sections that follow.
- To perform RF Planning, it is necessary to provide detailed information about the site and buildings by importing AutoCAD DXF™, AutoCAD DWG, JPEG, or GIF floor plan files of the buildings into WMS. As the floor plans are imported, one can modify them to add or remove RF obstacles. WMS includes a library of attenuators for building obstacles. The library includes doors, walls, ceilings, and other physical obstructions that can be selected. Attenuators can be defined by height, width, type of building material. WMS factors in the impact these objects have on how the radio frequency (RF) signals flow through a given site. If the network contains third-party APs, enter the information for these APs so that WMS takes the APs into account when calculating the placement (and optionally, the channel and power settings) of the Avaya APs.
- Auto-tune can help make WLAN installation easier by automatically selecting the channel and power setting. For Auto-RF to function correctly it is necessary to ensure that the AP's are properly located for best coverage based on anticipated use.

For models WLAN 2330, WLAN 2330A: Auto-RF in 802.11a mode supports channels 36, 40, 44, 48, 52, 56, 60, 64. Channels 149, 153, 157, 161, 165 are not being selected by the algorithm.

For models WLAN 2330B:

- US & Canada and other countries following US FCC (FCC Part 15.247, Subpart E rules) DFS requirements.
  - Auto-RF in 802.11a mode supports channels 36, 40, 44, 48.
  - Channels 52, 56, 60, 64, 149, 153, 157, 161, 165 are not being selected by the algorithm.
- Countries that do not require DFS equivalent support. Auto-RF supports channels 36, 40, 44, 48, 52, 56, 60, 64. Channels 149, 153, 157, 161, 165 are not being selected by the algorithm.

The recommended way to ensure proper AP coverage is by performing a site survey to determine the physical location of the AP's.

Avaya’s recommendation

Avaya recommends that Auto-RF cannot be used for the following:

- In 802.11a deployments where North American DFS regulation applies.
- Voice over WLAN network deployments
- Ekahau RTLS (Real Time Location Services) deployments.
- Mission critical data networks.
- Networks with more than 40 AP's.

AP’s need to be deployed with a minimum separation of **40db** (10feet is about 50db).

RF Auto-Tuning with Modeling

- As with the RF Auto-Tuning technique, it lets you set the auto tuning feature to adjust power and channel settings to provide RF signals to the coverage area for your users.
- Provides modeling information about your geographic location. By providing some information about your buildings and floors, you add enough details into WMS so that your can better visualize your network topology and support improved monitoring and locations services at your site. For example, you can manage your network overlaid on a floor plan, versus managing an abstract logical group of switches and APs.

WLAN 2300 Auto-RF deployments scenarios:

Once the site survey is completed and the AP's have been installed, Auto-RF can be turned on for the following scenarios:

- Locations with 12 AP's and less Auto-RF could be enabled freely, matching the WLAN 2360 maximum AP support.
- Location with 40 AP's and less Auto-RF could be used to select channels and power. Once the network has converged (see steps section 3.2) it is recommended to the channels be locked. Matching the 40 AP Licence for a WLAN 2380. Auto-RF for 802.11a should not be deployed where North American DFS regulations apply.
- For locations with more than 40 AP's it is **very strongly** recommended to use professional services to the plan for a static RF network

Which planning technique should I use?

The more detailed the network plan, the better able to manage and monitor the network. However, there are other requirements organizations should consider.

Table 9: RF Deployments

Deployment Type/Scale	Technique	When to Use	Pro's/Con's
Relatively simple user needs/ environment. Service use limited to data networking with tolerant applications (email etc.)	RF Auto-Tuning	If installing APs without consideration to: -blanket coverage -throughput -number of users	-Quick and easy -Reliability of latency sensitive applications is poor, such a VoWLAN. -No accurate rogue location services -No way to visualize your physical AP placement
Same deployment types as non-modeled RF Auto-tuning, however the customer requires some visualization and/or rogue location.	RF Auto-Tuning with Modeling	If you want to better monitor of your wireless network in terms of buildings, floors, or coverage areas over standard RF auto-tuning	-Quick and easy -Reliability of latency sensitive applications is poor, such a VoWLAN. -Lowest accuracy of rogue location services -basic visualization of your physical AP placement
Customers demanding rapid turn-around from sale to installation and which require latency sensitive application support (VoIP, etc.)	RF Auto-Tuning Snapshot	If you want to better monitor of your wireless network in terms of buildings, floors, or coverage areas over standard RF auto-tuning, and you are using latency sensitive applications. If there is not enough time to follow normal planning procedures.	-Allows system to determine acceptable RF settings for the particular deployment which can then be converted to a static configuration to improve stability and latency. -Settings captured at time of snapshot may not reflect true optimal configuration.

Table 9: RF Deployments

Well engineered networks requiring specific performance minimums and accurate rogue location services	RF Planning	When you want to provide customers or service staff a Work-Order including deployment details and a complete Bill of Materials If you want more accurate rogue locations services and planning with detailed network visualizations If you want to plan specific capacity requirements like per-user bandwidth, and number of users in a coverage area.	-Requires CAD file -CAD files need to be clean-up. Garbage-in, garbage-out. -Ability to define obstacles & attenuation considers cross-floor interference -Allows voice & data capacity calculations -Instead of you making a "best guess" as to how many APs you require for the desired coverage and where APs should be placed, WMS automatically calculates how many APs you need and where to place APs for optimal positioning. -You can generate a deployable work order to help installers place WSS switches and APs. -You automatically receive a deployable configuration that includes optimum power and channel settings. -More accurate monitoring options and network visualization based on the additional geographic modeling information loaded into WMS.

Table 9: RF Deployments

Well engineered networks with more challenging RF environments.	RF Planning with Ekahau Site survey	For the same reasons you would use RF planning, but when you want more accurate obstacle data for location services.	-Same pro's and cons as RF Planning -Used to calibrate/optimize an existing RF model by using "real-world" data with planning algorithms. Obstacle attenuations are adjusted based on the measured data. -Requires purchase of 3rd party application (Ekahau) -Requires significant time investment to get accurate data.
Very challenging RF environments without the need for rogue location or other visualization.	Manual Site Survey	If you want to perform planning with no building drawings or known inaccurate building drawings. Also used if there are known to be many interfering devices.	-Uses real-world measured data -Need to manually configure AP power & channel. -No Work-order -No accurate rogue location services -No way to visualize your physical AP placement -Requires large amounts of manual configuration to equal benefits of WMS Planning

System design – capacity

Most system capacity limitations and scalability information has already been defined in the System Scaling and Limitations Guidelines section. This section provides additional guidelines based upon best practices.

General guidelines

Be sure to plan for worst-case traffic capacity when working in environments that have intermittent high peaks or peak periods. Determine the peak traffic and peak traffic periods though measurement.

Design considerations

See the matrix below for specific design considerations:

Table 10: Design considerations

Area	Consideration	Type	Typical Design Guidelines
Total Sessions			
Clients	Types and Bandwidth	Data	Maximum: 802.11b= 5-6 Mbp/s peak 802.11g = 30 Mbp/s Protected Mode (802.11g in presence of 802.11b) = approx ½ of non-protected mode 802.11a= 30 Mbp/s Note: Guideline can vary greatly based on AP positioning and area RF properties such as interference and attenuation. Avaya recommends the use of a manual site survey to ensure proper bandwidth planning and AP positioning. Multicast streams can reduce available bandwidth dramatically for 802.11B or Protected Mode networks due to the amount of time-on-air for lower data rate packets. For aggregate multicast streaming bandwidth exceeding 512Kbps, 11a or pure 11g is recommended in production environments to ensure quality streaming.
Access Points	Sessions per AP	Data	Up to 50 Max up to 20 recommended with: No concurrent Multicast No other client types in use, like voice
		Voice (See Note)	SpectraLink: Maximum 10 calls / AP Vocera: Maximum 8 calls / AP
Security Switches Uplink throughput to Core		WSS 2350	100 Mbp/s minimum.
		WSS 2360 WSS 2361	Minimum - 100 Mbp/s for 1-12 APs Recommended - 200 Mbp/s for 4-12 APs
		WSS 2380	Minimum – 1 Gbp/s for 1-120 APs Recommended - 2 Gbp/s for 1-80 APs Recommended – 2-4 Gbp/s for 80-120 APs
		WSS 2382	Minimum - I Gbp/s for 1-128 APs Recommended - 2 Gbp/s for 1 - 128 APs
WMS	Number of servers required	1 per 1000 APs	1 server per 1000 APs

Note: Spectralink & Vocera count the number of calls differently. With Spectralink 10 calls would be considered between 5 Phones on a single Access Point. With Vocera 8 calls would be considered between 16 Badges on a single Access Point.

System design – redundancy & resiliency

Wireless LAN security switch

- Redundant power supplies
  - Wireless LAN Security Switch 2382 - Supports hot-swappable field upgradable power supplies.
  - Wireless LAN Security Switch 2380 – Supports two hot-swappable field upgradable power supplies.
  - Wireless LAN Security Switch 2361 – Supports two fixed internal power supplies.
  - Wireless LAN Security Switch 2360 – Supports one fixed internal power supply.
  - Wireless LAN Security Switch 2350 – Supports one external power supply.
- N+1 – switch cluster configuration

Single cluster

One approach to switch failover is to cluster the switches in N+1 redundancy groups in order to support both active and backup APs in the same cluster.

This greatly simplifies the configuration, management and also takes into account the fact that APs can failover from switch to switch in the cluster should they temporarily lose connectivity to their high bias switch during a network failure condition

Ensure that each wireless VLAN is present on at least 2 Security Switches in the cluster to prevent single switch failure from disabling any single VLAN. If possible, and if it is not configuration intensive, it is suggested to put all wireless VLANs on all switches, within a single cluster.

Multiple clusters

When you have multiple buildings that each have their own VLAN(s), is it recommended to use a separate cluster for each building in order to maintain availability of all VLAN's and eliminate unnecessary inter-building backbone traffic.

If you require seamless roaming by areas covered by multiple clusters, configure switches in each cluster into individual Mobility Domains and then add all applicable switches into an overarching Network Domain. DAP redundancy, RFdetect services and countermeasures are limited in scope to the individual Mobilty Domains but the Network Domain will permit seamless roaming between switches in separate Mobility Domains.

Three switch cluster example:

- No AP redundancy
  - A cluster of 3 MX 2382’s

  - Total Distributed AP capacity of 360
- N+1 redundancy
  - A cluster of 3 MX 2382’s
  - Total Distributed AP capacity of 240
  - Each Security Switch will be high bias for an even share of APs (80) and will provide failover connection for half of one of the other Security Switches distributed APs.

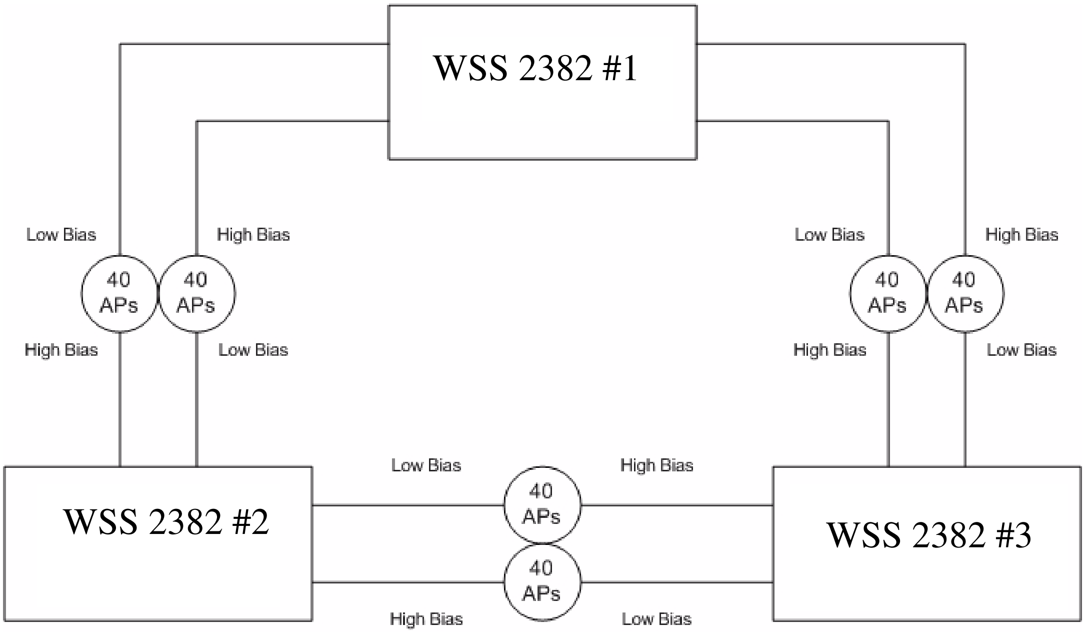

Perhaps the most commonly used redundancy feature in MSS today is the ability to use high and low bias configurations to support an N+1 redundancy scheme for controller hardware. The ability to weight boot preference towards some controllers, combined with support for more configured APs than the actual operational limit allow for rapid and automatic failover in the event of a single failure of a MX controller. Proper implementation of this redundancy scheme does require some planning and thought

The most common approach to planning an N+1 redundant configuration will start with an understanding of the amount of equipment required to satisfy the end customer's needs. This can be approached by determining the maximum capacity allowed when using a specific number of certain controllers or the number of APs required to meet service needs can be used to extrapolate how many controllers will be necessary.

The recommended method, which will be outlined in the example below, involves using all the installed controllers at reduced capacity. This leaves headroom for remaining controllers to pick up all lost high-bias DAPs in the event of a single MX failure.

Example Scenario:

Based on requirements defined by the customer, Site survey has shown that 360 APs will be needed for the customer's wireless network services. The customer has chosen the MX 2382 platform of MX controllers and desires to have N+1 redundancy. The MX 2382 will be licensed to support the full 128 DAP capacity which is their maximum.

Using the following formula the number of MX 2382 controllers necessary for N+1 redundancy can be determined:

# of controllers = (# of DAPs / capacity of controller) + 1 -OR-

4 = (360 / 128) + 1

With 4 MX 2382 controllers, each licensed for 128 DAPs, the customer can achieve N+1 redundancy. You can calculate the total possible capacity for an N+1 cluster with this formula:

Total DAP capacity = (# of controllers - 1) x capacity of controller -OR-

384 = (4 - 1) x 128

We can see from those results that 4 MX 2382 controllers with licensing to support 128 DAPs will provide slightly more redundant capacity than is strictly required by the customer, however this is the closest match.

The formula to calculate how many DAPs to configure as high-bias on each of the controllers is as follows:

# of high-bias DAPs per MX = (# of total DAPs / # of controllers) -OR-

90 = (360 / 4)

If we use the same formula using the maximum capacity of the cluster:

96 = (384 / 4)

We can see there will be 6 additional available licenses per switch which could be used for future expansion maintaining N+1 redundancy and which would not require re-calculating the spread of high-bias DAPs across the cluster; whereas capacity expansion that requires the addition of one or more controllers to the cluster will require re-calculation.

Now we know the number of controllers which will be required, as well as the number of high-bias DAPs which will be assigned to each of those controllers under normal operation. To achieve the redundancy desired we will have to add low-bias configurations for each of the DAPs, such that should their high-bias MX fail the available capacity on the remaining controllers in the cluster will pick up the dropped DAPs, spread out over the remaining switches in the cluster. In practice it is best to tackle this from the perspective of each individual controller.

The formula to determine the split of low-bias configurations to be spread across the cluster is:

# of low-bias DAPs divided amongst remaining WSSs per high-bias controller = # of high-bias DAPs per MX / (# of controllers - 1) -OR-

30 = 90 / (4 - 1)

For example:

MX 2382-A has high-bias configurations for DAPs 1-90. MX 2382-B has 91-180 etc.

Low-bias configurations should be added to MX 2382-B for DAPs 1-30, MX 2380-C would take 31-60 and MX 2382-D would provide low-bias backup for DAPs 61-90.

The same process is repeated for high-bias DAPs associated to MX 2382-B and so on.

Here is a simple diagram representing the high-bias (red line) and low-bias (blue line) configurations associated with DAPs 1-90. (This diagram could be extrapolated out to include all 360 DAPs, but the concept of concentrating the high-bias configuration and distributing the low-bias is better served by a simplified diagram):

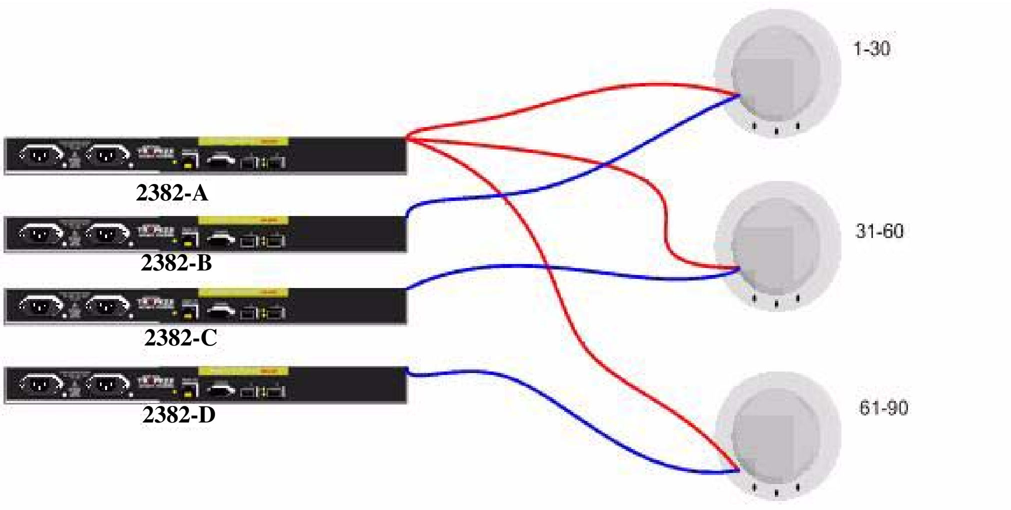

Wired network to WSS supported options

- Redundant uplink ports
  - Multi-Link Trunking (MLT)
  - EtherChannel
  - Split Multi-Link Trunking
- Redundant router protocols
  - VRRP
  - ESRP
  - HSRP
- L2 Loop detection
  - 802.1D Spanning-Tree Protocol (Untagged traffic)
  - PVST+ (802.1Q tagged traffic)

Security switch VLAN tunnel affinity

If multiple WLAN Security Switches are configured with the same VLAN, and the preference is for user traffic to use one Security Switch's uplink connection over another, the tunnel affinity on the preferred Security Switch can be increased.

This is typically used when one Security Switch has a higher speed connection to the core for a particular VLAN than the other Security Switch.

Access points

Access Points can support a wide variety of resiliency options. Redundancy for PoE, for data link connections and for Security Switch services can be provided to the AP.

Table 11: Typical AP redundancy types

Typical AP Redundancy Types	When to Use	PoE	Link	WSS
No AP redundancy	Not mission critical			
WLAN 2330A/2330B/Series 2332 Access Points Dual-Homed Direct Connections to a Single Security Switch	Smaller deployments (WSS 2360/2361)	X	X	
WLAN 2330A/2330B/Series 2332 Access Points Dual-Homed Direct Connections to Two Security Switches	Smaller deployments (WSS 2360/2361)	X	X	X
Single-Homed Distributed APs	Larger deployments when APs need to be more than 100 meters from the MX			X
Dual-Homed Distributed APs	Larger deployments when APs need to be more than 100 meters from the MX. Mission Critical usage	Opt	X	X

PoE redundancy

On WLAN 2330A/2330B/Series 2332 Access Points that have two Ethernet ports, PoE redundancy is provided by connecting both ports to PoE sources. PoE can come from a directly connected Security Switch, PoE Ethernet Switch or a PoE injector. Dual-homing support for PoE is automatically enabled when both WLAN 2330A Access Points Ethernet ports are connected.

Data link redundancy

Data link redundancy on the WLAN 2330A/2330B/Series 2332 Access Point is provided by connecting both Ethernet ports directly to one Security Switch, two Security Switches, an intermediate Ethernet switch, or a combination of Security Switch and Ethernet switch. If an intermediate Ethernet connection is used, a Distributed AP configuration on a Security Switch is required somewhere in the network. Dual-homing support for data link redundancy is automatically enabled when connecting both WLAN 2330A/2330B/Series 2332 Access Point Ethernet ports.

Security switch redundancy

Redundancy of Security Switch services is provided by dual-homing a WLAN 2330A/2330B/Series 2332 Access Point to two directly connected Security Switches; or by configuring a Distributed AP configuration either on two or more indirectly connected Security Switches, or on a combination of a directly connected Security Switch and one or more indirectly connected Security Switches. To provide Security Switch redundancy for a WLAN 2330 Access Point that has only one Ethernet port, configure a Distributed AP connection on two or more indirectly connected Security Switches.

Bias

On a Security Switch, configurations for APs have a bias (low or high) associated with them. The default is high. A Security Switch with high bias for an AP is preferred over a Security Switch with low bias for the AP.

In a failover situation, if an AP has more than 1 low bias connection then the Security Switch that has the greatest capacity to add more active APs is preferred.
For example, if one Security Switch has 50 active APs while another Security Switch has 60 active APs, and both Security Switches are capable of managing 80 active APs, the new AP uses the switch that has only 50 active APs.

AP selection of a WSS is sticky. After an AP selects a WSS to boot from, the AP continues to use that switch for its active data link even if another switch configured with high bias for the AP becomes available.

AP boot process

There are four methods that APs use to locate the Security Switches in their mobility domain. Following are the 4 methods:

- Static Boot-configuration
  - A DAP may be manually configured to store either IP address or DNS name values for an WSS switch from which it has to boot. If using the DNS name option, then an IP address for a DNS server has to be provided. DAPs has to be up and running on a system for these values to be written to persistent memory.
- DHCP (Option 43)
  - This option is recommended for most environments, especially medium to large environments.
  - Option 43 allows Distributed APs to find Security Switches across router boundaries, and provides the Network Administrator explicit control over which Security Switches, Distributed APs will initially contact upon boot-up.
  - When configuring option 43 for Distributed APs that are within a cluster (see switch design section) ensure that each cluster switch is included in the option 43 configuration. Placing the Security Switch with the high bias first in the option 43 configuration will slightly increase Distributed AP boot times, under normal conditions. For more information on configuration option 43 please see the "Avaya WLAN Security Switch 2300 Series Configuration Guide".
- Layer-2 broadcast
  - Requires APs to be on the same subnet as the Security Switch and is recommended for any environment where Distributed APs do not have to be placed across a router boundary which is typically smaller deployments, or evaluation systems.
- DNS
  - When use of option 43 is not possible, DNS can be used to allow Distributed APs to find Security Switches across router boundaries. This option also adds an additional failure point (the DNS server) to the boot process, which the other 2 options do not have. DNS is also not suitable for multiple Security Switch cluster environments due to the random nature of round robin DNS and loss of control associated with this approach.

System design – Network and Mobility Domains

Here is an illustration of the relationship between WSSs, mobility domains, and network domains. In this illustration mobility domain dependant features and information is exchanged between WSS1 and WSS2, as well as between WSS3 and WSS4. Only network domain level information (VLANs available for tunneling) is exchanged between mobility domain A and mobility domain B.

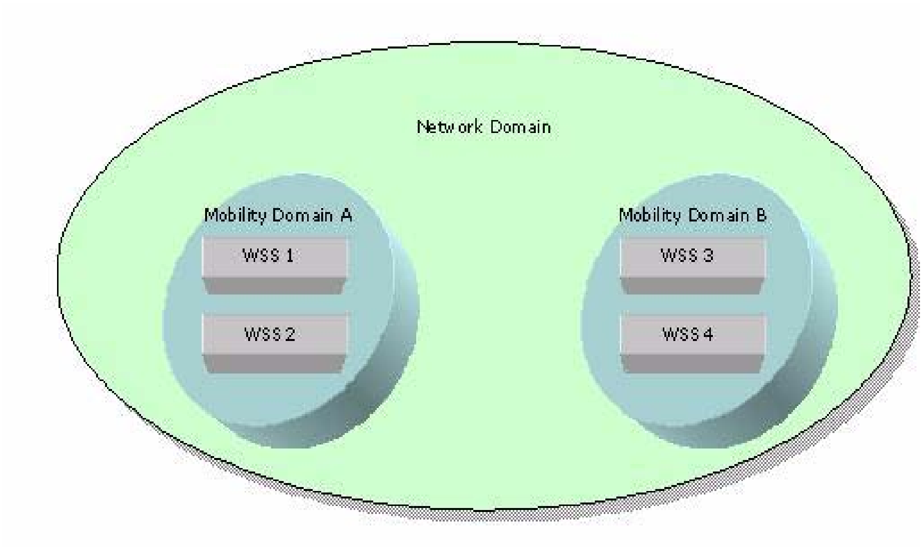

Table 12: Differences between Mobility Domains and Network Domains

Feature	Mobility Domain	Network Domain
AP failover (high-low bias)	Supported	Not supported
Subnet Roaming/VLAN tunneling	Supported	Supported
Inter-switch 802.11i Fast Roaming (Opportunistic PMK caching, requires WSS-to-WSS security enabled)	Supported	Not supported
Coordinated RF Countermeasures	Supported	Not supported
Switch and Client scaling limits	Lower	Higher

- The main consideration for which switches should be in the same mobility domain is AP booting and failover. Within a mobility domain each switch knows when every AP is configured and what it's bias is configured to be. This simplifies the AP boot process in that each AP can discover any member of a mobility domain in order to find it's best configured WSS. The mobility domain also knows when out-of-service WSSs return and will restore APs on low bias switches back to their high-bias switch when it recovers.

- When WSS-to-WSS security is enabled on a mobility domain it will allow PMKs for 802.11i fast roaming to be exchanged between all switches in the same mobility domain. This allows opportunistic caching to work between APS or different switches.
- As you increase the number of security switches in a mobility domain, it multiplies the amount of inter security switch communications when users authenticate or roam. These communications eventually reach a point where the scalability of the mobility domain breaks down, effectively limiting the total number of switches in any single mobility domain.
  For example, if you have 480 AP's across 5 WSS 2380s, is much better than across 24 WSS 2360s. To overcome the scaling limits in a single mobility domain split it into multiple mobility domains, all in the same network domain.
- Network Domains offer a limited subset of the services on a Mobilty Domain. The main function of a Network Domain is to allow seamless roaming between mobility domains when the roamed-to switch does not have the user VLAN locally configured. Because the Network Domain does not offer the full range of features which a Mobility Domain does it is possible to scale the number of switches in the domain much higher (500 WSS 2350 equivalent max.).
- If you have geographically separate areas that do not require seamless roaming, do not place WSS's in the same mobility domain. This is particularly important when you have large numbers of branch offices with single switches. In this case you should define them as stand-alone switches and not as members of a single mobility domain. This also holds true for Network Domain configuration as the primary benefit of the Network Domain is seamless roaming capability.
- Seed switches - Ensure that the seed switch in the Network or Mobility Domain is the highest CPU performance WLAN Security Switch in the domain.
  For example, you can make the WSS 2382 a seed switch when you have WSS 2360s and WSS 2382 in the same mobility domain.

System design – security

There can be many authentication and encryption combinations, and the table below list typical combination.

Typical security methods

Goal	Authentication and Encryption Method	Deployment Considerations
NO authentication at all, and minimal client configuration, broad compatibility including latency sensitive applications	Open, no encryption	NO authentication at all, and minimal client configuration, maximum compatibility
User name and password authentication and minimal client configuration, broad compatibility	Web Portal, no Encryption	User name and password authentication and minimal client configuration, maximum compatibility
User name and password authentication, basic data encryption, minimal client configuration, and broad compatibility	Web Portal, with Encryption	User name and password authentication, basic data encryption, minimal client configuration, and maximum compatibility
Stronger data encryption compatibility with latency sensitive applications	WPA-PSK	No per-user authentication, basic data encryption, minimal client configuration, and broad compatibility
Strongest user authentication and data encryption	WPA-802.1x	Strong, per-user authentication, strong data encryption, broad compatibility, and client configuration depends on authentication method used.
Stronger data encryption and compatibility with latency sensitive applications	WPA2-PSK (802.11i)	No per-user authentication Basic data encryption Limited compatibility with legacy devices Minimal client configuration
Strongest user authentication and data encryption	WPA2-802.1x (802.11i)	Strong, per-user authentication, strong data encryption, limited compatibility with legacy devices, and client configuration depends on authentication method used.

Encryption methods

Application	Method	Deployment Considerations
Only use if no other types are available	WEP	Industry standard WEP encryption has known vulnerabilities Widest compatibility
Widely available encryption that does not currently have known vulnerabilities	TKIP	Wide compatibility on older hardware
Security critical environments	AES	Strongest encryption currently available Requires hardware support

RF detection methods

- Level of IDS features needed?
  - Typical IDS features are already included in MSS.
  - For security-sensitive customers the Air Defense integration offers an additional feature set.
    - Best-in-class IDS features
    - "Air Defense Alerts sent to WMS-Unified WLAN and IDS alert system
    - "A WMS wizard can convert AP-2330A APs to Air Defense sensors if additional sensors are needed temporarily.
- Accurate rogue location required?
  - If so, must have planning data from the RF Planning technique provided by WMS.
- Countermeasures -See Matrix below

Table 13: Typical countermeasure techniques

Goal	Type	Description	Deployment Considerations
No counter measures and no AP performance degradation	None	No counter measures	-No counter measures +Does not degrade service
Best effort to contain rogue APs without affecting neighbors	Rogue	Attempt to only attack APs known to be present on the wired network without affecting neighbors	May not contain all rogue APs While APs performing countermeasure, AP performance is reduced Will not attack neighbor APs

Table 13: Typical countermeasure techniques

Best effort to contain rogue APs without affecting neighbors, and without reducing network performance	Rogue-with Sentry APs	Only attack APs known to be present on the wired network.	May not contain all rogue APs Increased cost While APs performing countermeasure, AP performance is reduced Will not attack neighbor APs Sentry APs tend to detect rogues quicker
Attack all rogues	All	Attacking all 802.11 APs not in the mobility domain	Will attack all APs on the network Will attack all APs on a neighboring network. While APs performing countermeasure, AP performance is reduced
Attack all rogues, without reducing network performance	All-with Sentry APs	Attacking all 802.11 APs not in the mobility domain	Will attack all APs on the network Increased cost Will attack all APs on the neighboring network. While APs performing countermeasure, AP performance is reduced Sentry APs tend to detect rogues quicker

Note. Countermeasure type “Rogue” is when countermeasures are configured to only attack APs which have been positively identified as being on the same network with the switches.

Countermeasure type “All” is when countermeasures are configured to attack all APs which are not on this mobility domain.

Sentry APs are dedicated APs only used for rogue detection and countermeasures. Sentry’s also provide more thorough rouge detection since they spend all their time dedicated to looking for rogues as opposed to also processing client traffic.

In dual-radio APs, 802.11a radios only attack 802.11a rogues, 802.11b/g radios only attack 802.11b/g rogues.

Authentication processing methods and recommended application

Application	Types	Description	Deployment Considerations
Environments that have small numbers of users not requiring centralized management or password synchronization, such as demo's or lab quick lab evaluations	Local	No radius server, database of users and certificates used for authentication are local to the switch.	No radius server required No password synchronization Limited centralized management of user names and passwords Limited scalability for the number of users Requires certificate configuration on the Security Switch Does not work with Windows XP "fast reconnect" OR machine authentication
Centralized user management and password synchronization where there is a slower WAN link between RADIUS server and the switch	Offload	All certificate processing is one by the Security Switch, and only the inner MSCHAPv2 request is passed to Radius.	Reduces number of packets exchanged between radius server and Security Switch Does not work with Windows XP "fast reconnect" Or machine authentication Only works with PEAP MSCHAPv2 Requires certificate configuration on the Security Switch
Most medium to large 802.1x deployments *	Pass-through	All EAP packets are passed-through to radius for processing	No certificate configuration on Security Switch Compatible with all EAP protocols Centralized use management and password synchronization for a wide variety of user backend types. Leverages fastest server technology that will be available on an ongoing basis for optimized EAP processing.

Implementation

Implementation – high level roll-out process

Listed below is a high level set of steps to perform prior to rolling out or upgrading any wireless network for medium to large environments:

- Build production similar test environment. Ensure test environment matches production environment as close as possible
- Stage small deployment with non-mission critical set of users on test environment
- Develop test plan to include:
    - all business goals
    - applications included
    - fail-over redundancy testing
- Performing testing on this small deployment
- If legacy wireless clients exist, make sure all types are tested as part of the initial staging
- If buying new wireless gear, purchase a small number of wireless client candidates to evaluate before purchasing in bulk.
- Adjust purchasing decision for new client or system design based upon testing results
- Develop a plan to address maintenance issues:
    - Maintenance windows
    - System change logs
    - Frequency of planned software upgrades

Implementation – optimization client performance

WLAN client performance

The top 3 considerations for obtaining optimal WLAN client performance include:

- WLAN client NIC hardware selection
- WLAN NIC drivers
- WLAN client configuration

WLAN client selection

The performance variability between various vendors of NIC hardware, operating systems, and supplicant software can be substantial.

WLAN client drivers

- Make sure the latest wireless client card drivers are installed. Drivers can make a significant difference.
- Ensure to reference the manufacturer's web site to be sure to get the latest driver; do not reference the CD that was shipped with the card.

802.1x supplicants

- Always use the latest version of the supplicant
  - For the Windows XP built-in supplicant, be sure to use XP service Pack 2+WPA-2 Hot Fix (KB # 893357), even if WPA-2 is not used.
- When deploying and configuring supplicants:
  - For Microsoft Windows XP clients, on a Windows Server 2003, Active directory domain, there are automatic configuration tools, in the Active Directory Wireless Group Policy. For more details, see Microsoft KB article 81123.
  - With the Juniper Networks clients the ability to create a pre-configured deployment package to ease installations and roll-outs exists. Consult Odyssey documentation for more details.

WLAN client cards

- While extensive compatibility testing with all hardware, driver, operating system and supplicant would not be feasible there is currently known performance issue leading to frequent client disconnects with the following cards:
  - Intel Centrino 2200BG - Symptoms lessened if you are using driver 9.0.4.27 or newer
  - Intel Centrino 2915 AG -- Symptoms lessened if you are using driver 9.0.4.27 or newer
- The following network adapters have been known to experience various compatibility related issues depending on software driver revisions:
  - Intel Centrino 3945 - Symptoms lessened if you are using driver 9.0.4.27 or newer. Setting Qos mode to SVP On the radio-profile may also help.

  - Pre-802.11N specification cards (including some Apple AirPort Extreme chipsets) - Optimum driver revisions will vary by manufacturer. Drivers dated March 2007 and later are generally recommended.

Client configuration

Roaming configuration

In the client drivers there is typically a configuration that allows how quickly clients can adjust to new APs when roaming. For very mobile users the client scans for new APs more frequently, and the configuration parameter on the driver should be set more aggressively.
For example: Scan_Valid_Interval (lower value more aggressive), or Roaming_Aggressiveness (named values) may be used a parameter names. Please consult the supplied client driver documentation for roaming options.

Power management

If users experience performance or client disconnects, users can try disabling any power management functions that may be enabled on the wireless card. Set this in the Advanced tab of Wireless LAN Adapter Properties.

Configured SSID's

Ensure there is only one possible configured SSID in one location. If two wireless networks are offering service in the same area and both of these networks are configured under the client's Microsoft DOT1X supplicant in "Preferred Networks", some clients tend to jump between both defined SSIDs. If only one of these two networks is defined, the same clients tend to stay connected to the single preferred network.

Multicast

In environments where multicast video quality is critical, restricting clients to either 2.4Ghz or 5.8Ghz band, as opposed to both, will reduce the amount of time the client spends scanning for APs off-channel, and will improve client video quality.

Simultaneous wired and wireless connections

Some client drivers have an option to automatically disable the wireless network when the client is connected to a wired Ethernet connection. This is a good feature and will cut-down on unnecessary wireless traffic.

If you do not understand driver settings, leave them at default.

Security switch configuration

Considerations for adjusting the 802.1x timers via WMS

- Quiet period timeout
  - The default setting of "60" will prevent a client who fails authentication/authorization from connecting for a period of 60 seconds. Lower this setting during the initial configuration and testing to avoid connectivity interruptions while configuring clients for the first time.
- Maximum requests
  - The default setting of "2" specifies the number of times an 802.1x packet is retransmitted before timing out. In noisy wireless environments, increasing this setting will make the system a little more flexible in the presence of a lousy wireless client.

- Retransmit timeout
    - The default setting of “5” specifies the number of seconds before the Security Switch will attempt to retransmit the initial EAP “Identity Request” packet. In environments where users have to manually enter usernames and passwords increasing this value will give the user more time to enter their credentials. In noisy environment where credentials are cached automatically lowering this value can reduce interruptions due to timeout.
- Supplicant timeout
    - The default setting of “30” specifies the number of seconds before the Security Switch will attempt to retransmit EAP packets after the initial “Identity Request”. In most environments this value should be reduced to 1 or 2 seconds to reduce delays introduced by packet-loss during authentication.

Long preamble

- Set radio-profile default preamble-length short.
- Older clients may require long preamble.
- Older wireless cards/drivers may experience connectivity problems if they do not understand short preamble.
    - Cards that use the Agere chipset, like the Dell TrueMobile 1150 and the Orinoco Gold card, require long-preamble or they will have connectivity problems.
    - The Symbol 6800 series terminal scanner may also require long preamble.

Avoid mixing ciphers

If the service-profile has both CIPHER-WEP104 and WPA-IE enabled, pick one or the other. Some cards cannot connect to SSID’s that advertise the ability to do both dynamic WEP and WPA at the same time.

Beaconing SSIDs

It is recommended that all SSIDs be beaconed to ensure maximum compatibility with wireless clients. Disabling of SSID beacons does not provide any tangible increase in security, it is purely cosmetic.

Proxy-ARP and no-broadcast

Due to the way in which broadcast packets are handled by 802.11 systems it is sometimes the case that the broadcast rate may exceed the transmit bandwidth/time-slice allocated to service these packets; which can in turn result in tail-drop. It is recommended that larger deployments enable both the proxy-arp and no-broadcast features on all service-profiles as a means to limit traffic that must be forwarded by this mechanism as much as possible.

WLAN client management

Client management strategies typically vary by the type of user environment that they will be deployed in.

Following are the 3 typical environments and management recommendations:

- Higher education (Colleges, Universities, etc.)
    - Since client selection is not controlled, and IT administration and configuration of clients is not centralized there are few management suggestions for these type of environments:
        - ❍ Perform regular end-user schedule configuration audits to improve end-user experience
        - ❍ Define and publish recommended configuration guidelines for typical system types.

  - Use authentication methods with broader compatibility
- Enterprise
  - Try to standardize the wireless clients on a known-good performing product.
  - Maintain standard images including testing drivers and configurations.
  - Implement push-technology for centralized configuration and driver management.
- Healthcare
  - A hybrid of Higher Education & Enterprise environments
    - COWS (Clients on Wheels) maintained similar to enterprise clients
    - Doctors that are not part of the staff are similar to college environments where there is no centralized administration or standardized client.

Implementation – security

Access control lists (ACLs)

ACL’s can be used to filter IP traffic based on IP protocol, TCP port number, TCP established state, UDP port number, or ICMP type and code. ACL’s can be configured per User, per VLAN, and per port; and can also be mapped to ingress (.in) traffic, egress (.out) traffic, or both for each mapping.

There are some limitations as to the number and complexity of ACLs and ACEs (individual rules that make up an ACL) supported.

- Each switch model has a maximum number of "mapped" ACE entries across the entire switch.
- Each time an ACL is mapped to a user/port/vlan/dap all of the ACES in that ACL count against the switch-wide maximum.
- Use VLAN or port ACLs when possible to conserve switch-wide ACE resources.
- If you are using the web-portal feature keep in mind that each un-authenticated web-portal user has the portal-acl mapped. Be sure to leave enough unallocated ACEs for the expected maximum number of un-authenticated web-portal users.

Table 14: ACEs per ACL

Model	# of ACEs Total	# of ACEs per ACL
WSS 2350	700	25
WSS 2360/2361	700	267
WSS 2380	2308	267
WSS 2382	2308	267

If providing guest access, it is highly recommended to create an ACL to limit the type of traffic guest users can access.

L2-restrict

The L2-restrict feature can be configured on a per-VLAN basis to limit network access to a defined set of destination MAC addresses. This is typically used to prevent users from communicating with any other devices on the same subnet other than the default gateway.

Restrict IP services for Guest users with a per-user ACL

If there are different classes of users on the same VLANs, you may also want to restrict which services they have access to. This can be accomplished by using the "Filter-ID" RADIUS attribute to dynamically map ACLs to specific users. You could use a filter-ID on the last-resort user to specify one set of ACLs, while mapping a different ACL (or no ACL at all) through MAC-Authentication for specific other sets of devices.

This example ACL specifically allows DHCP, DNS, HTTP, and HTTPS connections to be initiated from a client. Add permit statements for other protocols as appropriate, or deny statements to forbid specific internal networks.

```
# set security acl ip GuestInternet permit udp 0.0.0.0 255.255.255.255 eq 68 0.0.0.0
255.255.255.255 eq 67

# set security acl ip GuestInternet permit tcp 0.0.0.0 255.255.255.255 0.0.0.0
255.255.255.255 eq 80

# set security acl ip GuestInternet permit udp 0.0.0.0 255.255.255.255 0.0.0.0
255.255.255.255 eq 53

# set security acl ip GuestInternet permit tcp 0.0.0.0 255.255.255.255 0.0.0.0
255.255.255.255 eq 443

# set security acl ip GuestInternet deny 0.0.0.0 255.255.255.255

# commit security acl GuestInternet

# set user last-resort-GuestSSID attr filter-id GuestInternet.in
```

DHCP-Restrict

The DHCP-restrict feature, when enabled on a service-profile, will prevent a connected user from communicating over the network until the MX observes that the user system has received an IP address via DHCP. The MX will also prevent the user's system from communicating if it detects a change in that system's IP address that is not the result of further DHCP action. This can be useful on informal, open-access networks to prevent denial of service attacks related to duplicate IP addresses.

DAP Security and Encryption

To activate DAP security, the RSA fingerprint printed on the back of a DAP needs to be entered into WMS. DAP security controls the behavior of AP authentication by the switch, as well as encryption of WSS-to-AP control traffic. Encryption also provides improved CRC checking of TAPA packets, which can improve system stability in the presence of corrupted packets.

To avoid typing all of the RSA keys by hand, use the WMS alert resolution feature to automatically populate the fingerprint.

Table 15: DAP Security and Encryption

DAP Security Configuration	APs with RSA key	APs without RSA keys
DAP security = optional (default) No fingerprint configured	Encrypted	Unencrypted
DAP security = optional (default) Fingerprint configured	Encrypted, Authenticated	Non-operational
DAP security = required No fingerprint configured	Non-operational	Non-operational
DAP security = required Fingerprint configured	Encrypted, Authenticated	Non-operational
DAP security = none No fingerprint configured	Unencrypted	Unencrypted
DAP security = none Fingerprint configured	Unsupported	Unsupported

Secure MX to MX communications

Each switch in a mobility domain can be configured to authenticate and encrypt MX-to-MX control traffic. In addition to the security benefits, this also allows the Opportunistic PMK caching type of 802.11i fast roaming to operate across switches in the same mobility domain.

If you do not require these features, it is best to leave this feature disabled. The additional processing to encrypt the control traffic can reduce the scalability of the mobility domain.

Network users and guest user configuration recommendation

For security reasons, it is strongly recommended not to mix encrypted clear data on the same VLAN. There are two main reasons for this:

- Broadcast packets will go out on both the encrypted and clear data paths. This practice can reduce the overall security of the encrypted network.
- Assuming 802.1x in a normal network, if the same VLAN was used and clients are obtaining their IP address via DHCP, additional configuration would be necessary. Access lists would need to be, and can be, created such that the guests can have access to certain common network services, such as DNS. This would be necessary because it would be impossible to distinguish between guest and internal network users

Implementation – WMS

WMS RF planning tips

This area provides tips that are helpful when planning RF in WMS. It will cover the following broad sections:

- Planning Algorithm Overview
- Preparing the Floor for RF-Planning
- Planning Wireless Coverage
- General Drawing Tool Tips

Planning algorithm overview

WMS uses a mathematical model to model RF coverage that involves extending vectors around the transmitting device, and using a logarithmic path loss equation that approximates free space in an X, Y and Z axis. This equation has different constants, depending on the frequency being modeled (2.4 GHz or 5 GHz). Also taken into account are the optional capacity requirements (number of users, baseline association rate) for the area to be covered.

The computations takes into account defined RF obstacles. As a ray encounters an RF obstacle, that obstacle's attenuation is added to the path loss. When (TX Power - Path Loss) exceeds the expected receive capability for the given data rate, the end point for the ray is defined.

Compute and place APs

Two calculations are performed to determine the number of AP's required for each coverage area. One is based on capacity (traffic engineering) and the other is based on pure RF coverage (at a given data rate).

Capacity calculation for the number of AP's required is a straightforward math equation, drawing on several parameters: Baseline Association Rate, Number of Users, Throughput (Full duplex) per user, and a user traffic oversubscription ratio. This calculation also takes into account the protocol inefficiencies of CSMA/CA and contention. The result = #APcap.

Coverage calculation uses a proprietary iterative computational process combined with country-specific homologation information. The algorithm declares success when >90% of a grid of points within the coverage area (granular to 1 drawing unit) is covered by the contour at the baseline association rate with the least number of APs. This determines the number of APs, which = #APcov

The Number of APs from capacity (#APcap) and number of APs from coverage (#APcov) are compared, and the bigger count "wins". Placement re-uses the locations determined by the APcov algorithm.

Note: Using a "clean" RF model is imperative for best results. If you have many parallel RF obstacles that are close together, the placement algorithm will tend to place many more AP's required. So, even with the automatic clean layout mechanism in WMS, complex drawings demand additional pruning and isolation of single RF obstacles objects to keep the RF obstacle count as low as possible.

Compute optimal power

The initial “Compute and Place” step is performed using the maximum allowed power for the selected channel set in the defined regulatory domain. In “Compute Optimal Power”, optimal power can be computed for each AP, where transmit power is adjusted (up or down) to provide adequate coverage with minimum RF interference.

Also, this step is broken out so that the operator can manually change positions and counts of APs (add or remove) before the final power optimization is performed. This should typically happen often, given that an operator can interpret the floor plan and understand any cabling or other building constraints to avoid positioning problems.

Optimize AP Count

It applies to coverage areas that are based on coverage criteria - not capacity criteria. Coverage overlap is checked to see if any APs are not strictly required to provide coverage, that is. the APs around it are doing the job. If overlap is good enough the superfluous AP is removed when this option is checked.

Assign channels

The channel assignment algorithm assigns non-overlapping channels to neighboring APs from the selected channel set.You must choose the starting floor and the ending floor (in the downward direction) for multi-floor channel assignment. The algorithm takes predicted RSSI values between neighboring APs (including APs on different floors and 3rd party APs) and minimizes same-channel assignments between APs. You can control whether cross-floor attenuation is to be considered, and what 802.11 technology on which you want to perform the channel assignment. Note that it strictly uses predicted RSSI values for the imaginary "ray" that is drawn between two APs. So, it is possible that strange results can be seen if the exact path between the APs has many obstacles, but the areas around that path are relatively open. Again, the operator can change things if they feel it necessary.

Preparing the drawing for RF planning

An essential piece of RF planning is a floor layout or drawing. You can import a CAD DXF/DWG, JPG or GIF file into WMS to define the floor layout. Note that there’s a big difference in the usefulness of a drawing, depending on the file type.

- A DWG or DXF (AutoCAD) file is made up of vector graphics objects (i.e. lines), which can be subsequently converted into RF obstacles
- A JPG or GIF file is simply a raster graphics file (i.e. a screenshot or background), which is NOT made up of lines. Instead, it’s a (hopefully) scaled visual background of the floor layout that can be used to manually draw RF obstacles on top of it.
- Sometimes, there isn’t any file available at all. In this scenario, manually create RF obstacles.

Sometimes, the Visio file type is mentioned. Though Visio is not directly supported by WMS, Visio itself can export its drawings as DXF or JPG (depending on what method is used for RF obstacle creation in WMS).

DWG/DXF version support

DWG is the native binary format used by Autodesk’s popular AutoCAD software. The versions of DWG that can be imported into WMS are: R13, R14, R2000.

DXF is an ASCII-based interchange format that is designed for multi-vendor interoperability (much like RTF is for text documents). The versions of DXF that can be imported into WMS are: R12, R13, R14, R2000.

A DWG file, when converted into a DXF file, generates a file size of 3 times the size of the DWG file. This is a typical file size conversion: DXF File Size = 3 * DWG File Size.

Past experience has shown that DXF will tend to import more cleanly into WMS than DWG. However, exceptions to this have also been noted, so if possible, have both formats available to import in case one of the formats has import problems.

Preparing the drawing in popular CAD tools

WMS can be used with any size of imported drawing files. However, for optimal use, the drawings imported into WMS should be < 0.5MB for DWG files or < 1 MB for DXF files.

Drawing tools

A mouse with a scroll wheel – critical for easy zooming in CAD and WMS software

Drawing basics

A drawing typically contains “WORKSPACES” to contain objects. Typically, the following workspaces are found:

- Model (typically 1:1 scale to draw the objects)
- Layout/Paper (varies in units to fit the model onto a printable sheet of different sizes)

Note: It is the scaling units defined in Paper and other layouts that are NOT REQUIRED when considering using this drawing in WMS. WMS’s reader does not distinguish between these work spaces. All are read at the same time. So, it’s imperative to remove all the extra layouts.

Removing scaling units

AutoCAD

1. Delete all layouts/workspaces that are not required
2. Edit the properties of document
3. Ensure that Paper Space is 1” : 1” (Full Size)
4. Ensure that the scale type is same as that of the model space.

Review and prune the drawing

1. As mentioned before, remove all paper layouts. The last one cannot be removed, but simply remove the contents of the layout after it regenerates.
2. Check for externally referenced files.

   To keep drawings manageable and modular, drawings can link and reference other files. In some cases drawings are linked and referenced to other files. A sign that there is external references is when significant parts of the floor plan are missing, even with all the layers unfrozen and visible.

   If the externally referenced files are not used, in some cases these files may contain RF pertinent information, then obtain the files and put them in the same directory as the master file. Once done, use the sub-file for import into WMS.

WMS does not read from externally referenced files. If the information from the other files is required, use a "binding feature" to merge the external files into the master drawing.

3 Audit the CAD drawing

If possible use the audit function to find problems between objects and possibly fix them automatically. Consult your CAD documentation for information on feature usage.

4 Turn visible, unlock, and unfreeze all layers in the drawing (eventually) and delete unnecessary layers. Layers can be made visible or invisible (much like WMS). Locking a layer keeps it visible, but it's locked from any changes (read only). Freezing a layer both locks a layer and makes it invisible.It is suggested to review CAD drawing invisible layers, if it's believed that these layers will not add to the WMS RF calculation then remove them.

⚠ **Warning!** Some CAD software uses the Ctrl-A (Select All) function to select all the objects in the model space, regardless of layer status (invisible, locked, frozen). At this point all the invisible objects are unprotected. Exercise caution if deleting objects in this mode.

When completed, all layers should be modifiable and visible.

Grouped objects

Some drawings contain objects that have been grouped together. It can sometimes be the whole drawing. Grouped objects can span layers, they often appear when many objects are selected, even though only one object was initially selected. Here are some scenarios that have been encountered:

Complete sections of a drawing are grouped (sometimes the whole drawing). If the discrete objects for RF usage is needed, then "explode" the grouped object.

Grouped objects span layers, and may not be completely visible. Deleting one object of a group could delete all objects of that group, visible or not. The best way to spot a grouped object that spans layers is when the object is not selected normally when clicked on. Instead, a "selection" square ends up offset to the side of the object.

Don't explode everything, though.… sometimes it makes sense to leave certain groups alone, if they can be characterized as a single RF obstacle.

5 Purge all blocks, line types, and layers that are unused

To make things easier, remove any layers that do not depict the structure and materials of the building that would ultimately affect the RF characteristics of the deployment. .

6 Create new RF layers and pull walls, glass, and other RF-affecting objects from other layers in the new RF layers.

For a really clean drawing, it's recommended that new RF-specific layers be created and only desired objects from other layers are moved into these new layers. Usually, the desire is to end up with an RF-ExtWalls, RF-IntWalls, RF-Windows layer, where similar objects are grouped together. The possibly frustrating issue here is that typically only one of several redundant and parallel lines to represent the RF obstacle for WMS's model is needed. So, picking out a subset of the objects in the layer for transfer to the RF-specific layer may be manually intensive. The reward for this up front works is a very fast and clean conversion of your objects into RF obstacles in WMS.

7 Save the drawing into both a .dwg and .dxf format for import into WMS, just in case one particular format doesn't import well. File >> Save As allows a choice of the file format in which to save the CAD drawing. Usually R2000 format for either DWG or DXF files is selected.

Common CAD layer names:

- glaz – windows
- flor – floor-related stuff (e.g. flor-evtr = elevators, flor-hral = handrails, flor-strs = stairs)
- furn – furniture (e.g. furn-char = chairs)
- wall – walls (e.g. wall-clng = ceiling soffits/overhead)
- scol – steel column
- p-fixt or plbg-fixt – bathroom fixtures
- p-part – bathroom stall partitions
- ext – exterior
- int – interior
- anno – annotations (text)

RF planning with WMS

Importing the drawing into WMS

Import the drawing into WMS and observe the drawing. Perform the following, if necessary:

1 Remove white space around the drawing. Make sure the drawing is as close to the TOP-LEFT corner as possible. This will help in getting the entire floor plan in JPG snippets that are created for work orders. To perform this, do the following:

   a Select All

   b Group

   c Move the object to desired location. DO NOT move the drawing into invisible area in top or left direction. This is negative space and contents with negative coordinates will not be visible.

   d Ungroup

   e Fit To View.

2 Review the contents of the drawing and remove unwanted objects using WMS actions of deleting layer, delete objects, etc. This is unnecessary if the drawing was prepared in a CAD program.

Scale the drawing

It is very important that the drawing be scaled. When a file (DWG, DXF, JPG, GIF) is imported, no scale information is read from those files.

1 Verify the desired unit of scale (Feet or Meter) in the Building Setup page

In Edit Content, zoom into the drawing and select a door and use the scaling tool to correct it to the desired length. Typically office doors are 3 feet wide. However, if there is a known distance between any two points, it can be used to correct the scale of the drawing, as well.

Perform an automatic CAD layout cleanup

Using “Clean Layout” action is highly recommended for the following reasons:

- Removes parallel lines that might have been used to represent the same object. (For e.g. wall).
- Removes small objects that might not have been removed when reviewing the drawing in CAD tools.

The following objects are untouched when cleaning layout:

- It is an operation within the same layer. i.e.; if two objects are identical and one has to be deleted, it will be only if the two objects are in the same layer.
- A pre-existing RF-obstacle will not be touched.
- A Grouped object will not be touched. If the group contents are to be deleted, then ungroup the objects before performing Clean Layout.

Lines with arrowheads

Arrows depicting dimensions in the original drawing do come into WMS as separate objects. Some arrows might have Arrow heads. These grow in size when one zooms in/out from the drawing. Clean Layout DOES NOT delete such lines as they typically come in as part of the group. So, ungroup the objects where arrowheads are seen. Delete those lines with arrowheads.

Collate similar objects in one layer

It is recommended to put all objects intended to be similar in construction material to be in one layer. For e.g., if the drawing file had walls spread out in different layers and upon site-survey, they were found to be similar, it is better to put them in one layer so that assignment of RF attenuation can be performed in one shot.

Perform CAD cleanup to remove parallel lines

Typically, based on the drawing pencil chosen when the drawing file was created, a single object might be drawn with more than 1 line. When such an object is imported, it results in more than one object in WMS. To avoid the actual object now being defined as more than one obstacle, deletion of parallel lines within a certain distance is highly recommended.

The other way to achieve the same is to group all those lines into one object. However, this might not work well with polylines as the grouping of polylines causes planning tool to assume the bounds of the group to be an obstacle. This will cause incorrect results when viewing RF coverage. In some cases, this might be a perfectly logical way to simplify an obstacle. An example of this can be, 4 lines forming an office/conference room and one attenuation factor is desired for that entire area. Another example is where multiple lines are drawn to make a bigger line.

Note: Objects must not be RF Obstacles or Groups before Clean Layout is performed.

Removing RF obstacle information

The user can remove RF obstacle information by layer or by individual object by using the “Remove RF Obstacle Info” feature/button.

Coverage area shapes

As mentioned in the WMS Admin Guide, almost all possible shapes for a coverage area are possible. However, there are following restrictions to remember:

A shape where two sides intersect each other are not permitted:

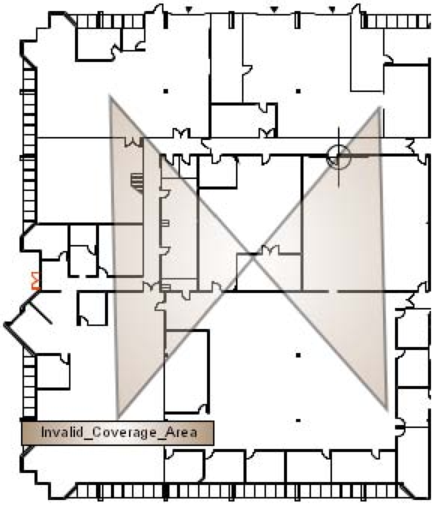

Shared Coverage Area where there is a partial intersection is not supported. The shared areas must either be identically placed and sized, or one area must be entirely contained inside another:

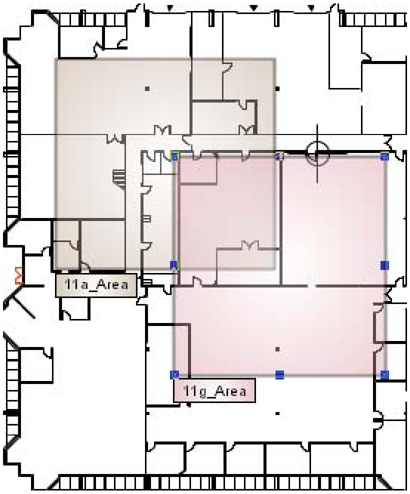

Note: For computational reasons, when a coverage area is drawn, it aligns to the grid to provide a whole number for width and height of the shape.

Shared coverage areas slightly moved

It is possible that after shared coverage areas are drawn, one of the coverage areas is slightly moved. This will cause RF planning tool to not compute for this coverage area, as it would fall into the restricted coverage area shapes. To align the two shapes manually is pretty tough. The easy way to do it is by changing X/Y coordinates:

1. Select Area1
2. Click on the Dimension Icon in the toolbar
3. Note the values and click OK
4. Select Area2
5. Click on the Dimension Icon in the toolbar.
6. Make sure the values shown in the dialog match what was noted in Step 3.
7. Click OK

Baseline association rate

Whether the AP computation is desired for capacity or coverage, please note that all computations are performed at the selected Baseline Association Rate.

If the AP count is found to be too excessive, the baseline association rate can be reduced. Do the following:

1. Check “Compute For Capacity”
2. Select the desired baseline association rate
3. Uncheck “Compute for Capacity”

For WMS 4.1 and above, the default values for association rate are: 11Mb/s for 802.11b/g and 36Mb/s for 802.11a.

Note: The baseline association rate affects the following areas in WMS:
AP Count computation and placement
AP Power optimization
RF Coverage
RSSI measurement visualization

AP count computation and placement

To fully utilize the features of the planning tool, the following conditions must be true:

a. No AP must be locked

- **b** All APs, if already present, must be of the same type as the constraint “Default AP Choice” in the coverage area

The important steps that happen in this entire operation are as follows:

1. For each coverage area, maximum RF cell radius is computed using the following information:
   - **a** Technology of coverage area
   - **b** Default AP choice
   - **c** Allowed channel set
   - **d** Max. TX. power allowed for that country on that allowed channel set
   - **e** Max. receiver sensitivity of the AP radio for that baseline association rate
2. Points within the coverage area are selected as “potential AP placement” points (the exact algorithm is secret sauce). At such a point, a potential AP is placed and the RF coverage is drawn to see if the entire coverage area is satisfied. Additional points are added as needed.
3. In the areas that are shared, the same “potential AP placement” points are used to see if it satisfies the other technology.
4. Once a potential AP placement point list is collated for each coverage area, RF coverage is envisioned at all points to see if any redundant AP placement points can be removed.
5. The end result is a set of points that ideally would cover the entire coverage area.
6. If APs are locked, then their location does not change. However, potential AP points closest to those locations are removed. APs are created and placed on the remaining locations.

➡ **Note:** For best results, APs must not be locked.

Factors that can impact computational time

There are lots of factors that can make the computational time a little long. Enough RAM must be provided if the floor plan is found to be complex and large.

The factors that impact the computational time are:

1. Size of the coverage area
2. Shape of the coverage area. More concave angles can take longer.
3. Whether the area is shared or not
4. Number of obstacles on the floor
5. RF attenuation of the obstacles on the floor
6. Type of AP selected as default choice
7. Selected Baseline Association Rate

Because there are a lot of variations and factors, an average time is not deterministic.

Power optimization

The only variables when power is optimized for a coverage area are:

- **a** The power for a radio

b. The starting power for a radio

All other characteristics of the radio are constants.

Factors that impact time for power optimization

The following factors impact the time it will take to optimize power for a given coverage area:

1. Technology
2. Starting Power
3. Maximum Power allowed for the channel set. This determines how many steps of power to permute on. WMS uses steps of 2dB.
4. Number of Radios (This parameter has high impact on time)
5. Number of RF Obstacles
6. RF Attenuation of Obstacles

Since there are numerous variations and factors, an average time is not deterministic.

Mesh, Bridging, and Local Switching

Overview

There are three key components to a mesh deployment, the Mesh Portal, Mesh Link, and Mesh AP. The Mesh portal is the AP that provides an uplink to wired access for the Mesh AP. The Mesh AP is the fully wireless AP. It accesses the wired networks using a Mesh Link to the Mesh Portal.

The bridge is simply a specific configuration of the mesh link where the traffic is flooded across the bridged mesh link. In the case of the bridge, the mesh link is a layer-2 transparent bridge with the Mesh Portal and the Mesh AP acting as bridge endpoints. It is assumed that both bridge endpoints are attached to wired networks. There is no requirement that there is a WSS on the Mesh AP side of the bridge.

The 2332 AP is the only AP that support Mesh or Bridging.

Figure 1 on page 52 shows how a client connects to a network using WLAN mesh services.

Figure 1. WLAN Mesh Service

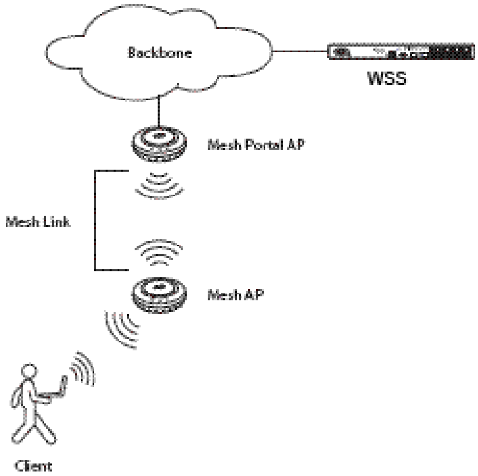

The critical differences between a Mesh AP and a wired AP are the initial configuration and authentication. A Mesh AP must be pre-configured prior to deployment with the Mesh link SSID, pre shared keys for authentication and mesh mode enabled. Optionally, the control channel timeout may be extended. The timeout should be increased depending on the length of the mesh-link.

The following commands are used for pre-configuring the Mesh AP:

set ap <*num*> **boot-config mesh mode** *<enable | disable> [ssid <ssid>]*

set ap <*num*> **boot-config mesh** *[psk-phrase <pass-phrase>] | [psk-raw <raw-pass>]*

Configures the AP to operate using a mesh link

mesh mode: enable or disable the use of the mesh link

ssid: specifies the SSID that should form the mesh link

mesh psk-raw: pre-shared key in raw form

mesh psk-phrase: pre-shared key pass phrase

clear dap boot-config will clear the static configuration stored on the AP

set ap <*num*> **time-out** <*seconds*>

Changes the TAPA control channel timeout on the AP (default: 10 sec)

When a Mesh AP boots it searches for the Mesh Link SSID and then gets authorized using the Mesh AP's highest or 2nd highest MAC address and then authenticates using the PSK information. The Mesh AP is authenticated with the Mesh Portal that advertising the Mesh Link SSID. The Mesh AP authenticates to the Mesh Portal as a client. If there are multiple Mesh Portal advertising the Mesh Link SSID it will choose the Mesh Portal with the strongest RSSI value. Once the authentication is complete the Mesh AP searches for a WSS using the same control packet exchanges as any other AP on the network. The Mesh Link is an authenticated encrypted radio link. Once the Mesh Link is established the Mesh AP will not switch to another mesh portal unless it looses contact with the original mesh portal.

The initial bridge configuration and establishing the Mesh Link is identical to Mesh. Once the link is established the service profile specifies that it is a bridge link. The Mesh Portal tells the Mesh AP that the link is a bridge link rather than a standard mesh link. The APs on both side of the bridge link must be dedicated to the bridge. Neither the Mesh Portal nor the Mesh AP that are part of the bridge link accept clients.

All APs participating in the mesh and bridge are counted against the WSS's AP license count.

Mesh Recommended Configuration

The following recommendations should provide the most stable Mesh.

- Dedicate one radio to client services and the other to the mesh service. Avaya recommends dedicating the 802.11a (radio 2) radio to mesh and the 802.11 b/g (radio 1) radio to clients.
- Dedicate the Mesh Portal to mesh services if there is a full client load through the Mesh AP.
- Limit the length of the mesh link to 3/8 mile or less when using 6.0.4 or earlier. WSS version 6.0.5 will support 1 mile links.
- Enable Local Switching on the Mesh Portal and all the Mesh APs. Although Local Switching is not related to the mesh technology in certain client usages it will improve your throughput.
- If the Mesh APs are "hidden nodes" to each other, RTS should be enabled on the radio-profile(s) used for the Mesh Links. If you are using highly directional antennas on the Mesh APs, you almost certainly have hidden nodes

Bridge Recommended Configuration

The following recommendations should provide the most stable bridge link:

- Limit the length of the mesh link to 3/8 mile or less when using WSS 6.0.4 or earlier. WSS version 6.0.5 and later has been tested to 1 mile.
- In point to multipoint deployments where the Mesh APs are "hidden nodes" to each other, RTS/CTS should be enabled on the Mesh Link radio profiles on all APs in order to avoid collisions. Most point to multipoint deployments would fall in this category unless omni-directional antennas are being used and all nodes within the point to multi-point tree are in range of each other.
- If you need the maximum possible throughput, avoid using point to multipoint deployments. Straight point-to-point links provide the highest possible throughput.
- Use a link-budget calculation tool to help determine the minimum link margin.
- Use narrow beam-width antennas whenever possible. A narrower beam will reduce interference from transmitters outside of the bridge link.

Mesh and Bridge Limitations and Restrictions

This is an initial release of the Mesh and Bridge features in WSS 6.0.

Mesh limitations in WSS 6.0:

- Each Mesh Portal can support at most 5 Mesh APs
- Each Mesh Portal support one Mesh Link deep
- There is no dual role Mesh Portal/Mesh AP support in WSS 6.0.
- Multi-hop mesh will be introduced in WSS 6.2.
- All Mesh APs using the same mesh SSID will use the same radio band
- The Mesh Portal may not be configured as a directly attached AP but must be configured as a DAP.
- Auto Tune (tx-power and channel) must be disabled
- Static configuration of channel and power is required in 6.0
- ActiveScan will not scan off-channel if left enabled Mesh Links.
- Auto-AP is not supported for Mesh Portals or Mesh APs
- Mesh cannot be configured using WebView. It is supported by the CLI and WMS

Bridge limitations in WSS 6.0 in addition to the above Mesh limitations:

- Bridged Mesh APs boot through the Mesh Portal even if there is a WSS on both sides of the bridge
- Mesh Portals and Mesh APs participating in a bridge link cannot accept clients
- Dual role Mesh Portals are not allowed. A Mesh Portal must support either the Bridge or a Mesh.

Mesh Performance

The maximum theoretical throughput of a client associated to a Mesh link is typically 10-15% lower that the performance of a client associated to a wired AP. This is due to the fact that a mesh client is traversing 2 air links with the subsequent packet loss.

Mesh and Bridge Hints and Tips

- The customer can monitor the Mesh AP associating with the WSS using the following commands:
    - show ap status terse
    - show sessions mesh
    - show rfdetect data
    - set trace sm level 5
    - set trace dot1x level 5
- Connect up a sniffer to the Mesh AP's Ethernet port - syslog packets with status are emitted during the bootup and association phase.
- Local Switching should be enabled on any Mesh AP and Mesh Portal accepting client traffic
- The Mesh AP side of the Bridge link must NOT be connected to the rest of the network to avoid loops. If a loop is required in the network topology enable spanning tree.
- An WSS may be present on the Mesh AP side of the Bridge link. The Mesh AP must not be configured on that WSS in order for the bridge to properly recover from outages.
- For each Mesh Link, use the highest gain antenna possible. Mesh Portal APs will generally need lower-gain omni or 180/120 degree antennas. Mesh APs can use the tighter-beam 18 degree antenna if you are not planning on adding additional hops later with WSS.
- Initially setting up the mesh link in the clear will facilitate debugging the link. Encryption can be enabled at any time.

Mesh Commands

Mesh Portal Command

- set service-profile <name> mesh mode <enable | disable>

Mesh AP Commands

- set ap <num> boot-config mesh mode <enable | disable> [ssid <ssid>]
- set ap <num> boot-config mesh [psk-phrase <pass-phrase>] | [psk-raw <raw-pass>]
    - Configures the AP to operate using a mesh link
        - mesh mode: enable or disable the use of the mesh link
        - ssid: specifies the SSID that should form the mesh link
        - mesh psk-raw: pre-shared key in raw form
        - mesh psk-phrase: pre-shared key pass phrase
- clear ap boot-config will clear the static configuration stored on the AP
- set ap <num> time-out <seconds>
    - Changes the TAPA control channel timeout on the AP (default: 10 s)
    - Set it to a longer value to survive WAN outages
- set ap <num> radio <num> link-calibration <enable | disable>
- show ap boot-configuration <n>
- show service-profile
- show ap status
- show ap status terse
- show ap mesh-links <num>
- show sessions mesh-ap verbose
- show sessions mesh-ap session-id

Bridge Commands

Mesh Portal (bridge) Commands

- set service-profile <name> mesh mode <enable | disable>
- set service-profile bridging <enable | disable>

Mesh AP (bridge) Commands

- No specific bridge commands - same as Mesh

Local Switching Description

Local switching, alternatively termed "Intelligent Switching", allows the customer to configure VLANs on APs where data traffic is switched directly to the wired network rather than being tunneled to the switch. The traditional overlay model and local switching can be supported simultaneously on the AP on a per VLAN basis. A VLAN is configured per AP for either overlay or local switching not both.

The AP switching path supports the following capabilities:

- Switching packets from the radio to the wired interface as a standard 802.3 packet
- 802.1q tagged packets (optional)
- Proxy ARP
- MAC or User based ACLs.

The following algorithm is followed for ACLs. The AP supports MAC or user based ACLs. If there is no MAC ACL but there is an ACL mapped to the client's VLAN, the AP will support the VLAN ACL. If there is no MAC or VLAN ACL defined but there is an ACL mapped to the client's DAP then the AP will support the DAP ACL.

Local switching is supported on the AP 2332 only.

Local switching is configured by creating a vlan-profile and then linking that vlan-profile to the locally switched AP.

The commands to configure an AP for local switching are:

- set vlan-profile <name> vlan <vlan-name> [tag <tag-value>]
    - name: name of the vlan-profile
    - vlan: name of the VLAN. There may be up to 128 VLANs for each vlan-profile. Execute this command iteratively to configure multiple VLANs for a vlan-profile.
    - tag: The optional 802.1q tag for the VLAN
- set ap <number> local-switching mode [enable|disable] vlan-profile <profile-name>
    - enable|disable: Enable or disables local switching; Default is disable.
    - vlan-profile: Specifies the name of the vlan-profile that is used to identify the VLANs connected to an AP. The default is no vlan-profile so there is one VLAN named "default" that is not tagged.

Local Switching Recommended Configuration

The following recommendation will provide the best use of Local Switching.

- If you want to use both traditional overlay and local-switching, every VLAN in a VLAN profile should exist on at least one WSS in the mobility domain to insure that you can seamlessly switch to overlay mode when the user's VLAN is not configured on their current AP.
- Use Identity Based Networking to assign users to VLANs by geography (proximity to wiring closets with the VLAN) in order to enable them to locally switch where they are likely to spend the most time.
- Use 802.1q tags to separate the AP management VLAN from any locally switched user VLANs. The management VLAN must be untagged, and each of the user VLANs needs to be tagged.
- Mix-and-match local switching and overlay on a per-application basis. If a specific application (e.g. VoIP!) can really benefit from local-switching, extend a VLAN for that application out to as many APs as possible with 802.1q tagging or physically connect as many APs as possible to the application LAN if one already exists (e.g. a VoIP VLAN).
- Locations where you are using local-switching and do not need seamless roaming (i.e. smaller branch offices) should not have the locally-switched VLAN created on the WSS. It should only be in the VLAN-Profile configuration for the APs.
- Configure Local Switching on all Mesh Portals and Mesh APs.
- Do not use Local Switching when the number of ACEs per ACL exceeds the limit for local switching.

Local Switching Limitations and Restrictions

- Local switching cannot be enabled on a VLAN that requires any of these features:
  - L2-restrict (VLAN)
  - IGMP snooping (VLAN)
  - DHCP-restrict (service profile -> VLAN)
  - Web Portal (service profile -> VLAN)
- ACLs on the AP are restricted to 16 ACEs. In 6.0.5 the limit will be increased to 25.
  - Do not locally switch VLAN with larger ACE lists.
- Local switching cannot be configured using WebView
  - Use the CLI or WMS
- Directly attached APs cannot be configured for local switching
  - Use the Distributed AP method on that port
- ACL changes are not dynamic. If an active ACL is modified, the stations must re-associate to pick up the new ACL
- VLAN Profiles changes are not dynamic. If a VLAN profile is changed from overlay to local switching or visa versa, all stations on the impacted APs must be terminated and forced to re-associate.

Local Switching Performance

The most noticeable performance impact from local switching is reduced latency on client traffic. With local switching the client will take the most direct path possible instead of being tunneled back to central controllers. The latency improvement will depend heavily on your network topology. The most dramatic improvements are in peer-to-peer traffic between wireless clients on the same VLAN, which is common in VOIP. Local switching will also improve the aggregate throughput of an AP, but to a much lesser degree. All control packets are still tunneled to the WSS.

Local Switching Hints and Tips

- You can view current VLAN information using the command "show ap vlan".

➡ **Note.** VLAN's are not created on the AP until they are needed.

- You can view the FDB table on the AP using the command, "show ap fdb"

Local Switching Commands

The commands to configure an AP for local switching are:

- set vlan-profile <name> vlan <vlan-name> [tag <tag-value>]
  - name: name of the vlan-profile
  - vlan: name of the VLAN. There may be up to 128 VLANs for each vlan-profile. Execute this command iteratively to configure multiple VLANs for a vlan-profile.
  - tag: The optional 802.1q tag for the VLAN
- set ap <number> local-switching mode [enable|disable] vlan-profile <profile-name>
  - enable|disable: Enable or disables local switching; Default is disable.
  - vlan-profile: Specifies the name of the vlan-provile that is used to identify the VLANs connected to an AP. The default is no vlan-profile so there is one VLAN named "default" that is not tagged.
- clear ap <number> local-switching vlan-profile
- show ap config
- clear vlan-profile <name>
- clear vlan-profile <name> vlan <name>
- show vlan-profile [<name>]
- show ap vlan <num>
- show ap fdb <num>
- show ap arp <num>

Appendix A - Identity Based Networking Concepts

Overview

In a wireless LAN system, a wireless client (sometimes called a "supplicant") will search the air for a wireless network to attach to. This client will then be asked by the wireless system to authenticate. The client may authenticate based on a client attribute such as the client's digital certificate or MAC address, or a user of the client may authenticate based on the user's digital certificate or other authentication credentials (MS-CHAPv2, PAP, MD5, etc.). Optionally an "unknown" client may be trivially authenticated as a guest.

Once the client/computer or user is authenticated (hereafter referred to as the "user"), the proper network authorization attributes are installed. These include the VLAN/subnet to which the user belongs, the area in which the user can roam and any ACL's or timeouts that may be needed. This authentication and authorization process usually involves a RADIUS server on the back-end, but may also be performed locally on the MX's.

Finally, if the user requires encryption services, dynamic keying material for both unicast and broadcast/multicast is generated. The unicast keying material is unique per-user and the broadcast/multicast keying material is unique per radio/VLAN/Encryption Type. In the case of dynamic WEP keying material, the keys may be automatically rotated on a periodic basis to help protect against the weak IV problems in WEP. This key generation can be done on a RADIUS server or locally on the MX's.

Once authenticated and encrypted, the wireless connection may now be used for the normal DHCP and NOS Login processes. Identity-Based networking insures that the user is placed in the proper VLAN/subnet so that as the user roams throughout the wireless domain he stays on the same subnet.

When a user roams across the wireless network, the Avaya Mobility System keeps track of the authentication and authorization credentials and authorizes roaming events without having to check back with the Web-based AAA server. This reduces the load on the Web-based AAA server and reduces the latency or a roaming event. Depending on client behavior, the user may either be re-authenticated or just re-keyed. In either case, this happens in the background with no user intervention required. Web-based AAA accounting follows a user as she roams and generates accounting updates for each roaming event.

Security primer

Security in 802.11 networking relies on the use of digital certificates, mutual authentication and encryption. To understand how all this works, you must first understand how Public Key / Private Key algorithms work. It's simple. Really, it is. Public Key / Private Key algorithms are used to provide encryption and authentication services. These "Keys" are analogous to passwords. Let's examine how this works.

With plain old username/password authentication, two parties have to know the same password: the person who is being authenticated and the authenticator. If both parties agree that the password being used is correct, the person is authenticated. This is just what happens when you log into your computer.

Likewise if two parties want to encrypt messages to send to each other they can use this same username/password concept. If both parties use the same encryption algorithm (DES for example) and both sides agree to use the same password then both sides can encrypt and decrypt messages to and from each other. This form of encryption is called "Symmetric Encryption" and the password being used is called a "shared secret" or "shared key" since both parties have to know and use the same password. In the case of encryption, the password is used as input to the encryption algorithm instead of being used to authenticate the user. But it's still just a matter of having two parties know the same password. By the way, Symmetric Encryption is used for all 802.11 encryption schemes (WEP/TKIP/AES).

The problem with the "Symmetric Encryption" model is that it's hard to keep "shared secrets" secret. After all, a "shared secret" is an oxymoron. Ideally you wouldn't share your secret, you'd have a real secret that only you knew.

This is where Public Key / Private Key algorithms come into play. The idea is that for any two parties that want to authenticate or encrypt messages, instead of "sharing a secret key" one side has a Public Key and the other side has a Private Key. These Public and Private keys are issued in pairs (key pairs). Any message that's encrypted with the Public Key can be read by whoever has the corresponding Private Key. Likewise any message encrypted with the Private Key can be read by someone with the corresponding Public Key. The most important attribute of these Public Key / Private Key pairs is that it's infeasible to generate one key from the other. That's what makes them so secure. The algorithms that support this are called Public Key algorithms (like RSA for example).

The real magic is that Public Keys can be freely distributed to anyone. There is nothing about a Public Key that's "secret". However the Private Key that corresponds to the Public Key is held in strict secrecy by the owner of the Public Key / Private Key pair. In fact, the Private Key is so secret that not even the user that owns it is likely to know what her Private Key is. It's generally a long string of seemingly random numbers stored in a secure repository within the operating system.

Public Key Algorithms are "Asymmetric Encryption" algorithms. The two keys are different, therefore it's "asymmetric". No shared secrets are required. But remember I said that all 802.11 encryption is done using "Symmetric Encryption" which requires sharing a secret. So why not use Asymmetric Algorithms for 802.11 encryption? Well, it turns out that Symmetric Algorithms are much faster and much more efficient that Public Key Algorithms. So in a perfect world what we'd like to do is use Public Key algorithms to authenticate parties to each other and to automatically generate really hard to guess shared secrets for the Symmetric Algorithms that perform 802.11 encryption. To do this requires the use of something called a "Digital Certificate".

Digital Certificates are a way of binding your user or computer's identity to your Public Key. There's nothing secret about a digital certificate. So your certificate can be sent to another party (called the "relying party" in crypto-speak). The other party now has your Public Key to use to send you an encrypted message. The other party can generate a random number, encrypt it and send it to you, then through a clever trick of cryptography both you and the other party can use that random number (and some other stuff you both know) to generate a shared secret key for the session's Symmetric Encryption. Note that every time you initiate a new session, a new random number will be used and a new shared secret will be generated. This makes it much harder to figure out the shared secret. Plus the shared secret that gets generated ends up being a really huge string that looks like random output. They are really hard to crack, unlike manually generated shared secrets which often end up being composed of words that are found in a dictionary. So Public Keys are useful for generating pseudo-random session keys. This is what we call "dynamic keying".

The clever trick I referred to above is to use one or more "hashing algorithms" which take a given input and produce a small "digest" output of it. Think of taking all the text from this paragraph and generating output of a dozen seemingly random characters from it. This digest is (in a perfect world) unique to the input, in other words a given input will always produce the same digest output. It is also computationally infeasible to determine the input from the digest output. Because of this, these hash functions are called "one-way hashes". So if each of the two parties that want to exchange encrypted messages takes a given input (a random number and other known stuff) and runs it through the same hash algorithm they will both end up with the same output. They can then take that output and run it through another round of the hash algorithm and they will both end up with new identical output. Crypto algorithms use hashes extensively to compute symmetric keying material. You can think of automatically generated symmetric encryption keying material as being a hash of a hash of a hash of a hash… (etc. ad nauseam). Examples of common hashing algorithms are MD5 and SHA.

We've seen that Public Key Algorithms can be used to encrypt messages but we haven't talked about how to use them for authentication. What if you want to prove to someone that you really are who you say you are and that the message you sent is authentic and un-altered? Turns out that you can hash your message (a packet for example), then encrypt the hash output using your Private Key. This results in a "digital signature" of the message. You tack your digital signature on to the end of the message. A party that receives your message can strip off the digital signature then perform the same hash on the remaining message. The party then decrypts the digital signature using your Public Key. If the decrypted signature and the hash of the message are the same, then the message is unaltered and definitely comes from the Private Key holder. But it's not quite enough to completely authenticate that it came from you and not a bad guy masquerading as you.

It is possible that someone might forge a certificate. In other words a bad guy could pretend to be you by creating a certificate that looks like it should be yours then attempting to use it to masquerade as you. Since the bad guy will have created his own Public Key / Private Key pair, the digital signature checking described above won't be that helpful. It will prove that the bad guy holds the proper Private Key that corresponds to the certificate he used, and it will prove that any message the bad guy sends you is unaltered but it won't prove who created the certificate and hence who the sender really is.

To insure that a certificate isn't forged, you need a way to validate it. For this you need the services of a Certificate Authority (CA). The CA is a trusted third party, usually a secure server in the Enterprise data center. The CA is charged with "signing" all user and server certificates. Like the digital signature example above, the CA signs a certificate by running a hash of the certificate contents, encrypting the hash with its Private Key (which creates a digital signature), then appending this signature to the end of the certificate. The CA also has a certificate of its own called a "CA cert". The CA cert contains the CA's Public key and can be freely distributed to all parties. Anyone that has the CA cert can then validate an incoming certificate to make sure that it really is from the claimed owner of it.

So in the case that a bad guy tries to forge a certificate and pretend to be you, he will always fail if all relying parties have the CA cert and check the CA signature of the bad guys cert. Since the bad guy can't forge a CA signature (he has no way of knowing the CA's Private Key), no one with a CA cert will believe that the bad guys certificate is valid.

Using a Certificate Authority to sign user and server certificates and validate them is the simplest form of Public Key Infrastructure (PKI). A certificate chain includes information about the CA or CA's involved in issuing certificates and the user or server certificates themselves. There are other PKI functions such as certificate revocation lists, cross-certification and certificate chaining that are used but are beyond the scope of a Security Primer.

Avaya does not require use of a Public Key Infrastructure (PKI), certificates can be "self-signed" instead. But it is a best practice to use a simple PKI infrastructure so that all certs can be CA-signed. It is expected that the use of digital certificates will be Enterprise Intranet-wide, not Internet wide. So there is no requirement for an external Certificate Authority (such as Verisign, Thawte, Entrust, etc.) instead you can use an in-house Enterprise Certificate Authority. This can be based on an existing Windows 2000 Server.

If you've enrolled in a certificate authority, you need a CA cert. So if you've received a certificate through a CSR process or a PKCS #12 object, then you need a CA certificate on the MX or RADIUS server too.

When the MX needs to communicate with WMS, Web View or an 802.1X wireless client, it asks the local MX certificate and key store for a private key to work with. If there is no private key available, the MX will not speak. If it gets a private key from the store then it asks for a corresponding certificate. If it is a self-signed certificate, then the MX will speak. If the certificate is not self-signed, then the MX looks for a CA certificate to validate the server certificate with. If there is no corresponding CA certificate, then the MX will not speak. If there is a corresponding CA certificate and the server certificate is validated (date still valid, signature okay), then the MX will speak.

Public-Key Cryptography Standards (PKCS)

When installing certificates on the MX, or when using WMS to install certificates on the MX, it is important to know about a few standards used for certificate management. The PKCS documents define various syntax and cryptography standards. They are sometimes called PKCS files or objects. This is really dry stuff. See the RSA Web site or the IETF RFC repository for the excruciating details.

From the point of view of managing the MSS, all you really need to know is that there are 3 commonly used PKCS file formats. They are used to manage the creation and distribution of certificates and keying material. They are usually stored in base64-encoded ASCII text files. They typically have extension like .pfx, .der or .cer. Descriptions of each follow.

PKCS #7

This is the Cryptographic Message Syntax Standard. This is a fancy way of saying that it is the format used for transporting a CA-signed certificate. If you generate a public key / private key pair and a certificate signing request from the MX, then the output returned to you from the CA that signs the request will be PKCS #7 output to be cut and pasted to the MX CLI.

PKCS #10

This is the Certificate Request Syntax Standard. This format is used to encapsulate certificate signing requests. When you generate a public key / private key pair and a certificate signing request on the MX, the blob that is output from the MX is in PKCS #10 format. You then either copy that blob to a file, or cut and paste the blob into the certificate signing request application on the CA.

PKCS #12

This is the Personal Information Exchange Syntax Standard. This format is used to encapsulate a CA-signed server certificate, CA certificate or chain and a public key / private key pair. A CA generates this file, then the MX imports it either through CLI or through WMS.

How RADIUS integrates with an existing Network Directory

It is common that Enterprises have existing directories used to authenticate users to network servers or applications. Prime examples of this are the Microsoft Active Directory, the Microsoft NT domain user database, and the Sun One (Netscape) LDAP Directory Server. It is common to want to leverage these databases for RADIUS authentication in order to avoid having to manage two different databases of users.

Most RADIUS servers have plug-ins that allow authentication to the directory back-end. It is common that the directory will contain both users and groups with users being assigned to one or more group. Most RADIUS servers allow for a mapping between the directory group and a profile on the RADIUS server, therefore the RADIUS server does not have to keep track of individual users. Instead, the RADIUS server passes user authentication credentials to the back-end directory, receives the group(s) to which the user belongs from the back-end directory and uses the group(s) to map to a specific RADIUS profile. The RADIUS profile contains the authorization attributes for that group, things like VLAN-Name, Mobility-Profile, Session-Timeout, Filter-ID, etc.

It is also possible to populate the users in the RADIUS local directory, but this is generally not recommended if you already have an existing network directory.

How RADIUS integrates with a Token Card back-end

Some enterprises use token card solutions for authentication. Examples of these solutions are the RSA/SecureID token card or the Secure Computing/Safeword Server. This form of authentication requires a token be sent from the client to the server in plain text (unencrypted). A token is simply a string of digits that the token card generates for the user. So for token card authentication to work in an 802.1X wireless LAN environment, and authentication protocol which can transport the token must be used. The token card itself comes in several forms, it may be a credit-card sized token generator, a small key fob, or a software shim that installs on the PC.

The two common choices today are Funk software's TTLS and Cisco's PEAP with EAP-GTC as the inner authentication method. EAP-GTC is the "Generic Token Card" EAP method and also allows both the username and password to be sent in the clear. It also allows for a conversation between the token card server and the client which may be needed for certain token card operations such as a re-sync. Today each of these mechanisms requires 3rd party client software.

Integrating a token card infrastructure requires the RADIUS server to be a client of the back-end token card server. For example, with the RSA SecureID solution there is a token card server called the Ace server. It is responsible for user authentication. Funk's Steel-Belted RADIUS has an optional module that enables it to proxy authentication to the Ace server. So Steel-Belted RADIUS becomes a client of the Ace server. This is transparent to the Avaya infrastructure.

All that is required is that the end-user client support an EAP method that supports the transport of token and that the RADIUS server can proxy the authentication requests to the back-end token card server.

How RADIUS integrates with a Smart Card Infrastructure

An optional client authentication technique is to use smart cards such as the Schlumberger CryptoFlex cards. Smart cards have their own certificate and key store. The smart card houses it's own key pair, client certificate and CA certificate. From Avaya's point of view, it is treated exactly the same way that EAP-TLS clients are treated. Smart cards are transparent to the MSS. See the instructions below for EAP-TLS, the same configuration is used for smart cards.

Elements of Identity-Based Networking

802.1X and EAP

802.1X is the primary IEEE standard for port-based network access control. The 802.1X standard, which is based largely on the IETF's Extensible Authentication Protocol (EAP), provides an authentication framework that supports a variety of protocols for authenticating and authorizing network access for wired or wireless users. These protocols are also known as EAP types or EAP methods. 802.1X defines EAPoL (EAP over LANs). EAPoL is spoken between the client (802.1X supplicant) and the MX (802.1X authenticator). EAP processing may also be performed on the RADIUS server in which case EAP-Messages are sent between the MX (802.1X authenticator) and the RADIUS server (802.1X authentication server).

All EAP types which are supported for wireless require a digital certificate on the MX. This is so the client can authenticate the MX and open and encrypted channel to the MX before any further authentication or authorization is performed. Some EAP types require certificates on both the client and the MX. In this case mutual authentication is done using the digital signatures in the certificates. A digital certificate on the MX also provides the necessary means to establish the dynamic encryption keying material used by WEP, TKIP or AES. Digital certificates may be installed in the MX through the MSS CLI or in some cases through WMS.

The EAP types supported by Avaya are EAP-MD5, EAP-TLS, PEAP and TTLS.

EAP-MD5 is for authentication only, it has no way to derive keying material, therefore it is only useful for wired-authenticated ports.

EAP-TLS requires client and server (MX or RADIUS) certificates. Transport Layer Security (TLS) is used to mutually authenticate and generate encryption keying material. This makes EAP-TLS the most secure option, but also the most difficult to administer. Since client certificates are required, a user can only use a device that has the user's certificate on it. This means that in an office environment, you cannot just borrow a co-workers PC or PDA and log into the wireless network with it. You would have to get your certificate onto it before you could use it. For some this would be a drawback, for others a security feature.

PEAP and TTLS are two-stage protocols. There is an inner authentication and an outer authentication. The first step is to perform outer authentication at which time the client authenticates the server (MX or RADIUS server) and establishes an encrypted tunnel to it. Once this encrypted tunnel exists, a secure inner authentication may be performed. Inner authentication protocols include MS-CHAP, PAP, TLS and others. Finally keying material is generated. Note that PEAP comes in two different flavors today, the Microsoft flavor which ships with Windows 2000 and Windows XP, and the Cisco flavor which requires Cisco client software. These two flavors are subtly different and can create interoperability problems. Needless to say, it is best to use what ships with the Microsoft client!

PEAP-MSCHAPv2 and TTLS with non-TLS inner authentication methods do not require client certificates. They only require server (MX or RADIUS) certificates. This means that in an office environment anyone can use anyone else's PC and login to the wireless network (assuming they know the correct password). The Wireless Security Switch is an 802.1X Authenticator and may optionally act as an 802.1X Authentication Server.

Web-based AAA

Web-based AAA stands for Authentication, Authorization and Accounting.

Authentication provides user identification and assurance that they are who they say they are, for example by user-name/password checking or a user challenge/response mechanism. Authentication may also be performed based on a client attribute such as the MAC addresses of the wireless client or a digital certificate bound to the client.

Authorization provides access control, for example VLAN membership, the area in which a user can roam, per-user access lists and timeout enforcement.

Accounting provides for collection and delivery of information used for billing, auditing, and reporting, for example user identities, connection start and stop times, the number of packets sent, and the number of bytes sent.

Authentication, Authorization and Accounting generally happen in order. A user must first be authenticated before she can be authorized to access the network. The Web-based AAA infrastructure on the MX has the ability to call on external mechanisms for Web-based AAA support. The mechanisms are called "Web-based AAA methods". In Avaya’s first release, the Web-based AAA methods supported are RADIUS and the MX's local database.