White Paper |

Improve Air Quality by Minimizing SSIDs: Using Role-Based Access to Increase Wi-Fi Application Performance

October 2010

ARUBA® networks

Table of Contents

Improve Air Quality by Minimizing SSIDs: Using Role-Based Access to Increase Wi-Fi Application Performance

1 Abstract

The role of 802.11 networks has shifted from what was once a nice to have to what is now deemed a mission critical component of an organization. The unprecedented growth in Wi-Fi devices, real-time applications and users connecting to the network introduces new security and quality assurance challenges.

The traditional approach is to use Virtual Access Points (VAPs) to deliver application QoS, wireless and access security to address these problems. Virtual Access Points are logical AP instances on the same physical access point that cater to the unique requirements of various user groups and encryption types. While it is easy to add SSIDs and VAPs as new services are provided on the WLAN, we will show that it can severely restrict the bandwidth available to users, having an adverse impact on application performance, and thus the number of VAPs configured on a WLAN should be minimized.

This document discusses the effects of multiple VAPs on the bandwidth of a WLAN network and makes recommendations for optimal deployments.

2 Introduction

A Virtual AP (VAP) is a logical entity that resides within a physical Access Point (AP). To a client, the VAP appears as an independent access point with its own unique SSID.

There are multiple approaches to implementing Virtual APs. One implementation uses a single BSSID and advertises all the SSIDs supported by the system on the same beacon. Some of the issues with this approach are:

- Incompatible with most 802.11 clients deployed
- Does not support different capability sets for each SSID

The de-facto industry standard is to use multiple BSSIDs, so each VAP has its own SSID and BSSID. Multiple beacons are used to advertise the SSIDs corresponding to the virtual APs configured. This solution however results in an increase in management traffic, which degrades in air quality and airtime availability. The remainder of this document discusses this architecture.

We will use the term 'Virtual AP' or 'VAP' to describe a BSSID/SSID set advertised by a physical access point.

3 What is a VAP?

Every VAP appears as an independent AP to the client. The VAPs emulate the operations of a physical AP at the MAC level. All wireless management traffic that would be transmitted by one physical AP is also transmitted by the VAP. For example, a single physical AP might broadcast 3 SSIDs (using virtual APs). This AP would also transmit the management traffic of 3 independent APs, one for each supported VAP.

4 Why do I need a VAP?

Virtual APs or VAPs are useful for several reasons. Some of the design considerations are for security reasons, others for application QoS, and some for ease of use. Listed below are some of the reasons why VAPs may be configured on a WLAN network –

- QoS – QoS is one of the primary drivers in a WiFi SSID design. Different devices have different QoS needs like voice, video, data, medical devices, etc. In a perfect world the devices tag the traffic with the right QoS levels. But in the real world, some devices don't. The tagging does not reflect the real QoS requirements and in this case, network architects often choose to ensure the right QoS levels by segregating the traffic completely using SSIDs and to correctly map them to the appropriate VLAN.

- Wireless Security – WiFi supports multiple encryption methods – open, WEP, 802.1x with WEP, WPA-PSK, WPA-Enterprise, WPA2-PSK and WPA2-Enterprise. Most devices now support WPA2 encryption mechanisms, but there are older applications and devices in production that only support lower secure encryption methods and in some cases no encryption method at all.

  Guest access. From a pure manageability perspective most companies prefer leaving their Guest SSIDs open.
  Some older WiFi scanners and voice clients still support only WEP or WPA/WPA2 PSK.

  Some network managers prefer not to mix different levels of encryption on a single SSID to ensure the sanctity of the encrypted traffic. This requires a different SSID for the guest users, and for dynamic key exchange encryption like WPA/WPA2 enterprise et cetera.

- Access Security – Some vendors' architectures must map each VAP to a VLAN and hence associate the VAPs to a VLAN. Users associating to the VAPs are placed onto the VLAN and policies applied on the VLANs are in turn applied on the users. This results in VLANs being used as more than just a broadcast containment mechanism, but also to secure traffic. Aruba does not recommend or require this architecture, but it is a common deployment strategy for other vendors' WLANs.

- WiFi settings – This issue is not a common one but it does exist. Although Wi-Fi Alliance certifications ensure a base level of interoperability, vendors' implementations of details of 802.11 standards can differ. This is occasionally significant for specific combinations of infrastructure and client, and can require certain counters and flags to be set or reset to ensure device interoperability. If the default wireless settings do not work for a device, setting up a separate VAP with adjusted settings can offer a solution.

5 SSID Explosion!!!

Adding SSIDs to an existing network is easy. Most WLANs support up to 64 VAPs per SSID and without proper design guidelines and/or tools to enable the designs, it is deceptively easy to add an SSID every time a new requirement or application is added to the network.

The following is an example of some typical network requirements:

- Dynamic WPA-2 encrypted secure network for the employees with different levels of access for sales, engineering and marketing
- Guest access for visitors
- Access for contractors
- Voice handsets
- Wireless tablet devices

- Special QoS requirements – WiFi pagers, medical devices,
- Barcode Scanners
- Legacy devices – old scanners or WiFi devices

5.1 Effects of VAPs on the channel bandwidth

Before adding SSIDs, consider the following limitations:

- The actual bandwidth supported by an 802.11 AP is constant and limited (11 Mbps for the 802.11b, 54 Mbps for 802.11g, 54 Mbps for 802.11a and upwards of 300Mbps for 802.11n) independent of the number of the VAPs
- Each VAP uses a significant fraction of this bandwidth for its management traffic. The bandwidth depends on the number of clients the AP can hear
- Since the bandwidth available per 802.11 channel is fixed and the bandwidth required for management traffic depends primarily on the number of VAPs the AP advertises, utilization of multiple VAPs results in a proportional decrease in the AP's net data capacity

This is further explained using the example below.

```
Net = Net bandwidth available on an AP (54 Mbps for a 802.11a/g AP and 300
Mbps for a 802.11n AP with MCS15)
Mgmt = Net bandwidth required per AP per SSID (per virtual AP)
VAP = Number of virtual APs configured
Data = Net data throughput available for data traffic
Data = Net – (Mgmt * VAP)
```

Thus the data capacity decreases as the number of VAPs advertised per AP increases. Large numbers of VAP definitions can result in very low data throughput especially in the 2.4GHz band. Depending on the client bandwidth requirements, this can also result in air traffic congestion, and higher noise floor, which impact application performance.

The net data throughput of an AP is also affected by the number of neighboring VAPs that the AP can hear on the same channel. These neighboring VAPs could belong to other coverage areas or WLAN systems but as long as they are in the receive range of the APs/clients, the data throughput of the cell (AP and its clients) will be affected.

Side Note: Other factors that exacerbate the effects of multiple virtual APs on the bandwidth include:

- *The RF band under consideration – 2.4GHz v/s the 5GHz bands*
- *The client distribution, # of 802.11b capable devices, 802.11g capable device, 802.11a and 802.11n capable devices*
- *The standards being used – 802.11b, 802.11a, 802.11n*
- *The number of clients*
- *The AP density*
- *The Tx Power of the APs*

5.2 Effects of virtual APs on 2.4GHz deployments

Consider the following 2.4GHz deployment scenario. The maximum number of non-overlapping channels available, in most countries, is 3 - channels 1, 6, 11.

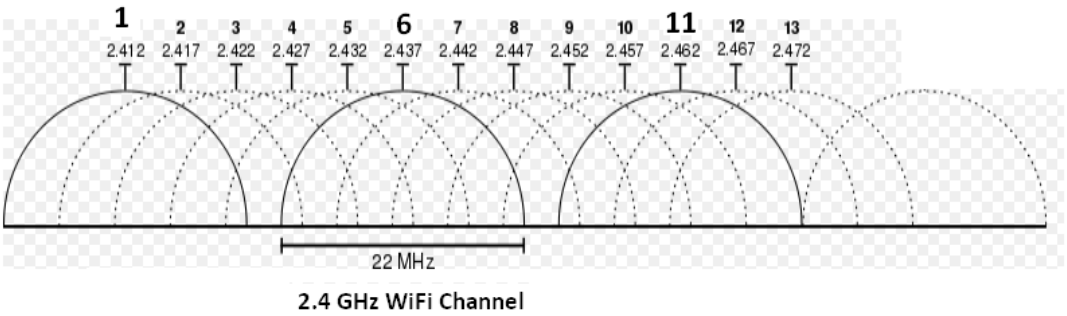

Figure 1: 2.4GHz Channel distribution

An ideal deployment for data and real time applications recommends placing APs at a distance of 30-75 feet from each other. In such a deployment, any 802.11 b/g client can hear at least 3 other APs on the same channel.

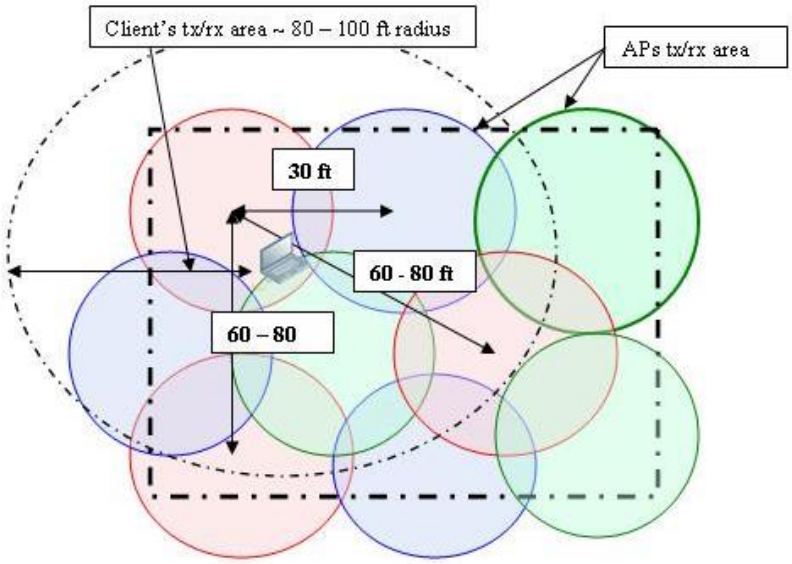

Figure 2: 2.4GHz AP deployment in a given coverage area. Channels used are 1,6,11 (color coded).

This is because the coverage area for an 2.4GHz environment is about 300 feet at maximum power, the 2.4GHz client's or AP's packets can be heard over 300 feet at lower traffic rates. In addition most clients also transmit at the highest tx-power levels resulting in large coverage areas.

*Larger coverage area **does not** equate to better performance. It is normally recommended to keep the coverage areas as small as needed to ensure less interference and collision in the neighboring cells*

In deployments of this size, there could be anywhere from 10 to 100 clients in the 2.4 GHz band. The following graph is based on the bandwidth calculations for the 802.11 management traffic for different number of stations and different numbers of virtual APs (per physical AP) when there are at least 3 APs in the receive range of each client. These calculations are based on beacon, probe request and probe response traffic alone. (Refer to the calculations in Appendix A).

Beacon and Probe Request/ Response traffic with the increase in BSSIDs and clients

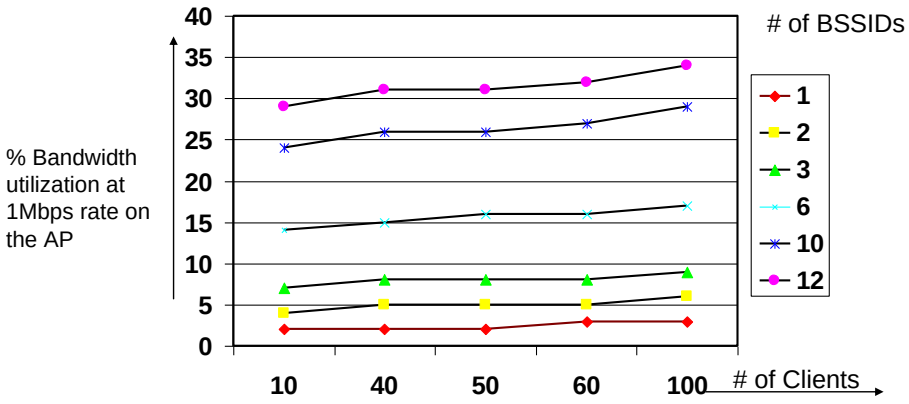

Figure 3: Effects of multiple BSSIDs on a 2.4GHz environment

From the graph, with 12 virtual APs (BSSIDs), 100 clients in a given coverage area and with 3 APs on the same channel, the management traffic takes almost 35% of the total bandwidth available to the AP using the default data rate of 1 Mbps for management frames. Beacons, probe requests and probe responses are transmitted at the 1Mbps data rate as per 802.11b/g standards. This is also true for 802.11n management traffic in the 2.4 GHz band. The bandwidth utilization for management traffic however is still well below 10% when the number of SSIDs is less than or equal to 3.

The effects are more pronounced in real world deployments in buildings with multiple floors and signals from neighboring office buildings bleeding into the coverage area. A client would now hear other APs on the same channel from neighboring WLAN deployments in addition to the APs on its own valid WLAN network. As a result the client could hear as many as 8 – 12 APs at any time. The bandwidth utilization for control traffic increases to 55% for 12 SSIDs at a data rate of 1Mbps assuming that there are at least 6 APs in the RF vicinity of each client.

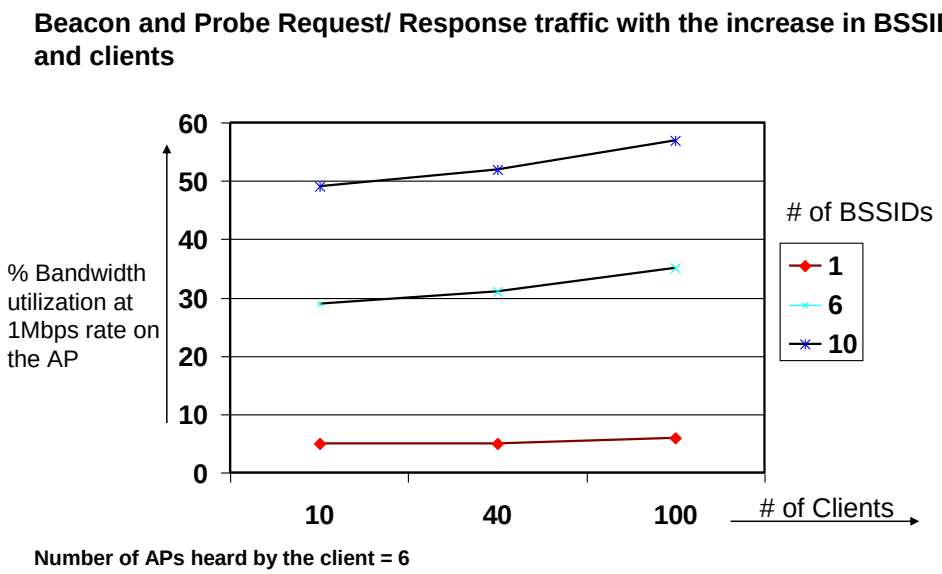

Figure 4: Effects of multiple BSSIDs in a 2.4GHz multi-floor environment

5.3 Effects of virtual APs on a 5GHz network

From the previous section it can be seen that multiple BSSIDs have a pronounced effect on any 2.4GHz channel. This is largely attributed to the fact that the 2.4GHz band offers limited non-overlapping RF channels while the coverage area for the 2.4GHz band is relatively large (around 300 feet).

This problem is not unique to the 2.4GHz band. On a per channel basis, the increase in the number of VAPs in the 5GHz band affects the channel throughput similarly. The 5GHz band offers some relief in that the number of non-overlapping channels is approximately 24. This greatly reduces the number of APs that can be heard by a client when used in pure 802.11a/n mode but with the following caveats:

- Not all of these channels are indoor channels
- Some of these channels require DFS
- With 802.11n and channel bonding, the number of available channels is halved

Note: *With channel bonding, the effects of VAP on the 5GHz spectrum would be quite similar to the 2.4GHz case as the number of independent channels would be limited to 4 or 11.*

Consider the case where there are 8 available 5GHz channels in use and no channel bonding:

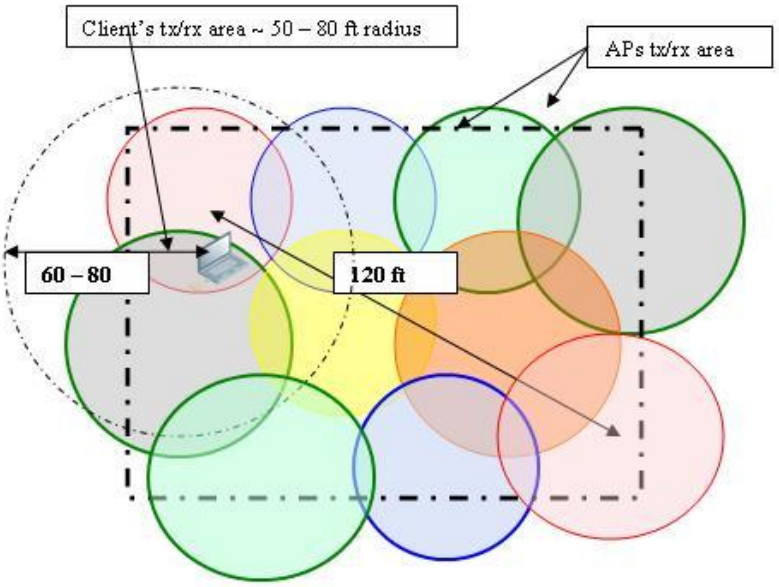

Figure 5: 802.11a AP deployment.

Note that for the same coverage area (shown in Figure 1) an 802.11a/n deployment can accommodate more APs on different channels than 802.11b/g/n, greatly reducing the possibility of the traffic from cells on the same channel bleeding over.

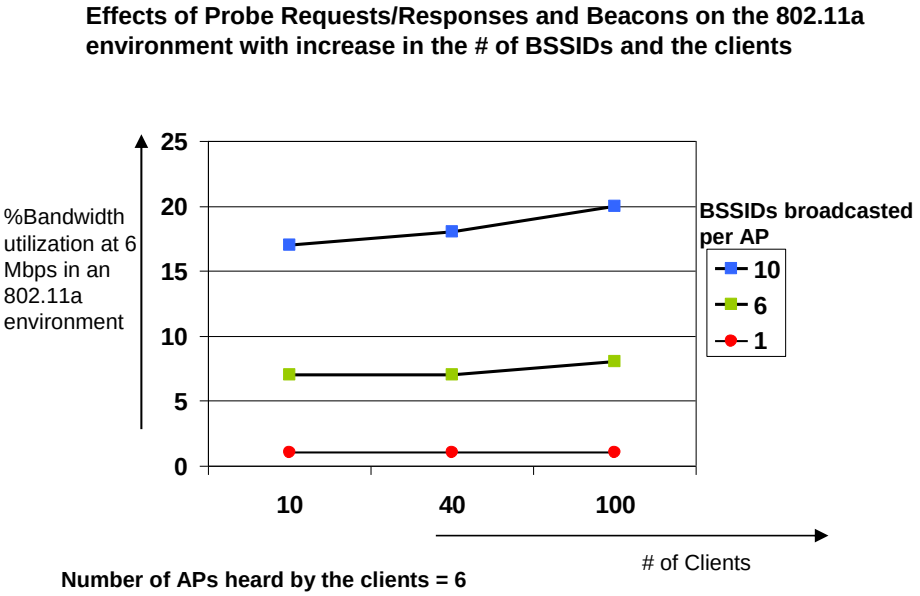

Figure 6: Effects of multiple BSSIDs in a 5 GHz environment

The loss in bandwidth may not be as profound as when using the 2.4GHz band or the 802.11n 40 MHz scenario (with channel bonding) but there is still a significant decrease in the net available bandwidth for data traffic.

6 When to use Virtual Access Points (VAPs)

A single SSID is sufficient to provide basic connectivity. A WLAN deployment, however, is seldom that simple. WLAN deployments are commonly required to support different types of devices from multiple vendors. In such cases:

- The devices support different authentication and encryption methods.
- Depending on the level of encryption supported, the devices have to be access-restricted to protect the integrity of the network and other wireless users.
- Different devices have different network access requirements. A WiFi phone needs limited access to the call servers and other phones whereas for a data device like a laptop the access required depends on the access privileges of the host using the system. The network access needs to be restricted to prevent excess access privileges.

6.1 Solutions

Two possible solutions that address these requirements are discussed below. While both solutions use VAPs to address the requirements of a heterogeneous network, the methodology used largely influences the security aspect of the solution and the number of VAPs defined.

Solution 1: Using a unique VAP for each device class and user class. In this case a unique VAP and SSID are defined for each encryption method and for each user class based on access privileges. Each of these SSIDs could optionally map to a unique VLAN on the wired network to restrict network access based on VLANs. The inherent problems with this solution are that:

- It requires the definition of too many VAPs and VLANs
- Each VAP definition increases the management traffic eventually choking the WLAN with management traffic
- Security is now enforced based on the VAP association and not the identity of the user accessing the network

Solution 2: Using VAPs for basic service separation and using firewalls to further segregate the users based on their access-privileges.

The advantage of this solution is two-fold:

- It restricts the number of VAPs defined to a bare minimum
- Since the access privileges are now based on the user/device identity and is firewall based, the network is secured from malicious attacks

This solution however requires the firewall capabilities to be integrated with the WLAN system.

Example:

Network Requirements:

- Support for visitors with no encryption enforced. These users would have no access to the intranet and will be able to access the Internet alone
- Voice handsets that support only WEP encryption and require specific RF settings

- Employee access with dynamic key exchange (WPA, WPA-2), advanced authentication like 802.1x, VPN and access based on their department - Sales, Marketing, Engineering, Administration

Solution 1

SSID	Encryption	Description
Guest	Open (No encryption)	Guest users can access the Internet and have no access to the intranet. This SSID is required since it is not recommended to use the same SSID for encrypted and non-encrypted users.
Voice	WEP shared keys	Voice needs limited access to the network (access to call servers only). Voice devices need to be on a different SSID if they have different DTIM requirements.
Sales Marketing Engineering Administration	WPA, WPA2 dynamic keys and using advanced auth methods like 802.11i, 802.1x, VPN	Access to the network limited based on the SSID the user associates to.

Solution 2

SSID	Encryption	Description
Guest	Open (No encryption)	Guest users can access the Internet and have no access to the intranet. This SSID is required since it is not recommended to use the same SSID for encrypted and non-encrypted users.
Voice	WEP shared keys	Voice needs limited access to the network (access to call servers only). Voice devices need to be on a different SSID if they have different DTIM requirements.
Employee	WPA, WPA2 dynamic keys and using advanced auth methods like 802.11i, 802.1x, VPN	Access to the network is limited by the authentication profile and not by the SSID.

6.2 Analysis of the Solutions

The difference between the two solutions might seem minimal but as shown in the earlier analysis, Solution 1 can consume significantly higher bandwidth than Solution 2. Solution 1 requires the definition of multiple additional SSIDs on the network, which results in an increase in the wireless management traffic and a decrease in the actual data throughput on the network. This approach can have a negative impact application performance, especially for latency sensitive applications like voice, video and medical devices. Additionally SSIDs are used for user classification and access rights policing. Thus users are assigned access rights not by

their identities but by their SSID association which could give a malicious spoofer privileged access into the network. The solution requires Employee A in the sales department to associate with the "Sales" SSID for the right network access privileges. Associating with the "Employee" SSID could result in Employee A gaining access to a privileged set of servers not accessible to the Sales user group. This is because the rights are assigned by the SSID and not Employee A's identity or authentication profile.

Solution 2 is the Aruba recommended solution. In this case the virtual APs are defined for basic service separation based on the radio configuration and user differentiation and access privileges are granted based on the individual user's identity and authentication profile. This approach minimizes the number of SSIDs by keeping wireless management traffic at an acceptable level, which results in better air quality, more bandwidth per AP, and increased application performance.

Employee A from the Engineering department and Employee B from the Sales department would both associate with the Employee SSID but the Aruba system would assign different them access privileges based on their identity and authentication profile. This ensures that users are always assigned the right access permissions depending on their identity.

When users associate with an SSID supported a weak encryption, their rights could be further limited to protect the integrity of the network.

7 Conclusion

Virtual APs address some of the basic wireless design requirements successfully only when used judiciously. The Virtual APs should be defined for basic service separation based on the radio configuration or device capabilities and not for user classification and access policing. Advanced and more secure methods like firewall definitions should be used to ensure that user groups are assigned the right access policies depending on their encryption and/or authentication methods. Access rights should be differentiated using firewall policies, which is more scalable and secure. Virtual APs should not be used to enforce security.

Judicious use of virtual APs helps improve and secure the connection on the wireless side by the encryption method with acceptable bandwidth loss. In conclusion, Virtual Access Points should not be used as the means to secure the network or classify users by their access rights but should be used to group users by their basic service sets and RF requirements.

8 Appendix A: Calculations

This section explains the process for computing the data used in this document.

Management Traffic Type

All calculations are based on the traffic generated by the clients and the APs in terms of probe requests, probe responses and beacons.

Client behavior model

On an average a normal WLAN client sends 2 probe requests per minute per channel. One of these is a broadcast with SSID set to the broadcast EESID and the other packet is sent with the ESSID set to the required SSID. The later is a broadcast but only the APs/Virtual APs with the corresponding SSID would respond.
The assumption made is that the client is pre-configured SSID as would be the case in an enterprise network.

AP behavior

The assumption made here is that the APs are configured to respond to broadcast probe requests.

Calculations

Num_{AP} = Number of APs that can hear the client or which the client hears.
Num_{VAP} = Number of virtual APs configured per AP. The APs will have a unique BSSID for each of these virtual APs.
Num_{C} = The number of clients in a given coverage area that hear Num_{AP} APs.

Every AP sends a beacon once every 100 milliseconds.
Number of beacons per AP per SSID = $10 * Num_{AP}$
Number of beacons per AP = $10 * Num_{AP} * Num_{VAP}$
Number of beacons per AP per minute = $60 * 10 * Num_{AP} * Num_{VAP}$

PBReq = Number of broadcast probe requests for Num_{C} clients per minute =
Num_{C} PBRes = Total Number of broadcast probe responses from Num_{AP}
Number of probe responses per client = $Num_{AP} * Num_{VAP}$
PBRes = $Num_{AP} * Num_{VAP} * Num_{C}$

PUReq = Total Number of probe requests from Num_{C}
PURes = Total number of probe responses for the PUReq from the client
= $Num_{AP} * Num_{C}$

Total number of packets from the client = PBReq + PUReq
Total number of packets from AP = Probe response + Beacons
= Beacons + PBRes + PURes

$= (60 * 10 * Num_{AP} * Num_{VAP}) + (Num_{AP} * Num_{VAP} * Num_{C}) + (Num_{AP} * Num_{C})$
$= (600 * Num_{AP} * Num_{VAP}) + (Num_{AP} * Num_{C}) (Num_{VAP} + 1)$

Total packets per minute = Packets from client + Packets from AP
$= (2 * Num_{C}) + (600 * Num_{AP} * Num_{VAP}) + (Num_{AP} * Num_{C}) (Num_{VAP} + 1)$

Packets per sec = PpS

$$\frac{(2 * Num_{C}) + (600 * Num_{AP} * Num_{VAP}) + (Num_{AP} * Num_{C}) (Num_{VAP} + 1)}{60}$$

$$PpS = \frac{(2 * NumC) + (600 * NumAP * NumVAP) + (NumAP * NumC)(NumVAP + 1)}{60}$$

Time to transmit

Considerations

- Beacons and probe requests / responses are transmitted at the lowest supported rates which would be 1 Mbps for 802.11b/g/n (2.4 GHz) and 6 Mbps for an 802.11a/n (5 GHz) network
- Assumption made is that probe requests, responses and beacons are of approximately 100 Bytes (since these calculations are used to provide a rough estimate of the bandwidth consumptions)
- Long preambles overheads and ACKs for unicast packets are not considered

Rate = The minimal rate at which these packets are transmitted (1 Mbps for 802.11 b/g/n and 6 Mbps for 802.11a/n)

Time to transmit 1 bit = $(1 / 2^{20})$

Time to transmit 100 bytes = $100 * 8 * 1 / 2^{20}$

Adding DIFs (inter packet interval)

$T_{pkt} = (100 * 8 * 1/ 2^{20}) + 50$ [50 microseconds is the DIFs time]

Time to transmit PpS number of packets =

$T = PpS * T_{pkt}$ microseconds

% bandwidth utilization when the AP is transmitting at rate R

$\%B = (T / 10^{6}) * 100$

$T = PpS * T_{pkt}$ microseconds

$\%B = (T / 10^{6}) * 100$

About Aruba Networks

Aruba is the global leader in distributed enterprise networks. Its award-winning portfolio of campus, branch/teleworker, and mobile solutions simplify operations and secure access to all corporate applications and services - regardless of the user's device, location, or network. This dramatically improves productivity and lowers capital and operational costs.

Listed on the NASDAQ and Russell 2000® Index, Aruba is based in Sunnyvale, California, and has operations throughout the Americas, Europe, Middle East, and Asia Pacific regions. To learn more, visit Aruba at http://www.arubanetworks.com. For real-time news updates follow Aruba on Twitter, Facebook, or the Green Island News Blog.

1344 Crossman Ave. Sunnyvale, CA 94089-1113
Tel. 408.227.4500 | Fax. 408.227.4550 | 1-866-55-ARUBA
info@arubanetworks.com | http://www.arubanetworks.com

WP_ Virtual-Access-Points_101011