RUCKUS™
an ARRIS company

Deploying Mesh

Best Practices Design Guide

Table of Contents

Copyright Notice and Proprietary Information

Destination Control Statement

Technical data contained in this publication may be subject to the export control laws of the United States of America. Disclosure to nationals of other countries contrary to United States law is prohibited. It is the reader's responsibility to determine the applicable regulations and to comply with them.

Disclaimer

THIS DOCUMENTATION AND ALL INFORMATION CONTAINED HEREIN ("MATERIAL") IS PROVIDED FOR GENERAL INFORMATION PURPOSES ONLY. RUCKUS AND ITS LICENSORS MAKE NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THE MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR THAT THE MATERIAL IS ERROR-FREE, ACCURATE OR RELIABLE. RUCKUS RESERVES THE RIGHT TO MAKE CHANGES OR UPDATES TO THE MATERIAL AT ANY TIME.

Limitation of Liability

IN NO EVENT, SHALL RUCKUS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL OR CONSEQUENTIAL DAMAGES, OR DAMAGES FOR LOSS OF PROFITS, REVENUE, DATA OR USE, INCURRED BY YOU OR ANY THIRD PARTY, WHETHER IN AN ACTION IN CONTRACT OR TORT, ARISING FROM YOUR ACCESS TO, OR USE OF, THE MATERIAL.

Ruckus Networks | 350 West Java Drive | Sunnyvale, CA 94089 USA | T: (650) 265-4200 | F: (408) 738-2065 ruckuswireless.com

Intended Audience

This document addresses factors and concerns related to deploying mesh. Many factors can affect both the initial design and final performance. These are considered here along with recommendations for an optimal design.

This document is written for and intended for use by technical engineers with some background in Wi-Fi design and 802.11/wireless engineering principles.

Overview

Need for mesh

In most cases, wireless LANs are seen as an additional connection method for clients to gain access to wired network resources. The AP provides this accessibility because it has both a wired connection and a radio. But what happens if there is no wired network connection available? Without direct, wired network access, can we still deploy wireless and gain some connectivity?

There are two popular methods of extending wired networks: wireless bridging and mesh. A wireless bridge is a simple connection that is either a Point-to-Point or Point-to-Multipoint connection. A bridge does not service clients. A wireless bridge is also considered a single "hop", e.g. bridges cannot connect in a wireless "daisy chain" of multiple hops.

Wireless mesh is an excellent way to provide wireless coverage in areas where wired AP connections are unfeasible. If power is available, a wireless mesh can be installed. Mesh APs work because client traffic is transmitted via a wireless link with another upstream AP. Traffic travels these links until it reaches an AP that is directly connected to the wired network.

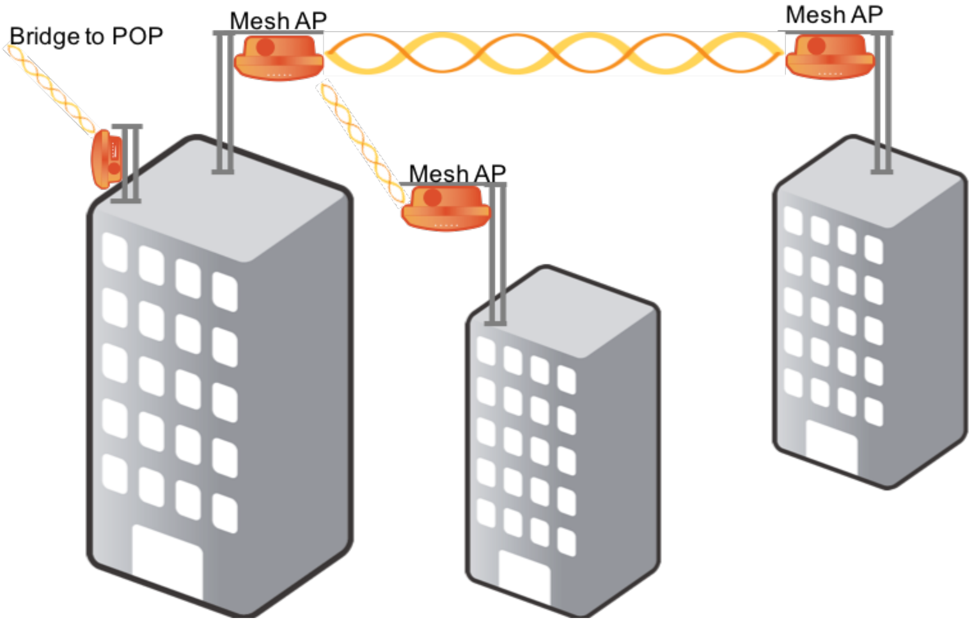

FIGURE 1: WIRELESS MESH NETWORKING

Only a few years ago, wireless mesh networks were considered an enterprise technology for spanning cityscapes. Today, wireless mesh is spanning your home—to get a good Wi-Fi signal to those gaming consoles in your basement. Although the home environment mesh is not a focus of this document, the basic principles of mesh performance would still apply.

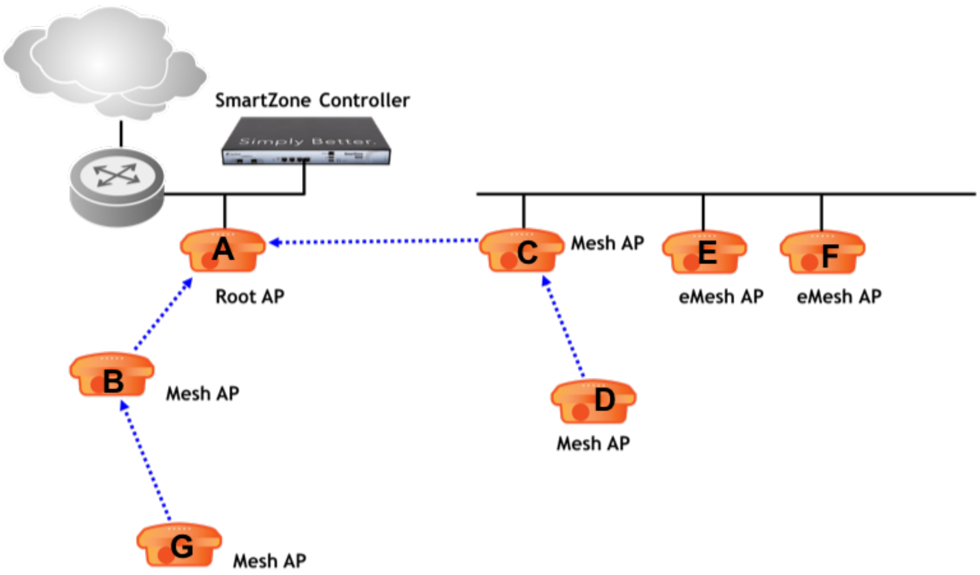

FIGURE 2: MESH ROLES

Devices within a wireless mesh network are divided into one or more of several roles. Using the reference in Figure 2, these roles are defined in the table below.

Name	Position in Figure 2	Role
Mesh Node	A, B, C, D, E, F, G	Any AP with mesh capability enabled
Root AP	A	Communicates with the SmartZone/ZoneDirector through an Ethernet interface
Mesh AP	B, C, D, G	Communicates with the SmartZone/ZoneDirector through another AP (via mesh uplink)
eMesh AP / eMAP	E, F	Mesh AP that communicates with the SmartZone/ZoneDirector via an Ethernet connection to another mesh node.
Mesh Link	A-B, A-C, B-G, C-D	Connection between 2 mesh nodes
Mesh Downlink	C to D	Client station to AP relationship
Mesh Uplink	D to C	AP to client relationship
Mesh Neighbor		Other mesh nodes that are visible to an AP

TABLE 1: MESH ROLES

Other key concepts include *mesh tree* and *hop count*. A mesh tree is the tree-like structure formed by interconnected mesh nodes. A mesh tree always has a root AP at the base of the structure. All other mesh nodes in a tree are downlinks from the root AP. The relationship between APs is either as an AP or a station. This refers to how the AP is handled by its upstream connection. A mesh AP has many characteristics of a station when it is downstream of another AP and may therefore be referred to that way.

The hop count refers to the number of links between the root AP and a specific mesh AP or client.[1] When designing a mesh network, it is recommended that no more than three hops be used. This ensures good overall performance within the mesh.

[1] The maximum hop count supported by Ruckus is 7.

A Wi-Fi access point (AP) is designed to connect wireless clients to a network. A mesh AP can also connect wireless clients. Unlike a wired AP however, a mesh AP does not have a direct connection to the wired backbone. It transports wireless client data through one of its radios to an upstream AP (with the exception of eMesh APs). The upstream AP can be another mesh AP or a wired (root) AP.

Ruckus-supported mesh APs

Almost all Ruckus APs support wireless meshing with SmartMesh intelligent mesh technology. Ruckus makes no distinction between a root AP and a mesh AP at the hardware level - any AP in a mesh can have any mesh role. A mesh node's role is determined by its connectivity and throughput compared to other mesh nodes rather than the AP model. While using different model APs in a mesh network, ensure the RF parameters like channel (indoor vs. outdoor), channel width are configured correctly for the APs to mesh successfully.

Radio Type	Compatible AP Models
802.11n	7781-CM
802.11ac wave 1	H500/R500/R600/R700/T300/T301
802.11ac wave 2	C110/H510/R510/T310/E510/M510 R610/R710/R720/T610/T710/T811-CM *M510 (support in future releases)
802.11ax	R730

TABLE 2: RUCKUS SUPPORTED MESH APS

A mesh AP can be single radio or dual-radio. In the case of dual-radio APs, in older releases the wireless mesh is always on the 5 GHz radio. But from SmartZoneOS 5 wireless mesh can now be on 2.4GHz or 5GHz radio. This is especially useful in countries like Israel where usage of 5GHz is restricted. Earlier these countries used a 2.4GHz-only model like 7352 but it is now end-of-sale. Similarly, on the ZoneDIrector APs, this support for mesh on either of the radios will be available from release 10.2.

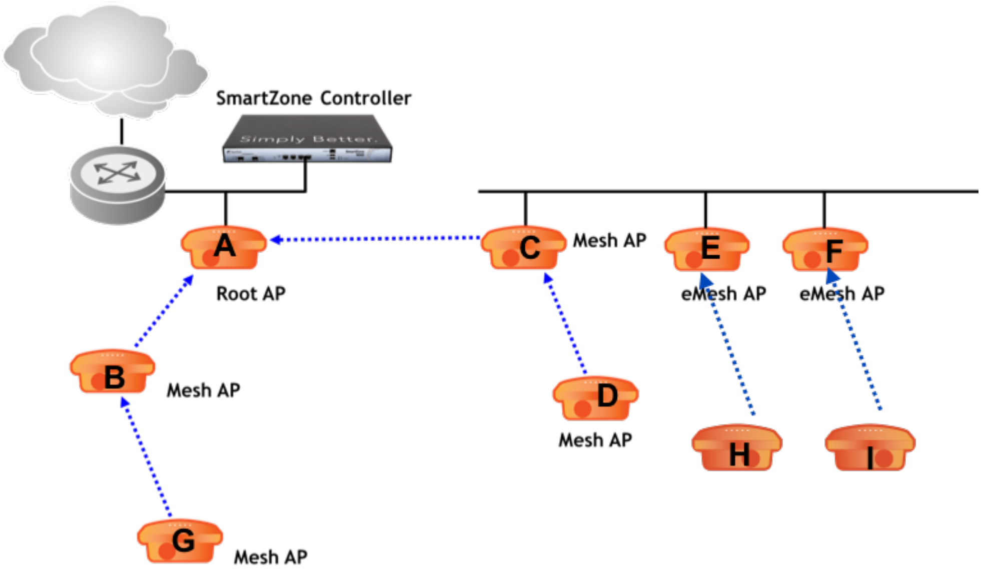

FIGURE 3: MESH TOPOLOGY

A wireless mesh must consist of APs on the same radio. For e.g. In Figure 2, all the mesh links A-B, A-C, B-G, C-D need to be either on 5GHz or 2.4GHz, you can't have A-B on 5GHz and A-C on 2.4GHz. But the mesh links that the eMAPs E and F form can be different from the A-B, A-C, B-G, C-D. E can have mesh links on 2.4GHz or 5GHz and independently F can have mesh links on either radios.

Supported Technologies

Ruckus Wireless supports the following mesh topologies:

- Standard SmartMesh
- eMesh / Hybrid mesh
- Mesh Bridging

Standard SmartMesh

The simplest configuration is a standard model that takes full advantage of automatic configuration and self-healing. With automatic SmartMesh, a mesh auto-forms to create an optimal uplink topology that is load-balanced across as many root APs as possible. In a standard configuration, each AP's mesh role is determined automatically.[2]

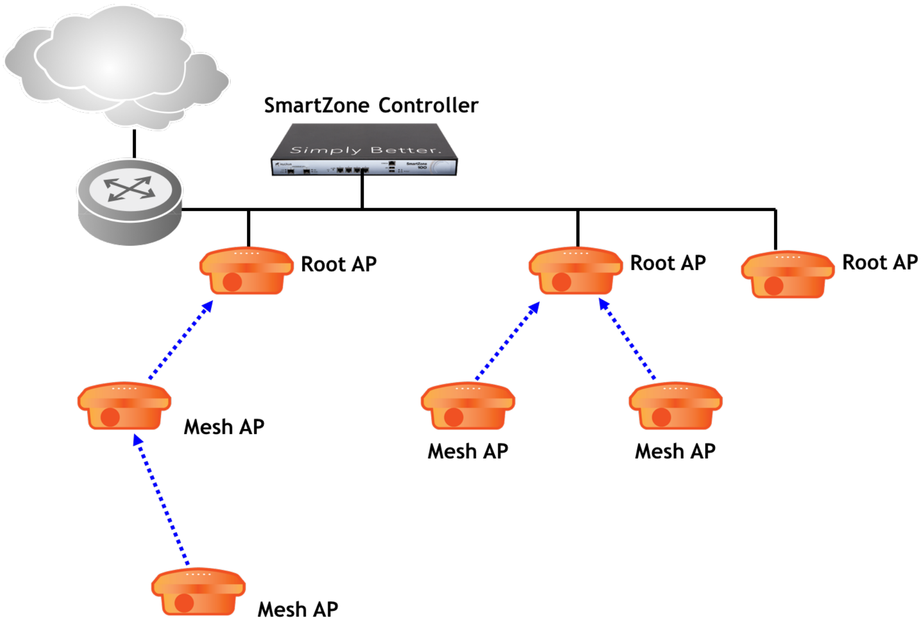

FIGURE 4: STANDARD SMARTMESH

Since a mesh role is automatically assigned, care should be taken during planning to make sure excessively long mesh hops are unlikely to occur. Mesh stability is also important: constant changes to the mesh (AP or Ethernet link goes/down up) will force topology reformation. Since the wireless network is offline during topology calculations, instability will impact overall performance and reliability.

A standard topology requires all APs use the same channel for the mesh. This allows APs to find each other in the case of mesh failures. Therefore, the most stable and highest performance channel should always be used. This can be configured manually or by SmartMesh automatically. In the first case, the mesh will not change channels in the case of excessive performance deterioration.

The standard Smart Mesh topology consists of a controller (a SmartZone 100 in Figure 4) and a number of Root APs and Mesh APs. In this topology, the controller and the upstream router are connected to the same wired LAN segment. You can extend the reach of your wireless network by forming and connecting multiple mesh trees. In this topology, all APs connected to the wired LAN are considered Root APs, and any AP not connected to the wired LAN is considered a Mesh AP.

That is, the Mesh APs do not have any Ethernet interfaces up. All network devices are connected to either the wired segment or through one of the mesh nodes.

Standard Mesh Advantages:

- Simplified design that takes advantage of self-forming and self-optimizing SmartMesh features
- Because of Ruckus' mesh algorithm, the path to the Root AP can change dynamically due to uplink node failure, signal strength changes, or the availability of a node with a better potential throughput.

Mesh Disadvantages:

- Hop count will affect latency. Assuming all mesh nodes are within recommended distances, latency will increase approximately 10-15ms per hop. Most data applications are not sensitive to minor latency increases but others, such as VoIP are very sensitive. A voice connection requires a minimum latency time from client to the core of 150ms. This will be a limiting factor in the number of hops in the mesh.

eMesh / Hybrid Mesh

Another type of network topology can be configured using the Hybrid Mesh (eMesh) concept. Ethernet-connected Mesh APs (eMesh AP or eMAP) enable the extension of wireless mesh functionality to a wired LAN segment. An eMAP uses a wired Ethernet link as its uplink rather than wireless. An eMAP is not considered a Root AP, despite the fact that it discovers the controller through its Ethernet port. Multiple eMAPs can be connected to a single Mesh AP to, for example, bridge a wired LAN segment inside a building to a wireless mesh outdoors.

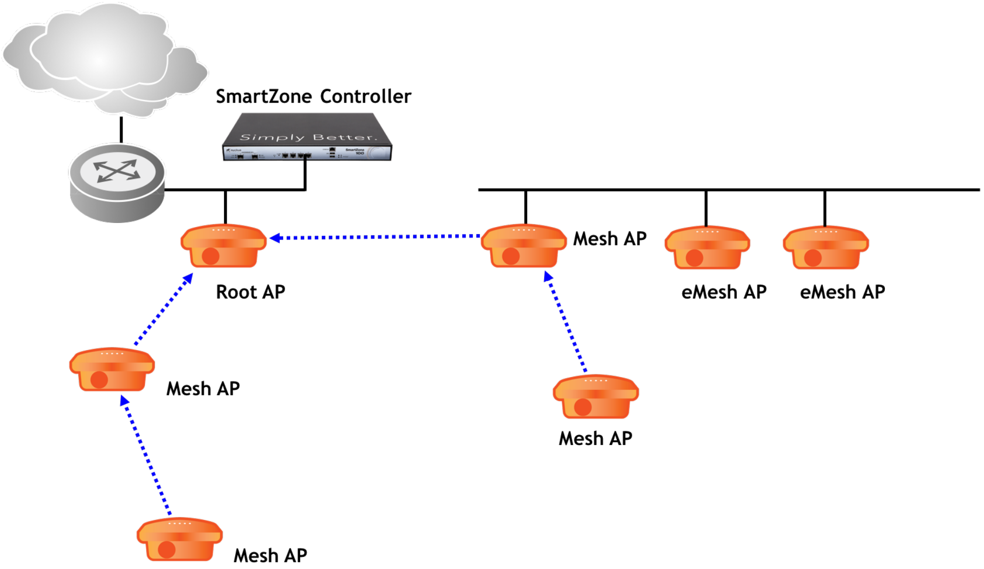

FIGURE 5: EMESH TOPOLOGY

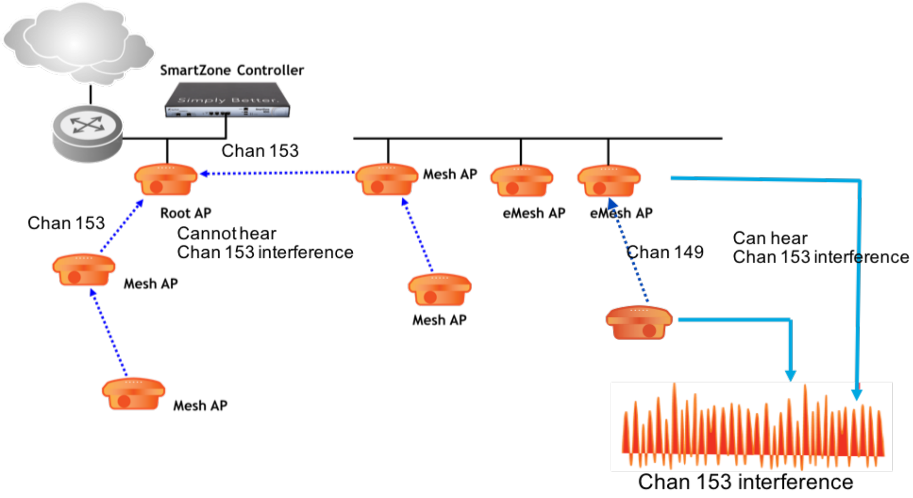

FIGURE 6: HYBRID MESH CHANNEL CHANGE

eMAP Mesh Advantages:

- In designing a mesh network, connecting an eMAP to a Mesh AP extends the Smart Mesh network without expending a wireless hop, and can be set on a different channel to take advantage of spectrum reuse.
- Very large networks are more likely to encounter RF interference on the selected mesh channel at some point. Hybrid mesh topologies allow the mesh beyond the eMAP to use a different channel. This makes the mesh more resilient and agile; particularly in cases where potential interference is expected. Note: The 7 hop maximum still applies.

Mesh Bridging

A wireless mesh can also be compared to a point-to-point or point-to-multipoint wireless bridge; both transport backbone data over wireless. A mesh network differs from a bridge in that it is not typically designed as a fixed installation – i.e. the mesh topology can change dynamically as needed. A mesh network also introduces some latency and throughput reduction after the first mesh hop. This is unlike a bridge, which is considered lossless and incurs no necessary throughput costs for a bridge link.

When to use mesh vs. bridge

Since both configurations can produce similar results, the deciding factor should be based on overall goals and future needs. The following are good rules to follow.

Use a bridge instead of wireless mesh when:

- Only need to connect two locations
- Link distance is greater than 300m-400m
- Low latency is required, e.g. voice
- Client access support is not required

Consider the use of a wireless mesh when:

- Need to connect multiple locations (point-to-multipoint, multi-point to multipoint)
- Link distances are short
- Client access is required
- LoS is difficult to achieve

In P300, bridge mesh is reused to great advantage as it brings in the flexibility of role selection, loop avoidance, uplink selection, and failover protection to point-to-point deployments.

Design Considerations Overview

Wireless Mesh Design Considerations

Every mesh network design must first start with a site survey. Only a site survey can identify the quality of the RF environment, potential sources of interference (both 802.11 and non-802.11), potential AP mounting sites, expected distances for each link, and potential Fresnel zone hazards. Without a good site survey (and follow up deployment validation), the network design will be GI/GO (Garbage In/ Garbage Out).

Any approach to mesh design should begin with an effort to wire every AP whenever possible. The rule of thumb is: *wire where you can and mesh where you must*. Only design an AP to act as a Mesh AP (MAP) when wiring it simply is not feasible. The number of 1-hop MAPs is called the fan count. The number of mesh links between a downstream MAP to the RAP is called the hop count. Combined, these form a mesh tree.

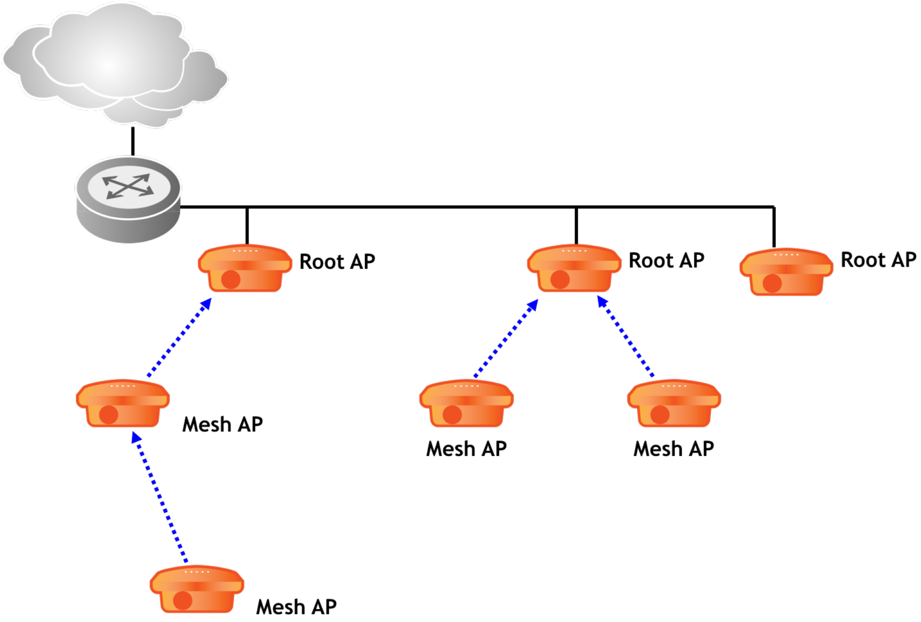

FIGURE 7: GENERIC MESH TREE EXAMPLE

The quality of the mesh network is a function of the width and depth of the mesh tree, the quality of the signal on each mesh link, and the volume of traffic being carried. The quality of the signal is affected by the distance between MAPs (or MAP and RAP), AP antenna type (omni, narrow, sector, or external antenna), and RF and other sources of interference. SNR is a key factor for mesh wireless link quality. However, traffic usage patterns drive airtime utilization at each AP, which may dramatically affect jitter and latency.

Mesh Performance

Regardless of the topology you choose, some general rules for optimal performance apply regardless of configuration differences. Care should be taken to follow these guidelines for location, mounting and hardware selection. Like all mesh networks, performance is determined in part by the number of hops. Overall mesh network performance is affected by the following:

Choosing the Right Equipment

The Ruckus solution portfolio includes a range of products geared towards specific applications. Selecting the right product is key to good design, scalability and performance. Selection depends on the design requirements, performance, and budget.

Radio Selection

One of the most important choices when designing a mesh is the AP hardware. As mentioned earlier, all mesh nodes must share the same radio type for the mesh.

If there is a lot of background RF interference, it may be a good idea to use a radio that is subject to less interference for the mesh. This is typically 5 GHz, which has more channels to choose from. Another alternative is to use a directional antenna, which does not hear as much since it is more focused. This might be enough to work around interference.

When to use 2.4 GHz	When to use 5 GHz
Little to no noise in the 2.4 GHz spectrum	Very noisy 2.4 GHz spectrum
Need to penetrate walls and/or foliage	Mostly Line of Sight (LoS)
Require longer distances	Shorter distances are achievable

TABLE 3: RADIO SELECTION CRITERIA

If a longer distance is required, a 2.4 GHz radio will give greater range. A 2.4 GHz radio is also a better choice if there are obstructions such as trees. This radio will penetrate obstacles much better than the 5 GHz radio.

Bridging with a Wireless Mesh

Some mesh topologies may require the use of wireless bridges as well as mesh nodes. Bridges are ideal for helping to span long distances within the mesh, preserving and injecting additional bandwidth into the mesh and creating more root APs, e.g. via a bridge backhaul rather than a wired connection can work the same as a direct Ethernet connection.

Mesh Link Distance

Distance between mesh nodes is a critical factor in overall performance. Signal quality and therefore throughput decrease as distance increases. A site survey is always recommended to determine the maximum distance between two mesh APs. This number takes into consideration factors such as local interference, signal quality, obstructions, etc. which are difficult to accurately model.

Although other factors may impact these numbers, if the mesh is used to offer client access, always plan for the weakest client (laptop, smartphone, etc.) The distance between two mesh nodes should not exceed two times (2x) the distance the client device can be from the AP before disconnecting. As a rough guide, consider that a handheld device (such as a phone) may be able to connect at up to 80m – 100m from the AP.

If there are no clients planned, a bridge may be a better option and is recommended.

Mesh Hops

The unique feature of a mesh network is the ability to backhaul APs through other mesh nodes. This provides greater flexibility in deployments. This flexibility does come with some considerations when planning the maximum number of mesh hops.

Tree Depth

In a purely mesh, dual-radio environment (no eMAP), Ruckus recommends no more than 1 or 2 mesh hops for high capacity or high throughput applications. For simple coverage extension and low bandwidth applications, a tree depth of up to 3 hops is acceptable provided there are good links between the mesh nodes.

Ruckus supports a maximum tree depth of 7: one root AP and up to 6 mesh APs in a single series. Realistically, few deployments should ever approach this number.

Throughput and Latency

Calculating the Root AP to Mesh AP Ratio

Generally, throughput available to a mesh node/client station can be estimated at 1/<hops-to-node> x Root radio throughput.

For example:

- For a 1st hop mesh node or a client that is connected to a Root AP, hop count to the node/client is 1. 1/1 = 100% of Root radio throughput.
- For a 2nd hop mesh node or a client that is 2 hops from a Root AP, hop count to the node/client is 2. 1/2 = 50% of Root radio throughput

For each hop through the same radio (no eMAP or client), the throughput is approximately 1/N where N is the hop number[3]. The following example assumes the connection rate is the same for each link. Links of unequal throughput will yield lower overall throughput.

[3] This is a limitation of all Wi-Fi mesh networks.

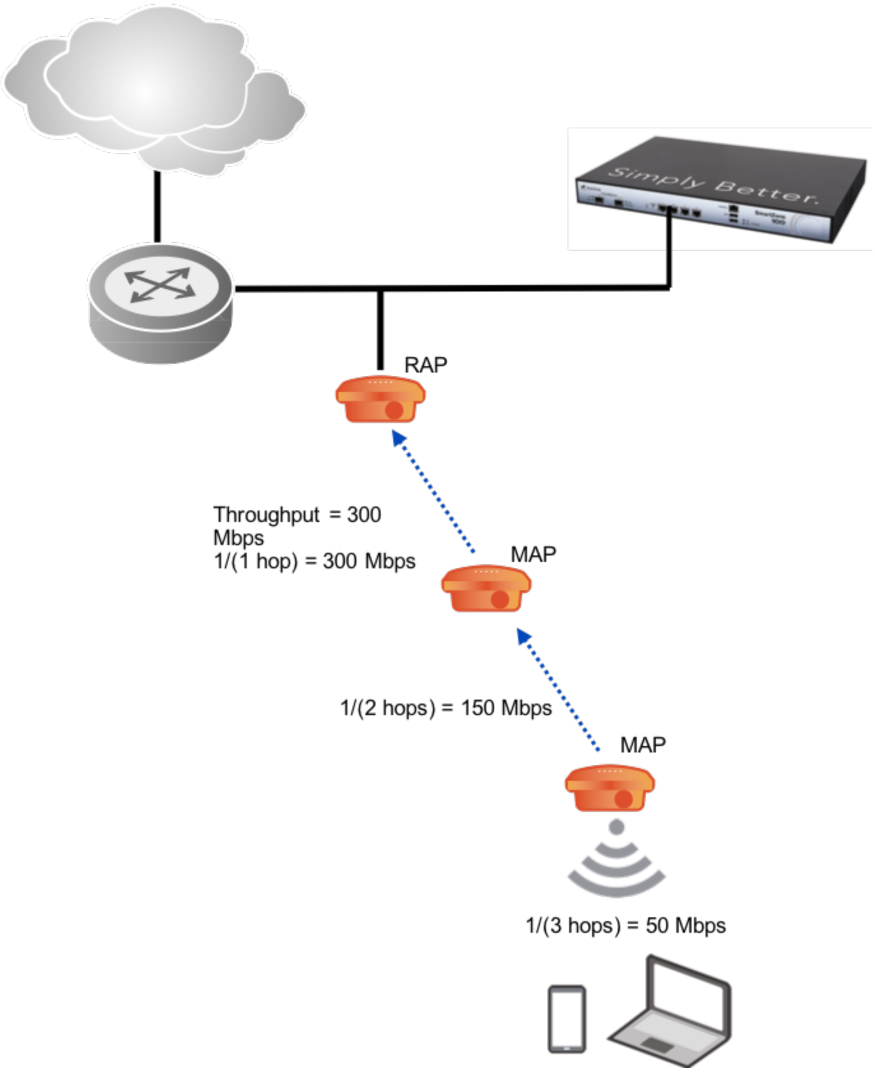

FIGURE 8: MULTI-HOP POTENTIAL THROUGHPUT EXAMPLE

Not only do hops affect throughput, they impact latency as well. Assuming all mesh nodes are within recommended distances, you can expect a latency increase of approximately 10-15ms per hop. This may or may not be a critical concern; it is largely determined by the application. For example, most data applications (web, email) are not sensitive to minor latency increases. Other applications, such as VoIP are very sensitive. A voice connection requires a maximum latency time (from the client to the core) of 150ms. So even though the voice application does not require a large amount of bandwidth, the latency requirement will constrain the ultimate number of hops within the mesh.

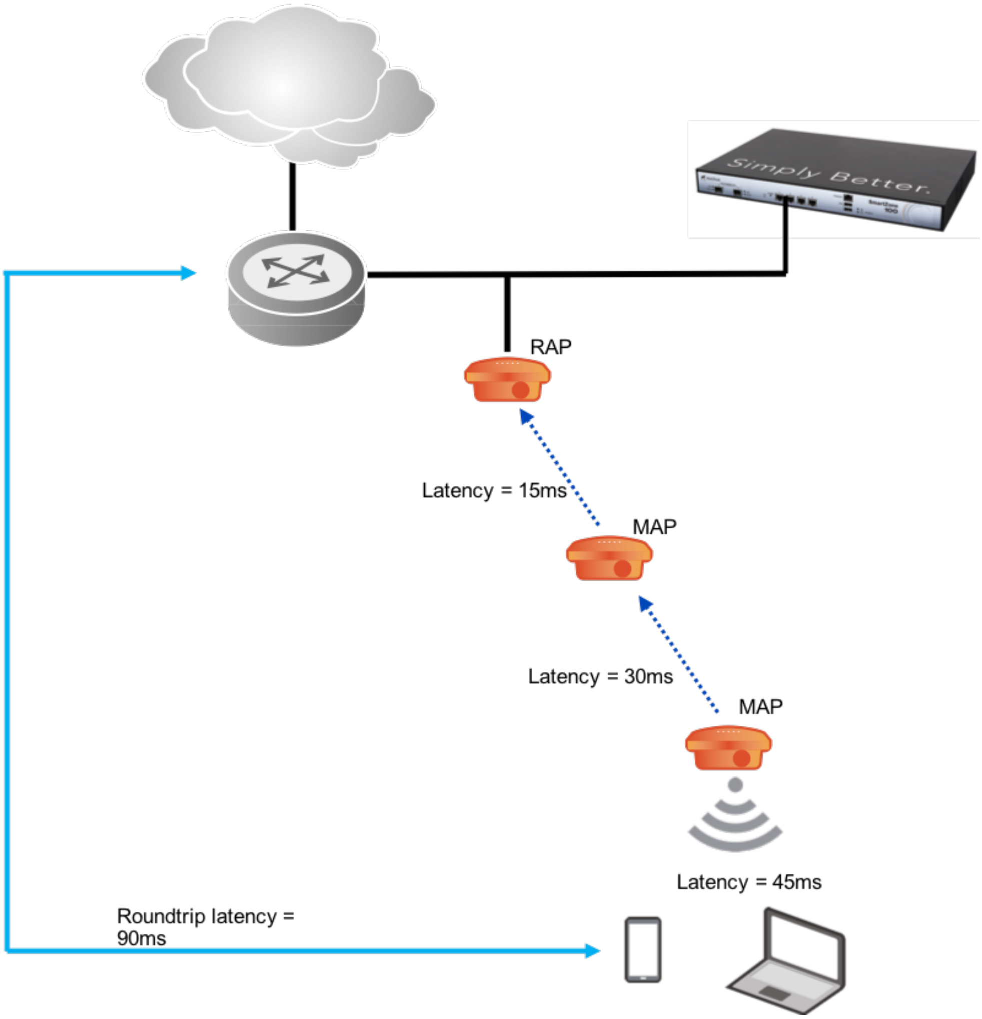

FIGURE 9: LATENCY THROUGH A MESH NETWORK

The above diagram shows latency as measured on the wireless network only. The backend network (wired, switches, routers, etc.) will also add some latency. This should be taken into consideration for latency-sensitive applications such as voice.

Packet Loss is a function of RF interference from sources such as 802.11 co-channel (CCI) and adjacent channel (ACI) devices and non-802.11 interference. As the number of hops increase, the RAP at the top of the tree may not avoid channels on which a MAP down the tree experiences RF interference. This sort of localized interference at the MAP could result in asymmetric throughput. Low throughput but high RSSI could also result due to external sources of interference.

Another key factor in determining the throughput (particularly in small cell deployments) is support for jumbo frames across the mesh or bridge link. Jumbo frame support is needed to prevent overhead due to fragmentation. All the AP models now support jumbo frames of up to 2290 bytes. This value is set by default for all mesh links.

Note: When deploying mesh on Unleashed the throughput hit on a per-hop basis would be the same. Latency and performance could be an issue depending on the amount of traffic etc. According to master election rule, the node with less mesh nodes has higher priority to be selected as Master in auto selection mode. This also implies the Master cannot take on the MAP role and has to only be a RAP. If the user wants to select a preferred Master, ensure that it has lesser MAP nodes connected to it. This behavior of the Mesh AP selection in Unleashed is intuitive as the Master performs most of the processing and is under heavy load in the network.

AP Mounting and Installation

An incorrectly aligned or positioned AP can result in signal loss and a decrease in coverage area. Since there is no reason to give up performance if possible, it's worth spending time considering how the AP is mounted as well as its orientation.

The optimal mounting orientation for an AP is dome down or facing the target coverage area. For example, if the coverage area is relatively horizontal, dome down is the preferred orientation. If coverage tends more towards vertical, mounting the AP dome facing out would give a better range. If mounting APs on a building, you may consider a small inward tilt. Most Ruckus APs now feature antenna arrays that self-optimize regardless of orientation. There may be times when client coverage does not match with optimal orientation for mesh backhaul. In these cases external antennas are recommended.

Fresnel Zones

Calculating and designing for the correct Fresnel zone is extremely important. Even an obstacle as small as the dome covering a light on a light pole can impinge on the Fresnel zone and impact performance. To avoid these kinds of issues, always plan for some clearance; when mounting to the side of a building, allow for horizontal clearance of about 1m. For rooftop installations, a vertical clearance of about 1m should also be used – or whatever is required to avoid obstruction of the first Fresnel zone.

Covering Hard to Reach Areas

The realities of physics limit just how far an RF signal will propagate. Many outdoor deployments will have Non-Line of Sight (NLoS) or Near-Line of Sight in the coverage area. Buildings, for example, can cause *RF shadows* where there is little or no coverage. These problems could be addressed by deploying enough APs to reduce RF shadows to a minimum.

External Antennas

Antenna selection can play a critical part in mesh networks. A high gain antenna can greatly increase the RF signal to another mesh node; resulting in a higher performance and more stable connection. A tighter beam width may also help where Fresnel zone blockage would occur with an omni-directional antenna.

The downside of a directional antenna however is loss of coverage. A 120° sector antenna will only be able to communicate with APs that happen to be within that zone of coverage. The coverage of an AP with a 30° beam width is restricted even further. This can reduce the number of potential uplinks for a mesh node if its primary uplink is lost.

Nearly any external antenna can be used with Ruckus products provided it meets the following criteria:

- Dual polarized
- N-type connectors
- Meets regulatory compliance standards

A number of antennas are also available as part of the Ruckus product catalogue.

Site Location

Any outdoor installation requires choosing the right location for the APs and bridges. Location is **highly** dependent on the geographic area. A more urban environment can have very different location options and restrictions than a rural project. In general, outdoor APs are mounted at least 5m from the ground. Any potential location must include:

- 24x7 power source (PoE or 12V)
- Installation point (pole, vertical/horizontal flat surface)
- Line of Sight (LoS) or Near LoS to the upstream APs
- LoS or Near LoS to the coverage area (clients)
- Accessibility (for installation and future touch)

Street Coverage

If coverage is primarily clients along streets, the simplest mounting locations are typically rooftops and/or streetlamps. Choice may be dependent on where power and connectivity is available. This type of pattern places APs in a staggered fashion down the street. Alternating mesh nodes with root nodes ensures there are enough APs, redundancy and capacity available at any given point.

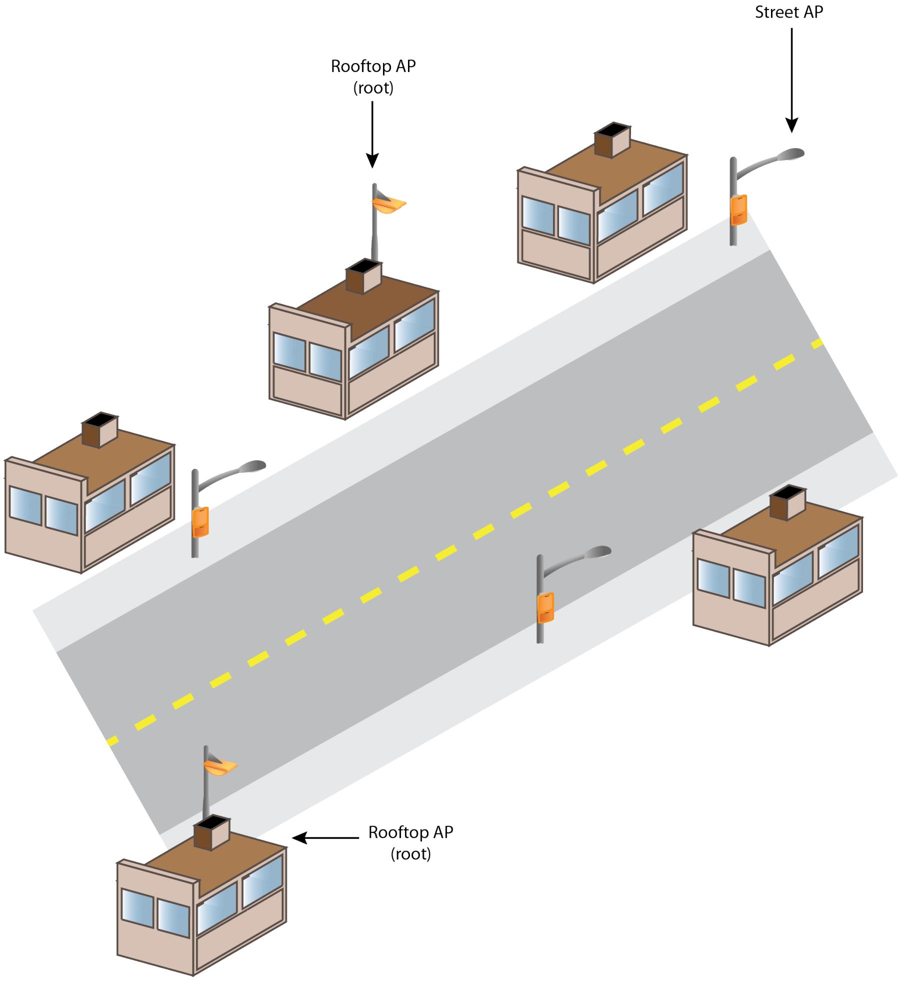

FIGURE 10: STAGGERED STREET COVERAGE

The diamond pattern model assumes coverage is possible from one side of the road to the other. If, for some reason, an AP on one side of the street cannot cover the opposite side, an alternate location or additional APs should be considered.

If coverage areas are separated by long stretches of road, a bridge link is a good way to span these distances without the cost of another mesh hop or the shorter distance limitation. It is not always possible to get location sites that exactly match this pattern. In this case look for possible alternatives:

- Try to find a different location that is nearby and has similar height/LoS/etc.
- Different antenna options, or mounting heights and orientation.

Do not try to use a site that is not near the ideal location unless you know it has the same characteristics (distance from other nodes, height, LoS). If you cannot guarantee the alternate location meets all requires, find a different location. If a good alternate solution is not available, change the design and use a different path through the area.

Line of Sight (LoS)

Although it has already been discussed, the importance of a clear line of sight (LoS) with no more than 60% of the first Fresnel zone obstructed cannot be stressed enough. Do not rely on memory or visual (eyeball) sighting. Make sure you know the exact distance between two locations, the exact mounting height and the height of any obstructions in between.

If a good LoS is not possible, consider raising the mounting height or try a different location. If you do choose to raise the mount height, make sure the link budgets (especially for clients) will allow for the extra distance. Mounting an AP higher seems like a good solution, but if you place it so high the weaker clients such as smartphones can't connect, it's not a good solution. Such a design error could cause more problems in the long run than it solves short term.

At times it makes sense to design for going low if you can't go high, for e.g. we have to get below the tree canopy because we can't get 80 – 100 feet in the air to get over the tree canopy.

Power

All APs and bridges require power. A unit can be powered by either Power over Ethernet (PoE) via a CAT5 cable or 12V from a direct current power source such as a battery. Selection of a site location must ensure there is enough power available to run the unit. In particular, there should be enough power to run the unit under all conditions and 100% of the time. Some locations may have power but not all of the time: for example, a park where lights are turned off at night. The same is also true of 12V: a solar charged battery must have sufficient storage and sunlight hours to provide continuous power.

Always consult the documentation for the particular model you are deploying and make sure there will be sufficient power. If using an AP with a large variance in potential power draw – for example a unit with an internal heater – make sure the highest possible power requirement is available.

Clients in a Mesh Environment

Capacity

The number of clients a mesh network can support is driven by the AP radios (single vs. dual) and the number of mesh hops. The client load on any mesh node is equivalent to the sum of all clients on downstream APs. Therefore, the more clients and/or hops, the higher the equivalent client load on upstream mesh nodes. Along with the number of clients, the client traffic is aggregated as well. This is an important point to remember; as many mesh throughput calculators do not include clients.

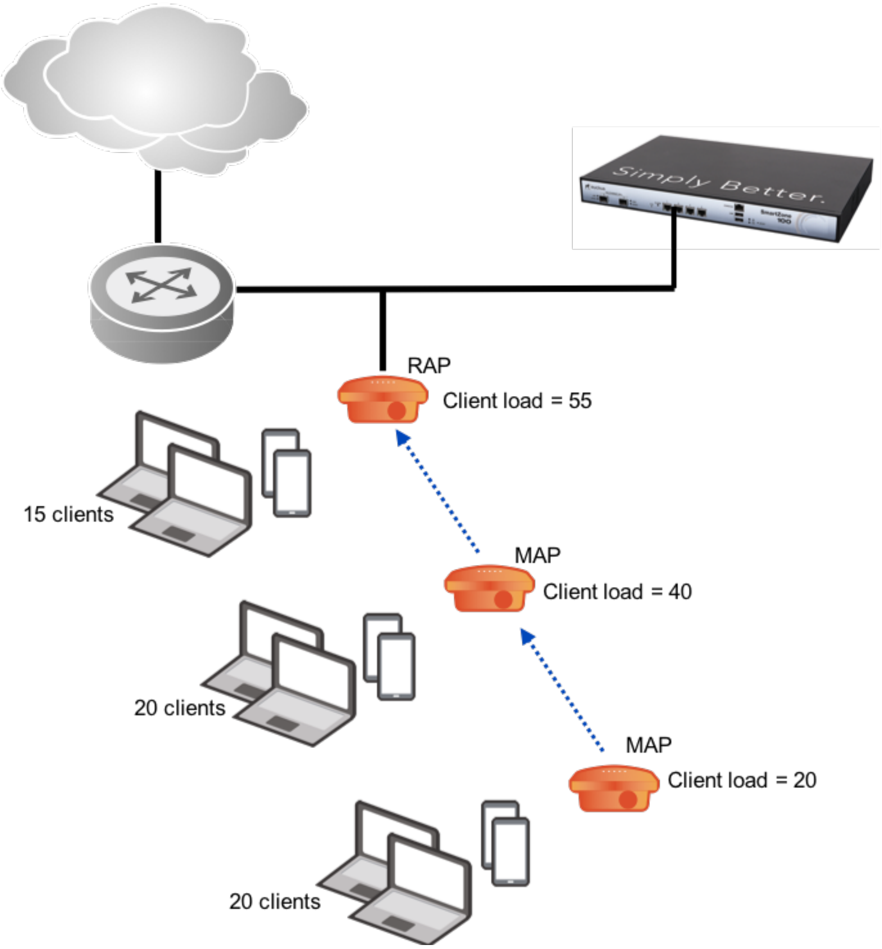

FIGURE 11: CLIENT LOADS IN A MESH

The Effect of Single vs. Dual-radio on Capacity

A common question is when to select one AP model over another. This usually breaks down into single vs. dual-radio. The answer is simple: clients drive everything. Clients – and their applications – determine everything from the minimum amount of bandwidth required to latency requirements. Both of these will restrict the mesh tree depth, choice of radios and minimum installation distance and SNR.

Data

Data applications such as web and email are not particularly constrained by bandwidth or latency. In this case, a single radio AP may be adequate to meet application needs. To ensure the best possible performance, a mesh node should be planned for every simultaneous 20-30 clients[4].

Video

Streaming video applications can vary widely in bandwidth requirements. The key factor for video is planning for the *minimum* throughput required. A 20 Mbps high-definition video stream will always require a minimum of 20 Mbps. If the available bandwidth is 40 Mbps, that is fine and doesn't impact the quality of the video. But if the available bandwidth ever drops below 20 Mbps, video quality will suffer.

Note: More details for video over mesh can be found here.

The zap[5] tool from Ruckus Wireless includes the ability to test a connection with video traffic and measure the minimum throughput available as a percentage of time. Video should have the minimum throughput available 99% of the time.

Voice

This is the most sensitive of all applications. Recommendations can vary by VoIP vendor, but in general a maximum of 150ms latency is recommended. Higher latency will result in dropped packets that can disrupt or disconnect active calls. Bandwidth for a voice call depends on the codec used, but in a campus environment using the G.711 codec is in the 64 Kbps range.[6]

Note: More details for voice best practices and design can be found here.

Client coverage

Some mesh networks are strictly backbone transport with little or no client access. The majority of mesh networks however are typically designed to support wireless client devices. If full coverage is required between mesh nodes, clients become a significant limiting factor on distance between APs. The reason is simple: client devices operate at greatly reduced power as compared to an AP. Therefore, each mesh node must be close enough to other APs such that the distance is no more than twice that of the client's ability to transmit. As a precaution, it is recommended this distance be somewhat less than twice to allow some overlap.

This distance will vary by client type. In general, laptops have a higher transmit power and greater range than devices such as phones. Even between devices that you would expect to have similar performance (like back-to-back models of 802.11ac phones), there is a significant difference in power and ranges. When planning wireless client support, always make a habit of designing to the lowest powered device[7].

Client Isolation

Some networks – in particular public access – require that end users not be able to communicate directly with each other or other local resources. This is known as client isolation.

There are several ways to isolate client traffic. For more information, please refer to the Global Client Isolation in the next section

[4] This number excludes connected, but inactive clients.
[5] Zap is available as an open source project. Pre-complied binaries are also available from Ruckus.
[6] A good VoIP bandwidth calculator can be found from Packetizer.
[7] An important corollary is the overall relationship between clients and wireless APs. Since it is so much more powerful than a client, APs can be run with reduced power and still support a large coverage area. This may be necessary if power reduction is required to reduce interference and noise.

Configuration Optimizations

This section describes various optional parameters available for optimizing the mesh configuration and how and when to configure them on the Ruckus equipment.

Global Client Isolation

This prevents clients from sending to ports other than their uplink ports. This prevents the forwarding of wireless client traffic between other wireless clients or APs. In the case of broadcast and multicast frames, the traffic is only allowed to go upstream. Unicast traffic to other wireless clients is dropped. The only exception is the APs that are uplinks to the mesh AP where the client is connected. In this case, the client traffic will be sent from the original AP via its uplink to the next mesh AP and so on until it reaches the root AP and the wired network.

This feature is disabled by default but may be enabled whenever there is no need to allow peer-to-peer wireless communications.

ZoneDirector

Global Client Isolation involves:

- Client isolation configured under the WLAN settings
- White Lists for permitted resources. A White List specifies certain devices in the local network a wireless client is allowed to reach. Clients attempting to spoof any of these devices will be blocked.

The following instructions configure client isolation and a white list:

1. Log on to the web UI and go to Configure->Access Control
2. Under the section titled Client Isolation White List client, click Create New
3. Under the Rules section, click Create New to define a resource
4. Enter the MAC and IP address for the device
5. Click Save
6. Configure to create rules until all have been added to the list
7. Click OK to save the list

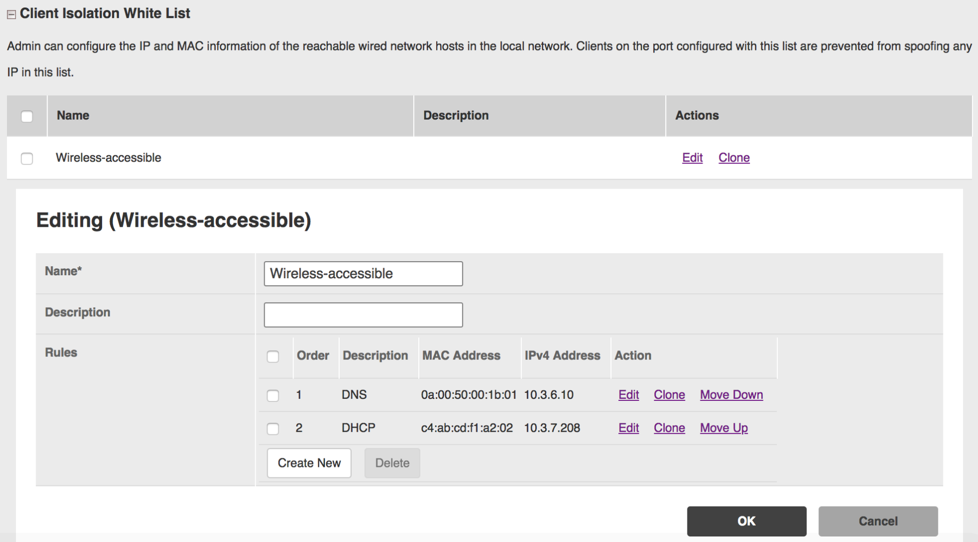

⊟ **Client Isolation White List**

Admin can configure the IP and MAC information of the reachable wired network hosts in the local network. Clients on the port configured with this list are prevented from spoofing any IP in this list.

☐	Name	Description	Actions
☐	Wireless-accessible		Edit Clone

Editing (Wireless-accessible)

Name*	Wireless-accessible
Description	
Rules	

☐	Order	Description	MAC Address	IPv4 Address	Action
☐	1	DNS	0a:00:50:00:1b:01	10.3.6.10	Edit Clone Move Down
☐	2	DHCP	c4:ab:cd:f1:a2:02	10.3.7.208	Edit Clone Move Up

FIGURE 12: ZD CLIENT ISOLATION WHITE LIST

The next step adds this white list to a WLAN.

8. Click Configure->WLANs
9. Click the Edit link next to the WLAN entry in the list
10. Check the Enable box under Wireless Client Isolation
11. Select the white list from the drop-down box underneath
12. Click OK to save the changes and apply the white list to the WLAN

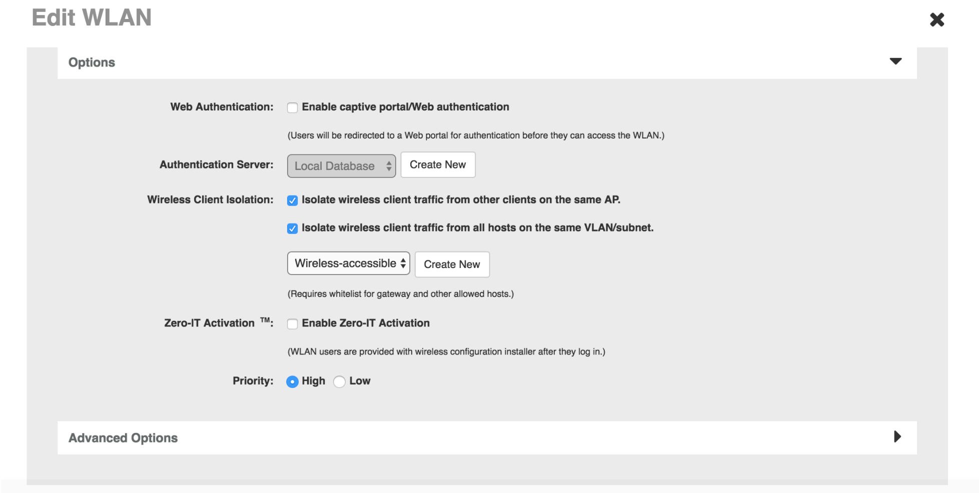

Edit WLAN ✖

Options

Web Authentication: ☐ Enable captive portal/Web authentication

(Users will be redirected to a Web portal for authentication before they can access the WLAN.)

Authentication Server: Local Database | Create New

Wireless Client Isolation: ☑ Isolate wireless client traffic from other clients on the same AP.

☑ Isolate wireless client traffic from all hosts on the same VLAN/subnet.

Wireless-accessible | Create New

(Requires whitelist for gateway and other allowed hosts.)

Zero-IT Activation ™: ☐ Enable Zero-IT Activation

(WLAN users are provided with wireless configuration installer after they log in.)

Priority: ● High ○ Low

Advanced Options

FIGURE 13: ZD WLAN WIRELESS CLIENT ISOLATION

This option also isolates clients on different SSIDs on the same AP from reaching each other.

SmartZoneOS

Similar to the ZoneDirector settings, global client isolation involves configuring the client isolation under WLAN settings and a White List for permitted resources.

The following instructions configure client isolation and a white list:

1. Log on to the web UI and go to Services & Profiles->Access Control
2. Under the tab titled "Client Isolation Whitelist", click on a zone under which you want to create this whitelist and click on "Create"
3. Enter the name of the whitelist and under the Client Entries box, click on "Create"
4. Enter the MAC and IP address for the device
5. Click OK to save the client entry
6. Create client entries until all have been added to the list
7. Click OK to save the list

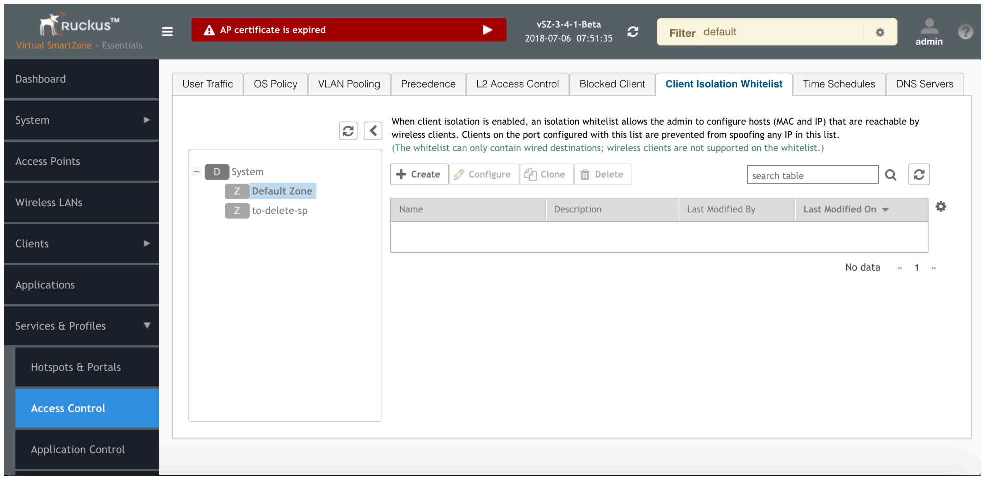

When client isolation is enabled, an isolation whitelist allows the admin to configure hosts (MAC and IP) that are reachable by wireless clients. Clients on the port configured with this list are prevented from spoofing any IP in this list.
(The whitelist can only contain wired destinations; wireless clients are not supported on the whitelist.)

Name	Description	Last Modified By	Last Modified On

FIGURE 14A: SZ CLIENT ISOLATION WHITE LIST

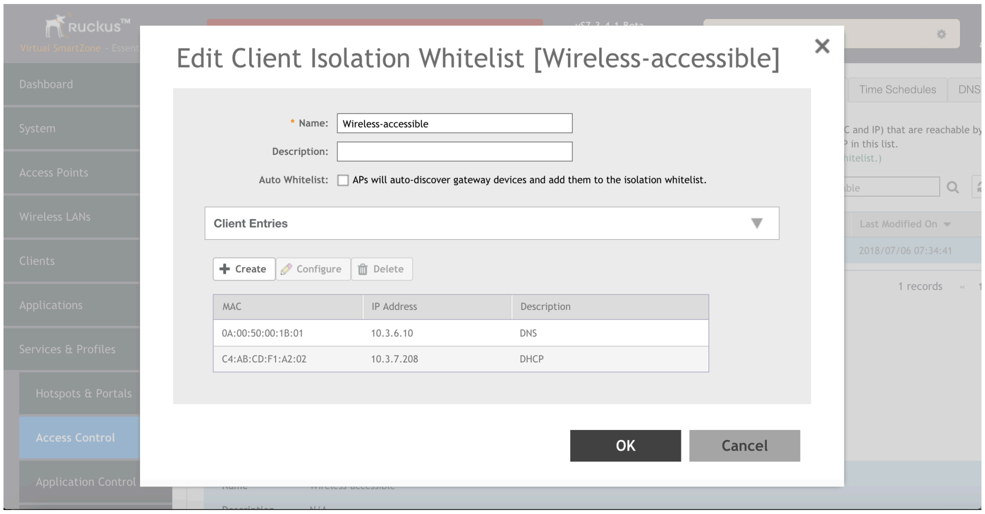

Edit Client Isolation Whitelist [Wireless-accessible]

* Name: Wireless-accessible
Description:
Auto Whitelist: □ APs will auto-discover gateway devices and add them to the isolation whitelist.

Client Entries

MAC	IP Address	Description
0A:00:50:00:1B:01	10.3.6.10	DNS
C4:AB:CD:F1:A2:02	10.3.7.208	DHCP

FIGURE 14B: SZ CLIENT ISOLATION WHITE LIST

The next step adds this white list to a WLAN.

1. Click on Wireless LANs
2. Click on the zone and then on the specific WLAN in which you want to enable Wireless Client Isolation
3. Click on Configure
4. Scroll down and then check Enable box beside Wireless Client Isolation
5. Select the white list from the drop-down box underneath
6. Click OK to save the changes and apply the white list to the WLAN

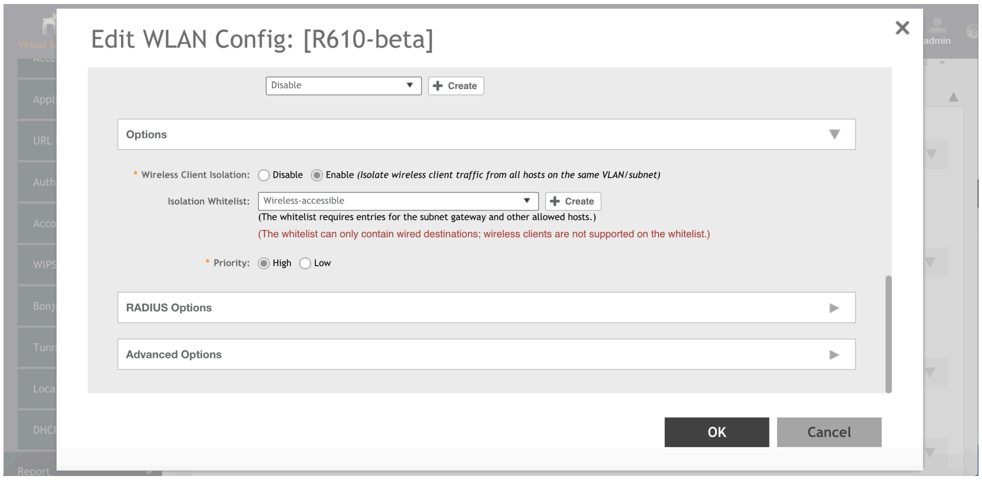

* Wireless Client Isolation: ○ Disable ● Enable *(Isolate wireless client traffic from all hosts on the same VLAN/subnet)*

Isolation Whitelist: Wireless-accessible

(The whitelist requires entries for the subnet gateway and other allowed hosts.)

(The whitelist can only contain wired destinations; wireless clients are not supported on the whitelist.)

* Priority: ● High ○ Low

FIGURE 15: SZ WLAN WIRELESS CLIENT ISOLATION

Mesh Topology Detection

If a mesh node fails, the topology must be reconfigured to self-heal around the failure. Depending on how many mesh nodes were deployed and how many other nodes are reachable, the new configuration might have more hops between some nodes than is desirable. The ZoneDirector can be configured to issue a warning message when a defined maximum number of hops is exceeded. Exceeding this number does not shutdown the mesh or change it; it simply informs the administrator that an event has occurred.

There are two types of warnings that may be configured. The maximum mesh hop count detection sends a warning if a mesh node exceeds the defined number of hops. The maximum downlink detection sends a warning if a mesh node has more mesh nodes connecting to it (downlink) than the defined number.

ZoneDirector

ZoneDirector may be configured as follows:

13. Log on to the web UI and go to Services & Profiles ->Mesh
14. Go to the section titled Mesh Topology Detection
15. Check the box for mesh hop count detection to configure the ZoneDirector to issue a warning when a node exceeds this number of hops
16. Check the box for mesh downlink detection to configure the ZoneDirector to issue a warning when a node exceeds this number of downlinks
17. Click Apply to save the changes

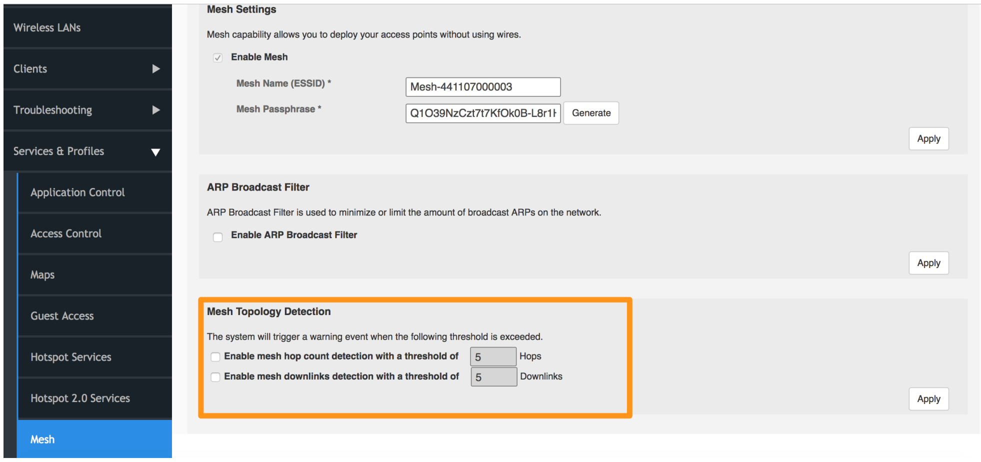

Mesh Settings

Mesh capability allows you to deploy your access points without using wires.

☑ **Enable Mesh**

Mesh Name (ESSID) * Mesh-441107000003

Mesh Passphrase * Q1O39NzCzt7t7KfOk0B-L8r1I Generate

Apply

ARP Broadcast Filter

ARP Broadcast Filter is used to minimize or limit the amount of broadcast ARPs on the network.

□ **Enable ARP Broadcast Filter**

Apply

Mesh Topology Detection

The system will trigger a warning event when the following threshold is exceeded.

□ **Enable mesh hop count detection with a threshold of** 5 Hops

□ **Enable mesh downlinks detection with a threshold of** 5 Downlinks

Apply

FIGURE 16: ZD MESH TOPOLOGY DETECTION

SmartZoneOS

The SmartZone architecture distributes more of the mesh functionality than the ZoneDirector does. And that's the reason mesh topology detection is not supported in SmartZone mesh.

Jumbo Frame Support

A key factor (particularly in small cell deployments) is support for jumbo frames across the mesh or bridge link. Jumbo frame support is needed to prevent fragmentation. The WLAN and mesh links typically default to 2290 bytes. The default on most ethernet interfaces is 1500 bytes. Different models support different max MTU sizes.

If jumbo frame support is required between two Ethernet interfaces on the same AP, the AP may be be configured via a CLI command as follows:

set mtu ethX

Where X is the Ethernet port number, e.g. eth0.

The following table shows the matrix of supported frames by AP model and software version:

Medium Type	R700	H500/R300	R500/R600/T300/T301/P300	R730/R720/R710/R610/T610/T710/T811-CM/7781-CM C110/H510/R510/T310/E510/M510
Ethernet	Default 1500 bytes Max 9578 bytes	Default 1500 bytes Max 1500 bytes	Default 1500 bytes Max 2290 bytes	Default 1500 bytes Max 9018 bytes
WLAN	2290 bytes	2290 bytes	2290 bytes	2290 bytes
WLAN Mesh	2290 bytes	2290 bytes	2290 bytes	2290 bytes

TABLE 4: JUMBO FRAME SUPPORT ON RUCKUS APS

Note: MTU size shown does not include the 18 byte header and Cyclic Redundancy Check (CRC).

Summary

Wireless mesh networks are an excellent way to provide wireless service in areas with limited in-place network access. By sending client traffic via a wireless backhaul link, the requirement for a wired network connection is removed allowing APs to be deployed almost anywhere.

Any wireless mesh network deployment should have the following issues addressed as fully and as early in the design process as possible:

- Mesh topology type (standard, eMesh, bridging)
- Radio and hardware selection
- AP mounting and distance between nodes
- Required throughput and latency
- Expected client performance

Taking these factors into account will greatly enhance mesh performance and resilience.

Troubleshooting

Verifying Mesh

Both the ZoneDirector and SmartZone network controller show the mesh topology on the Web UI.

SZ: Go to Access Points -> Zone and then click on "Mesh" tab in the right top corner to pick the mode.

ZD: Go to Access Points and click on "Mesh" tab in the right top corner to pick the mode.

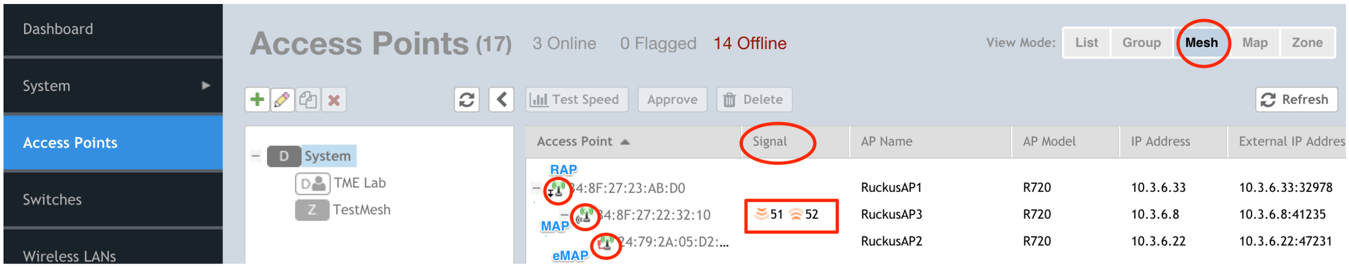

Access Points (17) 3 Online 0 Flagged 14 Offline

Access Point	Signal	AP Name	AP Model	IP Address	External IP Addres
34:8F:27:23:AB:D0		RuckusAP1	R720	10.3.6.33	10.3.6.33:32978
34:8F:27:22:32:10	51 52	RuckusAP3	R720	10.3.6.8	10.3.6.8:41235
24:79:2A:05:D2:...		RuckusAP2	R720	10.3.6.22	10.3.6.22:47231

Figure 17: SZ Mesh Verification

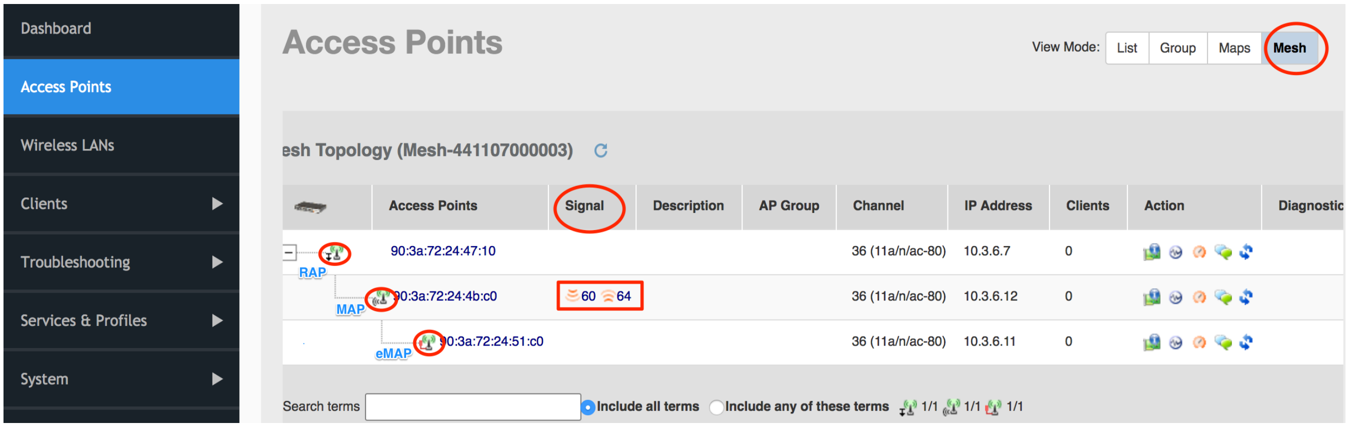

Access Points

esh Topology (Mesh-441107000003)

Access Points	Signal	Description	AP Group	Channel	IP Address	Clients	Action	Diagnostic
90:3a:72:24:47:10				36 (11a/n/ac-80)	10.3.6.7	0		
90:3a:72:24:4b:c0	60 64			36 (11a/n/ac-80)	10.3.6.12	0		
90:3a:72:24:51:c0				36 (11a/n/ac-80)	10.3.6.11	0		

Search terms ● Include all terms ○ Include any of these terms 1/1 1/1 1/1

Figure 18A: ZD Mesh Verification

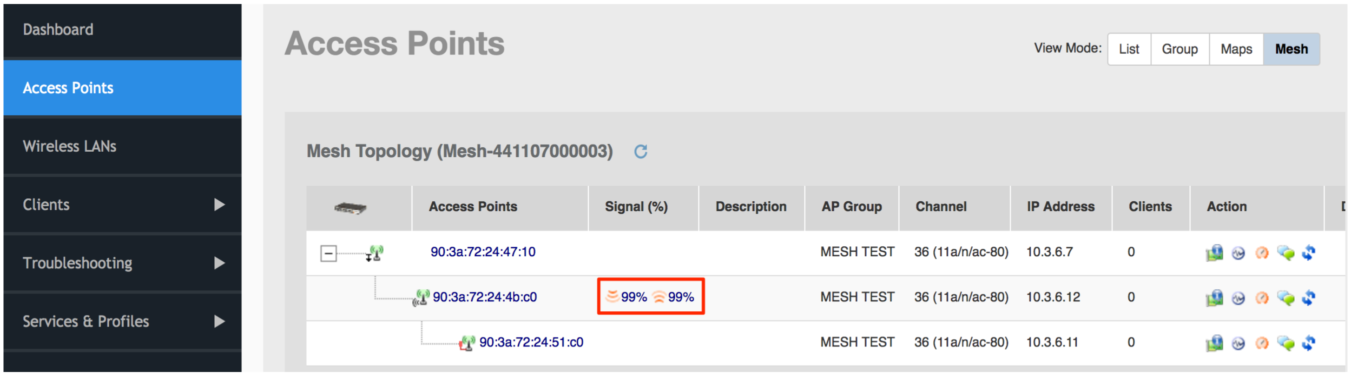

Access Points

Mesh Topology (Mesh-441107000003)

Access Points	Signal (%)	Description	AP Group	Channel	IP Address	Clients	Action
90:3a:72:24:47:10			MESH TEST	36 (11a/n/ac-80)	10.3.6.7	0	
90:3a:72:24:4b:c0	99% 99%		MESH TEST	36 (11a/n/ac-80)	10.3.6.12	0	
90:3a:72:24:51:c0			MESH TEST	36 (11a/n/ac-80)	10.3.6.11	0	

Figure 18B: ZD Mesh Verification

As seen in the figures 17 and 18, each type of mesh node has a special marker associated with it. There are many useful pieces of information to be gleaned from this page that we'll discuss in the next few sections.

Ruckus Signal % - SNR mapping

In Figure 17 on the SmartZone controller, the "Signal" at the MAP is displayed in SNR in dB. In Figure 18 on the ZoneDirector, the "Signal" can toggle between SNR in dB (Figure 18A) and Ruckus Signal % (Figure 19B) by clicking on the values. In Figure 18A, the first number (60) is the SNR from the RAP (90:3a:72:24:47:10) to the MAP (90:3a:72:24:4b:c0). The second number (64) is the SNR from the MAP to the RAP. The SNR is raw signal over the noise floor or noise generated by other devices on the same/adjacent channel. The higher the number the better the throughput and fade margin for when things go south.

Below is a mapping to understand the correlation between the Ruckus Signal % and SNR in dB.

Ruckus Signal %	SNR in dB
>45	99
45	99
43	94.05
41	89.1
39	84.15
37	79.2
35	74.25
33	69.3
31	64.35
29	59.4
27	54.45
25	49.5
23	44.55
21	39.6
19	34.65
17	29.7
15	24.75
13	19.8
11	14.85
9	9.9
7	4.95
5	0
3	0
1	0

TABLE 5: RUCKUS SIGNAL % - SNR MAPPING

The below table gives you a way to interpret the SNR values you see in the mesh network. An SNR of 25dB or higher is preferable.

SNR	Description
>40 dB	Excellent signal (5 bars); always associated; lightning fast
25 dB to 40 dB	Very good signal (3 to 4 bars); always associated; very fast
15 dB to 25 dB	Low signal (2 bars); always associated; usually fast
10 dB to 15 dB	Very low signal (1 bar); mostly associated; mostly slow
5 dB to 10 dB	No signal; not associated; no go

TABLE 6: INTERPRETING SNR VALUES

Test SpeedFlex for Mesh

Both the SmartZone network controller and the ZoneDirector have an option to use SpeedFlex to test the mesh link. This lets the customer test each link and then the complete mesh link to the controller or just in-between nodes. The result from this test gives an estimate of unused bandwidth. If there are more users on the network, then the test results will vary as the demand for backhaul will be used.

On the SmartZone click on the AP whose link you want to test and then hit the "Test Speed (Multi-hops)" button to start the test.

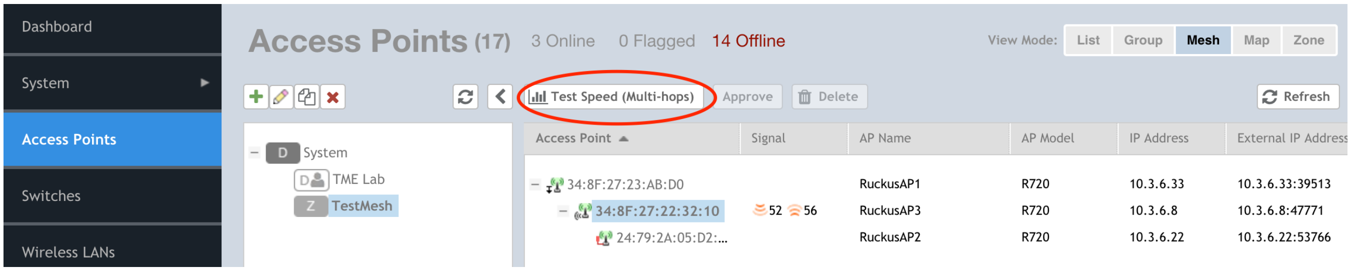

Access Point	Signal	AP Name	AP Model	IP Address	External IP Address
34:8F:27:23:AB:D0		RuckusAP1	R720	10.3.6.33	10.3.6.33:39513
34:8F:27:22:32:10	52 56	RuckusAP3	R720	10.3.6.8	10.3.6.8:47771
24:79:2A:05:D2:...		RuckusAP2	R720	10.3.6.22	10.3.6.22:53766

FIGURE 19A: SZ MESH SPEEDFLEX TEST

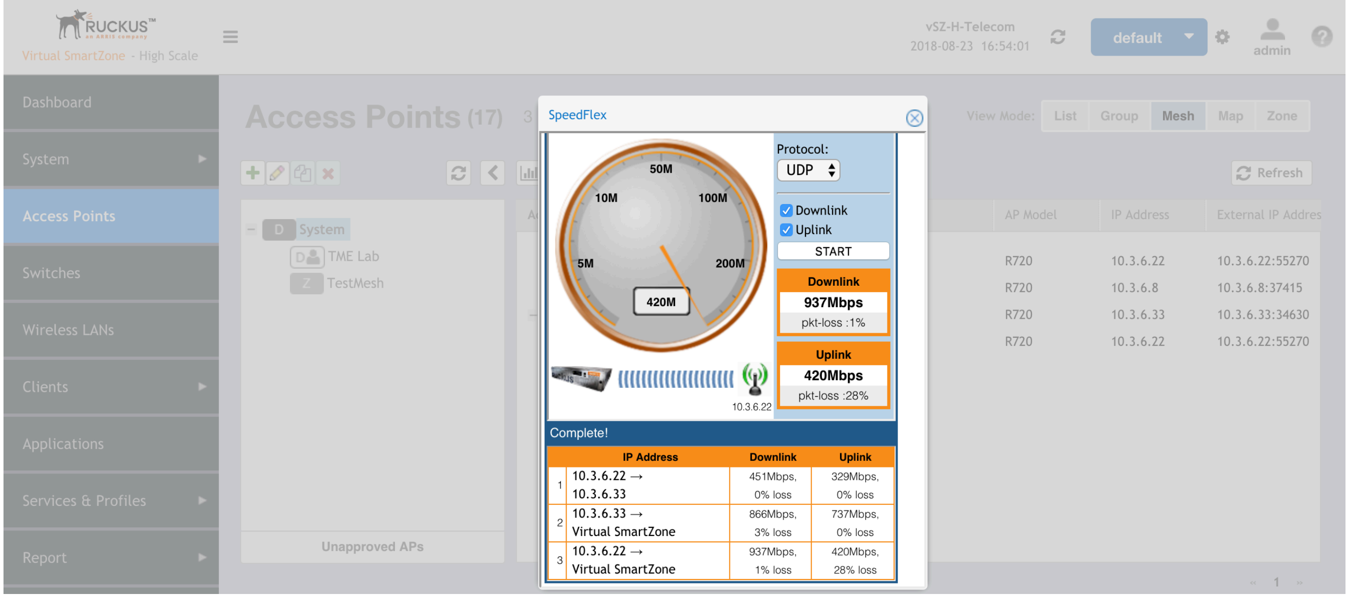

Protocol: UDP

☑ Downlink
☑ Uplink

START

Downlink
937Mbps
pkt-loss :1%

Uplink
420Mbps
pkt-loss :28%

Complete!

	IP Address	Downlink	Uplink
1	10.3.6.22 → 10.3.6.33	451Mbps, 0% loss	329Mbps, 0% loss
2	10.3.6.33 → Virtual SmartZone	866Mbps, 3% loss	737Mbps, 0% loss
3	10.3.6.22 → Virtual SmartZone	937Mbps, 1% loss	420Mbps, 28% loss

AP Model	IP Address	External IP Addres
R720	10.3.6.22	10.3.6.22:55270
R720	10.3.6.8	10.3.6.8:37415
R720	10.3.6.33	10.3.6.33:34630
R720	10.3.6.22	10.3.6.22:55270

Figure 19B: SZ Mesh SpeedFlex Test

On the ZoneDirector, click on the AP you want to test and then click on the SpeedFlex icon in orange to start the test.

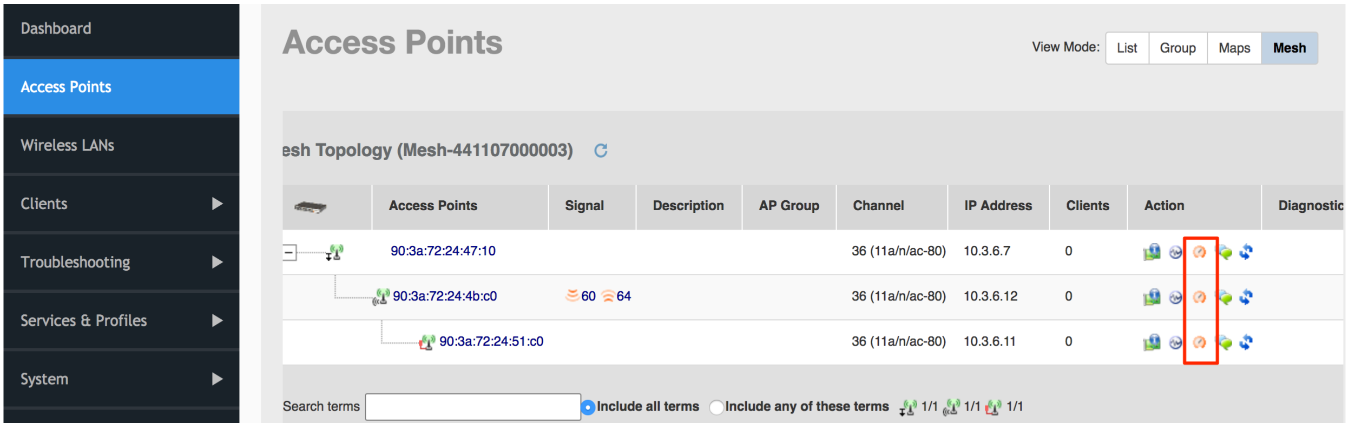

esh Topology (Mesh-441107000003)

	Access Points	Signal	Description	AP Group	Channel	IP Address	Clients	Action	Diagnostic
	90:3a:72:24:47:10				36 (11a/n/ac-80)	10.3.6.7	0		
	90:3a:72:24:4b:c0	60 64			36 (11a/n/ac-80)	10.3.6.12	0		
	90:3a:72:24:51:c0				36 (11a/n/ac-80)	10.3.6.11	0		

Search terms ● Include all terms ○ Include any of these terms 1/1 1/1 1/1

Figure 20A: ZD Mesh SpeedFlex Test

On clicking the orange icon against an AP, it takes you to the screen below where you start the test by hitting the “Multi-hops SpeedFlex” button.

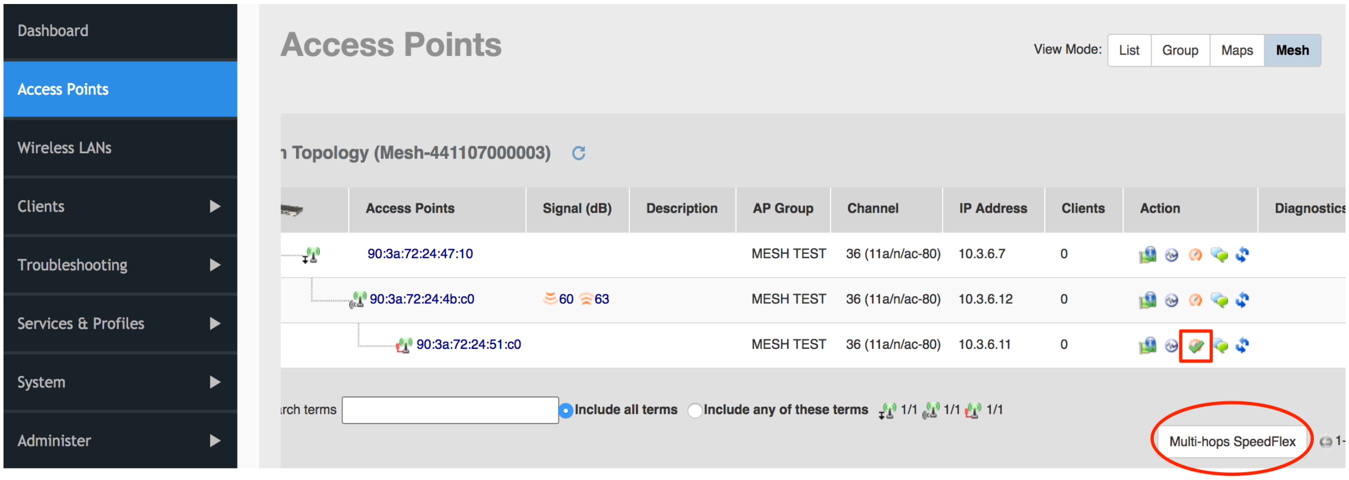

n Topology (Mesh-441107000003)

	Access Points	Signal (dB)	Description	AP Group	Channel	IP Address	Clients	Action	Diagnostics
	90:3a:72:24:47:10			MESH TEST	36 (11a/n/ac-80)	10.3.6.7	0		
	90:3a:72:24:4b:c0	60 63		MESH TEST	36 (11a/n/ac-80)	10.3.6.12	0		
	90:3a:72:24:51:c0			MESH TEST	36 (11a/n/ac-80)	10.3.6.11	0		

rch terms ● Include all terms ○ Include any of these terms 1/1 1/1 1/1

FIGURE 20B: ZD MESH SPEEDFLEX TEST

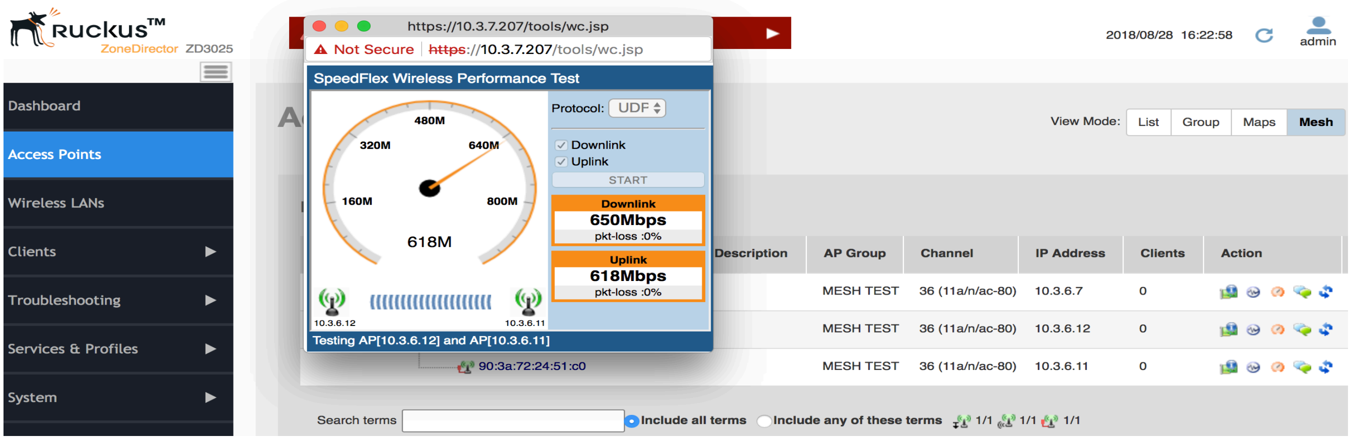

https://10.3.7.207/tools/wc.jsp

Not Secure | https://10.3.7.207/tools/wc.jsp

SpeedFlex Wireless Performance Test

Protocol: UDP

☑ Downlink
☑ Uplink

START

Downlink
650Mbps
pkt-loss :0%

Uplink
618Mbps
pkt-loss :0%

Testing AP[10.3.6.12] and AP[10.3.6.11]

	Description	AP Group	Channel	IP Address	Clients	Action
		MESH TEST	36 (11a/n/ac-80)	10.3.6.7	0	
		MESH TEST	36 (11a/n/ac-80)	10.3.6.12	0	
90:3a:72:24:51:c0		MESH TEST	36 (11a/n/ac-80)	10.3.6.11	0	

Search terms ● Include all terms ○ Include any of these terms 1/1 1/1 1/1

FIGURE 20C: ZD MESH SPEEDFLEX TEST

Note: SmartZone Monitoring Capabilities is discussed in great deal in the IP video over mesh document referenced in the last section.

Common Issues

Few of the common issues that we see in mesh networks are listed below. Understanding these issues and the possible root causes could help in providing the right guidance to the customer to mitigate their problems.

Root causes

1. Wire
   a. Bad Ethernet - Ethernet switches can misbehave, create loops and drop packets.
   b. Bad Backhaul – Backhaul outages and Root AP reboots wreak havoc in mesh networks.
   c. Bad DHCP - AP without IP can cause repeated join/leave causing other APs to become unstable.
2. Wireless
   a. Poor AP mounting locations – As seen in the factors affecting mesh performance, bad mounting location reduces signal strength and affects the mesh performance.

b. Interference - MAP experiences RF interference on certain channels, but RAP running ChannelFly does not always avoid such channels as it fails to detect interference down the mesh tree. This sort of localized interference at the MAP could result in asymmetric throughput. Low throughput but high RSSI – check for external sources of interference or if the link (distance between the mesh nodes) is too long. If no client access is needed, opting for a bridge over mesh for long distances is a good practice.

c. Stranded MAP – Could be due to a bad installation / misconfiguration

3. Client

   a. Fast roaming / Multiple simultaneous associations (make before break) – clients can associate with more than 1 AP at a time. We rely on the controller to notify the old AP, but that could take time. In the meanwhile, the client can send data packets to multiple APs on the same channel, causing the Linux kernel bridge to disengage. Misbehaving clients can cause infrastructure to act irrationally.

4. Others

   a. APs shipped with mixed SCG & ZD firmware - APs of different SCG & ZD FW versions mixed in the same network seems to cause interference.

   b. ChannelFly fast switching of channels at bootup causes mesh flaps.

   c. Other common mesh problems:

      i. Loops or suspected loops
      ii. Changing of role
      iii. MAP not joining controller quickly
      iv. APs join or lose connection to controller sporadically
      v. Misconfigured MTU on the Ethernet or mesh WLANs

Shell commands for debugging

Useful shell commands / meshd options

Meshd Option	Description
-h	Help
-hh	Reason codes help
-dn	Show neighbor list (first line is self)
-da	Show all neighbors (including non-mesh)
-dT	Show neighbor’s path to root list
-g	Concise meshd state info
-l	Full meshd state info
-m	Meshd debug level, see help (-h) -m 0xFFFFFFF8 (open all debug levels) -m 0xFFFFF7FE (default)
-U	Force reacquisition of uplink

TABLE 7: MESH DEBUG SHELL COMMANDS

Reading and understanding the mesh neighbor table

meshd -dn

The mesh neighbor table can be obtained from the **shell** command "***meshd -dn***" or the AP **CLI** command "***get mesh***".

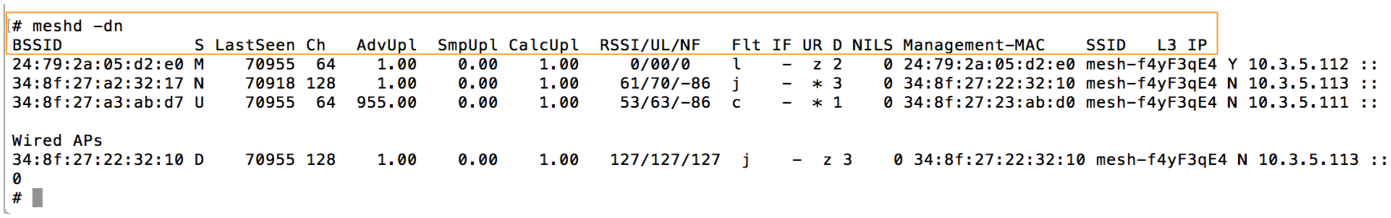

```
# meshd -dn
BSSID              S LastSeen Ch   AdvUpl  SmpUpl CalcUpl   RSSI/UL/NF   Flt IF UR D NILS Management-MAC    SSID   L3 IP
24:79:2a:05:d2:e0 M    70955  64    1.00    0.00    1.00      0/00/0    l     -  z 2    0 24:79:2a:05:d2:e0 mesh-f4yF3qE4 Y 10.3.5.112 ::
34:8f:27:a2:32:17 N    70918 128    1.00    0.00    1.00     61/70/-86  j     -  * 3    0 34:8f:27:22:32:10 mesh-f4yF3qE4 N 10.3.5.113 ::
34:8f:27:a3:ab:d7 U    70955  64  955.00    0.00    1.00     53/63/-86  c     -  * 1    0 34:8f:27:23:ab:d0 mesh-f4yF3qE4 N 10.3.5.111 ::

Wired APs
34:8f:27:22:32:10 D    70955 128    1.00    0.00    1.00    127/127/127  j     -  z 3    0 34:8f:27:22:32:10 mesh-f4yF3qE4 N 10.3.5.113 ::
0
# 
```

FIGURE 21: MESH -DN

Field	Description
BSSID	The advertised MAC address of the neighbor's VAP
S	Neighbor Mesh State (as the AP sees it)
LastSeen	Last time AP was seen in scan list (uptime in seconds)
Ch	Channel of neighbor
AdvUpl	Advertised Uplink – neighbor's uplink bandwidth metric (neighbor to its own uplink)
SmpUpl	Sample Uplink – estimated uplink bandwidth metric based on quality of link (me to neighbor)
CalcUpl	Calculate Uplink – estimate uplink bandwidth metric based on neighbor's uplink and the link between me and neighbor
RSSI UL NF	Downlink RSSI (AP's receive RSSI of neighbor) Uplink RSSI (neighbor's receive RSSI of me) Measured noise floor for the channel the AP is on. Its measured during background scan so results may vary
Flt	Reason AP won't uplink to that AP
IF	Island Filter (why neighbor is not connected)
UR	Uplink Reason (why did I choose the AP for uplink)
D	Mesh Depth (for RAP its 1)
NILS	Neighbor NOT in last n in-channel scan (criteria for AP to age out neighbor AP which no longer exists)
Management-MAC	Base MAC of neighbor
SSID	Neighbor advertised SSID
IP	IP address of neighbor

TABLE 8: MESH -DN

Lets further elaborate a few of the fields in Table 8.

Neighbor mesh state S

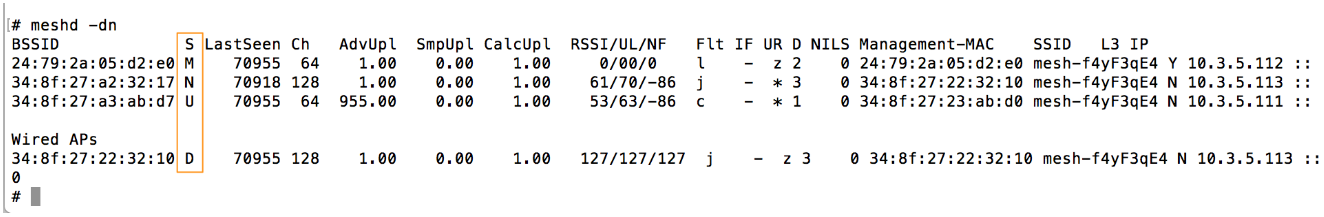

```
# meshd -dn
BSSID             S LastSeen Ch   AdvUpl  SmpUpl CalcUpl  RSSI/UL/NF   Flt IF UR D NILS Management-MAC    SSID   L3 IP
24:79:2a:05:d2:e0 M    70955  64    1.00    0.00    1.00     0/00/0    l    -  z 2    0 24:79:2a:05:d2:e0 mesh-f4yF3qE4 Y 10.3.5.112 ::
34:8f:27:a2:32:17 N    70918 128    1.00    0.00    1.00    61/70/-86  j    -  * 3    0 34:8f:27:22:32:10 mesh-f4yF3qE4 N 10.3.5.113 ::
34:8f:27:a3:ab:d7 U    70955  64  955.00    0.00    1.00    53/63/-86  c    -  * 1    0 34:8f:27:23:ab:d0 mesh-f4yF3qE4 N 10.3.5.111 ::

Wired APs
34:8f:27:22:32:10 D    70955 128    1.00    0.00    1.00   127/127/127  j    -  z 3    0 34:8f:27:22:32:10 mesh-f4yF3qE4 N 10.3.5.113 ::
0
# 
```

FIGURE 22: MESH -DN (S)

The first row in this command output is always the AP on which the command is executed, its own state.

Value	Description
R	RAP – AP is root (might not be advertising service)
I	Island – AP does not have an uplink
M	MAP – AP is connected to a wireless uplink
K	eMAP – AP is connected to a wired uplink
C	Consensus – AP is ensuring all eMAPs agree it has best uplink
W	Waiting – AP is trying to connect to uplink
O	Dormant – meshd is inactive

TABLE 9A: MESH -DN (S)

Following lines are the AP's (on which the command is executed) neighbors' states

Value	Description
N	Neighbor – depending on Flt might be used as uplink
U	Uplink – This is the neighbor I'm connected to (my uplink)
D	Downlink – These are the neighbors I'm offering service to (my downlinks)
P	Proposed – Proposed uplink, check consensus and try to connect
-	State unknown (unlikely)

TABLE 9B: MESH -DN (S)

Flt

Flt explains why neighbor is excluded from wireless uplink selection.

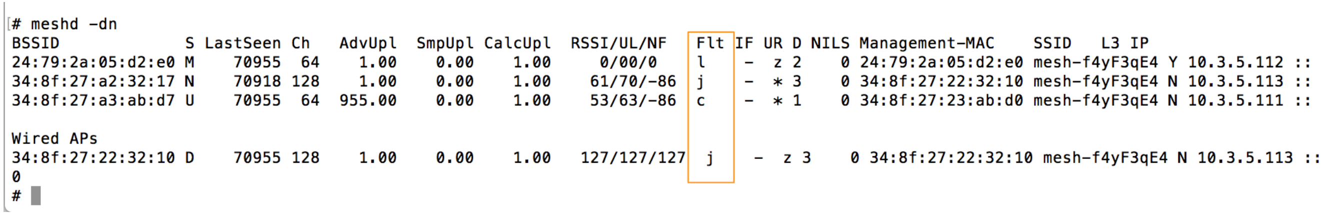

```
# meshd -dn
BSSID             S LastSeen Ch   AdvUpl  SmpUpl CalcUpl  RSSI/UL/NF   Flt IF UR D NILS Management-MAC    SSID   L3 IP
24:79:2a:05:d2:e0 M    70955  64    1.00    0.00    1.00     0/00/0    l    -  z 2    0 24:79:2a:05:d2:e0 mesh-f4yF3qE4 Y 10.3.5.112 ::
34:8f:27:a2:32:17 N    70918 128    1.00    0.00    1.00    61/70/-86  j    -  * 3    0 34:8f:27:22:32:10 mesh-f4yF3qE4 N 10.3.5.113 ::
34:8f:27:a3:ab:d7 U    70955  64  955.00    0.00    1.00    53/63/-86  c    -  * 1    0 34:8f:27:23:ab:d0 mesh-f4yF3qE4 N 10.3.5.111 ::

Wired APs
34:8f:27:22:32:10 D    70955 128    1.00    0.00    1.00   127/127/127  j    -  z 3    0 34:8f:27:22:32:10 mesh-f4yF3qE4 N 10.3.5.113 ::
0
# 
```

FIGURE 23: MESH -DN (FLT)

Value	Description
C	Connected to this AP (good state)
-	AP can be used as uplink
A	ACL Miss (neighbor is not in ACL list)
B	Neighbor not 'seen' enough times to allow connection to it
D	Max depth reached (neighbor has too many hops)
E	SSID mismatch (neighbor doesn't have the same mesh SSID)
F	Max fanout (too many APs connected to the neighbor)
G	Gateway/network mismatch (different subnet)
I	Missing Ruckus IE

TABLE 10: MESH -DN (FLT)

Island Filter IF

IF explains why neighbor has become stranded (not able to find uplink)

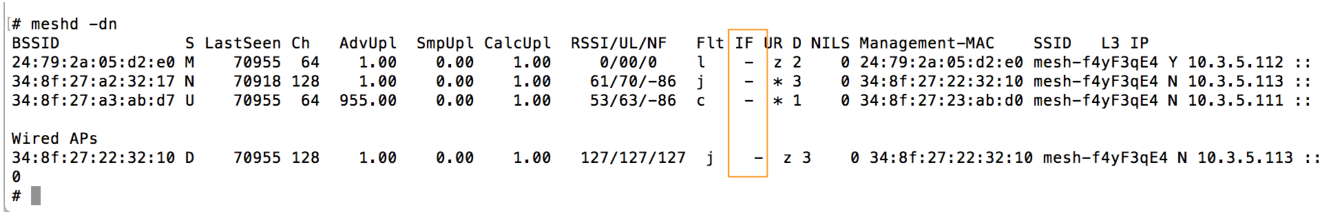

```
# meshd -dn
BSSID             S LastSeen Ch   AdvUpl  SmpUpl CalcUpl   RSSI/UL/NF   Flt IF UR D NILS Management-MAC    SSID   L3 IP
24:79:2a:05:d2:e0 M    70955  64    1.00    0.00    1.00     0/00/0     l    -  z 2    0 24:79:2a:05:d2:e0 mesh-f4yF3qE4 Y 10.3.5.112 ::
34:8f:27:a2:32:17 N    70918 128    1.00    0.00    1.00    61/70/-86   j    -  * 3    0 34:8f:27:22:32:10 mesh-f4yF3qE4 N 10.3.5.113 ::
34:8f:27:a3:ab:d7 U    70955  64  955.00    0.00    1.00    53/63/-86   c    -  * 1    0 34:8f:27:23:ab:d0 mesh-f4yF3qE4 N 10.3.5.111 ::

Wired APs
34:8f:27:22:32:10 D    70955 128    1.00    0.00    1.00   127/127/127   j    -   z 3    0 34:8f:27:22:32:10 mesh-f4yF3qE4 N 10.3.5.113 ::
0
# 
```

FIGURE 24: MESH -DN (IF)

Value	Description
C	Connected to this AP (good state)
-	AP can be used as uplink
A	ACL Miss (neighbor is not in ACL list)
B	Neighbor not 'seen' enough times to allow connection to it
D	Max depth reached (neighbor has too many hops)
E	SSID mismatch (neighbor doesn't have the same mesh SSID)
F	Max fanout (too many APs connected to the neighbor)
G	Gateway/network mismatch (different subnet)
I	Missing Ruckus IE

TABLE 11: MESH -DN (IF)

UR

UR (first row only) explains why the AP has changed its uplink status

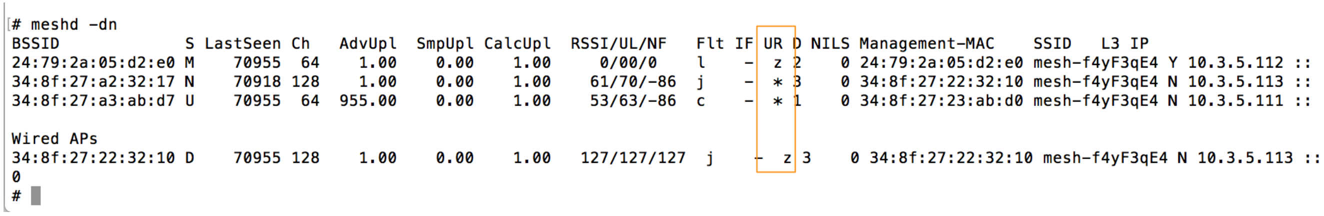

```
# meshd -dn
BSSID             S LastSeen Ch   AdvUpl  SmpUpl CalcUpl   RSSI/UL/NF    Flt  IF UR D NILS Management-MAC    SSID    L3 IP
24:79:2a:05:d2:e0 M    70955  64    1.00    0.00    1.00     0/00/0     l      -  z 2     0 24:79:2a:05:d2:e0 mesh-f4yF3qE4 Y 10.3.5.112 ::
34:8f:27:a2:32:17 N    70918 128    1.00    0.00    1.00    61/70/-86   j      -  * 3     0 34:8f:27:22:32:10 mesh-f4yF3qE4 N 10.3.5.113 ::
34:8f:27:a3:ab:d7 U    70955  64  955.00    0.00    1.00    53/63/-86   c      -  * 1     0 34:8f:27:23:ab:d0 mesh-f4yF3qE4 N 10.3.5.111 ::

Wired APs
34:8f:27:22:32:10 D    70955 128    1.00    0.00    1.00   127/127/127   j      -  z 3     0 34:8f:27:22:32:10 mesh-f4yF3qE4 N 10.3.5.113 ::
0
# 
```

FIGURE 25: MESH -DN (UR)

Value	Description
*	First time uplink (no reason)
T	Roam to path offering better throughput, was MAP
J	Roam to path offering better throughput, was eMAP
A	ACL enforced
Q	Manually triggered Reacquire
Y	Yielded to wired neighbor with better wireless uplink
Z	Previous wired parent disappeared
F	Lost connection to Gateway, RAP -> MAP failover
G	Lost connection to Gateway (implies Gateway previously detected)
X	Loop recovery, root previously detected on wire
E	Loop recovery, parent previously detected on wire
B	Loop recovery, Gateway previously detected

TABLE 12: MESH -DN (UR)

Mesh Logs

- Mesh log is now part of syslog (*logread*)
- To view only meshd logs, use '*mlogread*'
- To view only meshd logs and follow, use '*mlogread -f*'

Appendix A: Reference Architectures

Reference Architecture Option: SmartMesh Partitioning

There are times when it may be helpful to create multiple SmartMesh networks that work independently of each other. This can be done by using the concept of AP Zones in the SmartZone controller.

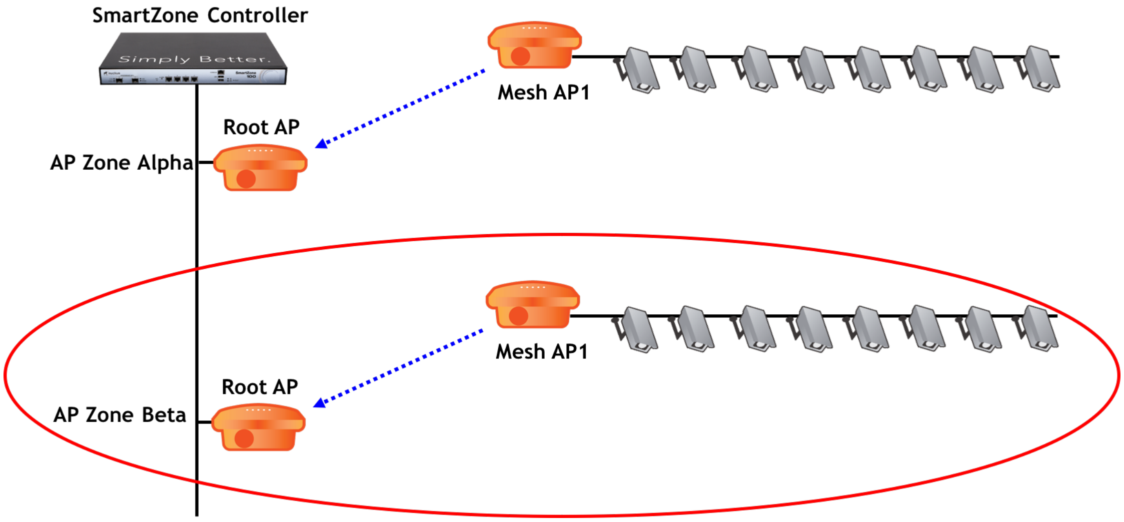

FIGURE 26: AP ZONE EXAMPLE FOR MESH PARTITIONING

Partitioning a Mesh Network

A limitation of mesh networks is the requirement to have all nodes of the network on the same channel. This limitation can be reduced by partitioning a single mesh topology into two or more. This have several benefits:

- Each additional channel represents more overall channel capacity
- Additional channels can help mitigate localized interference

The downside of partitioning a mesh network is mainly:

- Additional configuration and planning complexity
- If there is only one good channel available, additional mesh networks will suffer with lower quality links, impacting capacity

SmartZone Network Controller

The Ruckus SmartZone Network controller family offers features that can make the complexity of configuring and managing multiple mesh networks simpler. In particular, AP Zones can be used to create individual networks configurations that are simple to understand and monitor.

Use of AP Zones

When initially configuring a SmartZone system that manages mesh APs used for IP surveillance video transport, the first step is to determine how the mesh APs will be grouped and organized in the controller. It is recommended that AP zones be used to create logical groupings of APs. The use of AP zones is preferred over AP groups, and they are preferred to having all APs under one zone, or under one AP group.

When to Use AP Zones

Typically, there will be a commonality to the APs being grouped together. For example, one zone could be comprised of all of the APs within a geographic area, another zone could consist of all of the APs that a specific administrator has the job of managing.

Grouping all of the APs together that use the same wired backhaul could be the basis for organization as well. This could be all of the RAPs installed on a building and the MAPs that connect to them.

There is flexibility in how the APs can be grouped together, so there are no hard and fast rules. The only requirement is that all MAPs and the RAPs that they will be connected to be grouped into the same AP zone.

The first step is to determine how the APs will be grouped together. Once the AP groupings have been determined, the AP zones can be created. When the APs join the controller, they can be moved into the zones they have been assigned to.

How AP Zones are Created

With no prior controller configuration, on the vSZ-E and SZ100, APs will automatically be assigned to the Default Zone after they join the controller. On the vSZ-H and SZ300, after APs join the controller, they will appear in the Staging Zone.

For all controller models, several different methods can be used to either manually or automatically move APs into their assigned zones.

- Registration rules can be created to automatically move APs into their assigned zones when they join the controller for the first time based on one of several available rule types.
- The Import Batch Provisioning APs feature can also be used to automatically assign APs to zones based on identifying parameters entered in the pre-provisioning CSV file. The CSV file can be uploaded to the controller at any time to apply the AP to zone mapping.
- After an AP joins the controller, it can be moved manually to its assigned zone by clicking on it in the GUI and using the Move button to move it to its destination AP zone.
- The SmartZone API can also be used to move APs to specific zones.

Mesh Partitioning Advantages

By grouping APs together into a zone, one can achieve a more cohesive management experience. The following benefits are enabled with zone management:

- AP firmware version management
- AP firmware upgrade control
- Bulk AP configuration limited to the APs within a zone
- Bulk AP configuration through AP CLI scripts limited to the APs within a zone
- Streamlined map and dashboard viewing
- Simplified wired client, e.g. video camera, network monitoring
- Quick event filtering and viewing
- Simplified Syslog and SNMP management

While this list is not completely exhaustive, these are the primary benefits that an administrator will have by grouping APs into unique AP zones.

Note: Each AP Zone is managed independently of each other, which means that any changes to SSIDs or firmware on one AP Zone, will not affect any other zone.

Appendix B: Network Deployment Process

Network Deployment Process

A good network design process should always follow a best practices approach for determination of network needs, assessment of the local site, design, implementation, validation, and documentation. This document will touch on many parts of this cycle with regards to best practices for wireless mesh design, but it should not be considered a full and final substitution for the entire process. When it comes to Wi-Fi deployments, and outdoor deployments in particular, every installation is different and may have unique challenges that must be considered for a successful deployment.

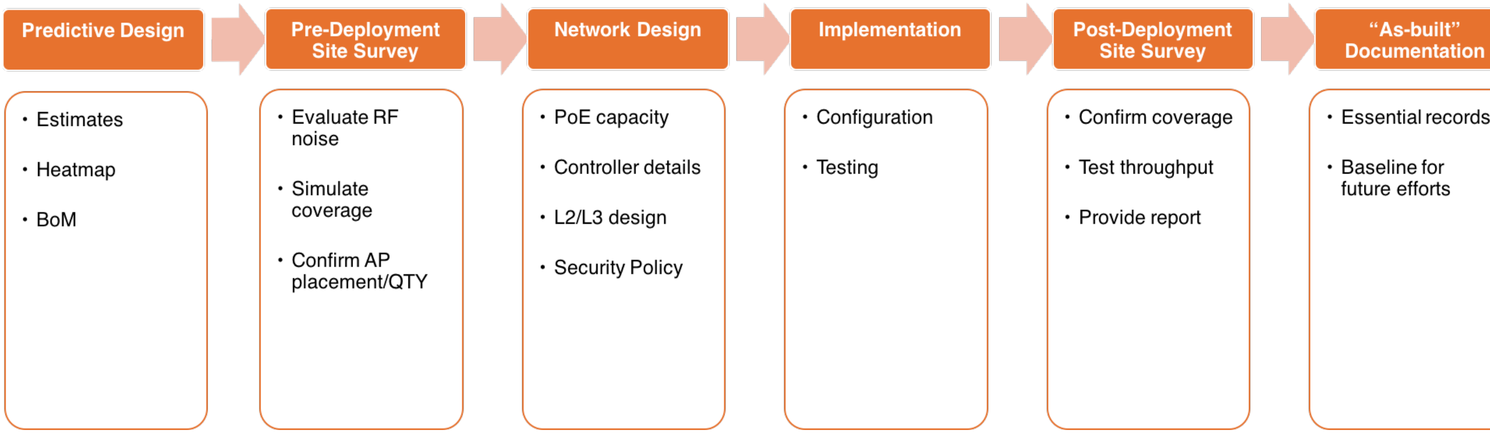

FIGURE 27: NETWORK DEPLOYMENT PROCESS

Why Wireless Deployments Fail

One of the top reasons **Wi-Fi deployments fail is because key questions where not asked and determined before the design process began.** Examples of this include:

- What is the goal of the deployment?
- What type of client device types will be used?
- What applications will be used and what are their minimum requirements?
- What type of RF environment is it?
- Estimated capacity required for day 1 installation? Capacity requirements 2-3 years from now?
- Is an accurate to-scale map available?
- What are the expectations and key performance indicators (KPIs)?
- What infrastructure is already available? Copper cabling? Fiber? Point-to-point bridges?
- What type of power is available?
- What are the available mounting options?

Few deployments will have the answers to all of these questions upfront, but they must be answered before a reasonable and workable WLAN design can be produced.

Another reason wireless deployments fail is because they do not consider hard limitations that may be inherent to certain designs.

Site Assessments

There are three parts to a thorough site survey: predictive modeling, on-site survey, and post-installation validation. Each of these steps are critical to deployment success and should be performed by an experienced Wi-Fi engineer using proper assessment tools.

Predictive Design

A predictive survey allows for an educated guess at the number of APs required. It is only an estimate, however. Predictive modeling is a good starting point to get the most out of an on-site site survey. **It is not a replacement for a site survey**.

Predictive modeling takes a to-scale map of an area and can be used to estimate factors such as coverage area, channel selection, etc. This can be helpful for a first-pass to identify possible AP locations and the number required. It is important to keep in mind however that a predictive model does not consider local conditions that an on-site survey does.

Indoor vs. Outdoor Prediction Tools

There are several different tools available to do Wi-Fi modeling and predictive analysis. Few however, support outdoor modeling. Before doing a predictive analysis, make sure the tool supports outdoor deployment modeling. It is not acceptable to use an indoor tool with something like a Google maps image. This will not provide sufficient or useful information since it does not consider parameters critical for outdoor designs.[8]

Predictive Modeling Checklist

A predictive model should be able to provide, at a minimum, the following information:

1. To-scale maps
2. Potential AP mounting locations
3. Estimated AP coverage heat maps with estimated RSSI and SNR values
4. AP channel assignments
5. Suggested number of APs based on coverage only

RF Signal Quality

The better the signal quality between mesh APs and their root APs, the higher the signal quality. This translates directly to better performance for camera video quality and reliability. A good site survey will determine these values. When doing the initial predictive model, at a minimum, the following values are recommended:

- Keep mesh APs as close as possible to root APs to maximize signal strength and signal-to-noise (SNR) values. This is typically 200m or less, although local conditions may alter this value
- Mesh APs should have an unobstructed Line of Sight (LoS) to at least one root AP node (two is recommended for redundancy in case of a node failure)
- If another AP will be used to provide access for other wireless devices, make sure this AP is mounted at least one meter away from the mesh AP
- Each mesh AP should have a minimum of -70 dBm signal strength to its root AP and a target SNR value of 20 dB or higher

Pre-deployment Site Survey

Only an on-site survey assessment can provide critical information such as whether or not there are limiting local factors that will change the predictive model. For example, a predictive model assumes there are no other Wi-Fi networks around that might interfere. Obviously, if the deployment site is in a heavily populated area such as a downtown, the RF environment is likely to be very busy. This will reduce capacity and must be taken into consideration *before* deployment.

[8] Tools such as iBwave offer outdoor analysis.

Survey Outcomes

An on-site survey should be able to provide, at a minimum, the following information:

1. To-scale maps with recommended AP locations
2. AP heat maps showing actual, measured coverage and recorded RSSI and SNR values
3. Per-channel spectrum assessment (how many and which channels are available and usable for this application)
4. Actual, measured throughput values at each AP location

It is important to note that an on-site survey will measure conditions as they exist at the time of the survey. Environmental conditions can change over time, e.g. trees grow and block LoS, other Wi-Fi networks are installed on the same channels, etc.

Challenges of Wireless Mesh to Network Designs

Although a site survey is recommended for any Wi-Fi network deployment, it is especially important for wireless mesh deployments. Because wireless mesh adds latency and jitter and reduces capacity as compared to a WLAN carried across wired APs, it is critical to understand what is possible and what is not. Environmental factors such as nearby Wi-Fi networks and transmitters, location of the AP, and physical objects in Line of Sight (LoS) all play a part.

Outdoor wireless networks have unique designs challenges compared with indoor designs:

- Lots of RF noise (channel optimization is more difficult)
- Foliage attenuation (especially evergreen trees), which may change seasonally and over the years
- Mounting, power, and backhaul can be difficult and expensive
- Fresnel zone issues on mesh and Point-to-Point (PtP) and Point-to-Multi-Point (PtMP) links

Calculations Behind the Network Design Process

AP Quantity

Once a site assessment has been performed, the next important task is to determine the quantity of APs required. This number will be determined by several factors, including:

- Coverage area
- Radio capacity
- Application requirements
- Distance required to get an acceptable signal quality

Radio Capacity

Airtime is Critical for Success

A common misconception of wireless networks is that the most important limitation of a network design is bandwidth. However, **for most wireless networks, the biggest limitation to scale is usually airtime utilization**. Airtime utilization is the percentage of time a device requires to transmit data. The busier a network is (higher airtime utilization) the less time there is for a device to transmit and the higher the chances are of interference. This parameter will vary as per the application (voice/video) being used over the wireless mesh network.

Estimating Airtime Utilization for a Device

Airtime utilization is heavily dependent on client behavior and local RF conditions which change constantly. It can be roughly estimated using the following formula:

```
Airtime Required (%) = (Throughput Required/Max Available Throughput) *100
```

It is important to note that Max. Available Throughput is the amount of throughput achievable by a device attached to the mesh AP, or goodput. This is very different from the 802.11 data rate or PHY.

Please note that for the last 2 factors - Application requirements and Distance required to get an acceptable signal quality, they are very specific to the application that's being run on the mesh network and cannot be generalized. This document gives you a general guideline. Please look at the specific documents in the next section for more details on video and voice applications.

Appendix C: Zero Touch Mesh

In the current mesh design to become mesh APs, APs need to have wired connection with controller to get provisioned for configurations including mesh configuration. Then the user can remove the wired connection and move the AP to the desired location. AP will find other APs to connect with using the pre-provisioned mesh configuration and form the mesh network.

Zero Touch Mesh is designed to reduce this manual procedure and allows the user to establish a mesh network in an easier manner. With Zero Touch Mesh the user can simply put the AP at the desired location, power it on. Then the AP automatically discovers the network, joins the controller and forms the mesh network. In short, the new AP can join the network via wireless without a wired connection.

It is used to simplify the mesh deployment, implement long distance wireless connections between networks, expand network coverage area, and reduce network deployment costs.

This feature is available from SmartZoneOS 5 and ZoneDirector 10.2 versions. On the SmartZone it can be enabled both at the domain and zone levels.

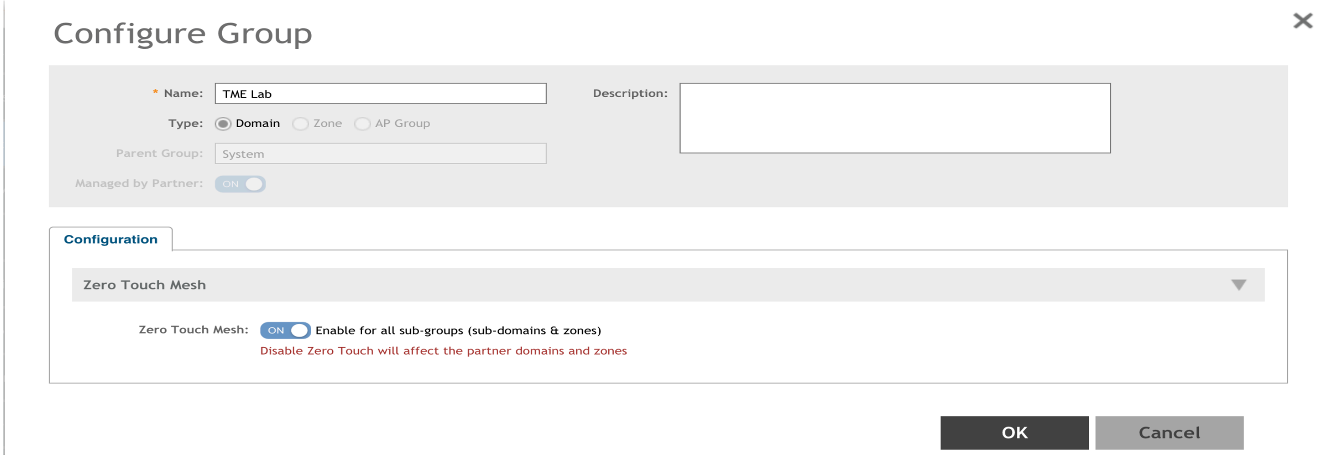

Configure Group

* Name: TME Lab

Description:

Type: ● Domain ○ Zone ○ AP Group

Parent Group: System

Managed by Partner: ON

Configuration

Zero Touch Mesh

Zero Touch Mesh: ON Enable for all sub-groups (sub-domains & zones)

Disable Zero Touch will affect the partner domains and zones

FIGURE 28A: ZERO TOUCH MESH – DOMAIN LEVEL

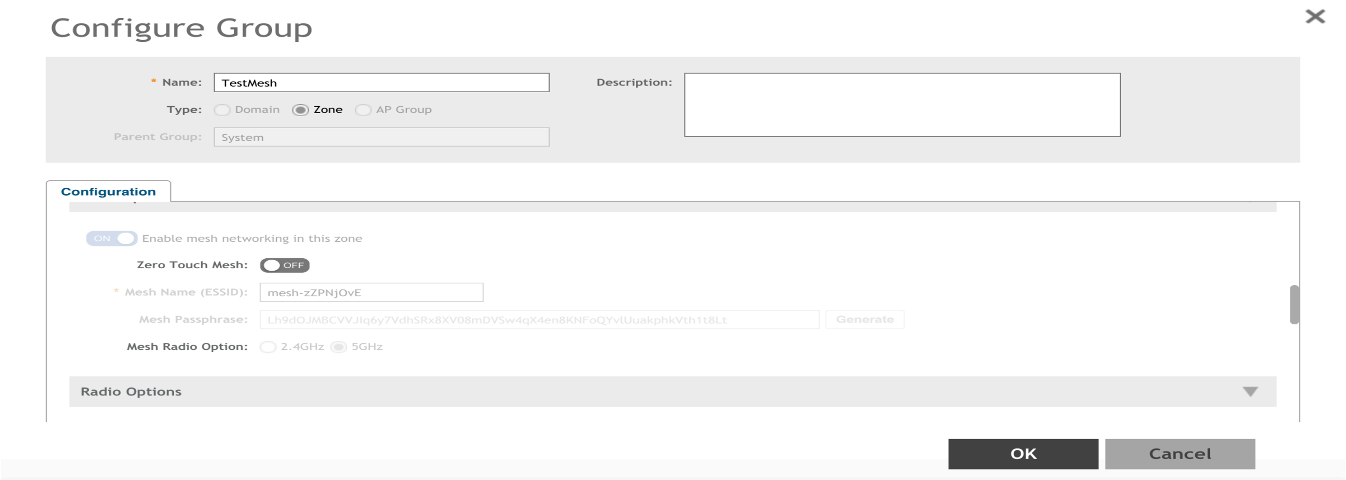

Configure Group

* Name: TestMesh

Description:

Type: ○ Domain ● Zone ○ AP Group

Parent Group: System

Configuration

ON Enable mesh networking in this zone

Zero Touch Mesh: OFF

* Mesh Name (ESSID): mesh-zZPNjOvE

Mesh Passphrase: Lh9dOJMBCVVJlq6y7VdhSRx8XV08mDVSw4qX4en8KNFoQYvlUuakphkVth1t8Lt

Mesh Radio Option: ○ 2.4GHz ● 5GHz

Radio Options

FIGURE 28B: ZERO TOUCH MESH – ZONE LEVEL

Appendix D: Additional Resources

IP Surveillance Video over a SmartMesh Network

BPDG IP video over Mesh - https://ruckuswireless.egnyte.com/dl/U4epDy6IDa/BPDG_IP_Video_Over_Mesh_Final.pdf_

Deploying Voice over IP for Wi-Fi

BPDG Voice over Wi-Fi - https://ruckuswireless.egnyte.com/dl/z5mtdBZa60/BPDG_Voice_over_IP.pdf_

About Ruckus Networks

Ruckus Networks enables organizations of all sizes to deliver great connectivity experiences. Ruckus delivers secure access networks to delight users while easing the IT burden, affordably. Organizations turn to Ruckus to make their networks simpler to manage and to better meet their users' expectations. For more information, visit www.ruckuswireless.com.

Ruckus Networks | 350 West Java Drive | Sunnyvale, CA 94089 USA | T: (650) 265-4200 | F: (408) 738-2065 ruckuswireless.com

About ARRIS

ARRIS International plc (NASDAQ: ARRS) is powering a smart, connected world. The company's leading hardware, software and services transform the way that people and businesses stay informed, entertained and connected. For more information, visit www.arris.com.

For the latest ARRIS news:

Check out our blog: ARRIS EVERYWHERE

Follow us on Twitter: @ARRIS