Traffic Management architectures and Service differentiators

Pronto vs. Typical WLANC architectures

In this document we describe the service deployment implications of two prevalent architectures typically found in the deployment of WLAN architectures. The primary solution under discussion is a typical WLANC architecture that utilizes WLAN controllers typically available in the market place coupled with the Pronto UniFi OSS. We also present an alternative architecture deployed using the Pronto UniFi OSS solution with the Pronto Hotzone Service Gateway (PSG) or Pronto Hotzone Service Controller (PSC).

Brief descriptions of the traffic management methods manifest in these solutions are provided here. Applicability of VLANs, SSIDs, and number of QoS levels, bandwidth partitions, and their impact on the service choices available to ISPs for metro Wi-Fi Hotzones and Wi- Fi Hotspots are discussed in this paper.

The intended audience is network planners and IT and service deployment architects of mixed-use networks, such as the ones found in Metro Wi-Fi networks that combine Public Access, Public Works, and Public Safety, and for Wi-Fi Hotspots for Enterprise Visitor Access or residential applications.

Scenario 1 - Generic WLANC

Considering a scenario where the Wi-Fi setup consists of a generic WLAN controller and back office is a simple AAA server, like RADIUS (Funk RADIUS etc) server, or the Pronto UniFi OSS.

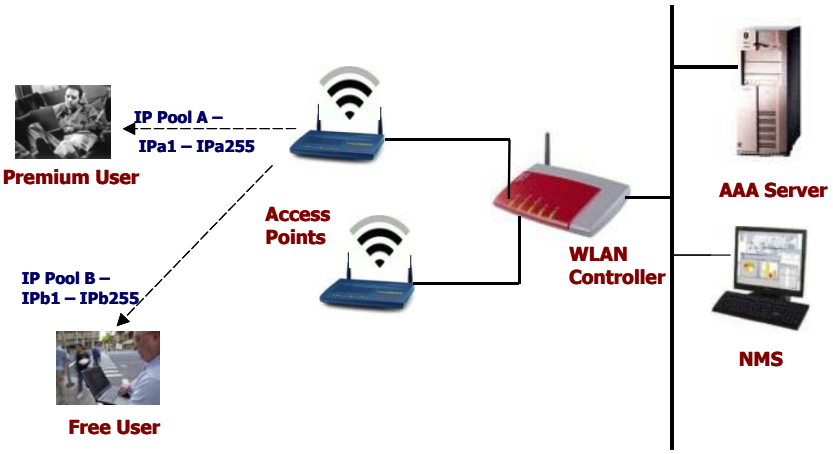

While setting up this solution, the WLAN controller is configured and separate pools are setup for different types of users. For example, Pool A for premium users, Pool B for free users, etc. IP addresses are assigned to these sets of pools, e.g. Pool A would have IP addresses range IPa1-IPa255, pool B would have IP addresses in the range Ipb1-Ipb255. WLAN controller divides WAN bandwidth into multiple bandwidth pools based on the IP address ranges.

When a user connects to access point, he is assigned a DHCP IP address from one of these pools. Based on which IP address he gets, he is assigned to a specific pool and he shares the bandwidth with other users in that pool. His user id/password is sent from the WLAN

controller to the AAA server-using RADIUS protocols. Once the user credentials are verified, he is assigned to the specific bandwidth pool. Once AAA server authenticates the user's credentials, the user is allowed to access the Internet, but the bandwidth pool is predetermined. So, the AAA server typically does not decide which Pool the user goes in, unless there are Vendor Specific Attributes, which are discussed in the following sections.

Bandwidth Queues

Some controllers can connect through multiple physical ports to multiple subnets in the network. This feature can be helpful when WLAN Solution operators want to confine multiple VLANs to separate subnets.

Many implementations support very few numbers of queues preconfigured in terms of allocated bandwidth on WAN, which have to be shared among all the users in a particular group connecting to the network. Further, these queues cannot share bandwidth with each other. For example, consider a WLAN controller, which can be preconfigured with 4 queues and these queues cannot share bandwidth from each other. So, while one can have 16 SSIDs in mesh AP, and they can be mapped to multiple VLANs, only 4 queues are available for these VLANs. So, those 16 VLANs would be mapped into these 4 queues, and would have to share the bandwidth.

In addition, once the user is assigned to one of the 4 QoS queues, all the VLANs (and all the users in the VLAN) contend for the same bandwidth. So, this kind of solution does not allow for per-user QoS. This imposes clear limitations on number of service plans that can be supported, and the guarantees that can be contracted to subscribers and other commercial users in the network. One can only separate different classes of users into separate Queues, but cannot do granulation beyond that.

8021.x

In case of 802.1X authentication, only 802.1X requests are allowed through a WLAN Controller until a decision is made to grant the access to the user. Throughout most of the 802.1X exchange, the WLAN controller ("the authenticator") is just a middleman, relaying EAP messages between the station ("the supplicant") and a RADIUS server ("the authentication server"). For example, the user station is asked to supply its identity, which the authenticator relays inside a RADIUS Access-Request. Based on the station's identity, the RADIUS server issues a RADIUS Access-Challenge, the content of which the authenticator relays to the station. And so on, until the RADIUS server makes a decision to accept or reject the access request. Once the user is authenticated, all LAN traffic can be relayed between the user station and the upstream network.

This 802.1X framework consolidates decision-making at the RADIUS server, so that ACLs no longer have to be individually configured into WLAN Controller/access point.

In 802.1x authentication also, similar method holds good, and since there are only limited number of queues, after authentication, all traffic would go into these queues only.

RADIUS VSAs

Vendor specific RADIUS attributes can be used to override the values already present in the wireless LAN profile. For example, in case of a Cisco WLAN controller, when the following

vendor-specific attributes are present in the RADIUS access accept message, the values override those present in the wireless LAN profile:

- QoS level
- 802.1p value
- VLAN interface name
- Access control list (ACL) name

The AAA server can return the VLAN number or name using the standard "RADIUS assigned VLAN name/number" feature defined in IETF RFC 2868 (RADIUS Attributes for Tunnel Protocol Support). To assign a wireless client to a particular VLAN, the AAA server can send the following attributes to the controller in the access accept message:

- IETF 64 (Tunnel Type): VLAN
- IETF 65 (Tunnel Medium Type): 802
- IETF 81 (Tunnel Private Group ID): VLAN # or VLAN Name String

Thus, if the AAA override is using RADIUS VSAs for assigning a particular user to a specific VLAN, it is essentially assigning a particular QoS level to the user. In this manner, the user can be moved to a specific bandwidth queue based on his Service Level Agreement registered at the AAA server or the Pronto UniFi OSS.

Scenario 2 – Pronto Service Controller/Gateway Architecture

In a Pronto based solution, set up will consist of Pronto Service Controller (PSC) and Pronto OSS. For the Pronto Service Controller to enforce QoS, it has to know what QoS to enforce for the specific IP address that is allocated to a user that is requesting access. i.e. VIP user gets QoS values (Qv1, Qv2,....) and a regular user gets QoS values Qr1, Qr2, etc.

These QoS values are associated with the user, not the IP address he gets from the mesh network. Thus, when a user logs in, the Pronto UniFi OSS communicates to the controllers what QoS values need to be enforced. Based on what the subscription/billing service level agreement information is associated with the user, the UniFi OSS will assign a service class to the user and it will get mapped to the QoS values. Thus, when a user enters the network, the Pronto WLAN Controller (via web services/RADIUS) sends the credentials to the Pronto UniFi OSS, and the UniFi OSS confirms to the WLAN Controller the (i) user's credentials and

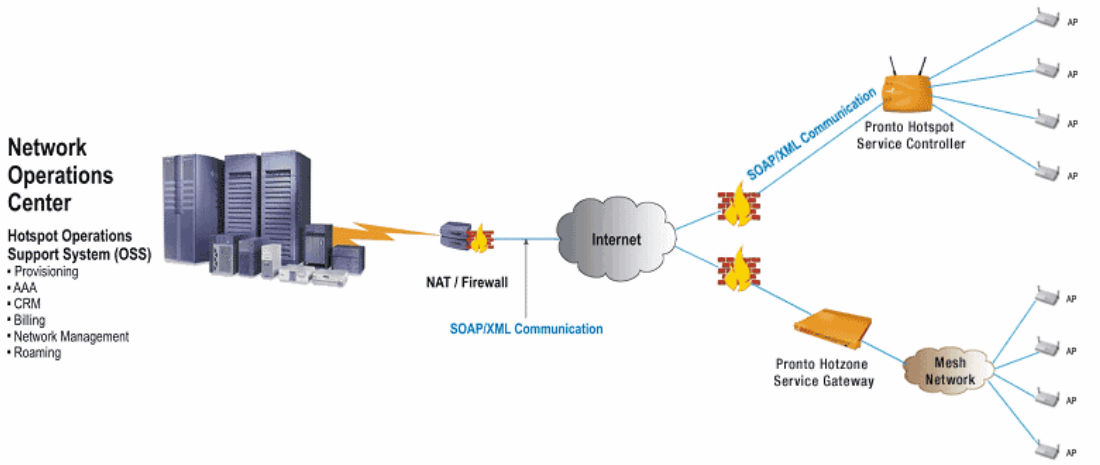

(ii) the QoS values that need to be enforced for the duration of the session. Traffic Management is implemented at the PSG/PSC and it includes:

- Bandwidth partitions
- SLA mapping to the defined partition

Bandwidth partitions can be of four types:

1) Bounded: A partition cannot borrow from any other partition. Thus, if a partition is bounded, then users of this partition are restricted by bandwidth allotted to it.
2) Unbounded: A partition can borrow from other partitions, subject to availability.
3) Isolated: A partition does not allow other partitions to borrow from it. Thus, if users of this type of partition were not using the bandwidth, then that bandwidth would go waste.
4) Shared (or not isolated): A partition allows other partitions to borrow from it. This would result in practically no wastage of the bandwidth when there is a demand for it.

This bandwidth partitioning done at the WISP level can be by percentages or actual bandwidth. This allows the system to apply different treatment strategies to different flows (session) of Internet access, e.g., a product plan can attach different bandwidth limits to applications/services being used. For specific services like emails (POP), browsing (http), download music (ftp) and talking to another remote user (VoIP), each of these services can be assigned their own SLAs (Bandwidth limits) simultaneously.

Pronto's OSS supports Quality of Service guarantees at the User level by allowing the Service provider to enforce SLAs on upstream and downstream bandwidth rates (minimum and maximum). The minimum rate defines a sustained level, and the maximum rate sets the peak level. Once these service plans are defined in the OSS, when a user subscribes to the service, he is assigned that QoS SLA. This QoS is associated with the user, not the PHC or the PHG. So, when the user attempts to login, he is assigned the QoS. Any number of SLAs can be defined in the system. The service provider on a PER PSC/PSG basis can set the values of the bandwidth rates.

All users that subscribed to the Service Plan are subject to the enforcement of this SLA upon successful login. When the available bandwidth on the PSC/PSG is committed to authenticated users, the OSS provides an option whether the next user that tries to login get his "SLA enforced". If the SLA is enforced, even a registered user is denied access to the network, since the PSG/PSC cannot meet the SLA commitments, as bandwidth is already committed to other users. If the SLA is not enforced, additional users are allowed on to the network.

Service Implications

Feature	Typical WLANC w Pronto UniFi OSS	PSC/PSG w Pronto UniFi OSS
Multiple VLANs	Yes	Yes
Multiple IP subnets within a single VLAN	No	No
Each VLAN has its own Bandwidth Partition	No	Yes

Bandwidth Priority Queues	Yes (0,1,2,3 based on Voice, Data, etc.)	Yes
Multiple Priority Queues within same VLAN	No	Yes
Share Bandwidth from other queues	No	Yes
Number of QoS levels	Fixed	Unlimited

Service Deployment choices

Metro Markets
Thus, in a typical deployment involving a typical WLAN Controller with Pronto UniFi OSS or PSC/PSG with Pronto UniFi OSS in a metro deployment, one would normally be able to support a mixed use environment by using multiple SSIDs within the mesh, some used for Public Safety, other SSID for Public Works, and one SSID for Captive Portal based Public Access etc.

In case of a typical WLAN Controller deployment, this would be possible by dividing users into multiple VLANs. These VLANs can be mapped to different SSIDs. Now, the users in these VLANs would be mapped into the available QoS/Bandwidth queues. Once the user is assigned to one of these queues, he/she will contend with other users in the same queue for bandwidth.

Thus, because of the fixed number of QoS levels and thus fixed number of bandwidth priority queues which cannot share bandwidth from each other a limited number of service plans can be implemented reducing the choices available for the ISP. In addition, since unused bandwidth cannot be shared between the queues, "ruthless preemption" for Public Safety traffic, and "drinking fountain" model for free-users becomes more difficult to implement.

Whereas, in a PSC/PSG deployment with Pronto UniFi OSS, QoS values are associated with the user, not the IP address he gets from the mesh network making a per-user QoS deployment possible. Users mapped to different VLANs and thus different SSIDs are subject to specific bandwidth partitions that can share bandwidth from each other thus differentiating the priority levels of users. For example, a Public Safety Employee user can use the bandwidth from a partition being used by a drinking fountain user if required.

Wi-Fi Hotspots
Similarly for regular Hotspots using a typical WLAN Controller, where there would be much fewer SSIDs, probably just one or two, mapping to users as Premium Internet Access users, VoIP users, enterprise users, etc, a fixed number of bandwidth queues will be available which will not be able to share bandwidth from each other. Further, users in a particular queue will contend for the available bandwidth, and each user cannot be guaranteed their individual bandwidth.

In a PSC/PSG scenario, however, these users can be placed in a different shareable bandwidth partitions thus allowing a Premium Internet access user to share bandwidth from a free user's partition. This allows for additional service options to be deployed at a hotspot or across the Wi-Fi network.