Layer 3 Roaming

Introduction

Large WLAN networks (for example, those found on large campuses) may require IP session roaming at layer 3 to enable application and session persistence while a mobile client roams across multiple VLANs. For example, when a user on a VoIP call roams between APs on different VLANs without layer 3 roaming, the user's session will be interrupted as the external server must re-establish communication with the client's new IP address. During this time, a VoIP call will noticeably drop for several seconds, providing a degraded user experience. In smaller networks, it may be possible to configure a flat network by placing all APs on the same VLAN.

However, on large networks filled with thousands of devices, configuring a flat architecture with a single native VLAN may be an undesirable network topology from a best practices perspective; it may also be challenging to configure legacy setups to conform to this architecture. A turnkey solution designed to enable seamless roaming across VLANs is therefore highly desirable when configuring a complex campus topology. Using Meraki's secure auto-tunneling technology, layer 3 roaming can be enabled using a Mobility Concentrator, allowing for bridging across multiple VLANs in a seamless and scalable fashion.

Typical Campus Architecture

Large campuses are often designed with a multi-VLAN architecture to segment broadcast traffic. Typically, network best practices dictate a one-to-one mapping of an IP subnet to a VLAN, e.g., client devices joining VLAN 10 will be assigned an IP address out of the subnet range 10.0.10.0/24. In this design, clients in different VLANs will receive IP addresses in different subnets via a DHCP server. Multi-VLAN architectures can vary to include multiple subnets within a building (e.g., one for each floor/area), or multiple subnets across a large site (e.g., one for each building/region in a large campus or enterprise environment).

As seen in the diagram below, the typical campus architecture has the core L3 switch connected to multiple L3 distribution switches (one per site), with each distribution switch then branching off to L2 access switches configured on different VLANs. In this fashion, each site is assigned a different VLAN to segregate traffic from different sites. Without an L3 roaming service, a client connected to an L2 access switch at Site A will not be able to seamlessly roam to a L2 access switch connected to Site B. Upon associating with an AP on Site B, the client would obtain a new IP address

from DHCP service running on the Site B scope. In addition, a particular route configuration or router NAT may also prevent clients from roaming, even if they do retain their original IP address.

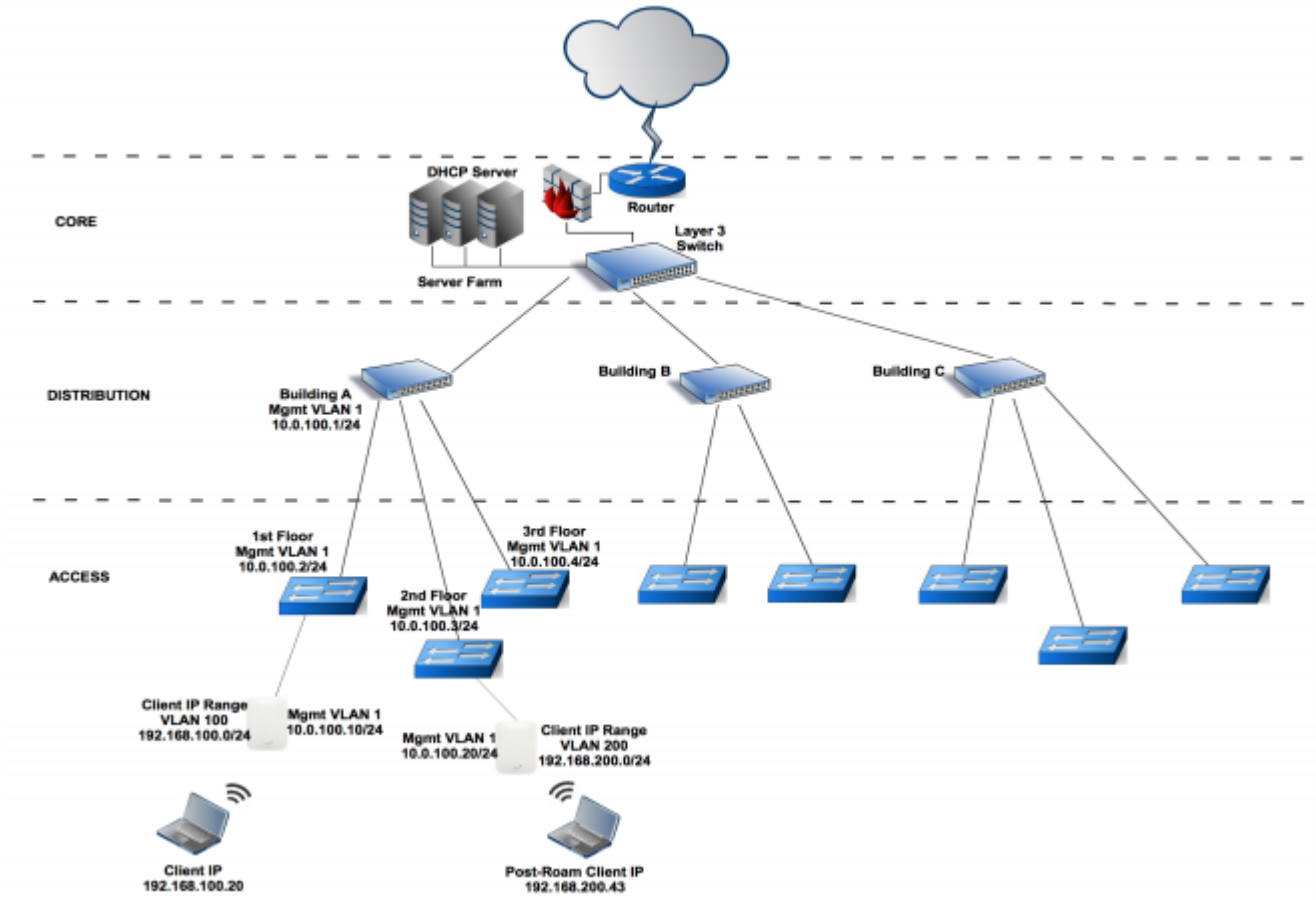

With layer 3 roaming, a client device must have a consistent IP address and subnet scope as it roams across multiple APs on different VLANs/subnets. Meraki's auto-tunnelling technology achieves this by creating a persistent tunnel between the L3 enabled APs and depending on the architecture a Mobility Concentrator. The two Layer 3 roaming architectures are discussed in detail below.

Distributed Layer 3 Roaming

Distributed Layer 3 Roaming maintains layer 3 connections for end devices as they roam across layer 3 boundaries without a concentrator. The first access point that a device connects to will become the anchor Access Point. The anchor access point informs all of the other Meraki Access Points within the network that it is the anchor for a particular client. Every subsequent roam to another access point will place the device/user on the VLAN that defined by the anchor AP.

Distributed Layer 3 roaming is very scalable because the Access Points are establishing connections with each other without the need for a concentrator. The target access point will lookup in the shared user database and contact the anchor access point. This communication does not traverse the Meraki Cloud and is a proprietary protocol for secure access point to access point communication.

A client's anchor AP will timeout after the client has left the network for 30 seconds.

VLAN Testing and Dynamic Configuration

The anchor access point runs a test to the target access point to determine if there is a shared layer 2 broadcast domain for every client serving VLAN. If there is a VLAN match on both access points, the target access point will configure the device for the VLAN without establishing a tunnel to the anchor. This test will dynamically configure the VLAN for the roaming device despite the VLAN that is configured for the target access point and the clients served by it. If the VLAN is not found on the target AP either because it is pruned on the upstream switchport or the Access Point is in a completely separated layer 3 network, the Tunneling method described below will be used.

Local VLAN testing and dynamic configuration is one method used to prevent all clients from tunneling to a single anchor AP. To prevent excess tunneling the layer 3 roaming algorithm determines that it is able to place the user on the same VLAN that the client was using on the anchor AP. The client in this case does a layer 2 roam as it would in bridge mode.

Tunneling

If necessary, the target Access Point will establish a tunnel to the anchor Access Point. Tunnels are established using Meraki-proprietary access point to access point communication. To load balance multiple tunnels amongst multiple APs, the tunneling selector will choose a random AP that has access to the original broadcast domain the client is roaming from. If the target AP detects a connectivity failure to the currently selected anchor AP, as a failover mechanism the target AP will choose a new anchor AP.

All APs must be able to communicate with each other via IP. This is required both for client data tunneling and for the distributed database. If a target access point is unable to communicate with the anchor access point the layer 3 roam will time out and the end device will be required to DHCP on the new VLAN.

Fast roaming protocols such as OKC and 802.11r are not currently supported with distributed layer 3 roaming. The best roaming performance will be using Layer 2 roaming with 802.11r.

Concentrator based Layer 3 Roaming

Any client that is connected to a layer 3 roaming enabled SSID is automatically bridged to the Meraki Mobility Concentrator. The Mobility Concentrator acts as a focal point to which all client traffic will be tunneled and anchored when the client moves between VLANs. In this fashion, any communication data directed towards a client by third party clients or servers will appear to originate at this central anchor. Any Meraki MX can act as a Concentrator, please refer to the MX sizing guides to determine the appropriate MX appliance for the expected users and traffic.

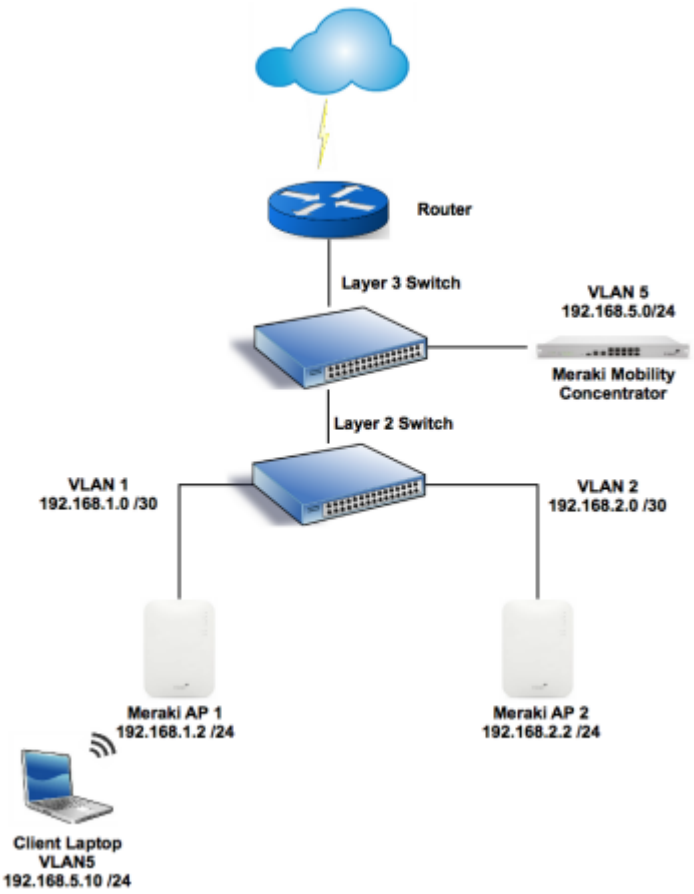

The diagram below shows the traffic flow for a particular flow within a campus environment using the Layer 3 roaming with concentrator.

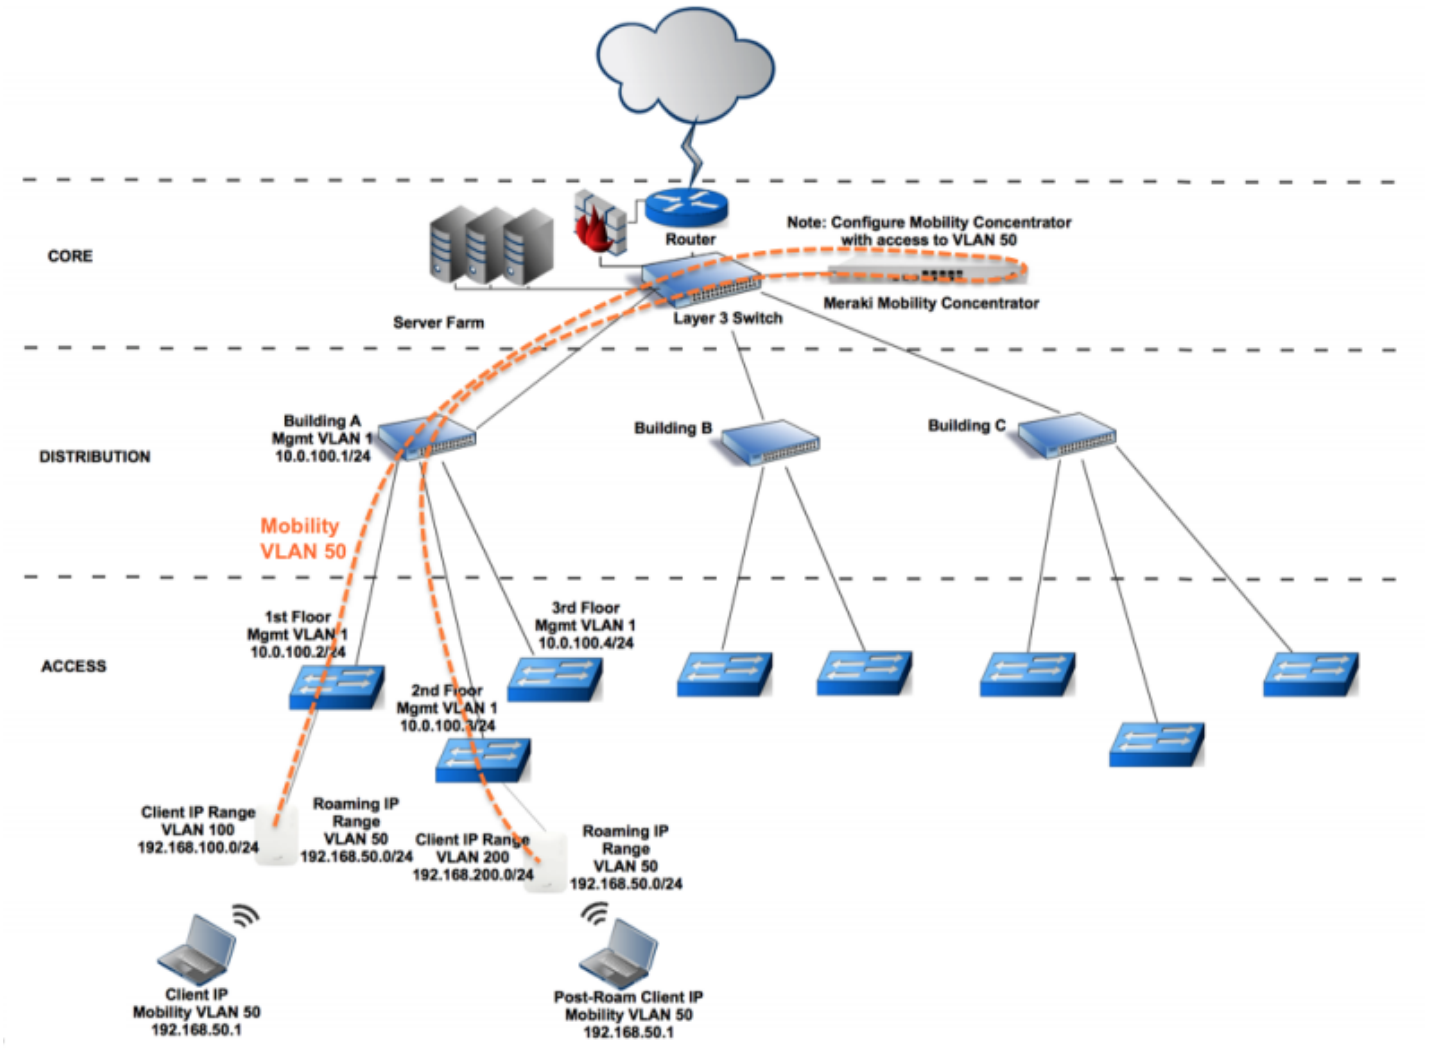
CORE
Router
Note: Configure Mobility Concentrator
with access to VLAN 50
Meraki Mobility Concentrator
Server Farm
Layer 3 Switch
DISTRIBUTION
Building A
Mgmt VLAN 1
10.0.100.1/24
Building B
Building C
Mobility
VLAN 50
ACCESS
1st Floor
Mgmt VLAN 1
10.0.100.2/24
3rd Floor
Mgmt VLAN 1
10.0.100.4/24
2nd Floor
Mgmt VLAN 1
10.0.100.3/24
Client IP Range
VLAN 100
192.168.100.0/24
Roaming IP
Range
VLAN 50
192.168.50.0/24
Client IP Range
VLAN 200
192.168.200.0/24
Roaming IP
Range
VLAN 50
192.168.50.0/24
Client IP
Mobility VLAN 50
192.168.50.1
Post-Roam Client IP
Mobility VLAN 50
192.168.50.1