Understand ARP Flooding and ARP Gleaning in ACI

Contents

Introduction

This document describes the use of Address Resolution Protocol (ARP) flooding and ARP gleaning in the Application Centric Infrastructure (ACI) fabric.

Understanding ARP Flooding

In Cisco ACI, there is an option to use the ARP flooding or disable it when needed. It is mandatory to know the fabric behavior regarding the ARP flooding so you can troubleshoot the Layer 2 issues.

If ARP flooding is enabled, ARP traffic is flooded inside the fabric as per regular ARP handling in traditional networks. ARP flooding is required when you need Gratuitous ARP (GARP) requests to update host ARP caches or router ARP caches. This is the case when an IP address can have a different MAC address (for example, with clustering of failover of load balancers and firewalls).

If ARP flooding is disabled, the fabric attempts to use unicast to send the ARP traffic to the destination. Therefore, a Layer 3 lookup occurs for the target IP address of the ARP packet. ARP behaves like a Layer 3 unicast packet until it reaches the destination leaf switch.

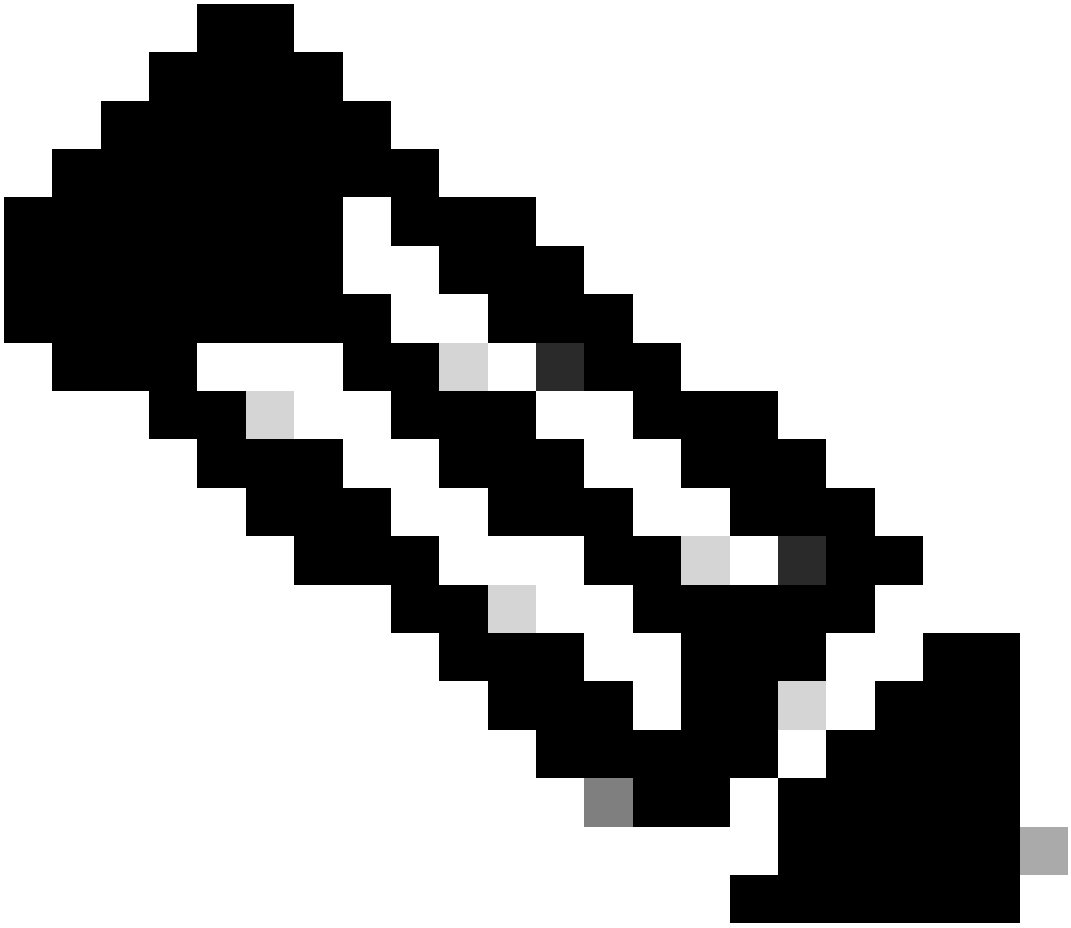

Note: Note that this option applies only if unicast routing is enabled on the bridge domain. If unicast routing is disabled, ARP flooding is implicitly enabled.

Next, you see a few use cases with respect to the use of ARP flooding.

Use Case 1. Endpoints are Learnt in ACI

This use case applies when both the endpoints are known to the leaf switch.

There is no role of ARP flooding in this scenario. The traffic is locally switched when the leaf switch knows its endpoint information. This behavior is the same, when one endpoint, such as H1, sends an ARP request towards the other (H2), and ARP flooding is disabled. Since the leaf switch knows where H2 is connected and checks the ARP target IP address (which is an H2 IP address), there is no need to flood the traffic or redirect it to the spine layer. Therefore, it sends the ARP request toward H2.

Regardless of the End-Point Group (EPG), Bridge domain, or Access/Encapsulation settings, If the endpoints are known to the leaf, they are treated in the same manner.

Example 1. Endpoints known to the fabric, working in the same EPG, Bridge domain, and Access/Encapsulation.

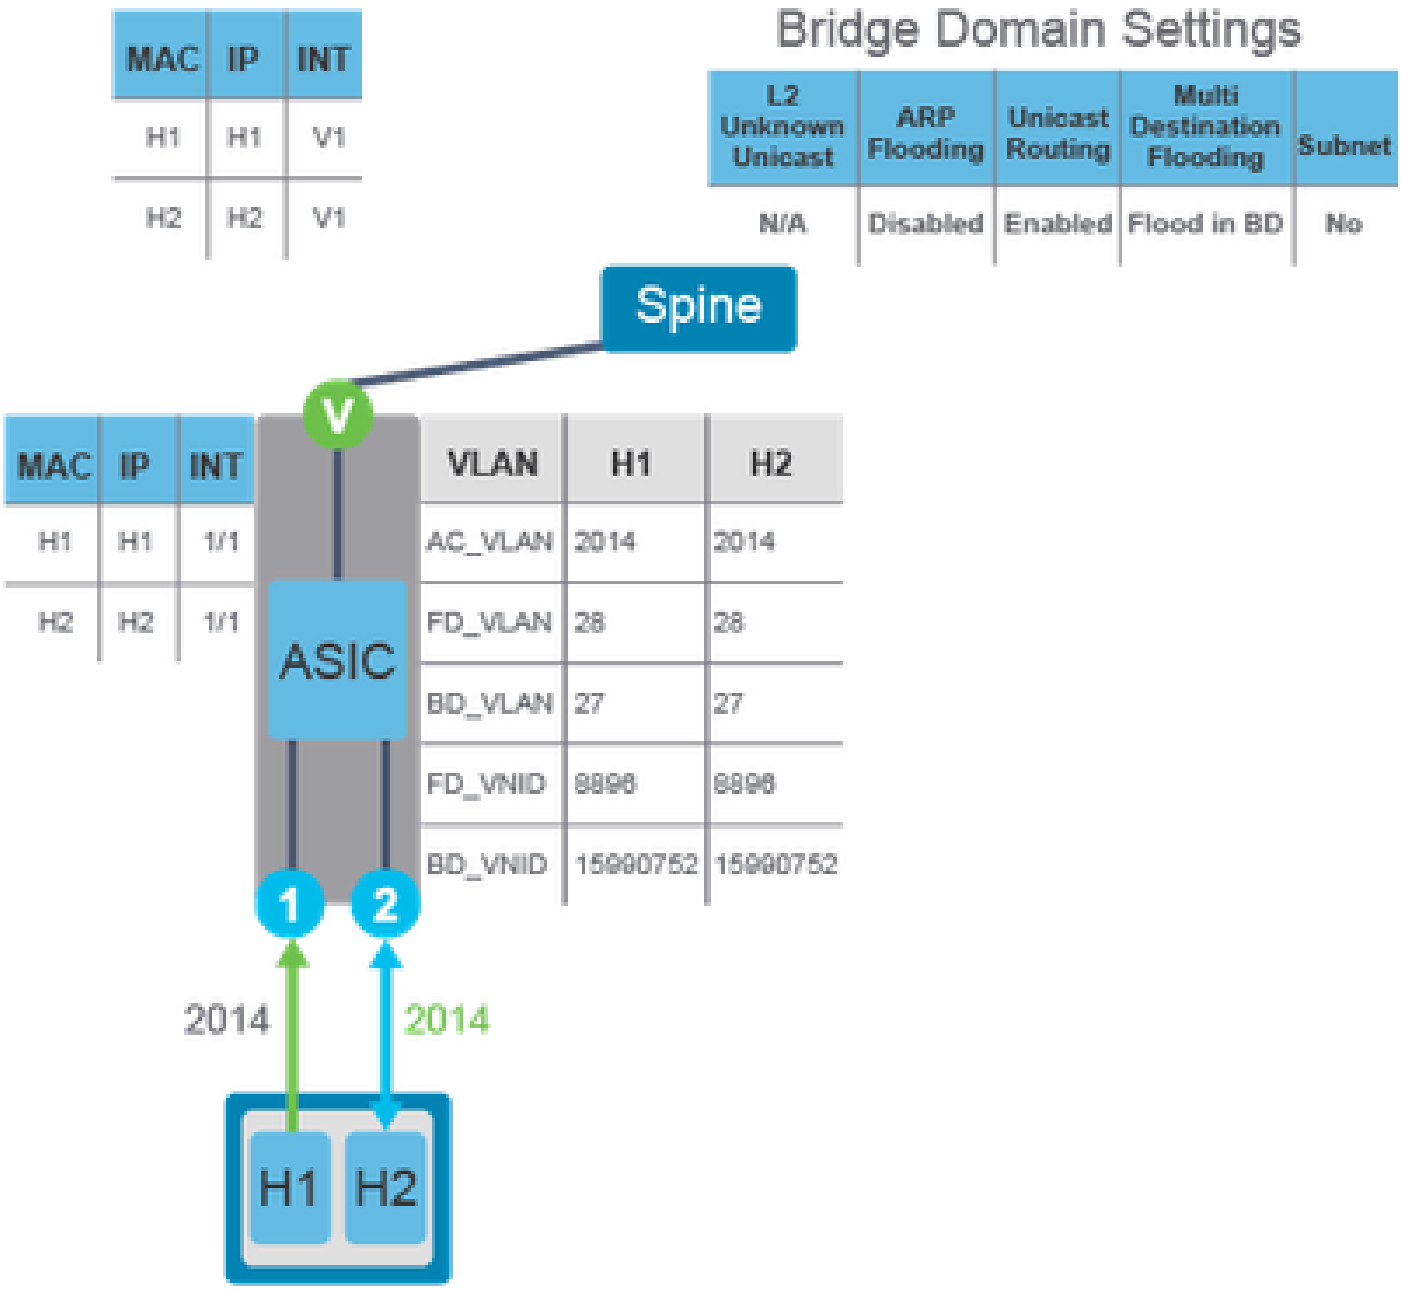

Example 2. Endpoints known to the fabric, working in the same EPG, Bridge domain but different Access/Encapsulation.

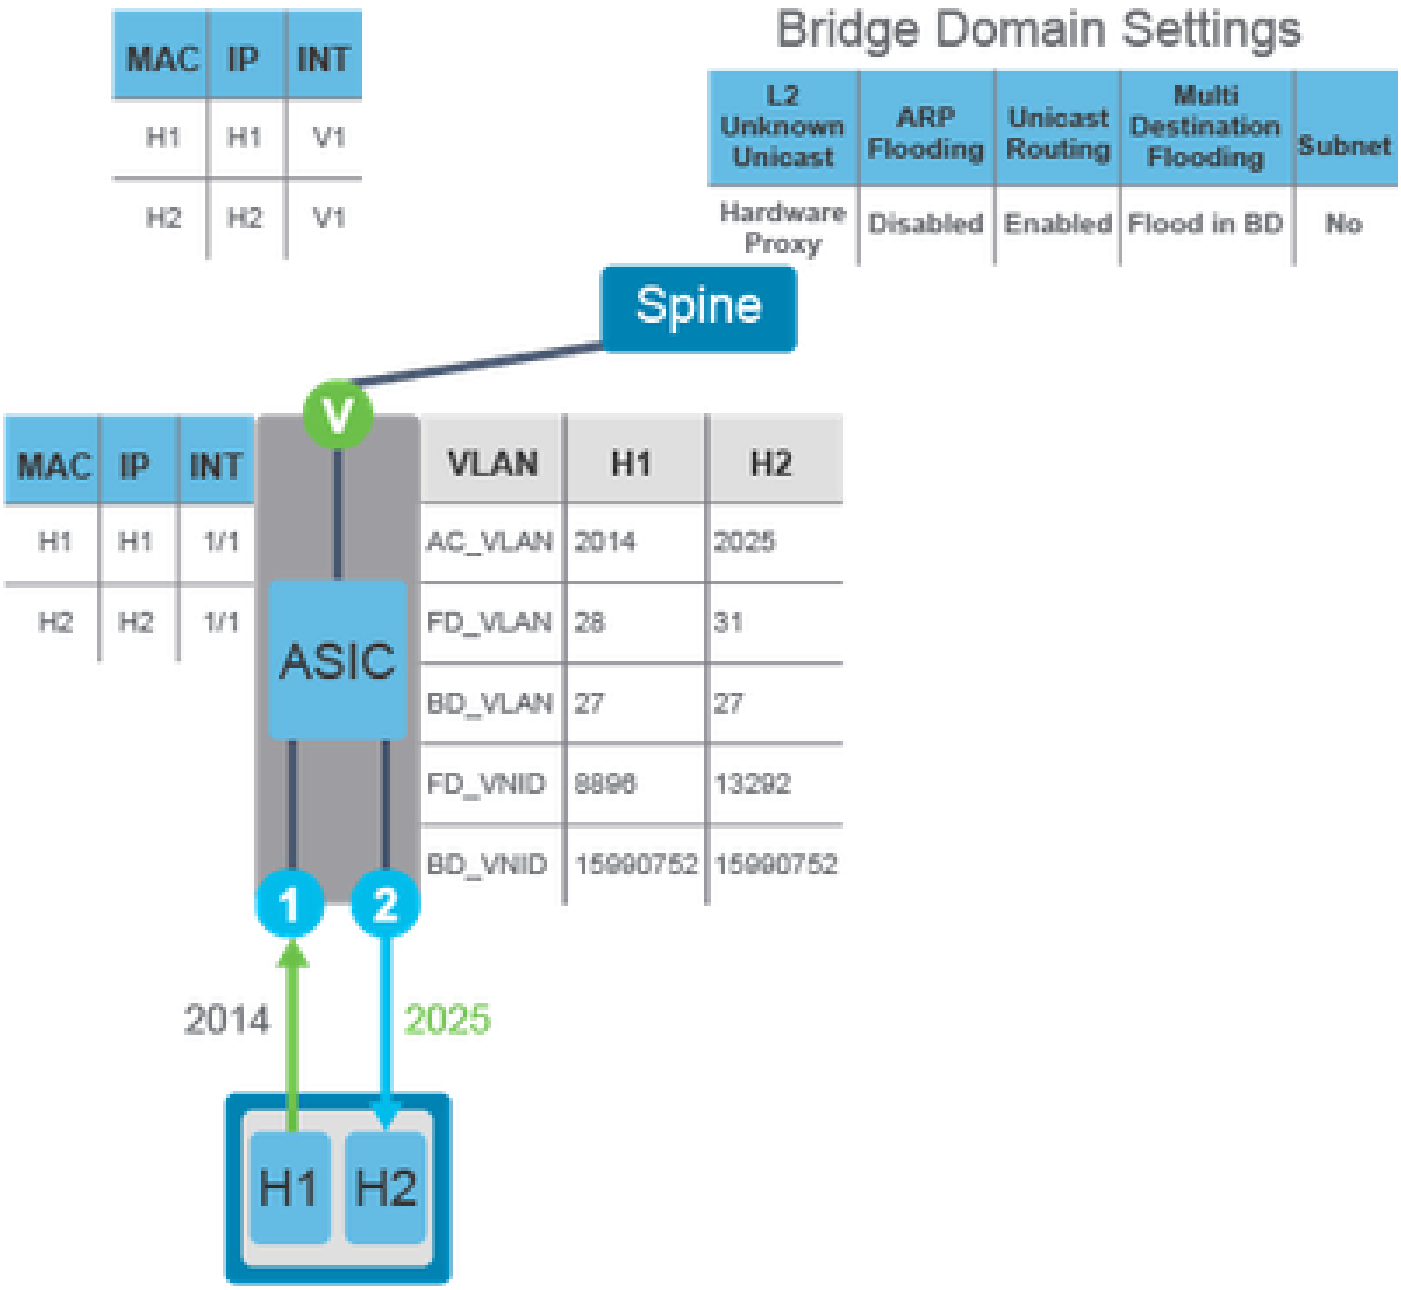

Example 3. Endpoints known to the fabric, working in different EPGs but the same Bridge Domain.

When ARP flooding is disabled and the endpoints are part of the different EPGs in the same bridge domain, while connected to the same leaf switch, the ARP traffic is locally routed if the leaf switch knows the ARP target IP address (unicast routing is enabled).

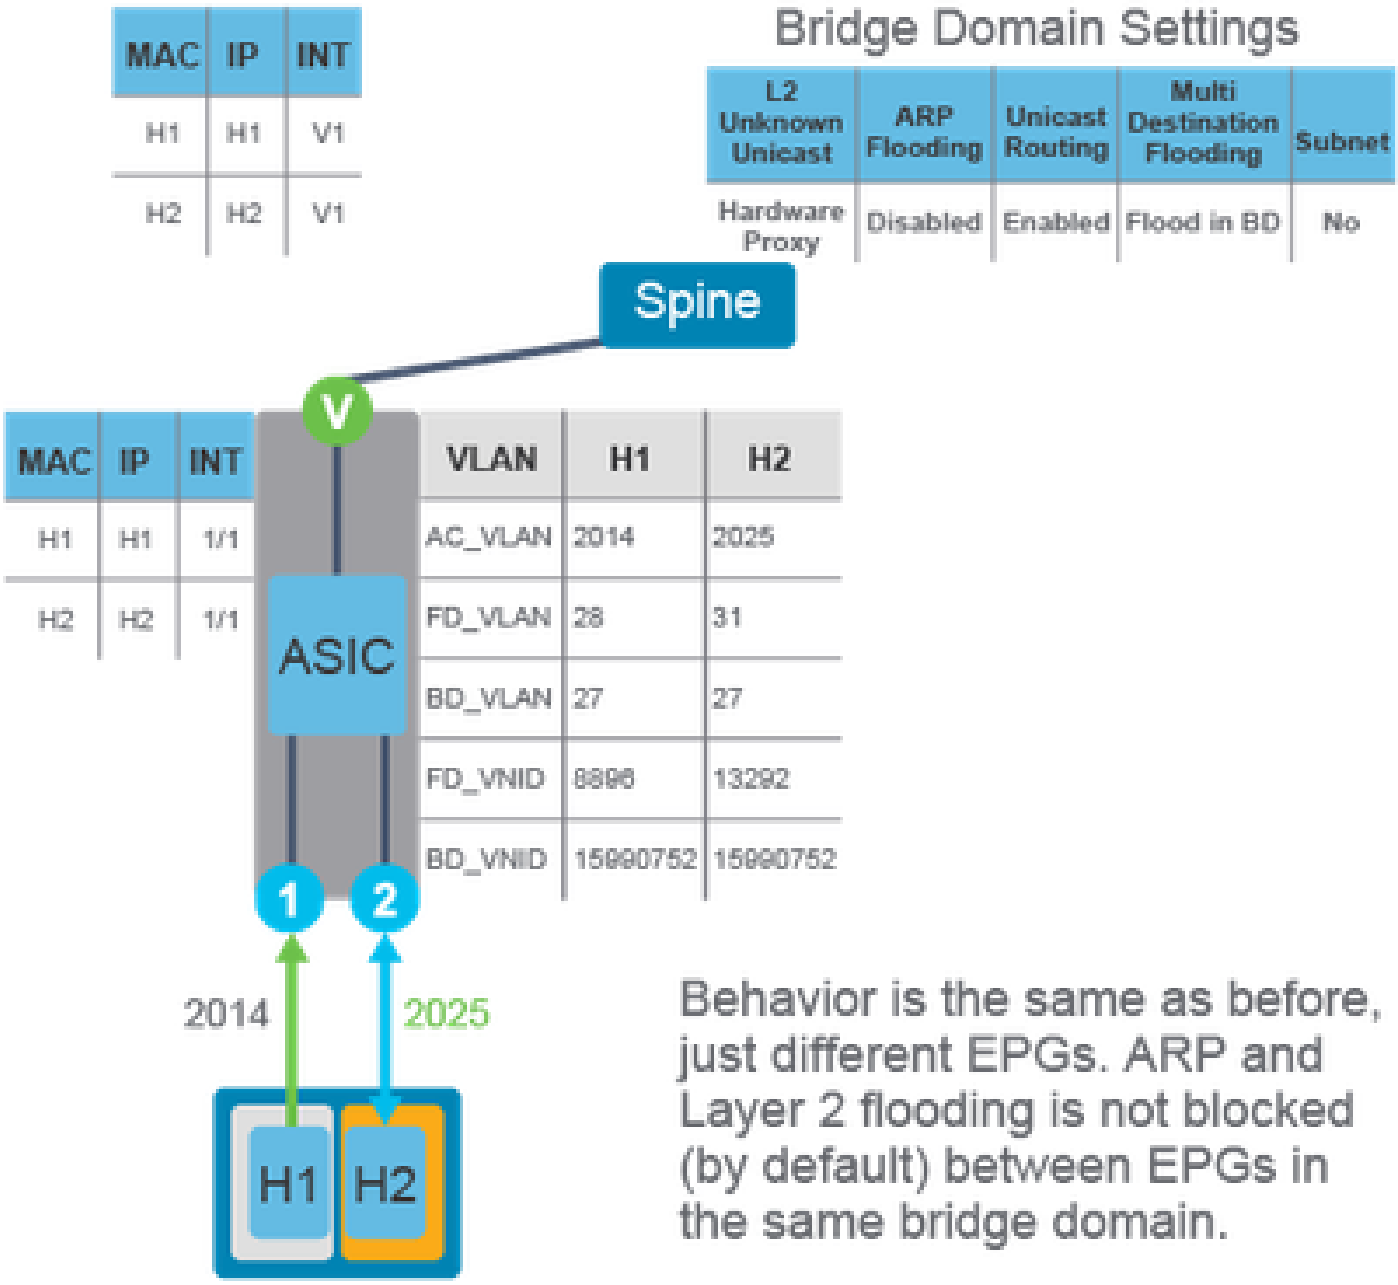

Use Case 2. Endpoints are Learnt in COOP

This use case applies when both the endpoints are connected to different leaf switches; present in the Cooperative Protocol (COOP) database of Spine Switch.

ARP request has to be forwarded across the fabric. The flow of ARP traffic from H1 to H3 is:

- H1 sends an ARP request for H3 using a broadcast destination MAC.
- The ACI attempts to use unicast forwarding in order to send the ARP request, so the local leaf switch checks the ARP target IP address, which is the H3 IP address. Since the local leaf switch does not know the IP address of the endpoint H3, it sends the ARP request to the spine switch for spine-proxy.
- The spine has the H3 information in the COOP database (unicast routing is enabled) and forwards the ARP request to the destination leaf switch across the fabric, which forwards it to H3. Once H3 receives the traffic, it replies to H1.

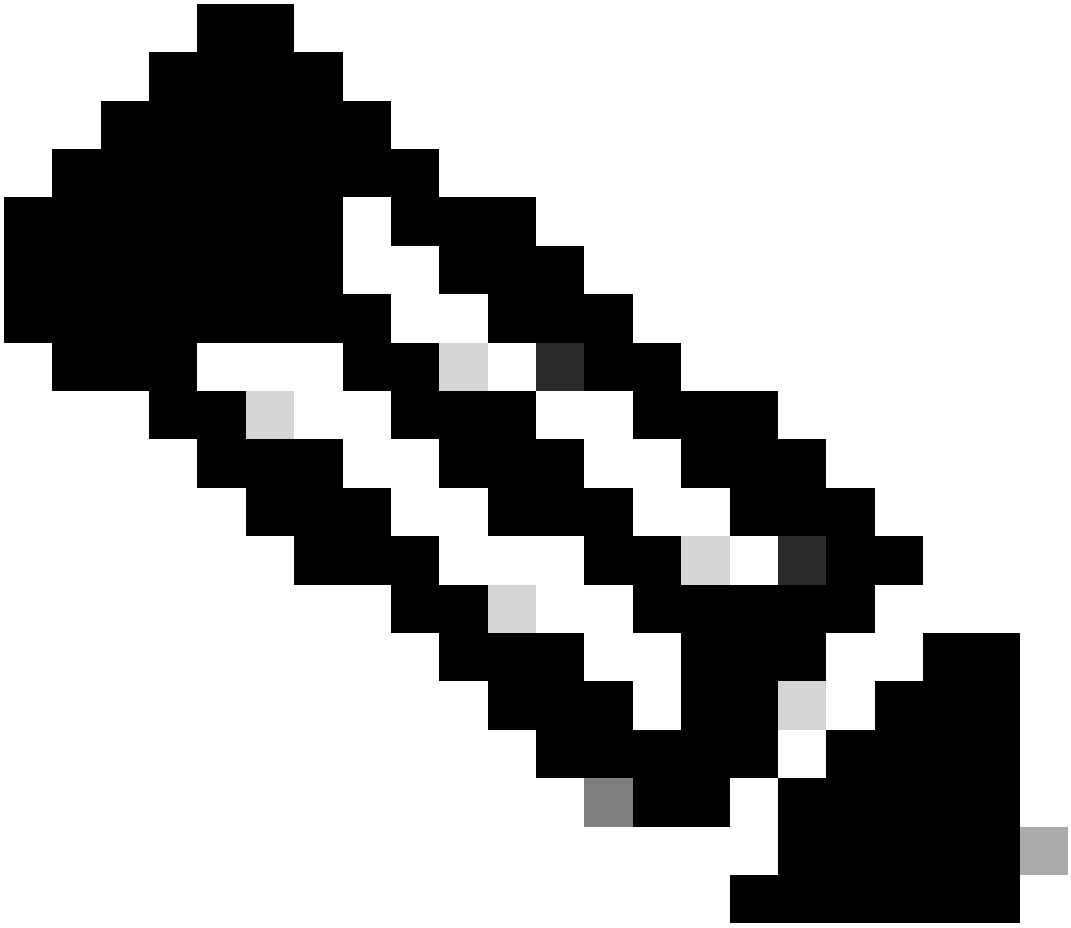

Note: The mentioned mechanism applies to all three scenarios.

Example 1. Endpoints known to the fabric, working in same EPG, Bridge domain, and Access/Encapsulation.

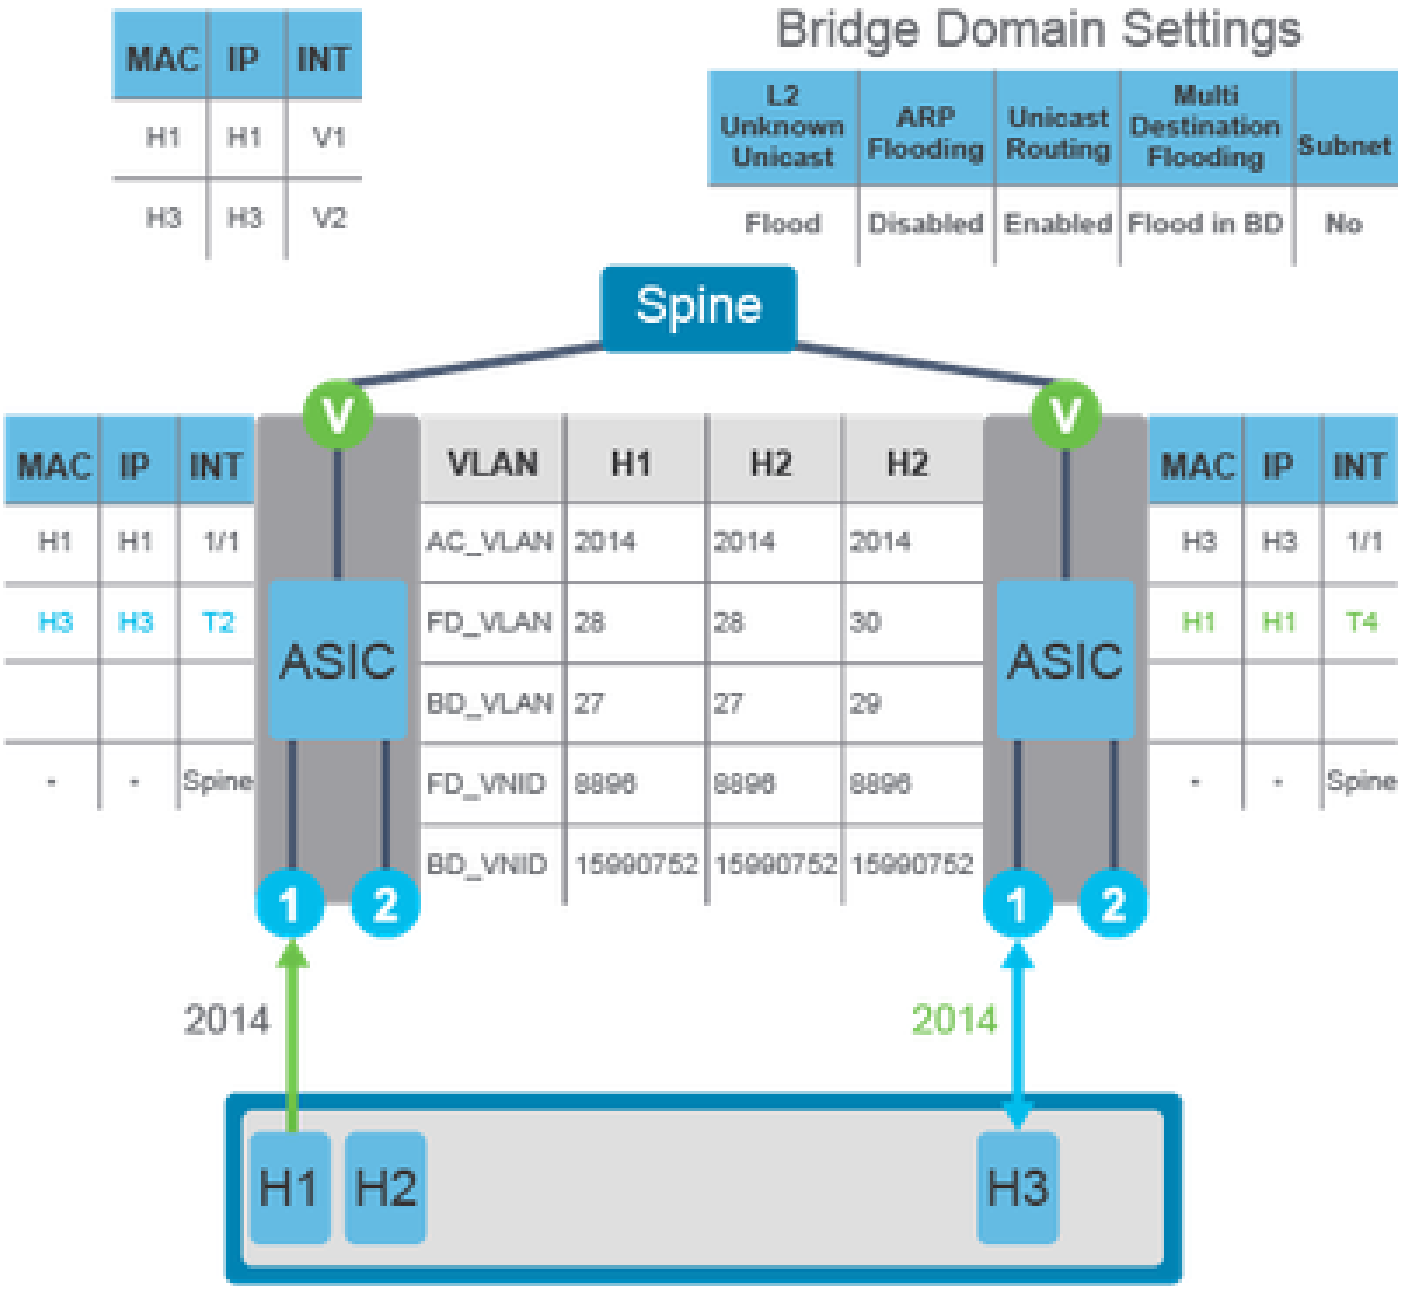

Example 2. Endpoints known to the fabric, working in the same EPG, Bridge domain but different Access/Encapsulation.

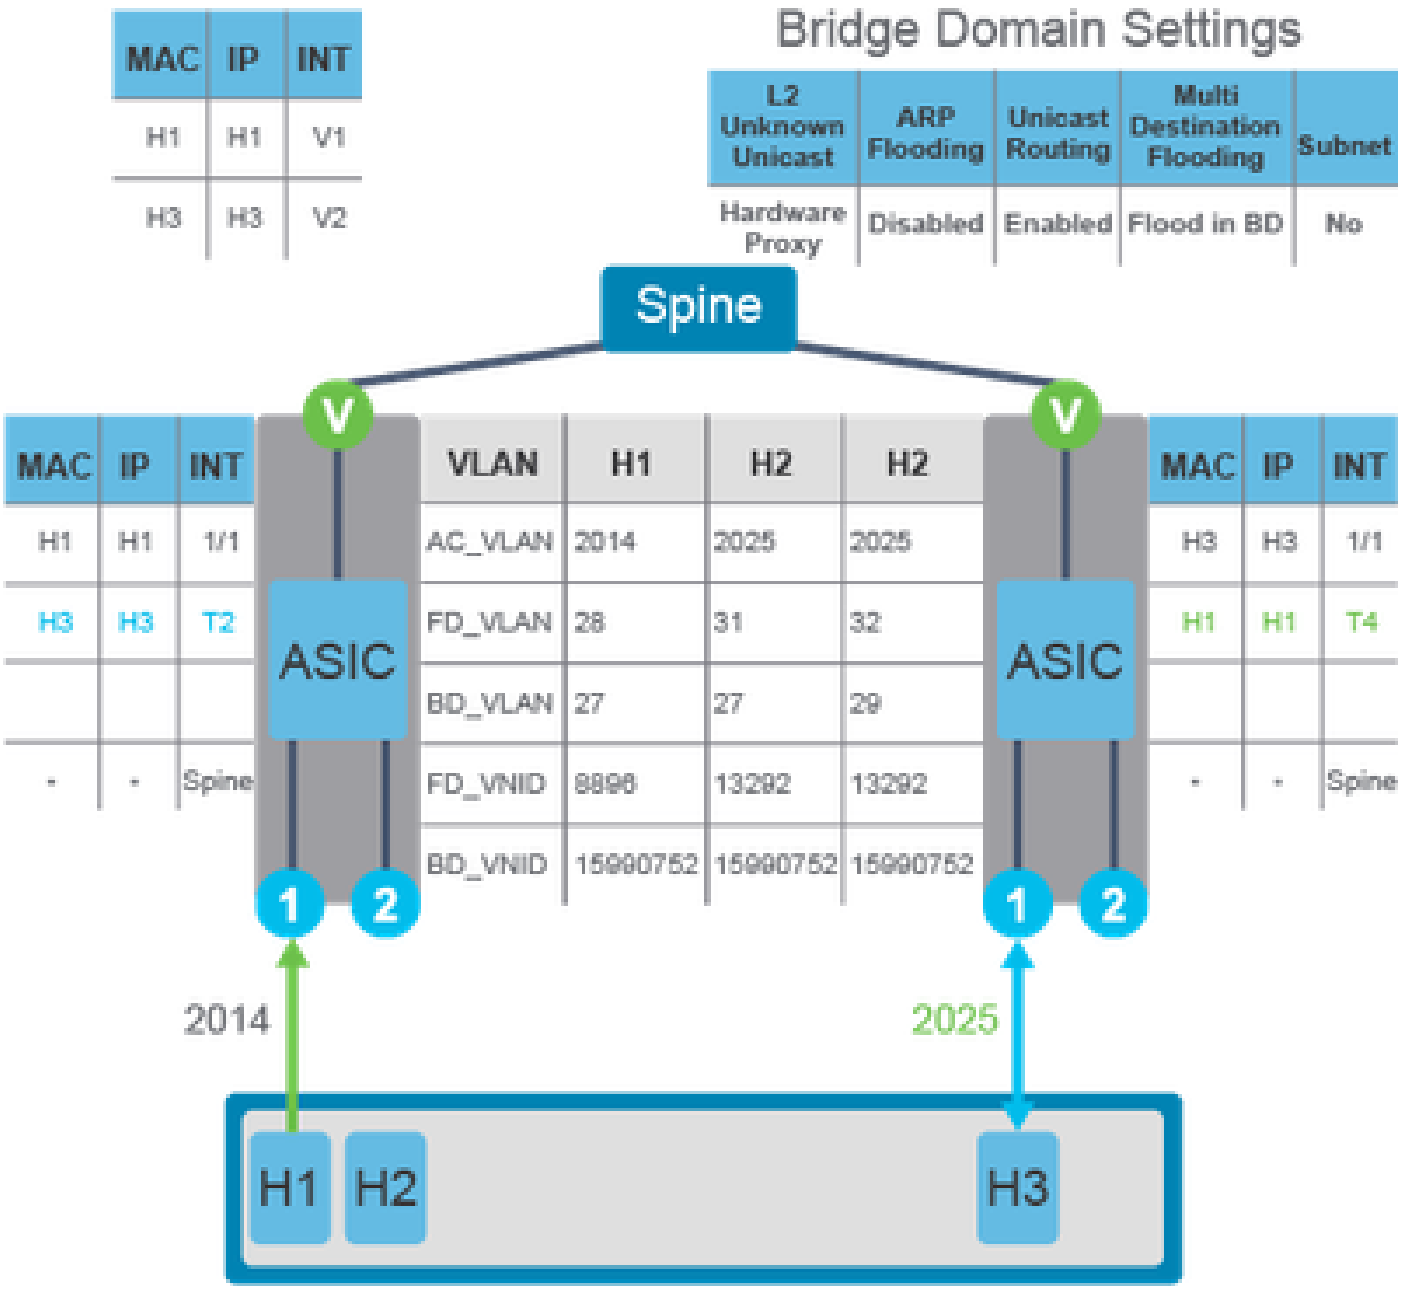

Example 3. Endpoints known to the fabric, working in different EPGs but the same Bridge Domain.

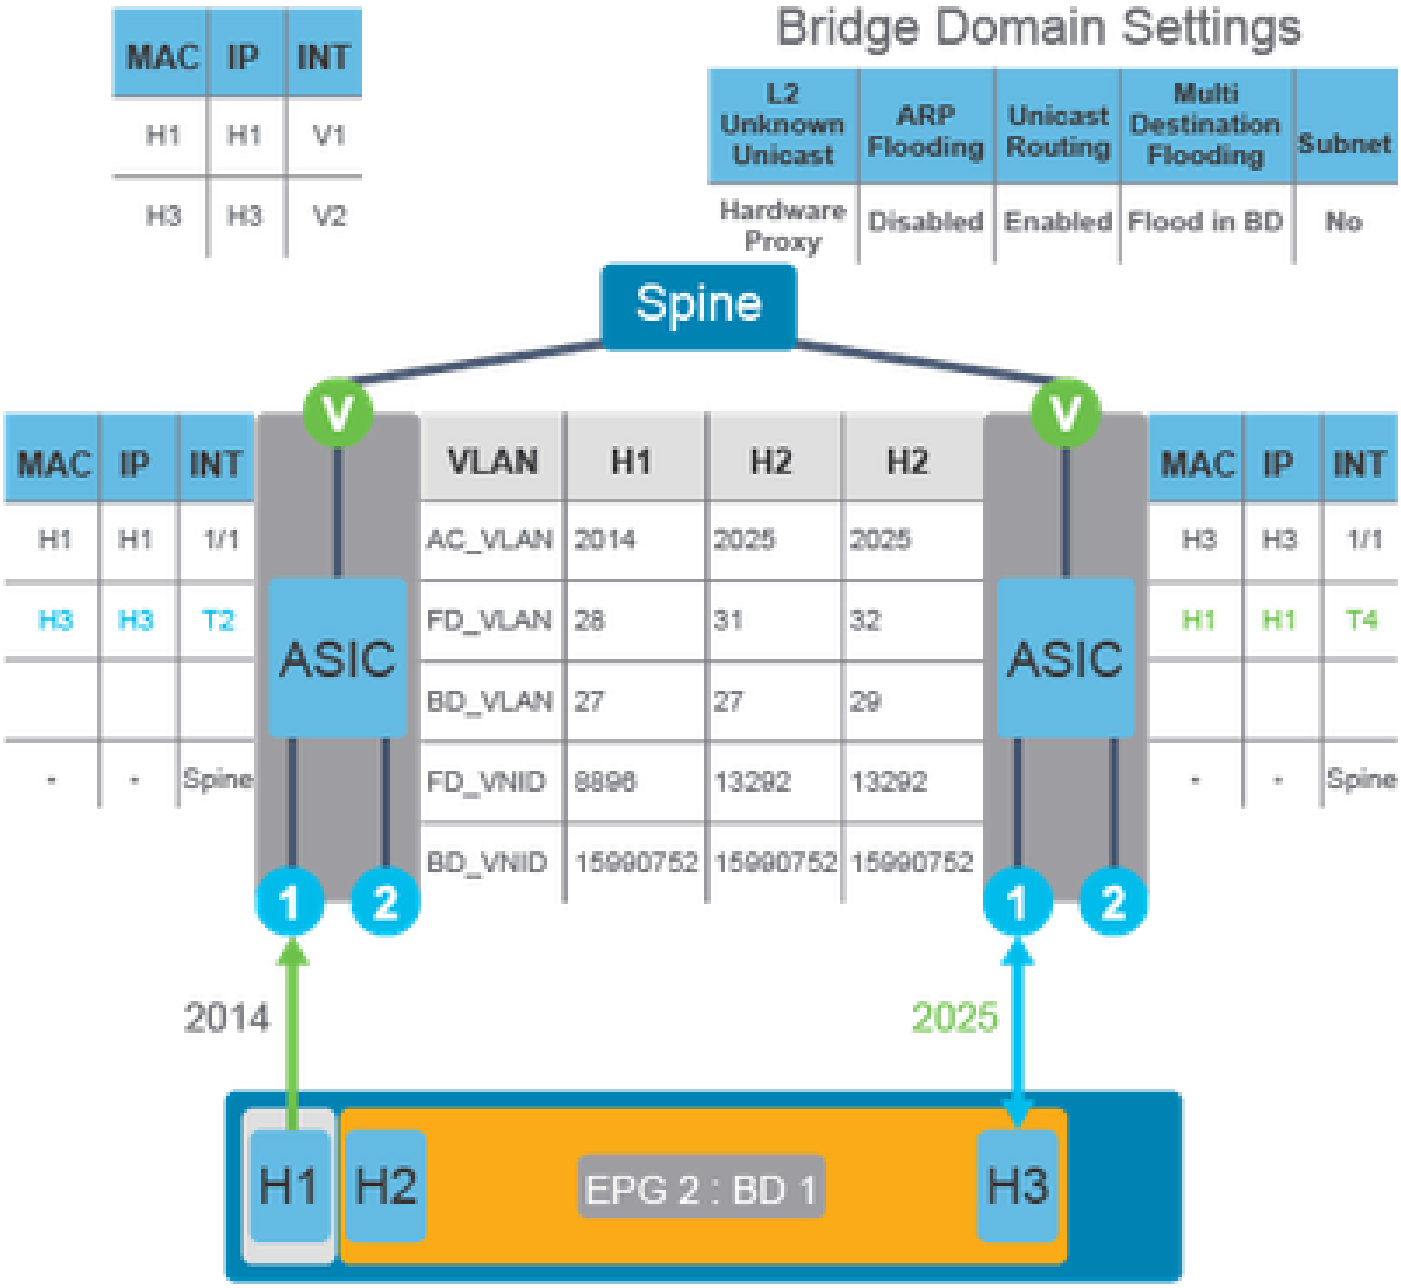

Use Case 3. Target IP Unknown, ARP Flood Disabled

This use case is applied when Ingress Leaf does not know the location of the target IP address (ARP flooding disabled, unicast routing enabled).

In a similar scenario, when ARP flooding is disabled and the ingress leaf does not know where the ARP target IP address is located, an ARP request is sent to the anycast spine-proxy Tunnel End-point (TEP) instead of flooding. The flow of ARP traffic from H1 to H2 is:

- H1 sends an ARP request for H2 using a broadcast destination MAC.
- The ACI attempts to use unicast forwarding to send the ARP request. The local leaf switch does not know the IP address of the endpoint H2 (the ARP target IP is unknown to the ingress leaf), so it sends the ARP request to the spine switch for spine proxy.
- Since H2 endpoint information is missing from the COOP database on the spine switch, the spine drops the original packet, but instead, it triggers ARP glean to detect the target IP, so that subsequent ARP requests are not dropped.

Example 1. Regardless of the EPG, Bridge domain, or Access/Encapsulation settings, the flow of ARP requests remains the same as mentioned earlier.

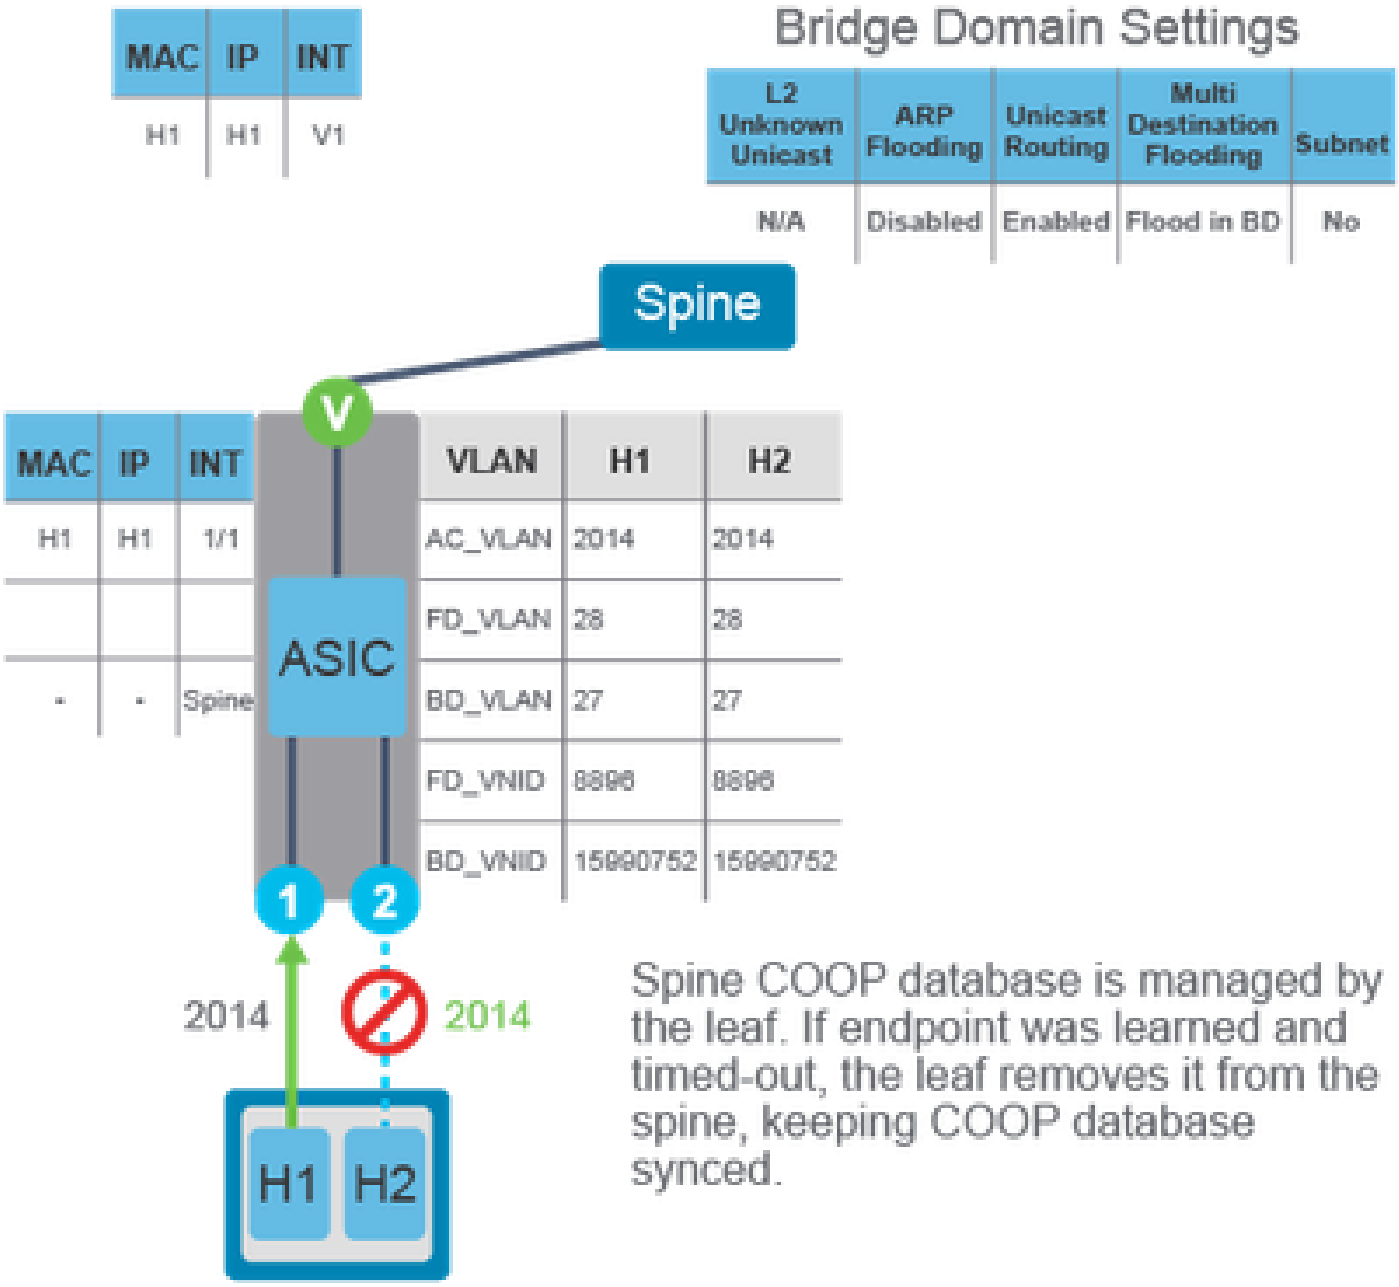

Use Case 4. Target IP Unknown, ARP Flood Enabled

This use case applies when Ingress Leaf does not know the location of the target IP address (ARP flooding enabled, unicast routing enabled).

If the ARP flooding is enabled in the bridge domain, the ARP request from H1 reaches H2 through flooding. The flow of ARP traffic from H1 to H2 is:

- H1 sends an ARP request for H2 using a broadcast destination MAC.
- The ARP request is flooded to all interfaces in the bridge domain. H2 receives the frame and replies, while it is learned in the fabric.

Example 1.

MAC	IP	INT
H1	H1	V1
H2	H2	V1

Bridge Domain Settings

L2 Unknown Unicast	ARP Flooding	Unicast Routing	Multi Destination Flooding	Subnet
N/A	Enabled	Enabled	Flood in BD	No

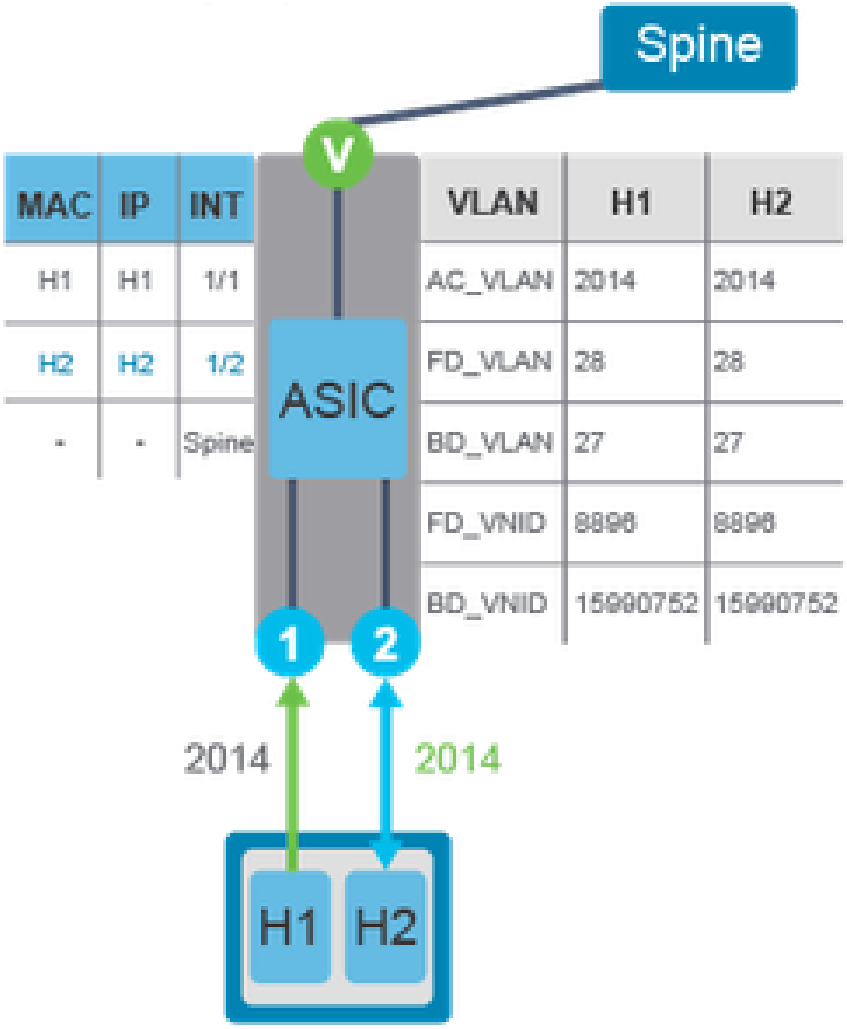

MAC	IP	INT
H1	H1	1/1
H2	H2	1/2
-	-	Spine

VLAN	H1	H2
AC_VLAN	2014	2014
FD_VLAN	28	28
BD_VLAN	27	27
FD_VNID	8898	8898
BD_VNID	15990752	15990752

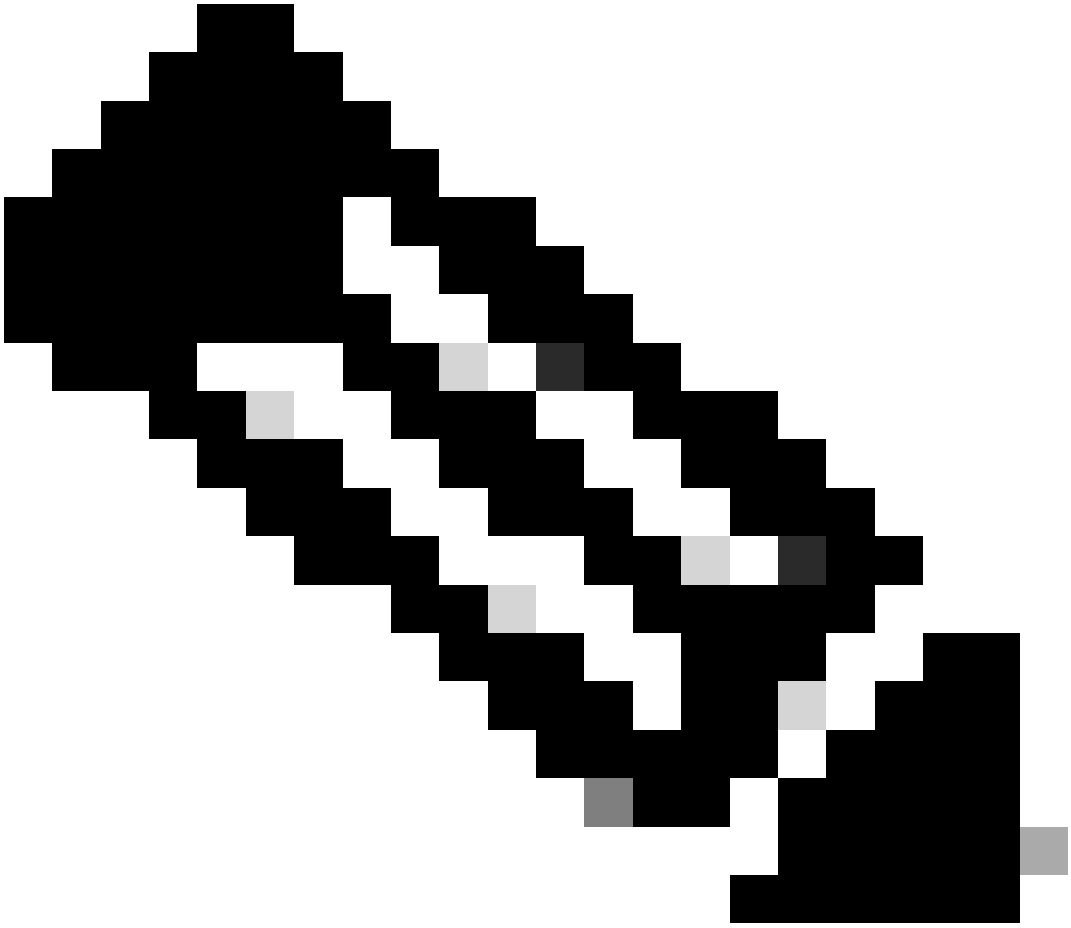

Note: The flood in encapsulation in Cisco ACI (bridge domain or EPG level) can be used to limit flooding traffic inside the bridge domain to a single encapsulation. When two EPGs share the same bridge domain and Flood in Encapsulation is enabled, the EPG flooding traffic does not reach the other EPG.

One of the benefits of enabling ARP flooding is to be able to detect a silent IP that moved from one location to another without notifying an ACI leaf. Because the ARP request is flooded within the bridge domain, even if the ACI leaf still thinks the IP is at the old location, the host with the silent IP responds appropriately so that the ACI leaf can update its entry accordingly.

If ARP flooding is disabled, the ACI leaf keeps forwarding the ARP request only to the old location until the IP endpoint ages out. On the other hand, the benefit of disabling ARP flooding is to be able to optimize traffic flow by sending the ARP request directly to the location of the target IP, assuming no endpoint moves without notifying its movement via GARP and such.

Use Case 5. Endpoints in Different EPGs and BDs

This use case is applied when endpoints are connected in different EPGs and different bridge domains.

When the endpoints are part of the different EPGs and different bridge domains, the traffic between them has to be routed. The flooding does not cross the bridge domains, including ARP flooding. So if H1 needs to communicate with H2, which is connected on the same leaf switch, the traffic is sent towards the default gateway MAC address, so ARP flooding is not relevant in this example.

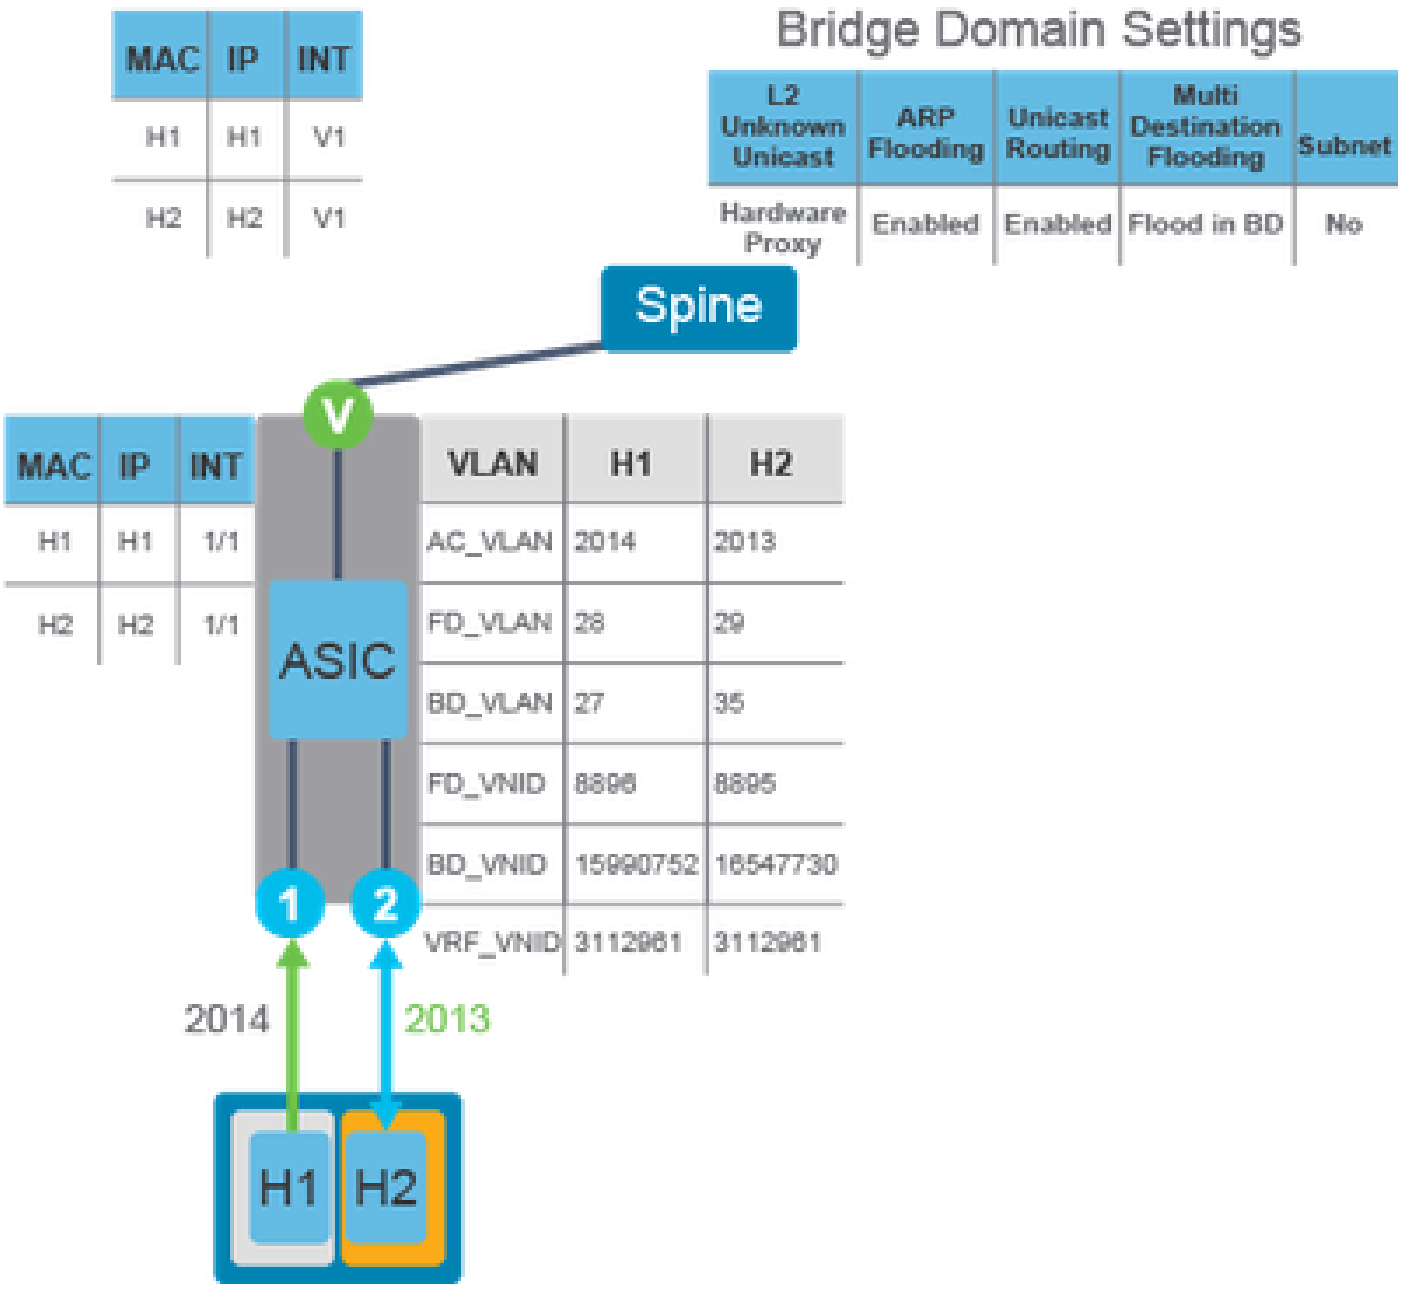

Understanding ARP Gleaning

Cisco ACI has several mechanisms to detect silent hosts, where an ACI leaf has not learned a local endpoint. ACI has some mechanisms to detect those silent hosts. For Layer 2 switched traffic to an unknown MAC, you can set the Layer 2 Unknown Unicast option under the Bridge Domain (BD) to flood, while for the ARP requests with a broadcast destination MAC, you can use the ARP flooding option under the bridge domain to control the flooding behavior. In addition, Cisco ACI uses ARP gleaning in order to send ARP requests to resolve the IP address of an endpoint that is yet to be learned (silent host detection).

With ARP gleaning, if the spine does not have information on where the destination of the ARP request is connected (the target IP is not in the COOP database), the fabric generates an ARP request that originated from the bridge domain Switch Virtual Interface (SVI) (pervasive gateway) IP address. This ARP request is

sent out to all the leaf nodes' edge interfaces part of the bridge domain. Also, ARP gleaning is triggered for (Layer 3) routed traffic regardless of configuration, such as ARP flooding, as long as the traffic is routed to an unknown IP.

ARP gleaning has a few requirements:

- IP address is used for forwarding (ARP requests with ARP flooding disabled, or traffic across subnets with ACI BD SVI as the gateway)
- Unicast routing enabled
- Subnet created under the bridge domain

Use Case 1. Target IP Unknown, ARP Flood Disabled

This use case applies when the target/destination endpoint is not known to the fabric (ARP flooding disabled).

When the endpoints are on different leaf switches, while part of the same EPG and bridge domain, and using the same VLAN access mapping, the ARP request (for example, from H1 to H3) has to be forwarded across the fabric. If H3 information is missing from the COOP database on the spine switch (silent host) and ARP flooding is disabled, ARP gleaning can be also utilized as depicted in this figure.

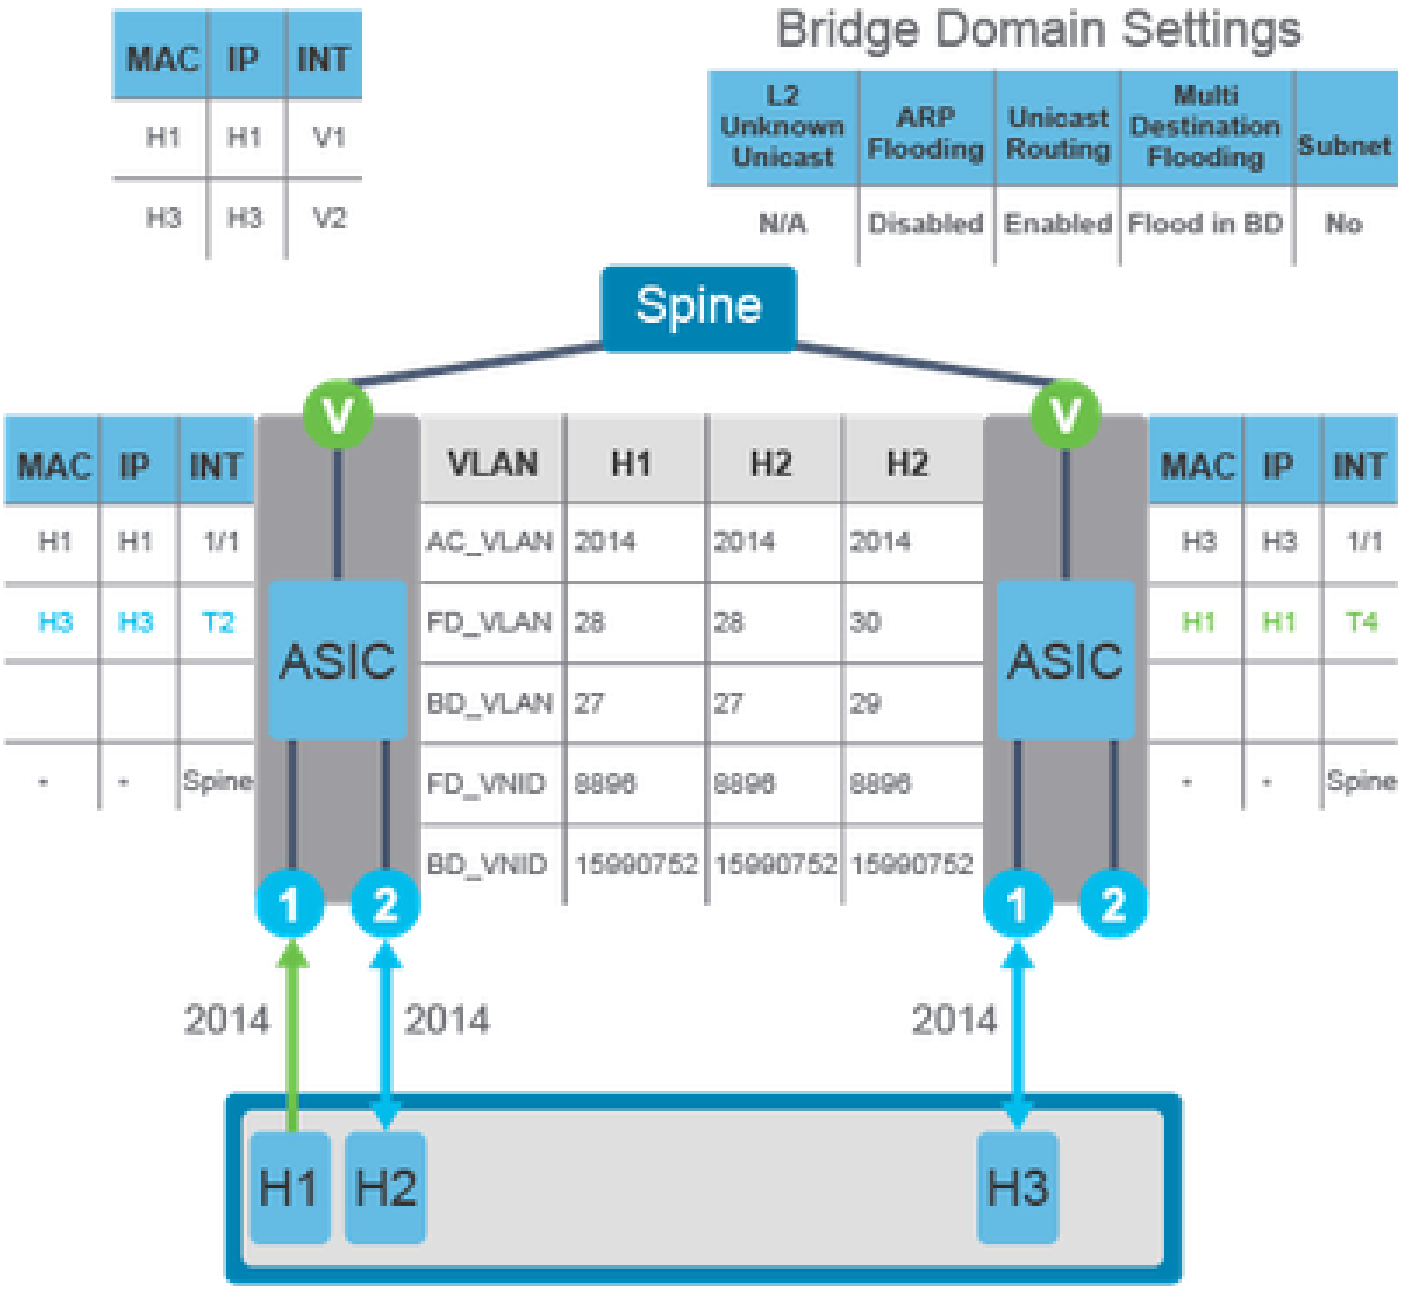

The flow of ARP traffic from H1 to H3 is:

- H1 sends an ARP request for H3 using a broadcast destination MAC.
- The ACI attempts to use unicast forwarding in order to send the ARP request, so the local leaf switch checks the ARP target IP address (H3 IP). Since the local leaf switch does not know the IP address of the endpoint H3, it sends the ARP request to the spine switch for spine-proxy.
- The H3 information is missing from the COOP database on the spine switch and triggers ARP gleaning using the pervasive gateway IP address as the source. This ARP request is flooded in the domain.
- H3 receives the ARP request and replies, while it is learned in the fabric.

Regardless of the EPG, Bridge domain, or Access/Encapsulation settings, the ARP glean feature works in the same way when two endpoints are trying to communicate with each other (irrespective of their connectivity to the same or different leaf switch inside the fabric).

Use Case 2. Endpoints in Different EPGs and BDs

This use case applies when endpoints are connected in different EPGs and bridge domains (ARP flooding enabled).

When the endpoints are part of the different EPGs and different bridge domains, the traffic between them has to be routed. The flooding does not cross the bridge domains, including ARP flooding that can be generated by ARP gleaning. So if H1 needs to communicate with H2, which is connected on the same leaf switch, the traffic is sent towards the default gateway MAC address, so ARP gleaning is not relevant in this example.