Advanced Application based Flow management and Firewall Services for new Age Enterprises

Tech Con 2025 Abstract 305

HPE Edge;

Advanced Application based Flow management and Firewall Services for new Age Enterprises

Abstract

Modern Enterprise networks are not just about providing connectivity. They all come with plenty of useful services that make the user experience seamless and network administration more robust and secure. Some of these useful services are Application Identification and Control through Deep Packet Inspection. It also includes Firewall services as well as Intrusion Detection and Prevention methods. These services should work seamlessly for a Roaming Client. One of the major requirement for a seamless network service for a Fast-Roaming client is the ability to manage all the flows and sessions across all the devices. As the scale of the network has gone enormously in the last few years, innovative methods are required for making the Fast-Roaming experience for any Client with any Role. This work proposes a smart and scalable method to solve this important problem in Enterprise Networks.

Problem Statement

Traditional firewalls offer perimeter security, controlling traffic entering and leaving your network. Distributed firewalls enable micro-segmentation, creating security zones at the individual device level for more granular control. Typically, in Campus networks EVPN is used to manage it smartly providing a better infrastructure to provide distributed firewall services. EVPN (Ethernet Virtual Private Network) combined with VxLAN (Virtual Extensible LAN) is a modern network architecture that offers significant benefits for campus and Data Center networks. This way of building the network provides enhanced scalability, flexibility, and efficiency in managing and operating large and complex network infrastructures. It uses BGP (Border Gateway Protocol) to distribute MAC addresses and IP routes, enabling efficient and scalable Layer 2 and Layer 3 VPN services. VxLAN encapsulates Layer 2 frames within UDP packets, allowing Layer 2 segments to be extended across Layer 3 networks. One of the primary benefits of EVPN networks is its Flexibility and Agility along with Fast Roaming for wireless clients. It supports seamless mobility for devices and applications across different segments of the campus network. The key to fast mobility is achieved by synchronizing MAC, Host and IP Prefixes across all the Network Nodes using BGP as the Control Protocol.

Flow Telemetry Application based Policies, Firewall policies and EVPN

Any modern Campus and Data Center network of today provides visibility and control for Flows and associated Applications for all the clients. This allows the Admin to have a fine-grained visibility and control for the connected clients, applications, and network services. Once the Application is identified, the admin can apply policies per for better Access Control and Quality of Services.

Fast Roaming and Application based visibility and control

Once a wireless client moves as shown in Figure 1 in a fast-roaming environment, the expectation is that the existing sessions and properties will stay unaffected. Application Identification Engine requires the first few packets, the client once moving to a new Switch will not have the Application Identification done at all. Consequently, any application specific policies will fail as well. One significant property of an IP flow that distinguishes it from other entities like Host, MAC and Routes is the sheer volume and transient nature of it. A single client (mac or IP) can open hundreds and thousands of connections. Similarly, a single VM (mac or IP) can host thousands of flows. The sheer transient nature and the huge volumes associated with flows make them unsuitable to be synchronized like other EVPN Control packets via BGP messages. Fast Roaming fails with Application Based Policies[5] without a proper solution to this problem. This work aims to solve this important problem in the new age of Enterprise networks.

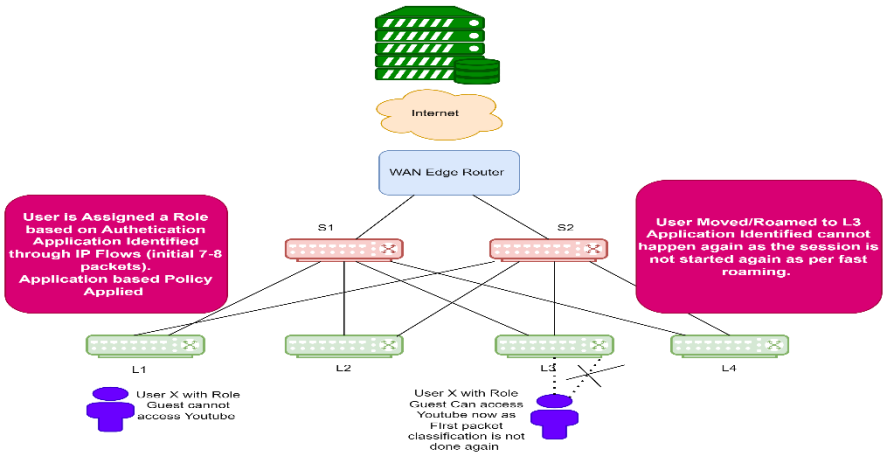

Figure 1: Fast Roaming with Application Based Policies

Our Solution

The solution that we propose here is to build a smart, low latency and scalable distributed flow management mechanism to exchange flow and its related information across the VTEPs (Switches) such that Application based Recognition and Control as well as other firewall services can happen seamlessly for a Roaming device. The steps involved in the solution is

1. Client Identification and Categorization – The first step is identification and categorization of clients. The categories of clients are

- **Wired Clients** – Wired Clients are the ones like IOT devices or Printers which are directly connected to the Switch Port. For those clients, there need not be any concept of fast roaming as the client will be physically disconnected and connected to a new port. All the existing sessions will be lost. Wired Clients are NOT considered for fast roaming.
- **Wireless Clients**- Among the wireless clients, multiple other characteristics are considered to build client clusters.
- Client VLANs- Check if the Client are part of roaming VLANs. Flows for client which are not part of roaming VLANs are not considered.
- Client Behavioural Characteristics- Based on historical data of the client, the roaming probability is built. This depends on the device type which is identified through device finger printing. In addition, the nature of the client is also considered. Multiple clients having which are using the same device type might be characterized as Roaming or Non-Roaming. This is done through behavioural characterization. Some Users are naturally roaming while others are stationary. Clients are characterized as roaming or non-roaming.

2. **Flow Synchronization**- IP flows from clients which are probable fast roaming clients are synchronized from the Client connected VTEPs to the Spine Switches. Note that in a typical VxLAN CLOS network, multiple VTEPs are connected to various Spine Switches. Typical leaf to spine ratio will be **m: n (m>n)**. The spines are of much higher capacity and can store way more flows than the VTEPs. As the number of roaming VLANs could be limited, each set of Spines will store flows for some of those VLANs. E.g. if there are 200 roaming VLANs and 8 Spines and 32 VTEPs then each Spine will store flows for 200/8=25 vlans. Out of these 200 vlans each VTEP will have approximately 50 VLANs stretched in them. Each of these VTEPs will synchronize their roaming client flows to the respective Spines. This will be a configuration in each VTEPS and their Spines. The flows are then pushed to the respective Leafs from the spines. This ensures that for Roaming Clients, the flows are synchronized to all the other VTEPs where it can move. MQTT[6](originally an initialism of MQ Telemetry Transport) is a lightweight, publish-subscribe, machine to machine network protocol for message queue/message queuing service.) based communication is used to send flows across Spine and Leafs. MQTT was chosen as it is lightweight, fast, and proven successful in multiple subscriber publisher environments.

3. Action on Client Move – When a client moves to a new VTEP, the existing flows of the client should be already present in the new VTEP if this was characterized correctly as a roaming client based on previous steps and flows are already PUSHED to it by the Spine Switch. This will achieve fast roaming as for the new client, Application based policies are rules are already set up properly. In case the flow is not present in new VTEP cache then the characterization of the client was wrong as a roaming client (in step 1). In this case, a slow path is executed where a lookup is executed to the old VTEP to get all the flows for the client. This is a PULL case. Note that most vendors in the world use the Slow path ONLY and our worst performance will be similar. The figure below shows the whole operation in detail for 2 roaming clients, one who is in VLAN X and one who is in VLAN Y. Note that if the client were not categorized as roaming then its flows would not get pushed to the Spine.

Impact of Spine Reboot – In case the Spines go for a reboot then the flows are pushed in bulk to the Spine once it is UP. In the meantime, if there is a roaming client then a slow path will be used for Flow data synchronization. We can also create a Backup Spine for every Active Spine for a set of VLANs. This would ensure that in case of spine reboot , the system never moves to slow path in case of a fast roaming . This requires that the MQTT Brokers in the Active and Backup Spines stay in sync, which might be little bit of an expensive operation. This will be evaluated

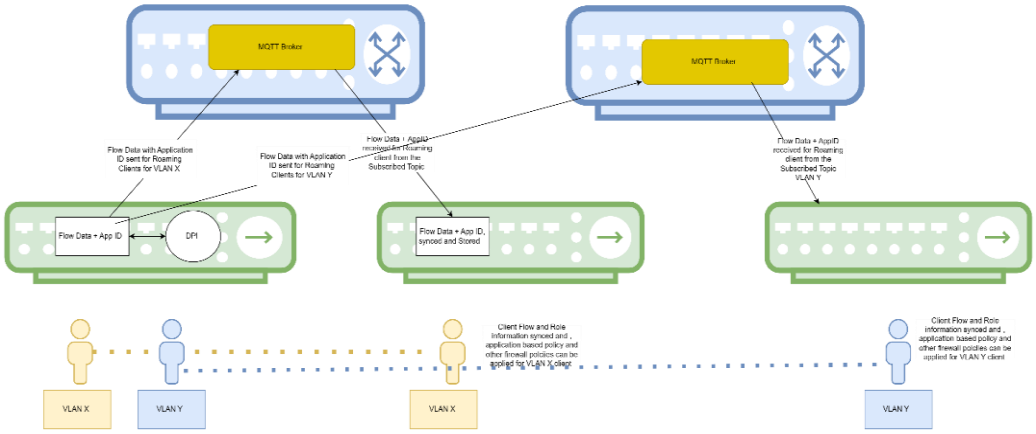

Figure 2: Fast Roaming with Firewall Policies

Business Value- In 2024, the total addressable market (TAM) for network security is forecast to amount to 35 billion U.S. dollars[1]. Software firewalls as well as the introduction of new security services such as IoT or DLP drives growth in the network security market. This market is expected to double in a few years. Not having a secured and flexible firewall and application policy engine for a Campus network would seriously harm Aruba in this segment. Aruba is bringing its next generation firewall-based line cards called King Fisher in 2026. The overall capacity of that system for flows and firewall services is higher than the current offering. This solution will be extremely useful once we get there.

Evidence the solution works

Flow Synchronization is achieved through running a MQTT based Brokers in each Spine Switches. The VTEPs selected were Aruba OS-CX 6300 which has Application Based Policies supported. Each Roaming VLAN is one of the Topic in the MQTT Broker. The clients which are characterized as Roaming Clients push the flow information to the respective Spines. We have evaluated the solution with 4 Roaming VLANs and 3 Spines with 2 VTEPs. The number of flows that were evaluated was around 10K flows with 100 clients. Out of that 30 percent of clients were characterized as Roaming. The solution worked very well as the scale of the Flows were increased.

Competitive Approaches

One alternate approach followed is for the new VTEP (to which the endpoint has moved) to request the old VTEP for client's flows upon seeing the move notification (PULL mode). The problem here is the latency between new VTEP detecting a client move and learning the client's active flows from the peer. For seamless roaming experience post enforcement of application-based policies, the PULL process must complete within 100ms. For a single client move scenario, this may be ok but when multiple clients move concurrently or the old and new VTEPs take longer to complete the transaction (e.g. CPU busy), it will be difficult to deterministically complete the process within 100ms.

Current Status

An MQTT based implementation of active flow synchronization is running on Aruba OS-CX 6300 access VTEPs. These VTEPs are the MQTT producers and subscribers of flow information associated with roaming clients. The Aruba OS-CX 8360 is the Spine and the MQTT broker pushing flow information from the producers to the subscribers. With the right classification of clients as 'roaming' or 'fixed' and assuming a reasonable 100 concurrent connections per client, the solution was able to scale for 100 clients and seamless roaming was achieved with application-based policies configured for their user-roles. Based on the analysis of the MQTT load and VTEP performance, this solution can comfortably handle a 10x increase in scale, i.e. 1000 concurrent roaming clients.

Next steps

The solution is being reviewed with PLMs and SEs and the plan of record is to add it to the FY'25 CX software roadmap. Some more refinements to the solution (e.g., filtering based on applications that a client is talking to as not all apps are latency sensitive and they can continue to be in PULL mode), support for hybrid PULL/PUSH modes, dynamic client-classification change (from roaming to fixed or vice-versa) are also being planned as part of the productization effort.

References

[1] Source: https://www.marketresearchfuture.com/reports/network-telemetry-market-8716
[2] https://technology.berkeley.edu/news/default-firewall-new-campus-networks
[3] https://www.checkpoint.com/cyber-hub/network-security/what-is-firewall/
[4] MQTT Broker - https://mqtt.org/
[5] https://www.arubanetworks.com/techdocs/Instant_41_Mobile/Advanced/Content/UG_files/AppRF/ConfACLRule.htm