VERSA NETWORKS

WHITE PAPER

Segmentation & Multitenancy: A Primer

Table of Contents

In this white paper, we will review the many types of segmentation that enterprises have traditionally utilized and explore the types and nature of multitenancy.

Traditional VLAN segmentation limitations

For years, enterprises viewed network segmentation as a necessity to accomplish network security and as a relief from limitations on Layer 2 domains. Virtual local area networks (VLANs) became the de facto standard for segmentation, providing a mechanism for segregating business units, zones, and security. However, VLANs only provide minimal security and separation within the average enterprise network. While it is true that a user on a given VLAN can't directly communicate or access information on another VLAN, the use of Denial of Service (DOS) attacks may impact another VLAN's traffic and communication. Also, given that a single switch probably houses more than one VLAN, compromise of that single switch would allow the user of one VLAN to gain access to the information on another VLAN. Similarly, compromise of the L3 device that is the gateway for those VLANs would provide a mechanism for a threat actor to gain access to the segmented data. (See figure 1). All of this stems from the fact that a VLAN or network segmentation only deals with Layer 2 or Layer 3 isolation, and does not deal with any true security separation nor any isolation of shared resources. Notice that in Figure 1, common Layer 2 switch and common Layer 3 router would be susceptible to DOS attacks meant to disrupt or perhaps allow elevated rights access from a given network segment.

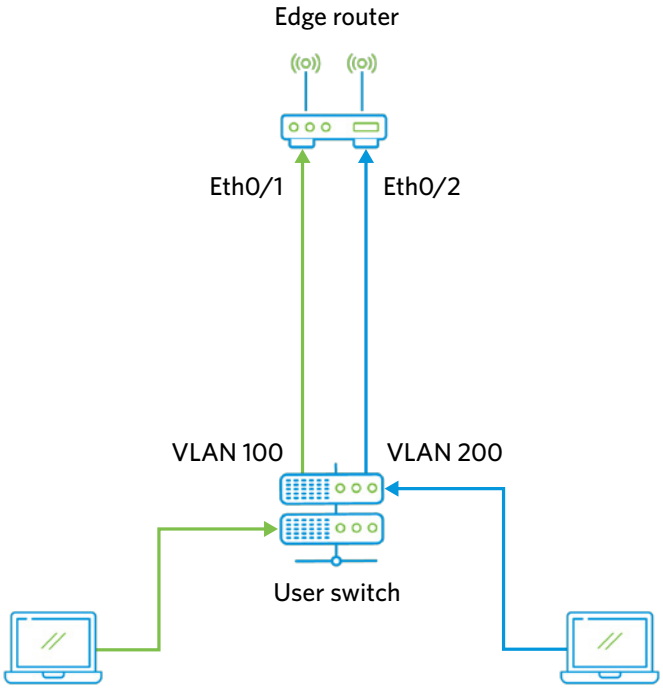

Figure 1 – The Common Layer 2 switch and Common Layer 3 router would be susceptible to DOS attacks.

Segmentation types

Segmentation comes in many different forms inside a network design. In addition to VLAN segmentation, discussed above, below we quickly review encryption, policy tags, virtual routing and forwarding (VRF), routing domains, virtual private networks (VPNs), network segmentation, zone concepts, and multitenancy as forms of segmentation.

Role-based access control (RBAC)

RBAC is intended to keep people from accessing the network elements within a network and only allow for levels of privilege associated with enterprise needs. However, there have been numerous Common Vulnerability and Exposure (CVE) announcements issued in the past decade where threat actors exploiting a vulnerability would gain elevated privilege to the systems. Sometimes, the access was granted without even needing initial access. So, while RBAC and other advanced authentication methods are good for protecting systems, obviously a more granular and layered approach is needed.

Policy-based tagging

Policy-based tagging, often referred to as a common group tag, is another segmentation mechanism that requires no actual segmentation of the data plane. In this model, packets are tagged based upon a policy as defined by the enterprise. This tag is then captured in a header format and utilized to traffic steer or apply policy based upon this tag. Many intent-based networking constructs utilize the policy tag or group tag to achieve the desired policyimplementation. This is an example of micro-segmentation. However, this does not provide any protection from anyone who has access to the data network from capturing the data from one group tag structure or another.

Encryption

Encryption is a powerful segmentation tool. This enables the end device to encrypt the traffic and segment it from the other traffic on the data network, such that someone who has access to the data network would not be able to effectively utilize the data that had been captured. All encrypted traffic has one major flaw – given enough memory, compute power and time, a threat actor can decrypt the traffic and utilize it. So, this type of segmentation becomes a probability problem. Utilization of a sufficiently complex encryption method and a small window for the data to be relevant would produce minimal risk for capture of the data by an unauthorized individual. Also, methods for encryption constantly need to evolve as the resources available to threat actors are always expanding over time.

Virtual routing and forwarding (VRF)

VRFs provide a mechanism to isolate the impact of a DOS attack on another aspect of the shared Layer 3 device. However, this protection is only limited to the routing and forwarding aspects between two segmentations. Unless the shared device has a method of segmenting the resources utilized by the network device, starvation of shared resources is still possible.

Routing domains

Another layer of segmentation can be the implementation of discreet routing domains. This way each of the VRFs would have their own routing and forwarding tables. This keeps one VRF from being able to adversely affect the other VRF via a routing issue. However, this still does not address the shared resource issue.

Virtual private networks (VPNs)

It is possible to accomplish actual segmentation by utilizing discrete Virtual Private Networks (VPNs). This concept normally, but not necessarily, requires discrete routing domains and discrete VRFs. By utilizing different VPN segmentation, each different VPN can have a different topology as there is no requirement for each VPN to connect to the same devices in the same way. In this manner, the segmentation would be complete with different security aspects. However, depending on how the security is implemented, security keys or certificates could be shared between VPN constructs. And as before, this still does not address the shared resource issue.

Segregation of control and data planes

Another layer of separation that should be considered is the segregation of the control plane and data planes. If the design does not segregate the data plane and control plane, then a DOS attack on the data plane could cause control plane loss. Control plane loss would cause irreparable harm to the enterprise. And, in fact, the design should have a multitenant control plane. This way no one tenant could cause another tenant to lose access to the control plane.

How to accomplish multitenancy

For many years, enterprises have considered multitenancy as the purview of the service provider networks. Enterprises traditionally would identify the need to carve shared resources into smaller chunks only if the intent was to resell the resources to customers. However, multitenancy benefits go beyond just this commercial aspect. The shared resource issue can be solved in multitenancy by assigning resource limits to each of the tenants and restricting the access to memory, bandwidth, CPU, and storage. RBAC controls would need to be architected in a way where access granted to a given tenant would not allow for access to any of the other resources not allocated to the tenant. Note in Figure 2, when a given tenant logs into the system, they are only able to see their resources as allocated by the system. Even in the case of a shared transport resource, the multitenant architecture allows for encapsulation of the tenant data in a manner where only the data that is pertinent to the tenant can be displayed or captured. In this way, the traffic from the other tenants is not able to be captured.

Figure 2 – Even if transport is shared, a proper multitenant architecture allows only the data that is pertinent to the tenant to be displayed or captured.

Orchestration multitenancy

Multitenancy can be accomplished in the orchestration platforms, the control plane, and the data plane. Many current systems allow for orchestration multitenancy. This level of multitenancy keeps the policies, configuration, logs and statistics segregated from those of the other tenants. But if there is no control plane segregation, then the control plane is shared by the many tenants and an orchestration lock could happen from a single tenant, adversely affecting the other tenants.

A fully multitenant system would take multitenancy to its most logical conclusion (see Figure 2). This would be a system where multitenancy was at the management level (see Figure 3), controller plane level, data plane level, and the analytics level. The system would be multitenant at the hub location and the edge device locations.

Multitenancy in the management plane

In the case of multitenancy in the management plane, each tenant:

- Sees both devices and their CPU/memory/HDD utilization.
- Will only see traffic that belongs to his ports and networks.
- Will only be able to configure its own policies, but will not be able to see configurations or statistics of the other tenants on the same devices.

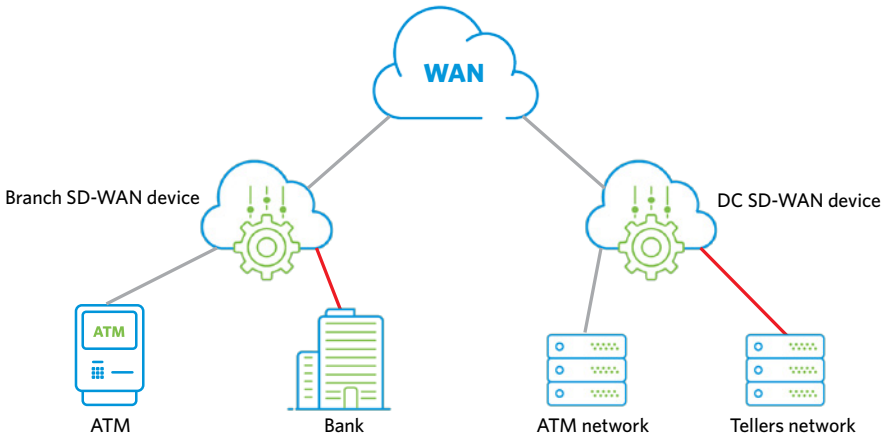

Figure 3 – Multitenancy at the management level.

Also, every multitenant device would have its resources allocated in a manner where no single tenant could cause issues with the whole environment so as not to adversely impact the other tenants. *(see Figure 4)*

Control Plane Multitenancy

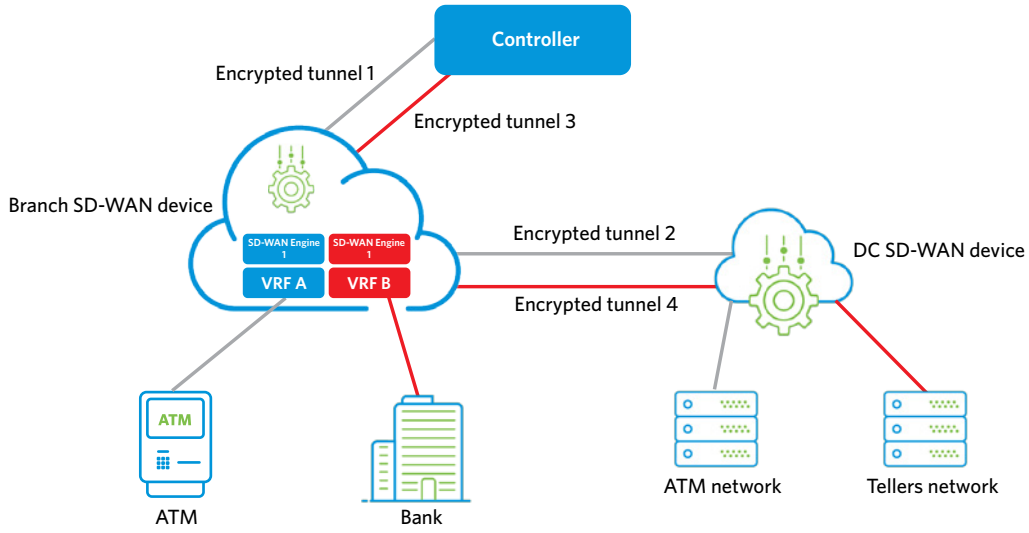

Figure 4 – Multitenancy in the control plane.

In control plane multitenancy, every tenant has its own:

- RBAC, logging, and statistics controls.
- Security posture, including distinct and discrete security keys or certificates for the control plane and separate ones for the data planes.
- Unique and discrete encryption algorithms *(see Figure 4)*.

Data plane multitenancy

In data plane multitenancy, each tenant:

- Has its own independently encrypted IPsec tunnels between SD-WAN devices. If any of the IPsec tunnels get compromised, other tenants are not affected.
- Will only see traffic that belongs to its ports and networks.
- Will only see its own ports and not the ports of other tenants on the device.
- Can configure only its own routing protocols, firewall rules and SD-WAN policies.

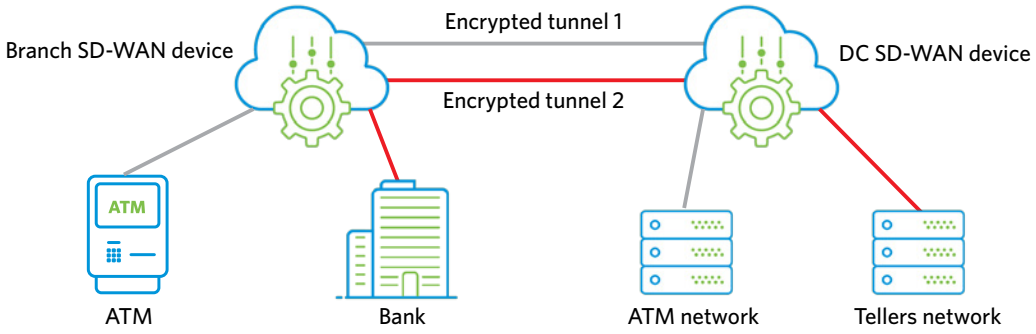

Figure 5 – Multitenancy in the data plane.

Each tenant would have its own distinct and discrete routing domains, VRFs, VLANs, zones and VPNs/topologies. *(See Figure 6)*

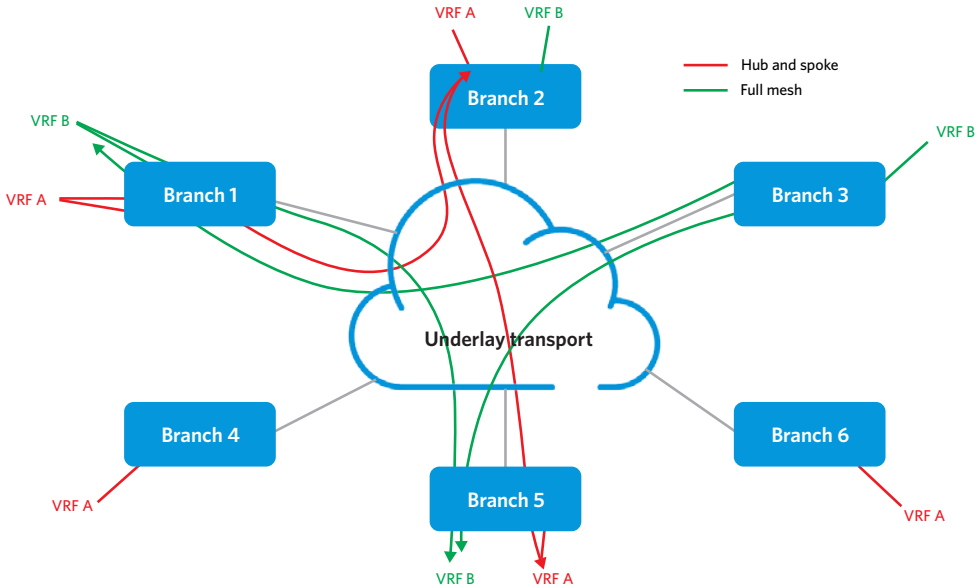

Figure 6 – Independent VPN overlays.

Each tenant also has its own policy tags, and these would not be shared by any of the other tenants.

Ideally, a true multitenant network should allow for multiple levels of multitenancy. This would allow for complex business logic to be implemented and for the system to be utilized for either service provider or reseller purposes, or for a very security-conscious enterprise.

About Versa Networks

Versa Networks, the leader in single-vendor Unified SASE platforms, delivers AI/ML-powered SSE and SD-WAN solutions. The platform provides networking and security with true multitenancy, and sophisticated analytics via the cloud, on-premises, or as a blended combination of both to meet SASE requirements for small to extremely large enterprises and service providers.

Thousands of customers globally with hundreds of thousands of sites and millions of users trust Versa with their mission critical networks and security. Versa Networks is privately held and funded by Sequoia Capital, Mayfield, Artis Ventures, Verizon Ventures, Comcast Ventures, BlackRock Inc., Liberty Global Ventures, Princeville Capital, RPS Ventures and Triangle Peak Partners. For more information, visit https://www.versa-networks.com or follow Versa Networks on X (Twitter) @versanetworks.

Versa Networks, Inc, 2550 Great America Way, Suite 350, Santa Clara, CA 95054
+1 408.385.7660 | info@versa-networks.com | www.versa-networks.com