VERSA NETWORKS

WHITE PAPER

Genuine Multi-tenancy

As enterprises look to glean more out of their business models and implement more security focused networks, Software Defined Networking (SDN), including Software Defined Wide Area Networking (SD-WAN) and Network Function Virtualization (NFV), provides for mechanisms to modernize infrastructure and bring a more business focus to traffic engineering.

In this White paper, we will explore the many aspects of segmentation that enterprises have utilized, and we will explore Multi-Tenancy and discuss why this is no longer a Service Provider construct.

Segmentation comes in many different forms inside a Network Design. VLAN segmentation, Encryption, Policy Tags, Virtual Routing and Forwarding (VRF), Routing Domains, Virtual Private Networks (VPNs), Network Segmentation, Zone Concepts, and Multi-Tenancy are all forms of segmentation.

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. It provided a mechanism for segregating Business units, zones, and security.

However, VLANs provide a minimal set of security and separation within the average Enterprise network. While it is true that a user on a given VLAN can not directly communicate nor access information on the other VLAN, the use of Denial of Service (DOS) attacks may cause impact to the other VLANs traffic and communication. Also, given that a single switch probably houses both VLANs, compromise of that single switch would allow the user of one VLAN to gain access to the information on the other VLAN. The same compromise holds true for the L3 device that is the gateway for both VLANs. Compromise of that device would provide mechanisms for a threat actor to gain access to the segmented data. (See figure 1). All of this stems from the fact that 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 right access from a given network segment.

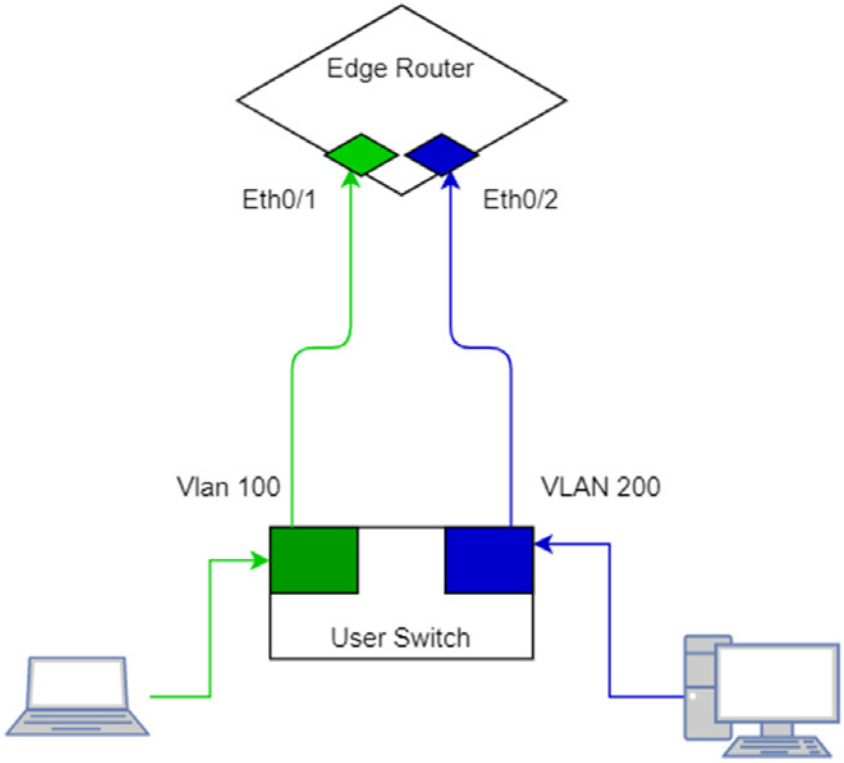

Figure 1

Role Based Access Control (RBAC) is intended to keep people from accessing the network elements within a network and only allow for levels of privilege associated with Enterprise need. However, there have been numerous Common Vulnerability and Exposure (CVE) announcements have been 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 more granular and layered approach is needed.

Policy based tagging, often referred 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 policy implementation. 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 and another.

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, enough compute power and enough time, a threat actor can decrypt the traffic and utilize it. So, this type or segmentation

becomes a Probability problem. Utilization of a sufficiently complex encryption method and a small window for the data to be relevant would produce a minimal exposure to capture of data by an unauthorized individual. Also, methods for encryption constantly need to evolve as the resources available to threat actors is exponentially expanding with time.

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 the 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.

Another layer of Segmentation can be 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.

Continuing there could be actual segmentation by utilizing discreet Virtual Private Networks (VPNs). This concept normally, but not necessarily, requires discreet Routing Domains and discreet 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.

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 Multi-Tenant Control plane. This way no one tenant could cause another tenant to lose access to the control plane.

For many years, Enterprises have considered Multi-Tenancy as the purview of the Service Provider networks. Enterprises could identify the need to carve shared resources into smaller chunks if the intent was to resell the resources to customers. However, Multi-Tenancy benefits are far more than just a commercial aspect. The shared resource issue can be solved in Multi-Tenancy 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 Multi-Tenant architecture allows for encapsulation of the Tenant data in a manner where only that data which 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.

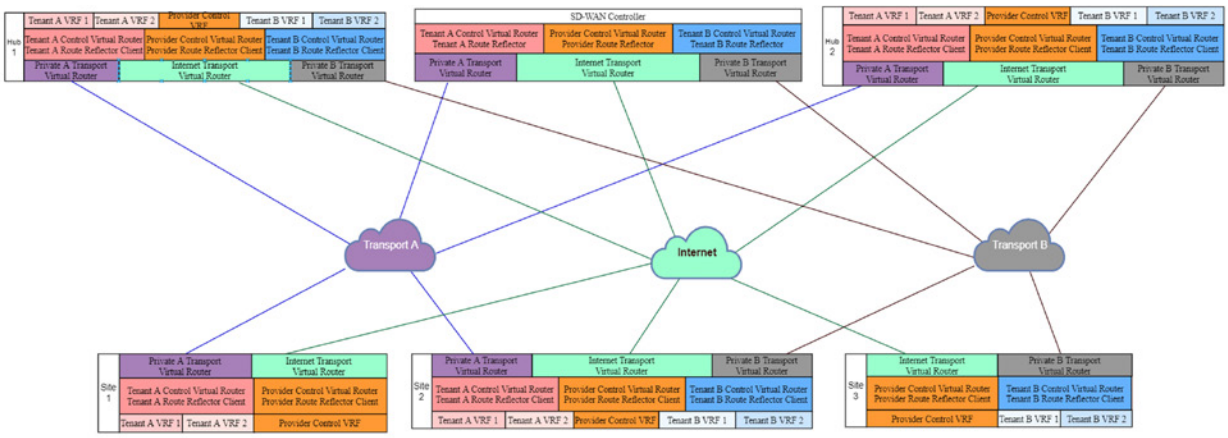

Figure 2

Multi-Tenancy can be accomplished in the Orchestration platforms, the Control plane, and the Data plane. Many of the current systems allow for Orchestration Multi-tenancy. This level of Multi-tenancy keeps the policies and configuration and the logs and statistics segregated from that 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 result from a single tenant and adversely affect the ability of the other tenants.

A fully Multi-tenant system would take Multi-Tenancy to its most logical conclusion. (See Figure 2) This would be a system where Multi-tenancy was at the Management level (see Figure 3), Controller Plane, Data Plane, and the Analytics Level. The system would be Multi-Tenant at the Hub location and the Edge Device locations.

Multitenancy in Management Plane

- Each tenant will see both devices and their CPU/memory/HDD utilization
- Each tenant will only see traffic that belongs to his ports and networks
- Each tenant will only be able to configure its own polices but will not be able to see configurations/statistics of the other tenants on the same devices

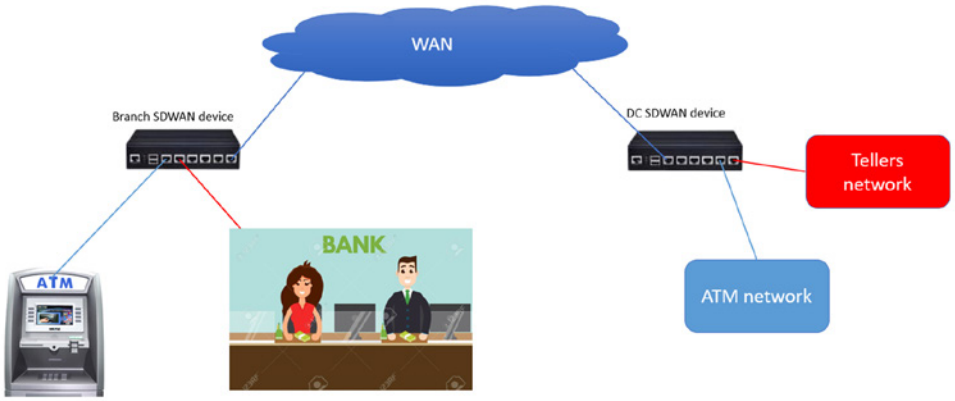

Figure 3

Also, every multi-tenant 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

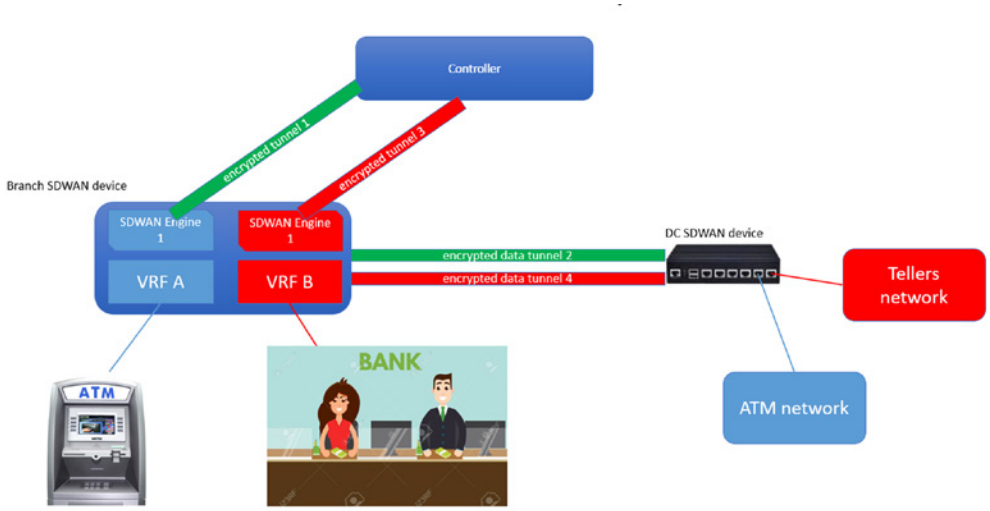

Figure 4

Every Tenant would have its own RBAC, Logging, and Statistic controls. Every Tenant would have its own security posture, including distinct and discreet Security Keys or certificates for Control Plane and separate ones for the Data Planes. Each Tenant would have its own unique and discreet Encryption algorithms. *(see Figure 4)*

Data Plane Multitenancy

- Each tenant will have is own independently encrypted ipsec tunnels between SD-WAN devices. If any of the ipsec tunnels gets compromised other tenants are not affected
- Each tenant will only see traffic that belongs to his ports and networks
- Each tenant will only see his own ports and not the ports of other tenants on the device
- Each tenant can configure only its own routing protocols, firewall rules and SD-WAN polices

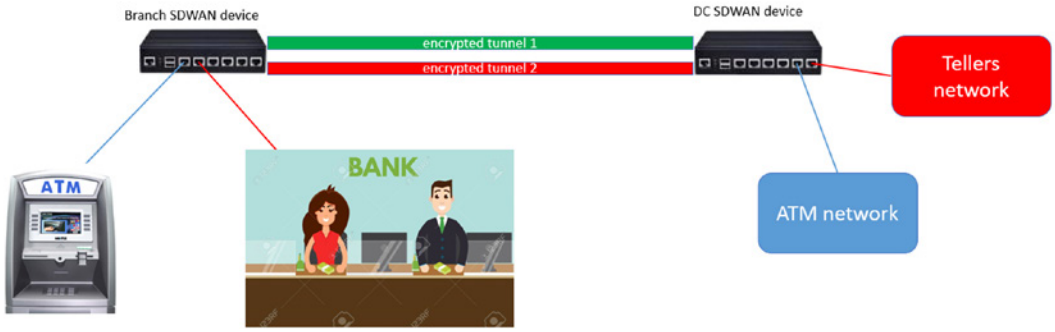

Figure 5

Each tenant would have its own distinct and discreet Routing Domains, VRFs, VLANs, Zones and VPNs/Topologies. *(See Figure 6)*

Independent VPN Overlays

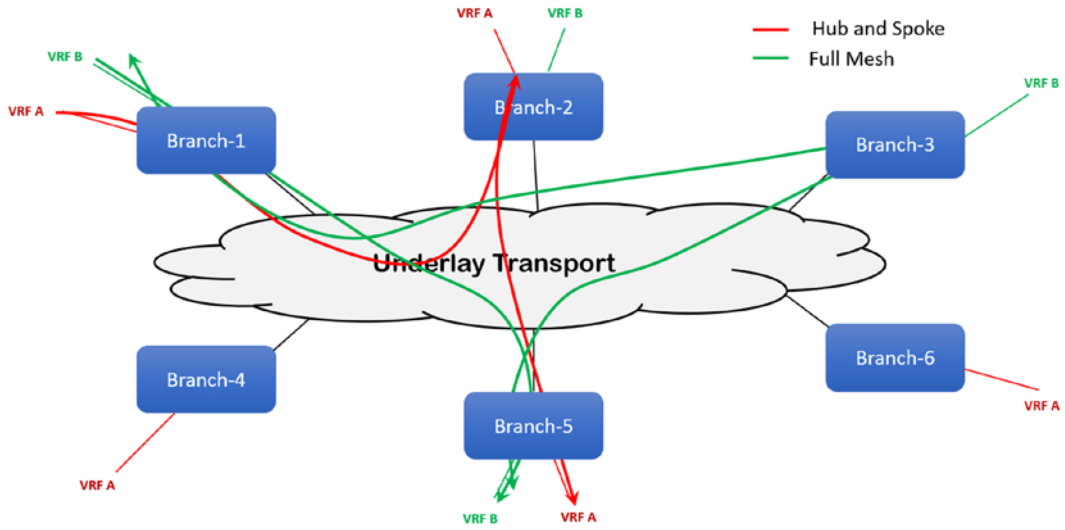

Figure 6

Each Tenant would have its own Policy tags, and these would not be shared by any of the other tenants.

Ideally, a true Multi-Tenant network would allow for multiple levels of Multi-tenancy. This would allow for complex business logic to be implemented and allow the system to be utilized for either Service Provider/Reseller purposes, or a very security conscious Enterprise.

Multi-Tenancy

SD-WAN Management Plane Multitenancy

- Independent RBAC for each tenant
- Users of a tenant can see only devices of that particular tenant only

SD-WAN Data Plane Multitenancy

- Routing tables separation
- Each tenant can have up to 1024 VRFs
- Data Plane independently encrypted tunnel between SD-WAN devices
- Independent instances of a routing table, BGP instances, OSPF instances, etc.

SD-WAN Control Plane Multitenancy

- Independent SD-WAN engines for each tenant
- Independently encrypted secure tunnels with Controllers for each tenant
- Independent topologies for each tenant

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