Network Segmentation for a More Secure LAN

ATLANTIC

Computing Technology Corporation

Network Segmentation at Layer-2

Version 1.0

Sponsored by:

Thank This Guy

[1]In 1988, a student at Cornell, Robert Tappan Morris, introduced a worm into the Internet. The worm took advantage of peer computer weaknesses in applications on DEC VAXes, Sun

workstations and a few other Unix platforms. Morris didn't mean to cripple the network, but he wanted his worm to go everywhere, albeit discreetly. Early versions of the worm would test to see whether a host was already infected,

[1] Image, Trevor Blackwell, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=5131967

before re-infecting. But, Morris recognized that this was an exploitable method to shut the attack down. So, he modified the worm for a modest-seeming infection rate of 14%. Taken over time, 14% applied repeatedly to the same hosts turned the worm into a significant denial of service attack, and much of the Internet was crippled for days.

Us vs. Them

Morris' experiment marked the end of innocence; before the worm, the Internet had been a playground. In the 90s, networks dressed in fig leaves--access lists, proxy and stateful firewalls. The Internet slipped into dualism--us versus them, inside versus outside, heaven versus earth. Each organization's network was clad against the world.

But, the worst security problems turned out to be inside jobs, emanating from machines on the LAN. They were delivered through email, device drivers or software updates. Compromised machines are not detectable at the firewall until they try to phone home, and even then the traffic may be concealed in sanctioned connections and still go undetected. Internal

attacks can include probes for SMB shares, brute force password cracking, and internal vulnerability searches. Exploits can go on for months without notice.

How do we secure the internal network? It is possible to deploy network flow taps across the LAN. This can be expensive. It is possible, but burdensome, to extend the firewall to every device via host agents. One could also cut the network into segments and put a layer-3 firewall in the middle. The result will be a coarse partitioning that can neither adapt nor support mobility. You'll hate it.

This paper talks about a relatively simple and very effective way to segment the local area network at layer-2. The partitioning is transparent and flexible. It can adapt and change even as people are using it. You may be able to stage it on equipment that you already own. Our discussion will demonstrate the concepts on switching and identity products from HPE Aruba Networks.

A segmented layer-2 network

You know that you can run multiple layer-3 networks over a layer-2 VLAN. Did you also know that one layer-3 network can span multiple VLANS? We have become accustomed to associating layer-3 networks with VLANs but they're different animals. A VLAN is a virtual LAN--a virtual wire. Two VLANs are two virtual wires.

Below, we have a DHCP server on one LAN or VLAN and a potential client on another. But, they cannot talk because they are not connected.

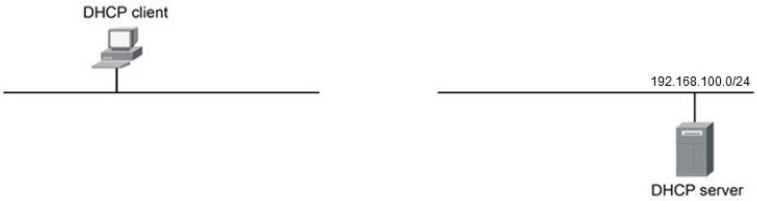

I can place a jumper between the networks. Now, the host on the left can reach the DHCP server on the right and obtain an address. In the picture below, VLAN 101 and VLAN 1, together, make up *one* layer-3 network with and address space of 192.168.100.0/24.

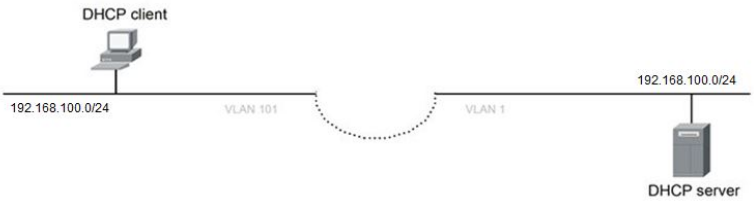

By replacing the jumper with a layer-2 firewall, we can segment the network without need to re-address anything. There are still two VLANs, but just one layer-3 network.

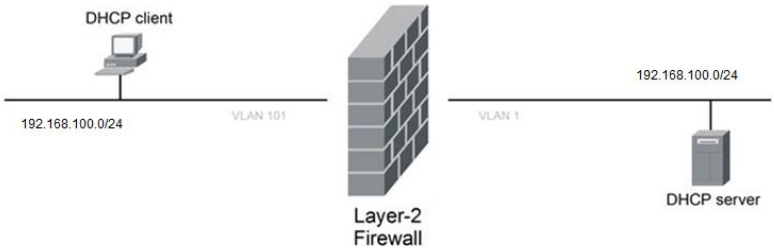

The traffic on the left and the traffic on the right are part of the same layer-3 network,. But, in order for the host on the left to talk to the server on the right, traffic will have to traverse a layer-2 firewall.

Here's an important point: even though a firewall may be filtering at layer-2, the rules can be written for layer-3. So, for example, this firewall may have a rule that rate limits SMB traffic, and rules that block RDP or portmapper access. It may also contain rules for layer-2 traffic, including protocols like ARP, UPnP or non-IP networking.

Configurable segmentation

Let's say that the following depicts ports on a layer-2 switch. The white ports are in VLAN 1; the red port is in VLAN 101. Whether a host is in VLAN 101 or VLAN 1 is determined by the access VLAN designation on the switch port. This can be changed on-the-fly; a host can be in VLAN 1 at this moment and in VLAN 101 the next.

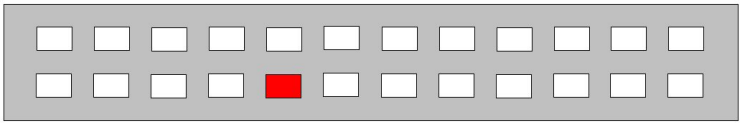

Below, we show a laboratory configuration using a Fortigate[2] as the layer-2 firewall and an Aruba CX switch hosting VLAN 1 and VLAN 101.

[2] One can fashion an L2 firewall one from a generic edge firewall by configuring interfaces for L2 traffic. Fortinet calls L2 capability ISFW (internal segmentation firewall). Palo Alto calls it VWIRE. Watchguard calls it LanBridge. This firewall will be able to filter traffic based on L3 firewall rules, L2 rules (some) and IDS/IPS UTM capabilities.

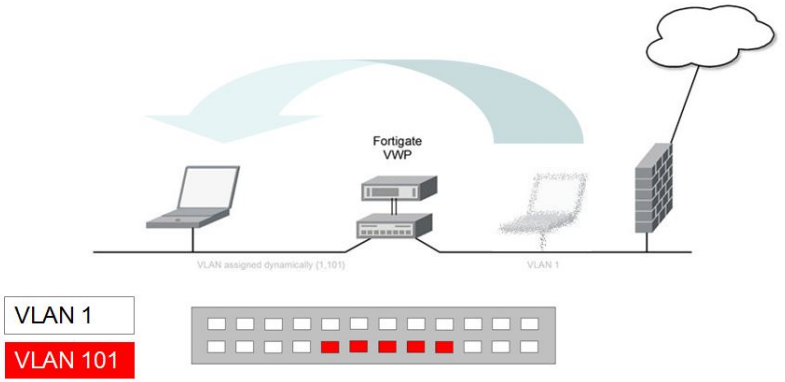

A host's position on the network—to the left or right of the layer-2 firewall—is determined by its VLAN membership, and this can change without affecting the user experience for permitted network traffic.[3] The arrow depicts the logical movement of a host from VLAN 1 to VLAN 101 by changing the port's access VLAN. One moment a user is in VLAN 1 and in the next moment, in VLAN 101. Save for a quick hiccup, the user will not be aware of the change.

Enterprise Deployment

A typical enterprise switching environment is spine/leaf, with multiple VLANs trunked down to access switches in multiple closets. Segmenting at layer-2, none of this would change. However, there will be additional access VLANs extended to each switch.

[3] Note: need to filter BPDU packets where VLANs meet the firewall. Otherwise, switch may see spanning-tree issue.

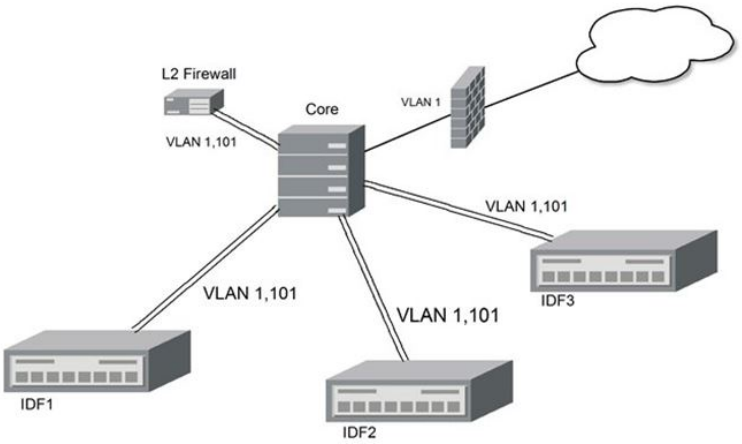

This picture shows VLANs 1 and 101 from our previous examples trunked from closets to the core switch and extended to a common internal layer-2 firewall. In order to avoid routing between VLANs 1 and 101 in the core, we configure separate virtual routing and forwarding domains (VRFs) in the core switch, as supported in the Aruba CX line and in some other manufacturer's switches.

Dynamically Configurable Segmentation

VLAN assignment automation is the only missing piece. It will require some scripting. Based on telemetry, we can create a whole-WAN segmentation solution. To automate inspection of the network, we could employ network taps to gather and report flow data in closets. Or, more economically, we can use VLAN assignment automation to examine different sections of the WAN in a round-robin basis. By this, we don't

need a whole-enterprise L2 data collection strategy; we can time-share our L2 firewall or IDS for telemetry at the core.

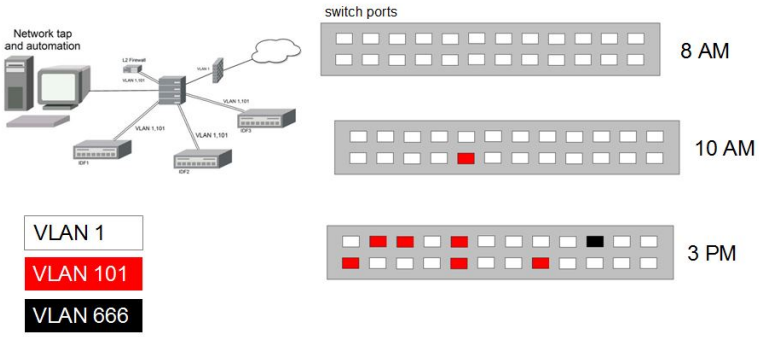

Since it is possible to switch a user's VLAN without interrupting their session, detection of unexpected SMB or email traffic, as examples, could result in a VLAN re-assignment that rate-limits or prohibits traffic, plus an email to the network administrator.

Aruba Networks Additional Capabilities

HPE Aruba Networks wired networking technology adds components that make internal network segmentation at layer-2 even more powerful:

1) CX Switch ports can choose VLANs by profiling devices with CDP or LLDP
2) ClearPass can choose VLANs for switch ports based on user attributes, and in some cases issue a CoA to change VLANs in mid-session.
3) The Aruba CX10000 is a TOR switch that can be used as a core switch both to enforce traffic at L2 and to issue client telemetry for making VLAN choices.
4) Dynamic Segmentation

Let's look at each in turn:

1) Access ports on Aruba AOS-CX switches can dynamically choose a VLAN based on LLDP or CDP-retrieved client information. When a device is plugged in, the switch can profile it locally. We use this capability to automatically configure trunk groups for APs and access VLANs for desktops. But, it can also be used to place a device into a untrusted VLAN (our VLAN 101 in the discussion above) so that we can monitor the client at the same time we allow it onto the network. This is similar to scanning, but without the agent or the attendant delay. When we're happy with its behavior, it can go into a trusted VLAN, unfettered.

2) Profiling and NAC are among ClearPass's many capabilities. If we were to employ 802.1x and at an access port, VLAN assignment could be made by ClearPass and changed later via a RADIUS Change of Authroization.

3) The Aruba CX10000 switch includes two AMD Pensando data processing units. This switch provides wire-speed security and telemetry in the core. With a CX10000 in an enterprise configuration, we wouldn't need a separate L2 firewall tromboned off the core; the core *would be* the segmentation firewall.

CX10000

4) Aruba networks wireless controllers support a wired network feature known as *Dynamic Segmentation*. In conjunction with switch-based profiling or ClearPass, an Aruba controller can build a layer-2 tunnel from a switch port to somewhere else—into the core or across the world. The traffic is carried in an encrypted layer-2 tunnel at layer-3. This could extend across the WAN. An example of how this could be used: one could sit in a remote office and connect to their home access VLAN in the main office as if they were sitting there.

Kevin Dowd
Atlantic Computing
atlantic.com

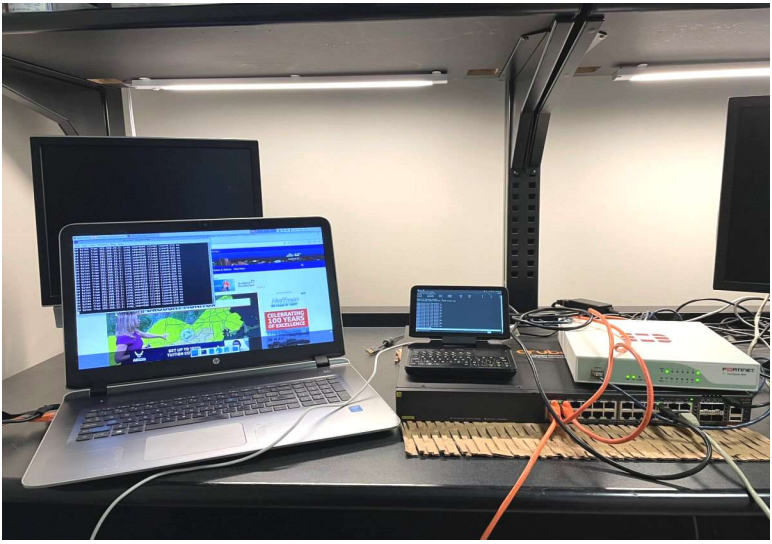

Laboratory testing with Aruba CX3000 and FortiGate

sponsored by: