

HIGHLIGHTS:

- Automates Quantum Tenant Network VLAN provisioning on Arista top-of-rack or leaf switches.
- The first OpenStack solution to allow hardware devices such as Arista switches to interoperate with the Open vSwitch plugin.
- Truly multi-vendor architecture allowing hardware devices from other vendors to co-exist with Arista switches and the Open vSwitch plugin.

OVERVIEW

OpenStack is being increasingly deployed in private cloud environments today. In many such deployments, Open vSwitch (OVS) is being chosen as the virtual switch technology of choice for server hypervisors. Though the Open vSwitch Quantum plugin supports both GRE tunnels and VLAN based tenant isolation, performance issues with GRE tunnels has led to the widespread use of VLANs for network isolation and segmentation.


Unfortunately, the current plugin model in Quantum only permits the deployment of either the OVS plugin or one of the several supported hardware device plugins, but not both at the same time. Using the OVS plugin requires the network operator to manually trunk/un-trunk VLANs as Quantum networks are created/destroyed. Clearly neither approach is desirable.

Arista's solution remedies this by contributing software to OpenStack that enables the Open vSwitch plugin to interface with a hardware device layer into which hardware drivers from different vendors can plug-in. An Arista driver is provided which interacts with the Arista switches to automate VLAN provisioning as needed. By deploying this solution, Arista switches automatically learn about Quantum Tenant networks as they are created/destroyed and can automatically manage the trunk VLAN lists on each trunk port leading to the compute nodes, thus providing tight integration between network and compute provisioning.

This solution also provides enhanced visibility into the network topology and how the virtual Quantum networks map onto the physical data center. It is possible to issue show commands on the Arista switches and gather information about how OpenStack VM instances are placed on different compute hosts and how they interconnect to different Quantum networks.

TECHNICAL INFORMATION

This section describes the solution in more detail. As the figure on the following page shows, the Quantum Network Controller node runs an enhanced OVS plugin which interfaces with the Arista driver through a public interface. The Arista driver communicates using the EOS API (eAPI) with a configured Arista spine switch which orchestrates all the racks in the OpenStack Cloud.



The figure above shows two OpenStack tenant networks (Red Network and Blue Network) for which corresponding Quantum networks have been provisioned. The OVS plugin decides on the VLAN allocation for these tenant networks - Red VLAN and Blue VLAN above. These VLAN choices are conveyed by the OVS plugin to the Arista driver.

The Arista driver in turn uses eAPI to communicate these VLANs to the Arista spine switch in addition to communicating to the compute host information about the particular VM that has been launched. By combining these pieces of information, the Arista spine switch is able to configure the trunk VLAN lists on the appropriate interfaces on the leaf switches (TOR1 above).

The communication between the Arista spine switch and the leaf switches leverages the same distributed system technology that is used for sharing state within processes in EOS. As VM's are terminated and relaunched on other compute nodes, or as new tenant networks are created, the Arista Quantum hardware driver keeps the spine switch updated, which in turn updates the trunk VLANs on appropriate interfaces.

This architecture is also resilient with respect to hardware/software failures. If the spine switch were to be restarted, or the communication between the Quantum network controller and spine switch were to be briefly interrupted, the Arista Quantum hardware driver resyncs its state to the spine switch when a connection is reestablished. If the leaf switches were to be restarted, the spine switch automatically reprograms the trunk lists on them when they come back online.

HARDWARE DRIVER API

We have added the following API between the OVS plugin and hardware device drivers. When the OVS plugin is loaded, it will load one or more hardware device drivers (based on configuration) and invoke the following API when tenant networks are created, VM's are launched, etc.

- `create_network(network_id)`
Configures an L2 isolated network on the vendor hardware for given tenant using the configured segmentation_type.
- `plug_host(network_id, segmentation_id, host_id)`
Adds a host to the network.
- `unplug_host(network_id, segmentation_id, host_id)`
Removes a host from the network.
- `delete_network(network_id)`
Removes the previously provisioned network segmentation from the switches

The Arista driver implementing this interface will be provided with the OpenStack Quantum distribution. The API is open and allows other hardware device vendors to implement similar drivers to create a multi-vendor solution. Arista has open-sourced the driver and the entire framework back to the OpenStack community.

Arista is currently in the process of getting this API standardized and pushed into the Havana release of OpenStack. However the Arista solution will be available from our website for early adopters and should be usable with OpenStack Folsom/Grizzly.

CONCLUSION

The Arista OpenStack solution automates the provisioning of VLANs on Arista switches to work in concert with the OVS plugin. This is in contrast to the other OpenStack plugins available today, which allow deployment of either the virtual switch plugin or the hardware device plugin, but not both. The architecture implemented is open and vendor agnostic. Providing this capability within the OpenStack solution is a significant achievement for integration with physical network devices. The availability of the drivers published back into the OpenStack community by Arista Networks can lead to truly multi-vendor solutions for VLAN provisioning automation within OpenStack environments.

ARISTA

Santa Clara – Corporate Headquarters

5470 Great America Parkway
Santa Clara, CA 95054

Phone: 408-547-5500
Fax: 408-538-8920

Email: info@aristanetworks.com

San Francisco – R&D and Sales Office

1390 Market Street
Suite 800
San Francisco, CA 94102

India – R&D Office

Eastland Citadel
102, 2nd Floor, Hosur Road
Madiwala Check Post
Bangalore - 560 095

Vancouver – R&D Office

Suite 350, 3505 Gilmore Wayt
Burnaby, British Columbia
Canada V5G 4X5

Ireland – International Headquarters

Hartnett Enterprise Acceleration Moylish Park
Limerick, Ireland

Singapore – APAC Administrative Office

9 Temasek Boulevard
29-01, Suntec Tower Two
Singapore 038989

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

© Copyright 2013 Arista Networks, Inc.

Information contained herein is subject to change without notice. Arista and the Arista logo are trademarks of Arista Networks Inc .