

EVE VLAN support - switch network instances

Background

Several enterprise users expressed desire about using VLANs. Gopi started this document based on the his discussion with one such enterprise user.

The first step requirement (the topic of this document) is to focus on application traffic on a switch network instance, where some application instances will use trunk ports (for example, a firewall of virtual router VNF), and other application instances will be using an access port configured to be on one VLAN.

The switch network instance may or may not have an external port. If it does, such a port will by default be a trunk port.

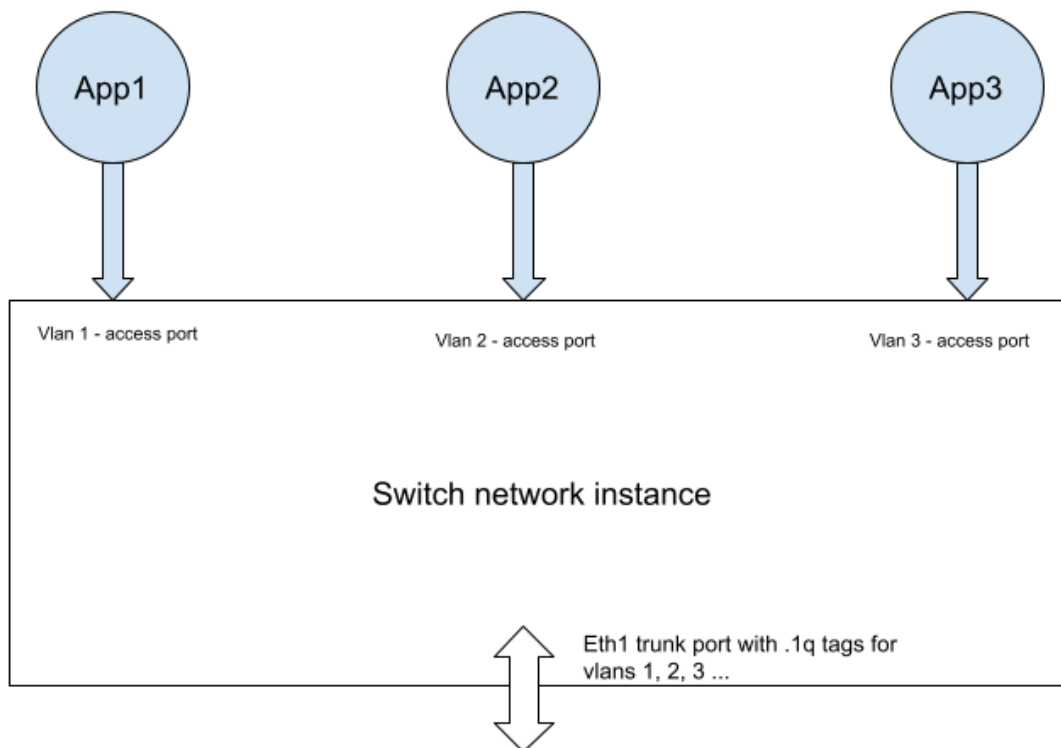
Note that the current EVE implementation of a switch network instance is completely transparent to the Ethernet types and payload, hence carries 802.1Q tags transparently. Thus in effect, today all the attachments to a switch network instance are trunk ports (allowing all VLANs and also allowing untagged packets).

Out of scope

Using VLAN (and/or bonds/LAGs) for EVE management traffic is out of scope for this document. That would require expanding either the model (to carry L2 networking config) and/or expanding the configuration of adapters in the UI (to allow the definition of VLANs and/or LAGs) for an edge node.

Impact

- Single new field added to app instance network adapter
 - UX question about busy screen?
 - When we later add jumbo-frame support there will be yet one more field on the same screen
 - UI validation: Either unspecified (meaning zero in API), or number between 2 and 4093.
- Passthrough of that field through zedcloud
- (Thus new fields in zedcloud and EVE APIs.)
- New behaviors in EVE to configure the access VLAN
- Test and documentation impact; get the news out somehow?
 - Larger question about workflow vs. object
- API transition? vlan=0 means old behavior
- Likely that we need to add more policy??



Assumptions:

- 1) Apps are configured with their virtual NICs made part of the desired vlan (UI/UX needs to be figured).

2) EVE does not care how these apps get their IP addresses. And EVE does not have to support providing them with IP addresses. Basically no EVE IPAM support for switch network instances.

3) The trunk port of the switch network instance is connected to another managed switch that understands VLANs.

4) The trunk port of the switch network instance is connected to another device (may be a switch as above) that can then trunk these packets to a router/firewall. No EVE involvement needed

in routing packets between VLANs or to the router/firewall.

All EVE does is recognize that an app packet comes from an access port (with vlan) and then trunk the packet out with 802.1q vlan header.

Additional configuration

In the first phase it is sufficient to add a "access vlan" field in the UI for the application instance network adapter configuration (where Sahil just added a Mac Address field). This is only needed for a switch network instance. Allowed values between 2 and 4093.

This means carrying that field in the controller and EVE APIs.

In a subsequent phase we might be asked to add VLAN filters on the trunk port (which might be a list of VLAN ranges such as "1-10,20,30-39" if that matches how network gear configures trunk ports.)

If we want to specify VLAN filters on the uplink port(s), then we need a place to put such configuration. That most likely falls in the out of scope category above.

Implementation notes

EVE will use the access VLAN to set up a port-based VLAN for the vif handed to the app instance. If the field is zero (aka missing in protobuf), no VLAN will be configured.

If we later add VLAN range filtering on the trunk ports, then we need to configure VLAN filtering in Linux. When VLAN filtering is configured, if the API specifies no access VLAN and no VLAN ranges, then for backwards compatibility we must assume that the port should be configured to allow all VLANs 2-4093.