

## Topic

- Network
  - Why not use FOAM everywhere?

*You don't need to expose it to experimenters - just call FOAM instead of calling FlowVisor. You don't even need to write specs or use the GENI AM API to call FOAM, we are adding a pluggable front-end API to FOAM, so we could work out a nice ExoGENI plugin to minimize the work you needed to do. If you write directly to FlowVisor then you'll have to track FlowVisor changes and add support for new features to avoid getting out of sync with the rest of GENI, which seems like an unnecessary ongoing maintenance issue.*
  - Why not run pure OpenFlow and slice on VLAN in FlowVisor w/translation at the rack edge?

*Running in a hybrid mode creates limitations on what the openflow datapath is capable of, particularly in relation to VLAN tagging and QoS*
  - How is IP space managed within the rack environment - can experimenters request more / specific IP space?
  - The OpenFlow control channel looks to be extremely throughput constrained.

*The connection from the head node to the OpenFlow switch appears to be at most 1Gbps (shared), and the aggregate throughput of the (current) data plane connections is at least 420Gbps. Even if experimenters didn't do bad things, only having 0.25% of the data plane traffic sourced as packet-ins seems overly optimistic.*
  - Does the switch not support the ENQUEUE action at all, or does it just not support all the openflow packet-queue structures?
- Rack Configuration
  - Is there an IPMI connection from the head node to the management switch? If so I think that makes for 45 management switch ports used.
  - I am concerned that the head node is under provisioned for all the services it needs to run - 12GB of ram seems low.
  - How is the head node configured - do the services run in their own VMs, or do they need to co-exist on the same OS instance?
  - PDUs are also useful for remote management if a node gets completely bricked (such that IPMI is useless) - I would think that the marginal cost would be more than worth it.
- Resources
  - Why not allow arbitrary bare-metal images? Is this any more dangerous than arbitrary VM images?
  - Where is the storage for the running instances - on the worker nodes?
  - What are the average IOPS available for each VM on a fully loaded (max running VMs) worker node?