OPNFV SFC Brahmaputra

Release Planning

Brady Johnson
June 26, 2015

Detailed Requirements in later slides

Project Box

Minimum viable requirements
(Cant release without these)

- Deploy a complete SFC solution
  - Create BGS/Genesis/Fuel scripts
- SFC ODL and OpenStack Coordination
  - SFC needs info from OpenStack about the Service Functions VMs

Want but at risk requirements
(Like to have, but high risk)

-

Working plan requirements
(Minimum Viable, low risk, but high effort needed)

- Clustered ODL SFC deployment
- VNFFG Requirement compliance

Outplan requirements
(Future requirements not in this release.)

- Support multiple Compute nodes, 1 SFF each

OPNFV

Risks Assessment

Description	Likelihood	Impact	Trigger	Mitigation	Contingency	Owner	Status
OVS NSH patch is not formalized	Medium	High	IETF NSH spec not formalized	Work with OVS upstream	Don't release SFC or use different encap approach	Brady	Ongoing
ODL SFC clustering does not work	Medium	Low	ODL cluster testing fails	Perform ODL clustering tests	Don't use clustering	Brady	Ongoing

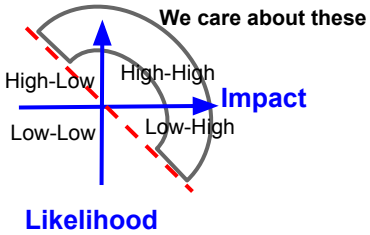

Minimal Viable *Detailed* Requirements

Taken from the “Project Box” on previous slides

The OPNFV SFC Brahmaputra Release ***cannot*** release without these:

Notice:
This is still a draft

- Deploy a complete SFC solution
  - Create BGS/Genesis/Fuel scripts
    - Control node with 2 SFs and the SFF in the host OVS
      - This depends on the SFC ODL/OpenStack coordination (which SFs to create?)
    - Deploy a Compute Node with ODL SFC, OpenStack, and an injector
    - Need to distinguish between initial setup versus rebooting once everything has been setup
    - Needs to allow for multiple Compute Nodes: In the future this will need to be supported

- SFC ODL and OpenStack Coordination
  - SFC needs the following info from OpenStack for each Service Function (SF) VM:
    - IP Address
    - encapsulation details (VxLAN, NSH enabled)
    - OVS switch and port the SF is connected to
  - SF creation alternatives (the proactive approach will most likely be the way to go)
    - **Reactive approach**: SFC can query OpenStack for SF VM info, but we need to know which VMs to query
      - For example, it may be that 10 VMs are started, but only 2 of them are SFs, which ones do we query?
    - **Proactive approach**: SFC receives the JSON SF config, and for each SF, request OpenStack to create it and return the necessary info. Then SFC will have to complete the SF config with the retrieved VM info.
      - This is only needed for initial SF creation.