NFVIaaS Architectural Framework for Policy Based Resource Placement and Scheduling

draft-krishnan-nfvrg-policy-based-rm-nfviaas-01

IETF 91

Ram (Ramki) Krishnan & Norival Figueira - Brocade Communications

Dilip Krishnaswamy - IBM Research

Diego Lopez - Telefonica I+D

Steven Wright - AT&T

Tim Hinrichs – VMware

Introduction

- What is NFVIaaS?
  - A third party offers NFV infra as a service to SPs; key reasons
  - Internal separation of concerns in the SP
  - Infrastructure sharing by SPs
  - Meet performance requirements (e.g., latency or throughput) - locations where the customer SP does not have physical data center presence
  - Meet regulatory requirements
  - And more …
- NFVIaaS Challenges
  - VNF's stringent service level specifications (SLS) required by customer SPs.
  - NFV PoPs often have capacity, energy and other constraints.

NFVIaaS Architectural Framework
Policy-based resource placement/scheduling

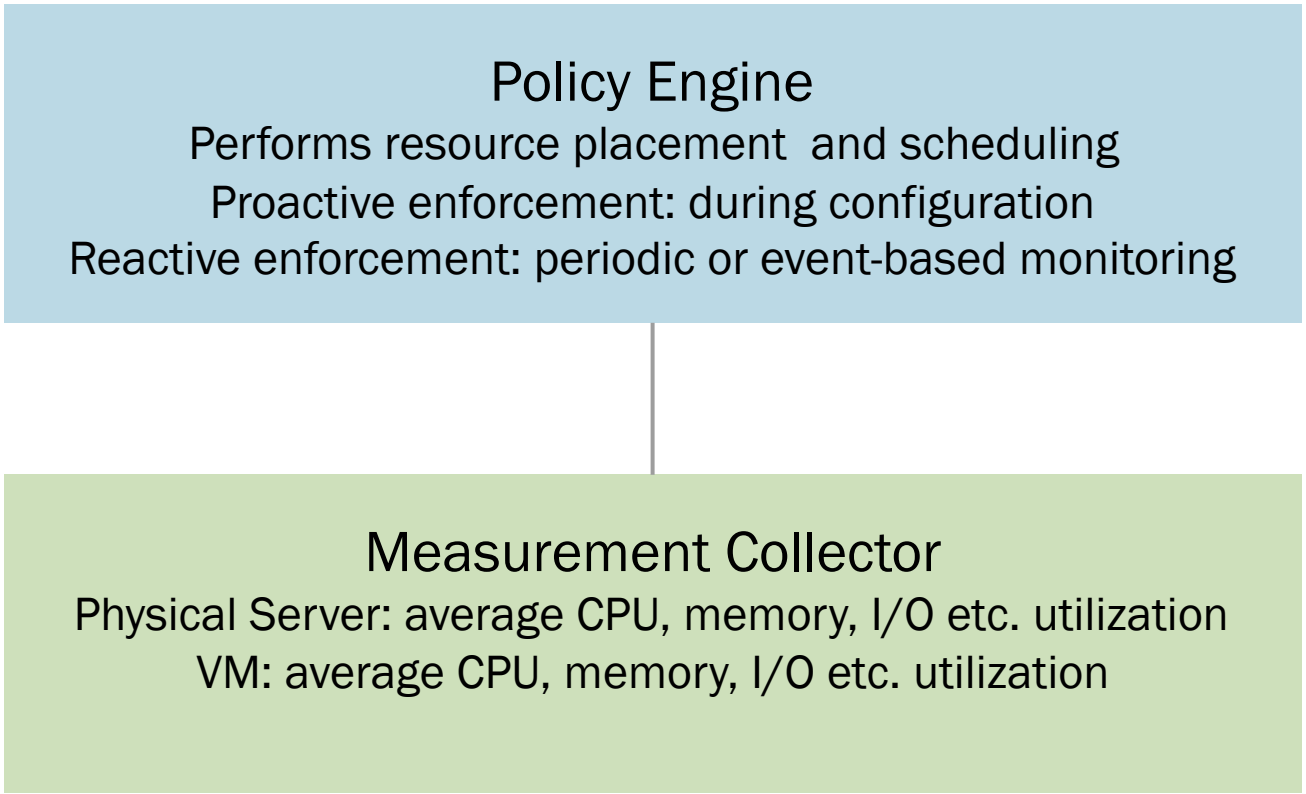

System Analysis in OpenStack Framework (1)

- Policy Engine -- OpenStack Congress
- Measurement Collector -- OpenStack Celiometer
- Exemplary mini NFV PoP configuration
  - Multiple physical server types – same or different manufacturer
  - Multiple NFVIaaS instance types – an instance is a VM
- Exemplary NFV policy
  - Global policy across multiple sub-systems
  - Description: For physical servers of type 1, there can be at most only one physical server with average overall utilization less than 50%.
  - Objective: Address the energy efficiency requirements specified by ETSI NFV by ensuring that servers are not kept under low utilization (Note 1).

Note 1: Servers have a non-linear power profile and exhibit relatively higher power wastage at lower utilization. For example, in the active idle state as much as 30% of peak power is consumed.

System Analysis in OpenStack Framework (2)

- Implementation Summary
  - Policy expressed in Datalog policy language by OpenStack Congress
  - Monitor OpenStack Congress periodically or based on an event (e.g. customer instance addition/deletion) average physical server utilization for server type 1 by querying OpenStack Celiometer
  - On detecting policy violation, OpenStack Congress uses constraint based placement techniques to find the new optimized placement(s) for physical server type 1 to address the policy violation.
  - OpenStack Congress performs NFVIaaS instance (VM) live migration to addressed the policy violation.

- Related Work
  - A related proof of concept in ETSI NFV on placement and scheduling - http://nfvwiki.etsi.org/index.php?title=Constraint_based_Placement_and_Scheduling_for_NFV/Cloud_Systems

Next Steps

- Current architectural framework maps to existing OpenStack modules
- Draft Progression
  - Examine the need for a separate OpenStack module for placement and scheduling
  - Policy Engine <-> Measurement Collector API information model definition