

# NFV Real-time Analytics and Orchestration: Use Cases and Architectural Framework

draft-krishnan-nfvrg-real-time-analytics-orch

NFVRG IETF91 – Honolulu

Diego R. Lopez – Telefónica  
Ramki Krishnan – Brocade  
Dilip Krishnaswamy – IBM  
Steven Wright – AT&T  
Norival Figueira – Brocade  
Asif Qamar – Evolv

# The Scenario

- The NFV promise combines
  - Elastic and resilient service provisioning
  - Optimized infrastructure usage
  - Operational simplicity
- A highly dynamic, changing environment
  - An integrated resource view
  - A holistic service orchestration
- An ideal field for applying real-time analytics across multiple sub-systems
  - Even more a need than a feature

# The Goal

- Outline an architecture framework for the integration of real-time analytics in NFV orchestration
  - At all levels: VNF, VIM, VNFM, NFVO
- Explore the applicability of machine learning techniques for NFV orchestration
  - In line with related efforts around SDN
- Consider initial use cases
  - Resource usage and security

# The Use Cases

- Optimizing resource usage
  - Relying on current usage monitoring
- Managing noisy neighbors
  - Detect and correct interference by analyzing correlated usage of common resources
  - Learn correlated patterns to avoid them at deployment
- Detecting inconsistent configuration
  - Analyze behavior patterns against expected ones
- Achieving QoX
  - Identify patterns that translate in better (perceived) service by any entity in the network

# The Framework

- A set of distributed probes collecting data
  - At the infrastructure, the VNFs and the endpoints
  - Suitable to be dynamically deployed or activated
  - An aggregation layer to avoid event flooding
  - Event-based, avoid polling where possible
  - Could leverage IPFIX, LMAP, Ganglia etc.
- An analytics application able to generate recommendations to the orchestrator(s)
  - APIs at the VNF, VIM, VNFM, NFVO...
  - Reactive and predictive
- Machine learning in the loop
  - At the application
  - At the orchestration elements

# Summary and Next Steps

- First reference framework defined, including machine learning considerations
  - Refine it
  - Accommodate different autonomy levels and the need for coordination among them
  - API information model definition
  - Mapping to OpenStack, OpenDaylight, OPNFV...
- Initial use cases identified
  - Explore the full range
  - Systematize practices
- A couple of practical projects on their way