



2016 中国云计算技术大会  
Cloud Computing Technology Conference 2016

CSDN

# OpenStack and NFV

April 2016



## Web 2.0

- Jeopardizing CSP' s revenue
- OTT exploding, customer affinity eroding
  - new compute platform, new challenges



## ARPU Down

ARPU is falling due to:

- Increasing churn
- Growing customer acquisition and retention costs
- Limited ability to monetize increasing and outlier utilization due to Net Neutrality regulation



## Agility

Not able to add VAS\* fast enough  
Need Service Agility  
While reducing OPEX and CAPEX?

## Traditional Network Function


### SERVICE PROVIDER APPLICATIONS



### VIRTUAL TRADITIONAL NETWORK APPLIANCES



## Network Function Virtualization (NFV)

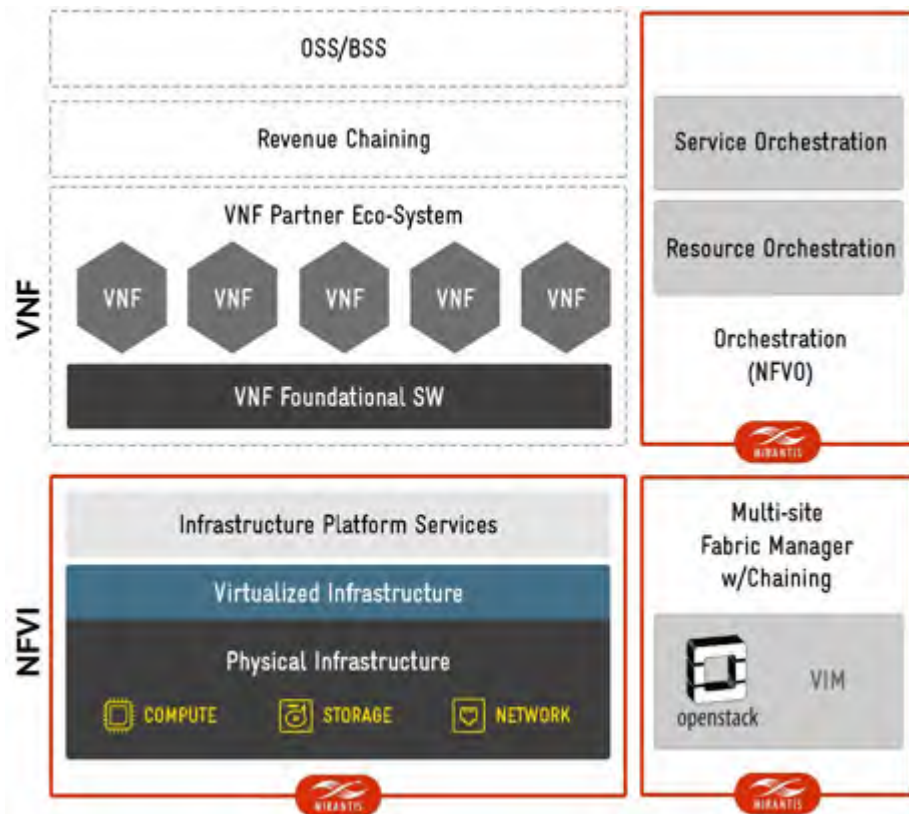


### OPENSTACK



## Benefits

- Improve service delivery cost structures
- Accelerate service innovation and creation
- Customer wallet share
- New customer segments



## Mirantis OpenStack Web-scale Distribution

- #1 Purpose-built OpenStack installer
- Hardened, support, HA
- Optimized for NFV

## Management and Orchestration (MANO)

Orchestration = Service Orch (SO) + Resource Orch (RO) Make Murano de-facto RO, and part of OSM

Murano & Application integration

## NFV Infrastructure

### Data Plane Acceleration

- High performance OVS with SRIOV
- OVS with DPDK support in-progress

Today  
(Siloed)



Transformation

Key Enablers:

- Virtualization & x86 compute
- Automation
- Multi-tenancy

Tomorrow  
(Cloudified)



Incumbents

Cloud era (Pure

Gen

Vendors





## Open

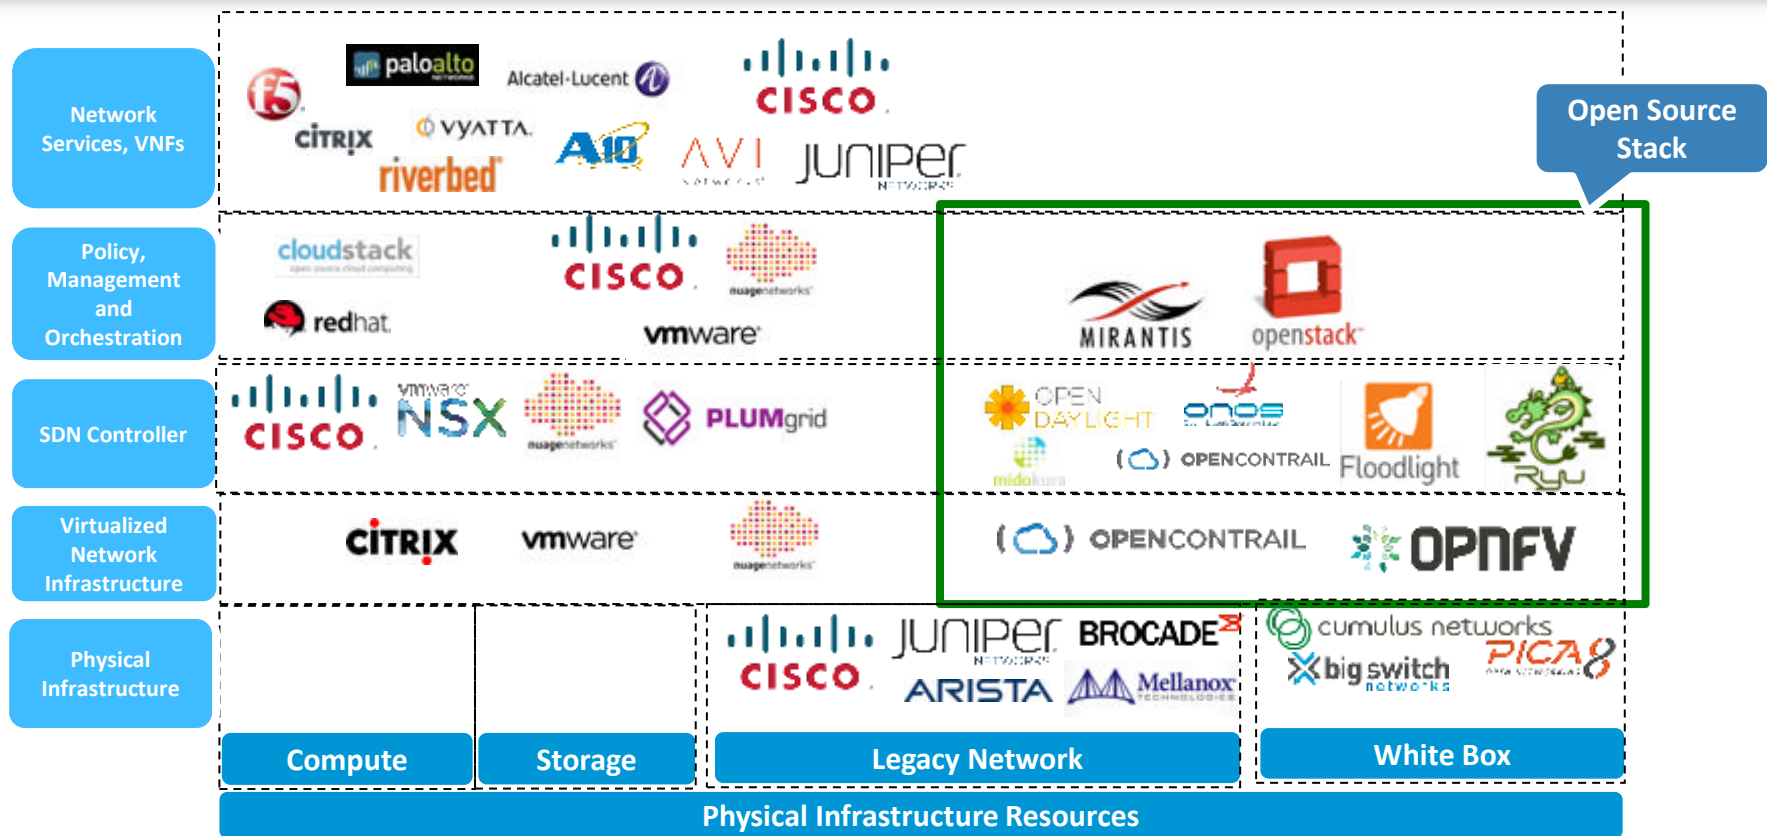
- Open community approach to developing NFV platform
- ETSI-NFV compliant
- 100% Open Source, No vendor-lock in

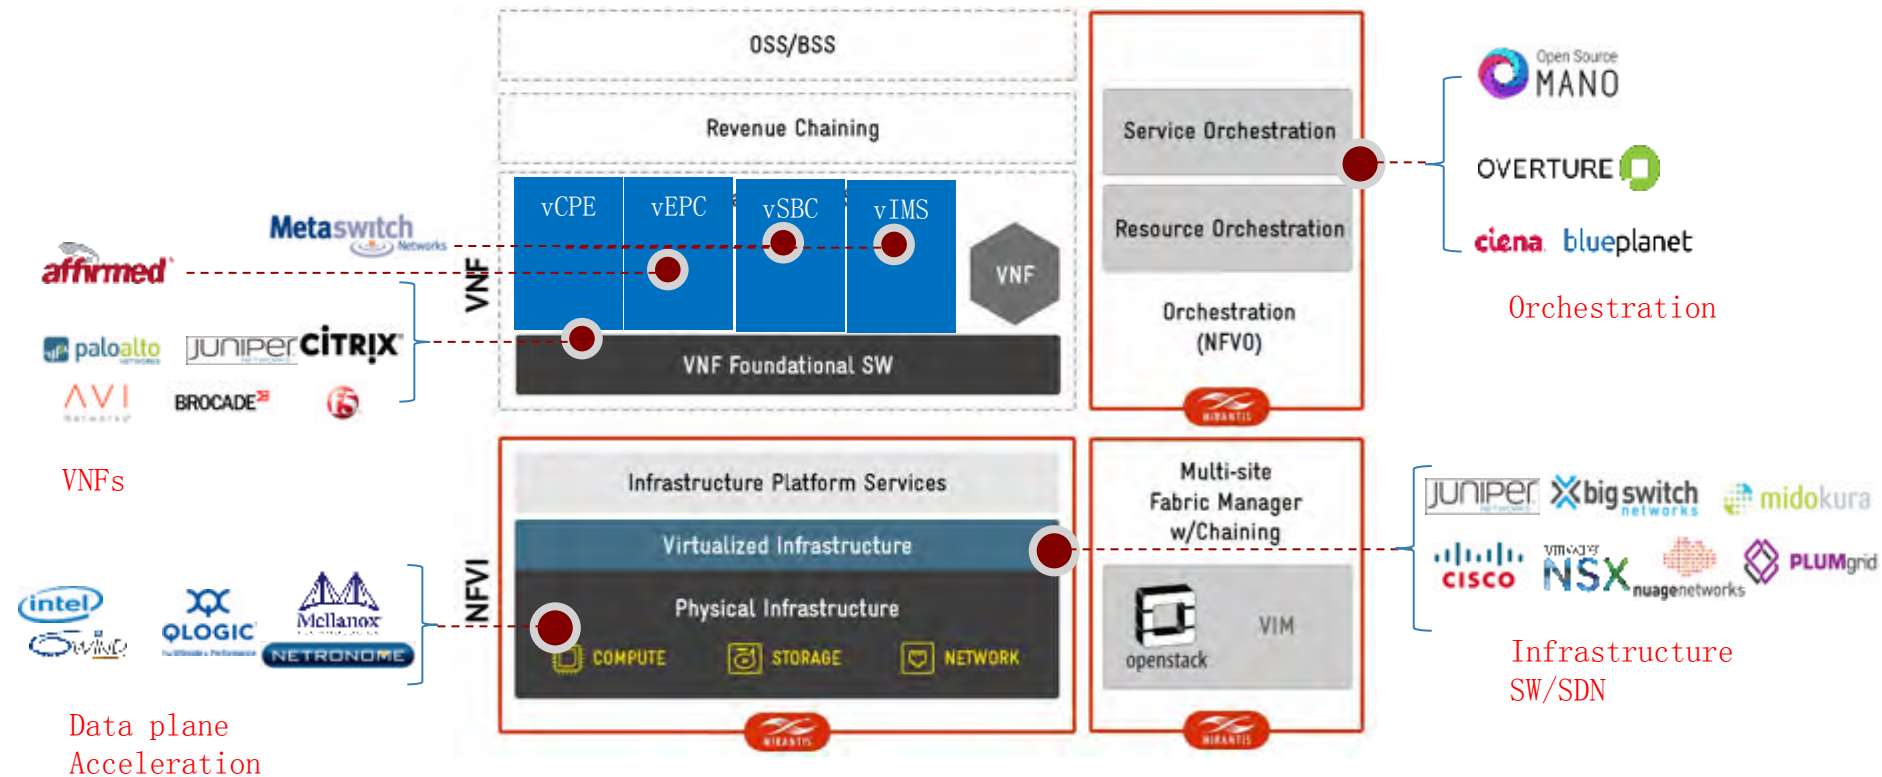
## Unified

- One Cloud for NFV, IoT, and Enterprise IT workloads
- Best-of-breed, validated solutions from large partner ecosystem
- Purpose-engineered by Mirantis: the leading OpenStack contributor

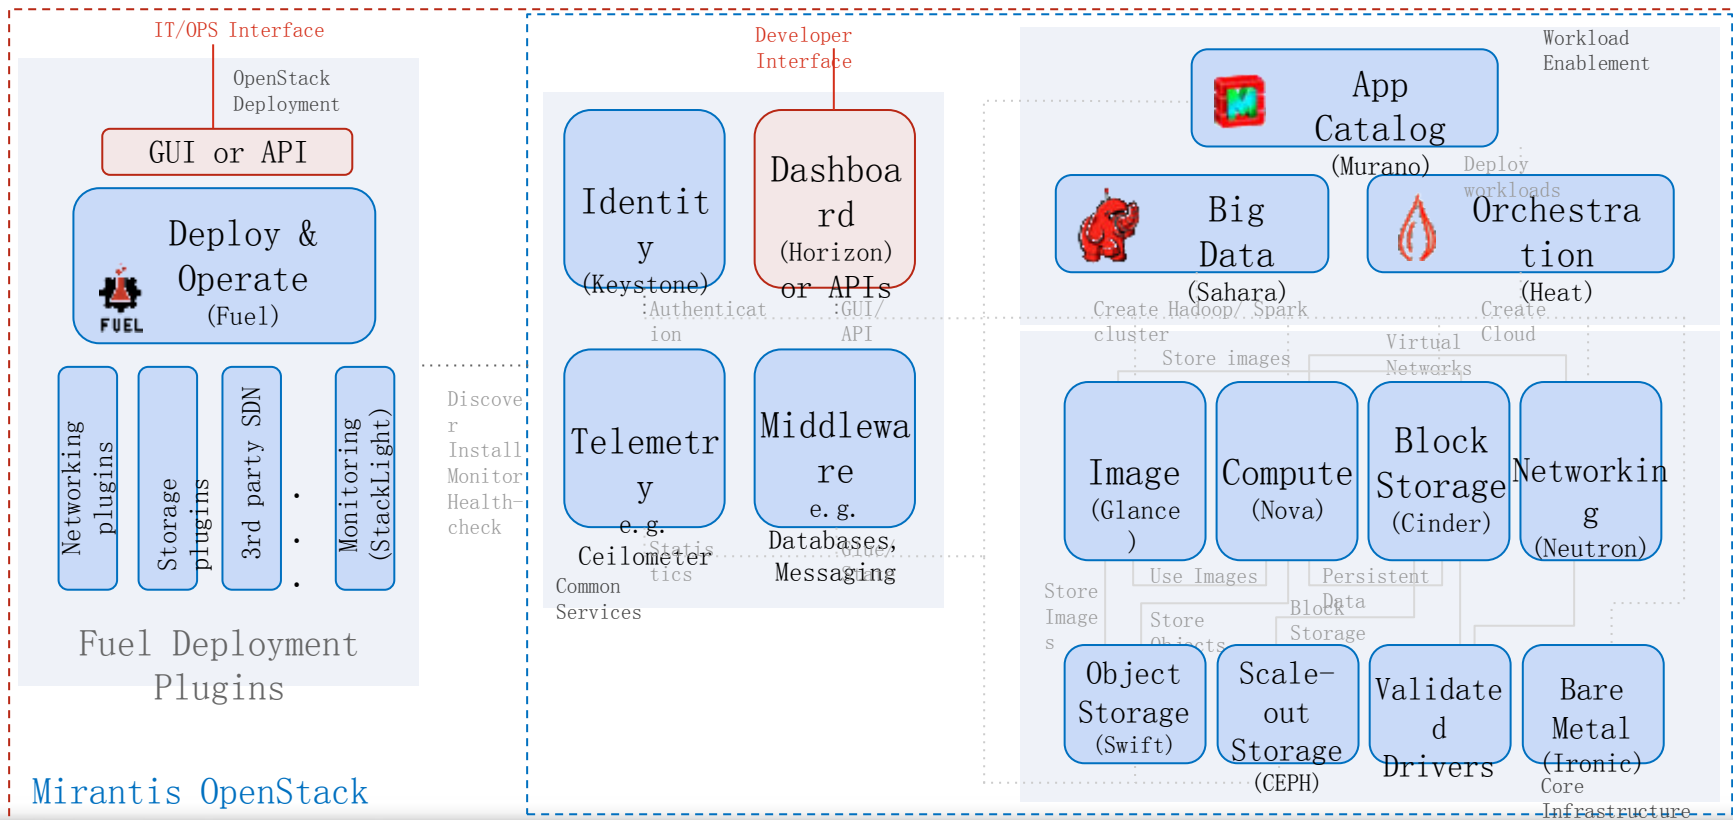
## Agile

- Software-defined service definition
- Scalable and Extensible Architecture









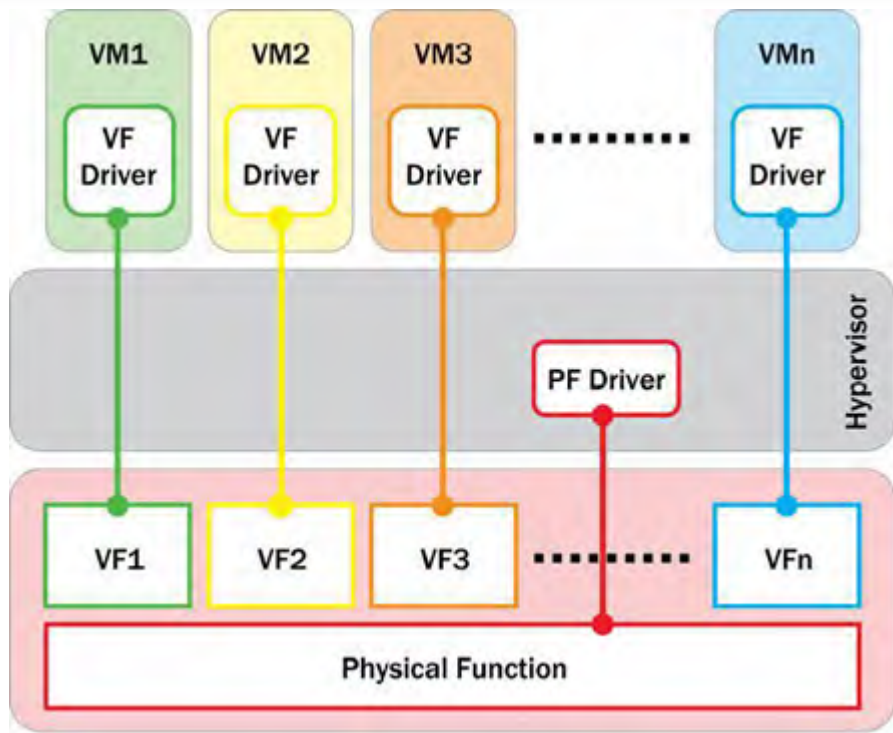
- Scalable Cloud
- High Performance Low Latency Compute
- Accelerated Virtual Network
- Monitoring and Service Assurance

- Multi Region
  - Shared Nothing - Global
  - Shared Critical Components – Country
  - Shared Infra - Metro
- Cluster Scale
  - Out of the Box – 200
  - Buildable to – 1000
  - Mirantis and Co-Dev partners roadmap
    - Out of the box: 500, 1000, 5000

- CPU and Memory Pinning
  - Guaranties Resources
  - Limit Response Jitter
  - Eliminates Noisy Neighbor
- NUMA aware scheduling
  - Optimizes IO and Memory Transfers
- Huge Page Support
  - Enhances Virtualized Performance
  - NUMA and IO-NUMA aware

- Preserve Openness, Provide Choices
- Hardware Acceleration – SR-IOV
  - Ultra-low latency
  - Very core efficient
  - Basic Multi-tenancy
  - Interface dependent
- Software Acceleration - DPDK
  - Flexible
  - Low Latency
  - Advanced network encapsulation
  - Advanced features: Service Chaining, Security
- QOS





Number of VFs depend on NIC  
Intel Niantic up to 64.

- NIC divides itself into PF and VFs
  - PF control interface
  - VF virtual functions
    - VF has at least one RX/TX queue pair
  - VFs are given to VM as PCI pass-through devices
  - Packet classifier on a NIC is used to steer traffic
- Advantages
  - DMA into Guest, no host CPU is used to move packets
  - No vSwitch latency
  - Multi-tenancy – VLANs
  - Anti-spoofing support
- Limitations
  - No VXLAN in current NIC generation
  - No Firewall (Need switch ACL integration)
  - Cannot run OVS on VFs (promisq mode is required)
    - Can run vRouter/vRouter-DPDK on the VF

- Packet rates
  - 64 byte – 13.95 mpps (test tool limitation)
  - 128 byte or greater - wirespeed
  - Multi-tenancy – VLANs
  - Anti-spoofing support
- Jitter
  - 64 byte - .01 mpps
  - 128 byte or greater – too low to measure

- Intel (Supported by Mirantis)
  - Niantic
  - Forteville
- Mellanox (Supported by Mellanox)
  - Connect X3
  - Connect X3pro
- Qlogic (Supported by Qlogic)
  - 3xxx
- Netronome (future)
  - OVS acceleration
  - vRouter acceleration

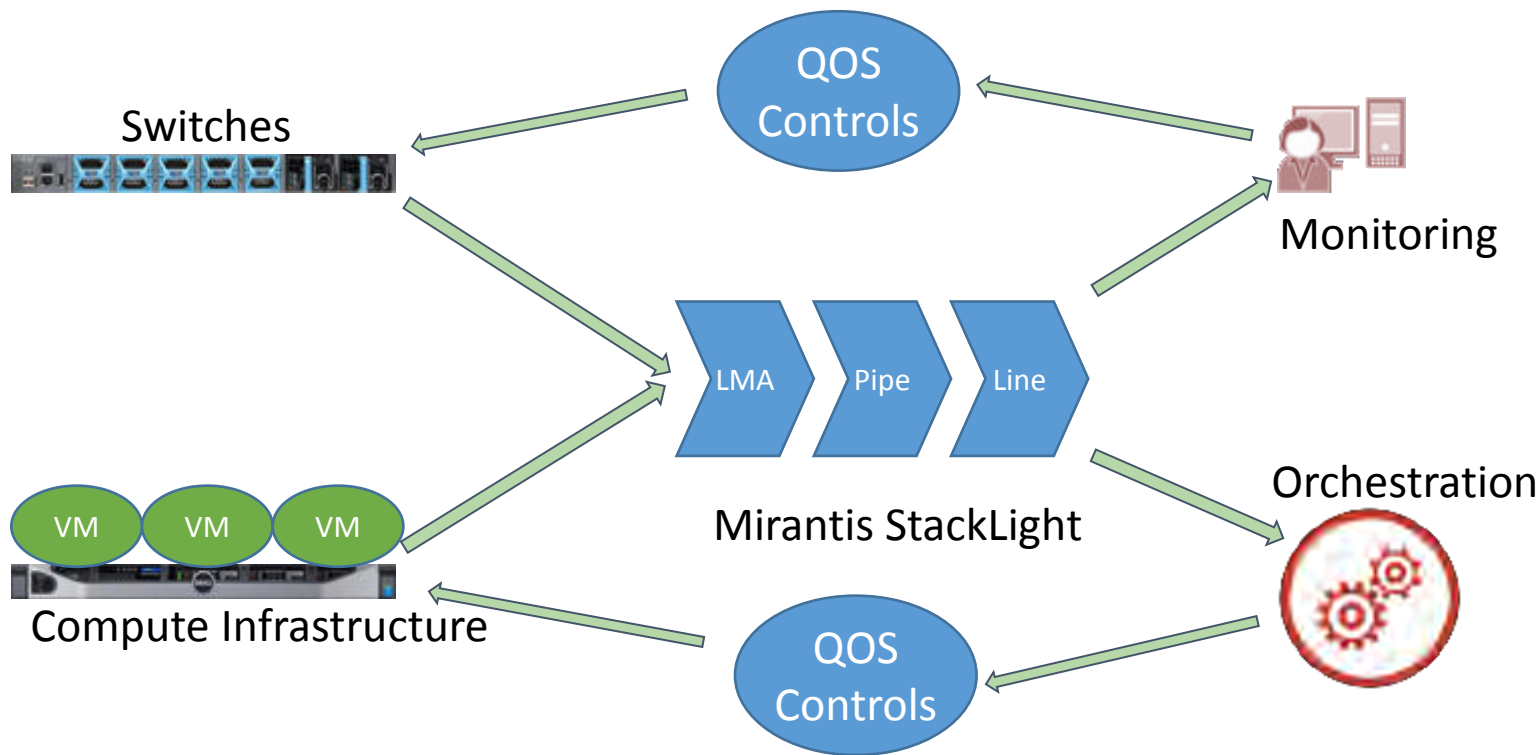
- OVS/DPDK
  - Open Source
  - Fast Innovation Cycle
  - Large Community
- vRouter/DPDK
  - Flexible
  - Powerful feature set
- 6Wind Virtual Accelerator
  - Ultrahigh performance
  - Good host and VNF solution



- Community effort
- OVS 2.4 or greater
- Bypass Kernel
  - 3 mpps per core
  - Significant improvement in progress
- Current limitations
  - VLAN only (MOS 9, VXLAN MOS 10)
  - Security groups under development

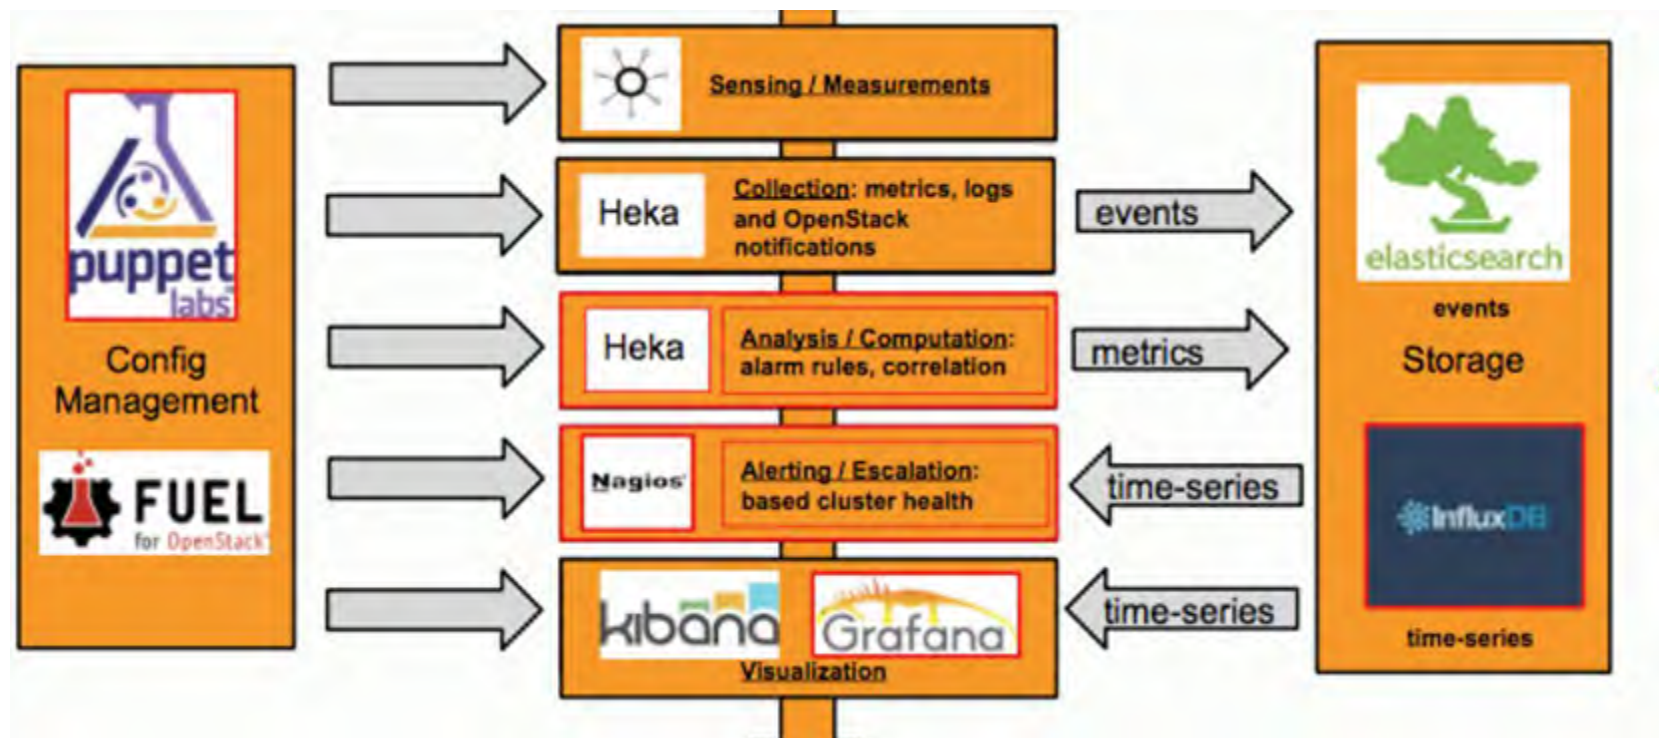
- OpenContrail/Juniper effort
- Contrail 3.0 or greater
- Bypass Kernel
  - 2 mpps per core
- Full Contrail features supported

- Commercial product
- Integrated with MOS 7 or greater
- Hooks existing OVS
  - 6 mpps per core
- Features are extensible



- Flexible collection infrastructure
  - Collectd
  - Heka
- Heka - Integrated Scalable Transport
  - Logs and Sample data
  - Real Time Alerts
- Flexible, Scalable, Efficient Storage
  - InfluxDB
  - ElasticSearch
- Effective Presentation of Data
  - Graphana
  - Kibana
- Stand-alone monitoring or feed into other systems
  - Lots of connectors are available
  - New ones can be easily written
  - Good way to provide service assurance feed to Orchestrators





# Questions

Thank You

for your time