





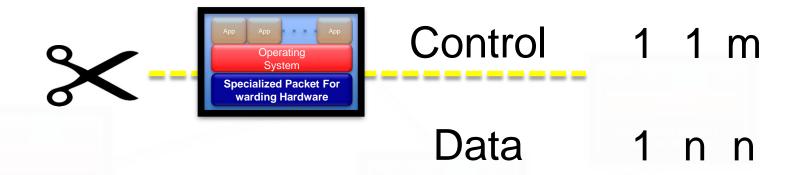
SONA: ONOS SDN Controller based OpenStack/Kubernetes Network Management Solution Trellis: Multi-Purpose Leaf-Spine Fabric Solution

Feb 22, 2019

Sangho Shin

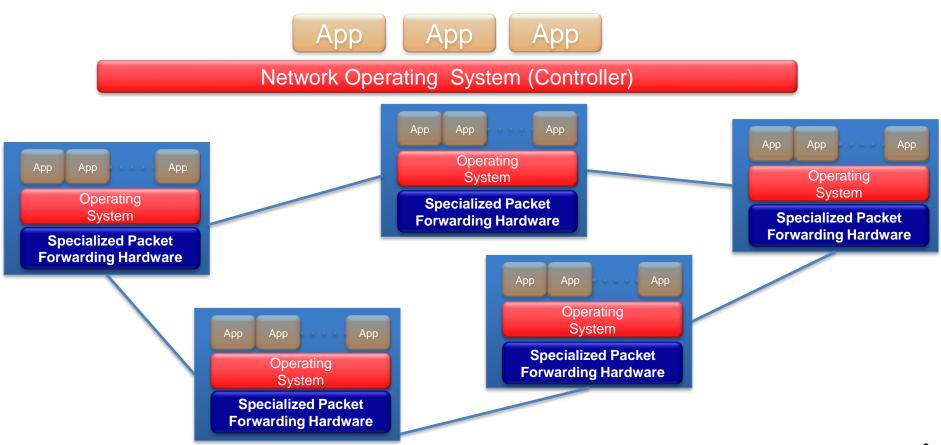
Introduction to SDN & ONOS

Software-Defined Network (SDN) (1/2)



Other aspects of SDN follow

Software-Defined Network (SDN) (2/2)



SDN Evolution and ONF



- Non-profit, carrier and vendor neutral
- Provide technical shepherding, core team
- **Build community**
- Many organizations supports















Demonstrations

Platform Development





2007 - Creation of SDN Concept





































And Beyond











2012 - Define SDN





ONOS Architecture (1/2)

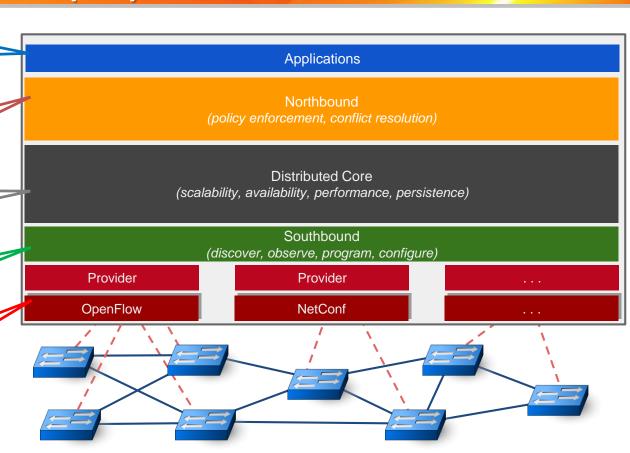
Contains user applications E.g., reactive forwarding, ProxyARP, segment routing, SDN-IP, etc.

Transfer network info to app layer Provide management interface for controlling lower layer component

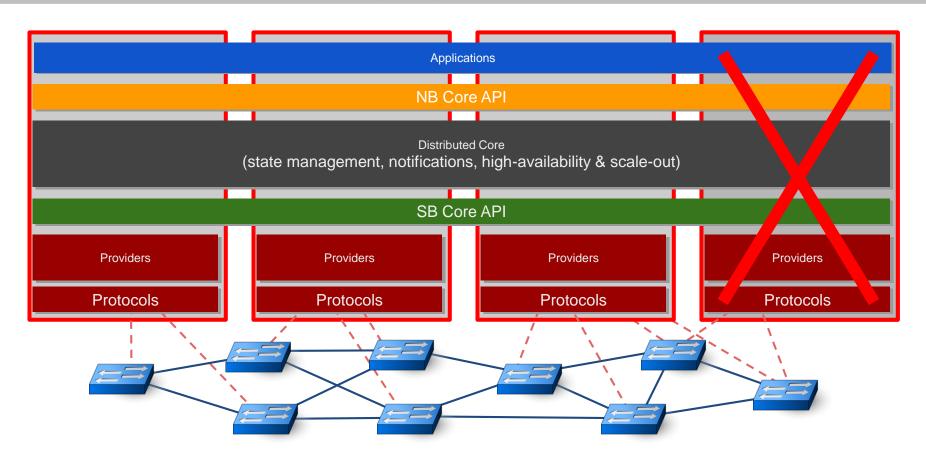
Contains many core features Provide distributed clustering func. for supporting HA and scalability

Provide an abstracted interface for controlling the network infrastructure

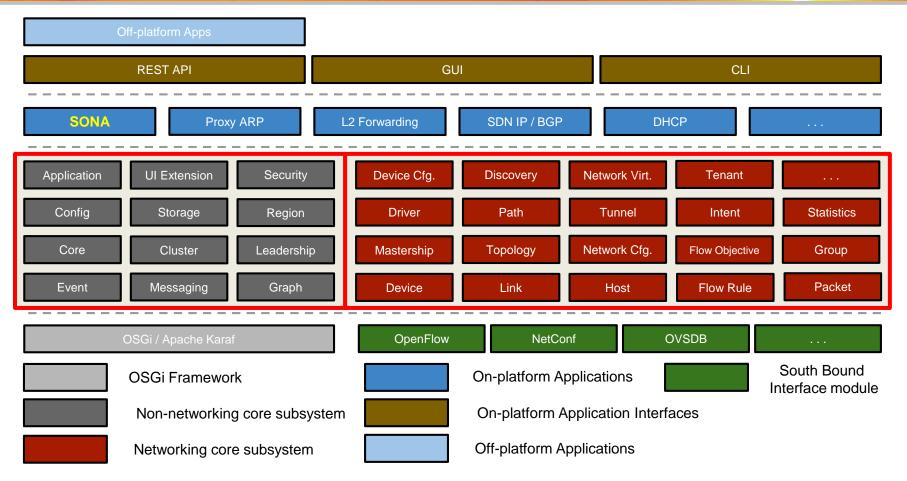
Network protocol implementation for managing network elements E.g., OpenFlow, NetConf



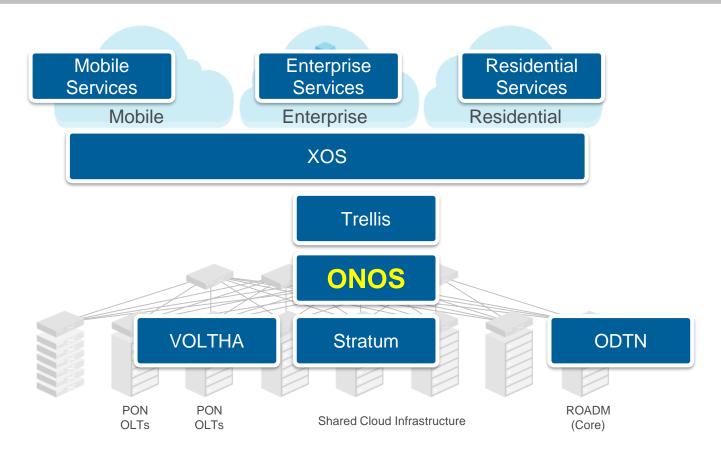
ONOS Architecture (2/2)



ONOS Subsystems (Services)



ONOS Project at ONF



ONOS Community

Partners

Collaborators





ONOS Release History



Q4/14 Avocet

Base Architecture



Q1/15 Blackbird
Performance



Q2/15 Cardinal

ONS Use Cases SDN-IP Packet Optical R-CORD



Q3/15 Drake

ONF ATRIUM Secure Mode ONOS VxLAN Device Configuration



Q4/15 Emu OPNFV SONA AARNET

KREONET-S



Q1/16 Falcon

ONS Use Cases {A, E, M} CORD Disaggregated ROADM Global R&E Deployment



CPMan Apps
Intents using Flow Objectives
P4 DEMO support
YANG tool chain



Q3/16 Hummingbird

RabbitMQ, Kafka Message YANG NBI, SBI CODECs ACTN Traffic Engineering



Q4/16 bis

BUCK Build Tool
Trellis Fabric enhancement
LISP SBI support, REST Client,
FatTree simulator



Q1/17 Junco

TL1 SBI support
Virtualization support
Regionalization support
Dynamic conf. enhancement



Q2/17 Kingfisher

YANG Tools 2.0
OpenFlow 1.4 support
Intent F/W improment
vRouter, OpenROADM support



OpenFlow 1.5 SBI gRPC NBI support P4 runtime initial support



Q4/17 Magpie

Topo2 initial support More switch driver support



Q1/18 Nightingale

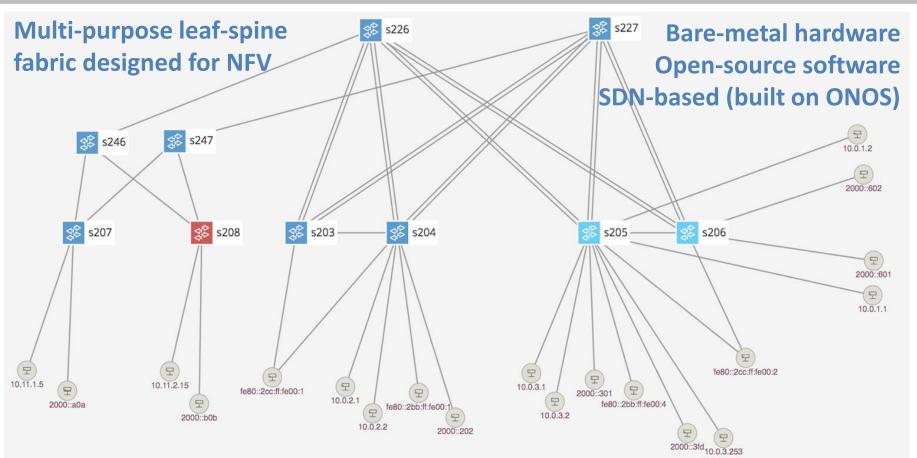
ISSU initial support Trellis enhancement (T3) P4 support enhancment



Q2/18 Owl Coming soon...

Introduction to Trellis

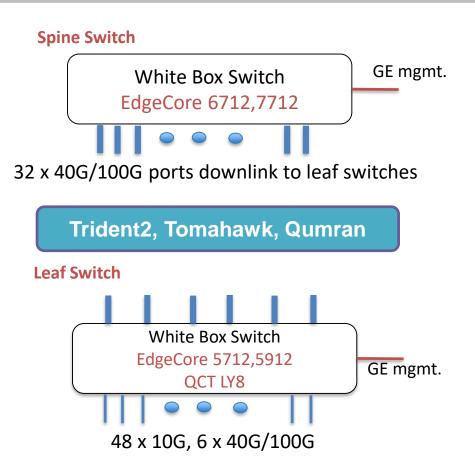
Trellis Overview



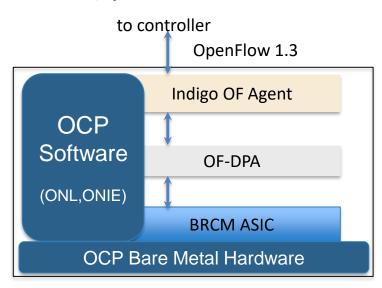
Trellis Features

- Bridging with Access & Trunk VLANs (within a rack)
- Routing (inter-rack)
 - IPv4 & IPv6 Unicast routing with MPLS Segment-Routing
 - IPv4 & IPv6 Multicast routing
- Dual-homing for compute-nodes and external routers
- Multi-stage fabrics (2 layers of spines)
- vRouter entire fabric behaves as a single router
 - BGP (v4/v6) support for external (upstream) connectivity
 - Static routes, route blackholing
 - DHCP L3 relay (IPv4/v6)
- MPLS Pseudowires
- QinQ termination
- T3 Trellis Troubleshooting Tool
- ASIC Support
 - Broadcom Qumran, Tomahawk, Trident2 switches from EdgeCore & QCT
 - Preliminary support for Cavium Xpliant switches and P4-based Tofino switches

White-Box = Bare-metal hw + Open-Source sw



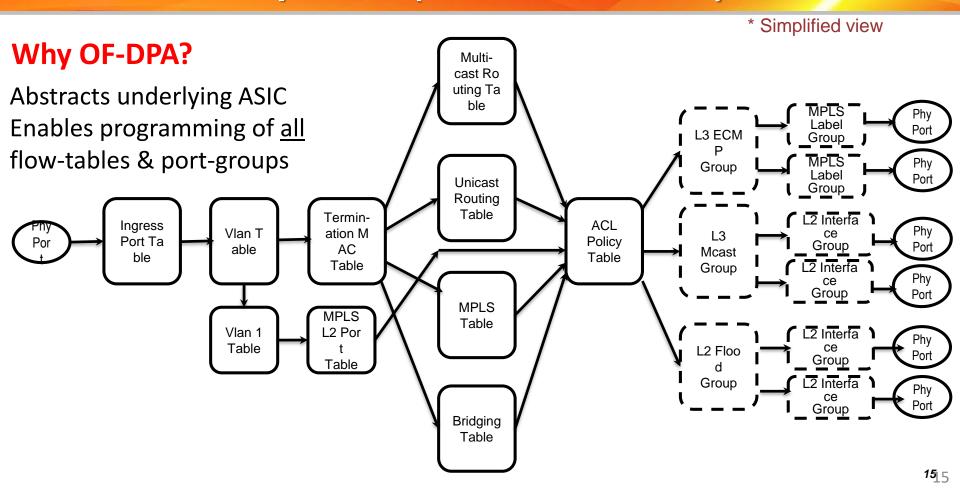
Leaf/Spine Switch Software Stack



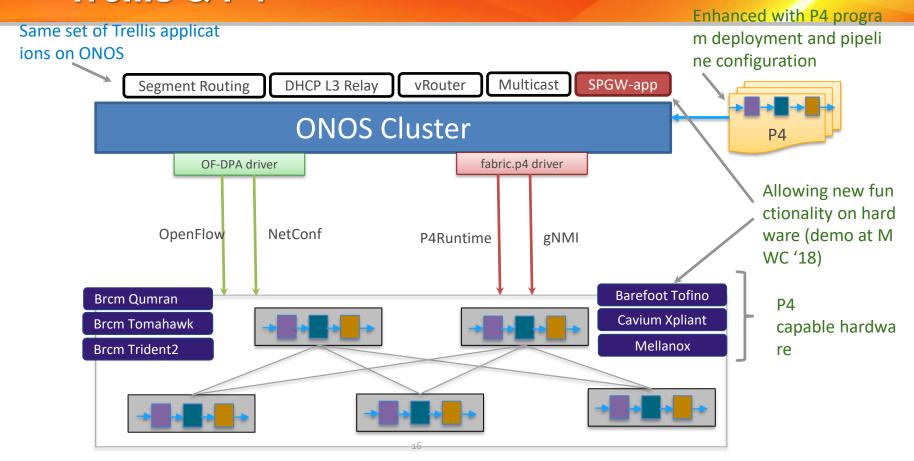
OCP: Open Compute Project ONL: Open Network Linux

ONIE: Open Network Install Environment BRCM: Broadcom Merchant Silicon ASICs OF-DPA: OpenFlow Datapath Abstraction

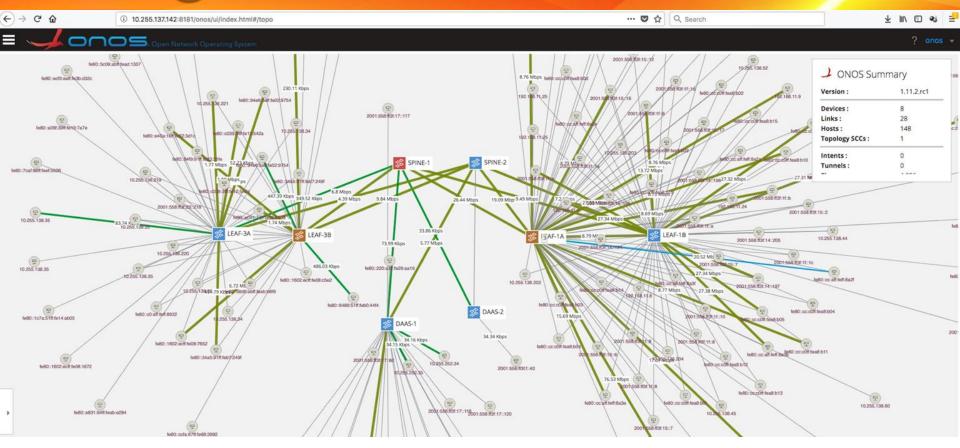
Fabric ASIC Pipeline* (BRCM's OF-DPA)



Trellis & P4



Trellis @ Comcast

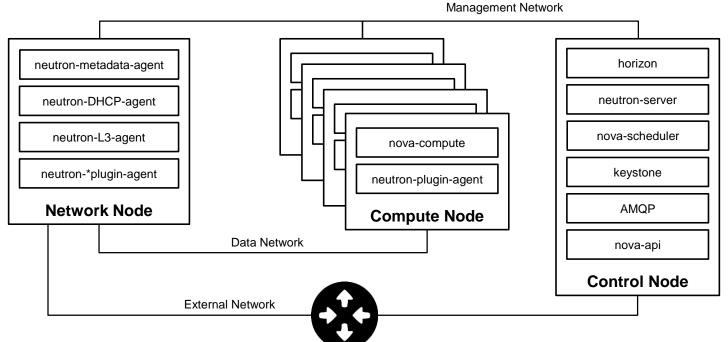


Introduction to SONA

Why SONA?

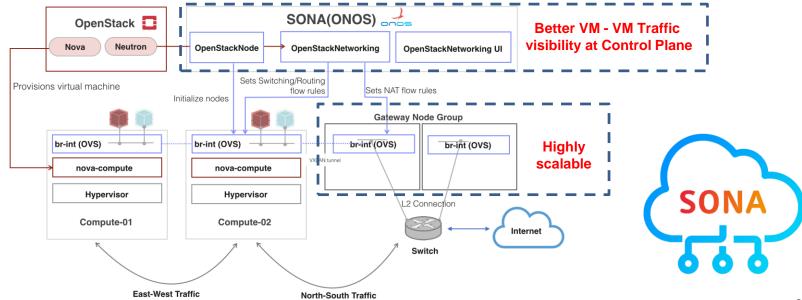
Limitation of Neutron network

- Limited visibility of VM traffic
- Limited scalability of network node



SONA (Simplified Overlay Networking Architecture)

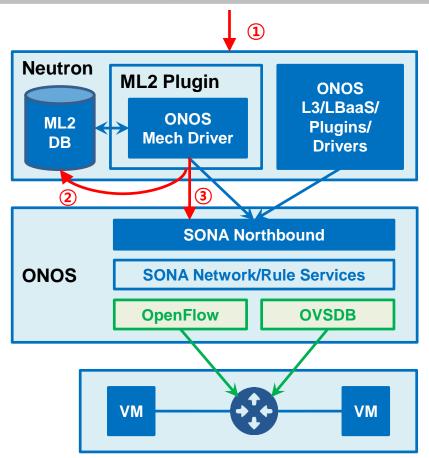
- SONA: Overlay Network Management Solution for SDDC
 - ONOS based Virtual Network Management solution (support VxLAN, VLAN, FLAT)
 - Empowered by SDN controller, a better replacement of neutron, scalable gateway
 - Fully compatible with OpenStack (mitaka, newton, ocata, pike, queens)

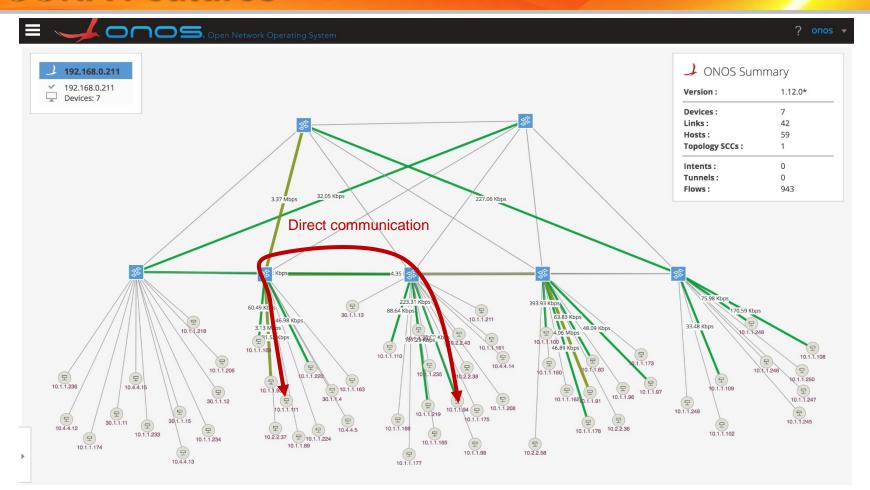


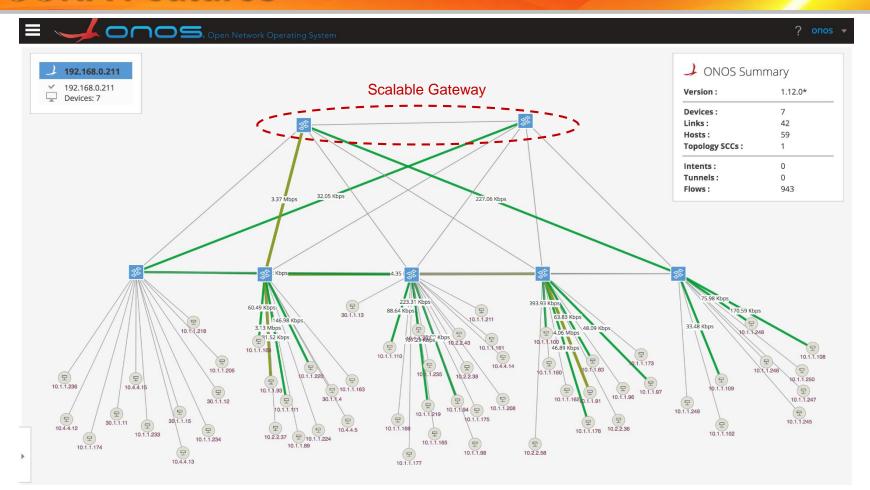
SONA (Simplified Overlay Networking Architecture)

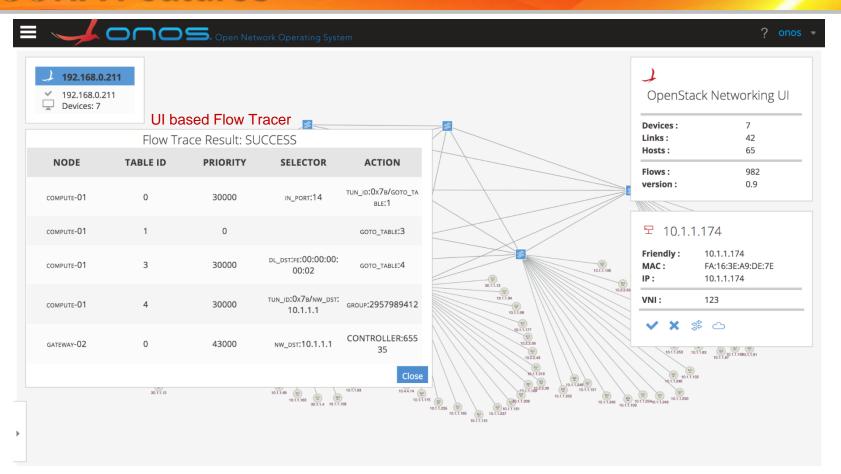
Integration with OpenStack

- OpenStack neutron
 - Plugin: modular layer 2 plugin
- networking-onos
 - ONOS L3 plugin
 - Drivers for LBaaS, FWaaS, etc.
- SONA
 - Northbound interacts with networking-onos
 - https://github.com/openstack/networking-onos
 - Southbound protocol
 - OpenFlow: install/uninstall flow rules
 - OVSDB: configure OpenvSwitch
 - » Add/delete virtual port
 - » Create/delete bridges (e.g., br-int, etc.)



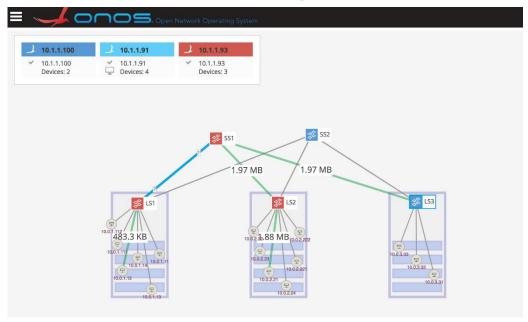






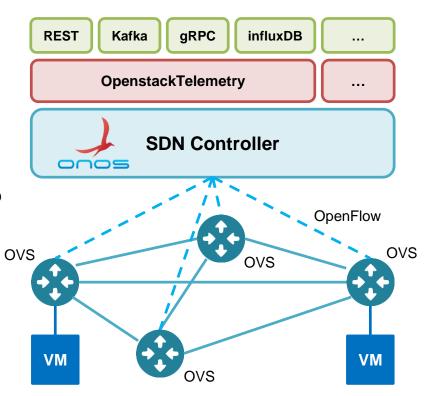
SONA Fabric

- Pure OpenFlow based Leaf-Spine Fabric Solution
- Supports ECMP, Failure detection & auto recovery
- Physical + Virtual Network Integration



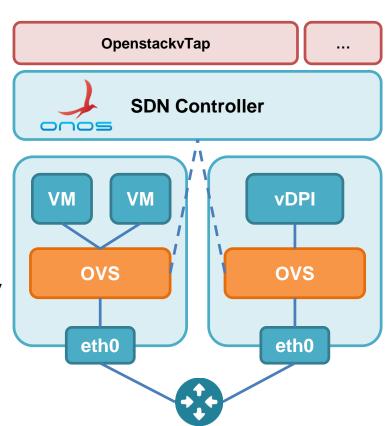
vFlow Statistics

- Collect VM to VM real-time flow statistic
- Stats collection is realized using **OpenFlow** standards protocol (no extra overhead!)
- Seamless integration with monitoring systems through various NBIs
 - REST, Kafka, gRPC, influxDB, etc.
- Realized through OpenstackTelemetry app
- No additional software installations are required at OpenStack side
- No additional hardware installations are required at compute/control node
- Open source!

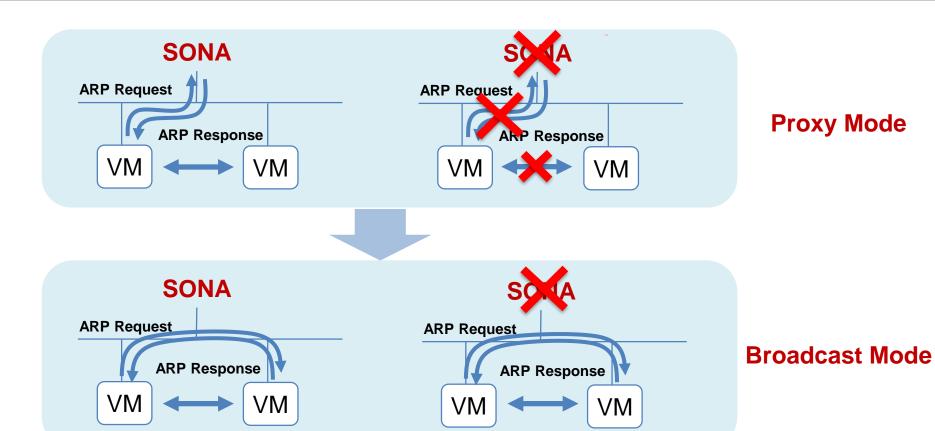


vTap

- Mirror VM to VM real-time traffic
- Leverage OVS's traffic mirroring feature
- Two traffic mirroring schemes
 - Port-based: specific to OVS
 - Flow-based: uses OpenFlow group table
- Realized through OpenstackTelemetry app
- No additional software installations are required at OpenStack side
- Further **improve** the mirroring performance by leveraging data plane acceleration technology
- Open source!

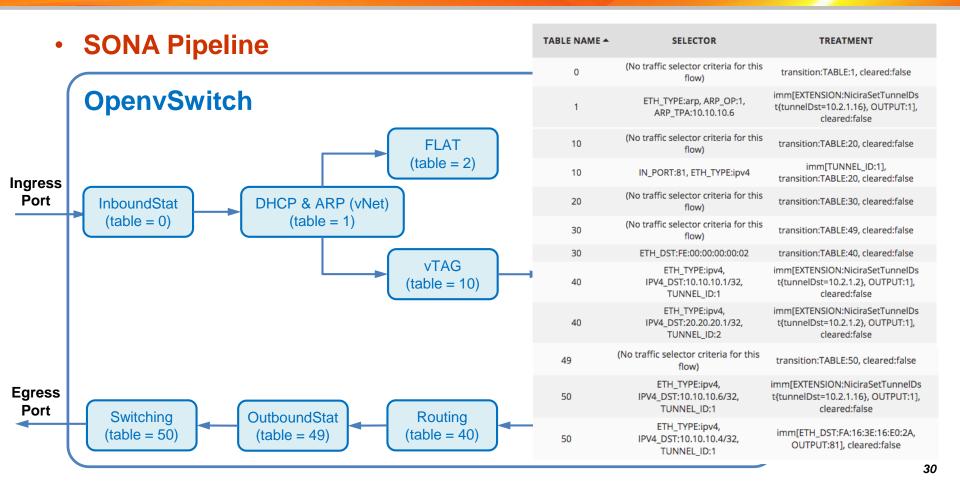


SONA Failover



Kubernetes Support Kuryr-Neutron kubernetes kubernetes openstack. Kuryr CNI **SONA** OVS Pods Container VM **SONA Fabric Spine** Leaf **Tenant A Tenant B Tenant C Tenant D** VM VM VM VM VM VM VM

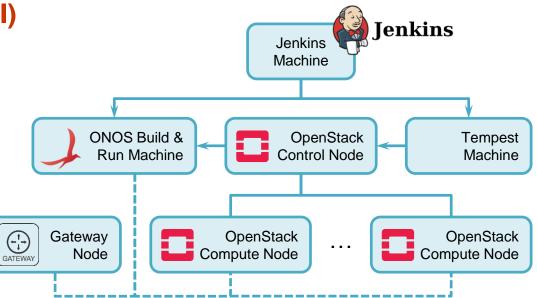
How does SONA Process Packets?



SONA CI/CD

Continuous Integration (CI)

- Fetch latest SONA source
- Build against stable ONOS
- Run unit test
- Package & deploy SONA
- Run integration test
- Notify the CI result via slack
- Deliver SONA container

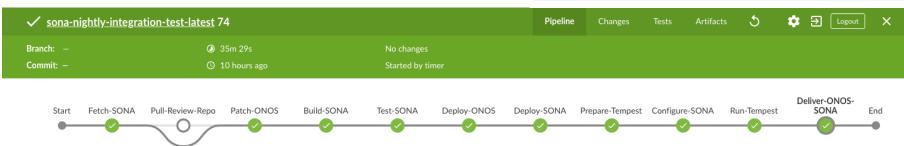


SONA CI/CD

Integration Test

- Initialize environment to spawn ONOS cluster
- Tempest basic test
 - OpenStack API test
 - OpenStack scenario test
- ONOS failure test
 - Terminate ONOS nodes, run tempest
- SONA app failure test





Open Source Contribution

Open Source Strategy

- 100% open source
- 136 commits were upstreamed in 2018 (2018.01 ~ now)
 - https://gerrit.onosproject.org/#/q/project:onos+branch:master+topic:sona

Helps from Community

- More tests and feedback from community
- Code contributions are always welcomed :)

Wiki

https://wiki.onosproject.org/display/ONOS/SONA%3A+DC+Network+Virtualization

Slack Channel

– #sonaproject @ onosproject.slack.com