OCP SUMMIT

March 20-21
2018
San Jose, CA

OPEN. FOR BUSINESS.

SONiC - Programmability, Extensibility and Beyond

David A. Maltz

Distinguished Engineer

Microsoft Azure Networking

Application & Management tools

docker
CANONICAL
Mellanox TECHNOLOGIES
Alibaba Group
apstra
ARISTA
Metaswitch Networks
Microsoft
Tencent 腾讯

SONiC [Software For Open Networking in the Cloud]

Switch

DELL
ARISTA
Mellanox TECHNOLOGIES
DELTA
Inventec
Quanta Computer
Edge-corE NETWORKS
ingrasys®
CISCO
Celestica™
ALPHA Alpha Networks Inc.

SAI [Switch Abstraction Interface]

Silicon/ASIC

BAREFOOT NETWORKS
BROADCOM
Mellanox TECHNOLOGIES
nephos
CISCO
CAVIUM
centec networks
MARVELL®
Innovium™

SONiC Recap – Containerized Open Source NOS

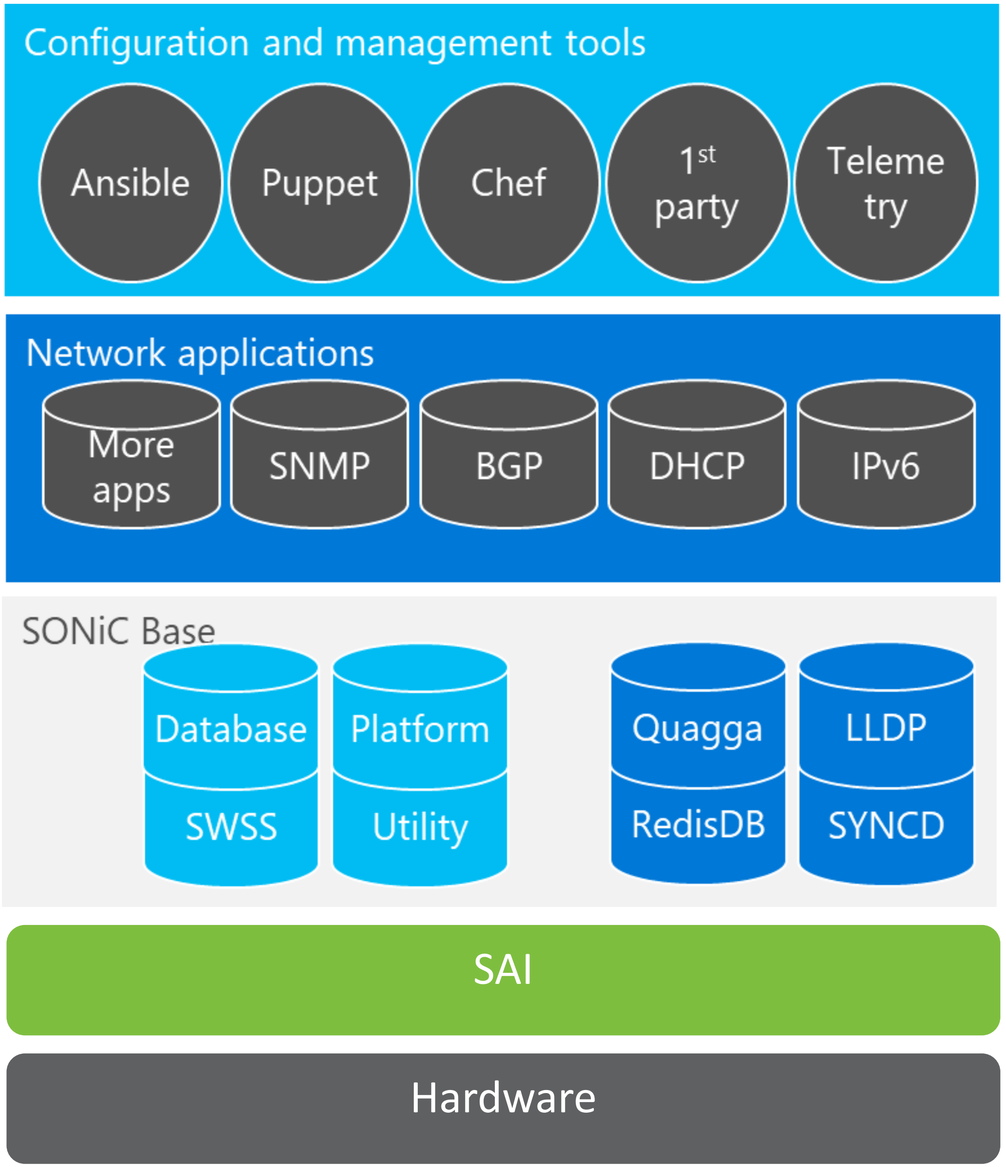

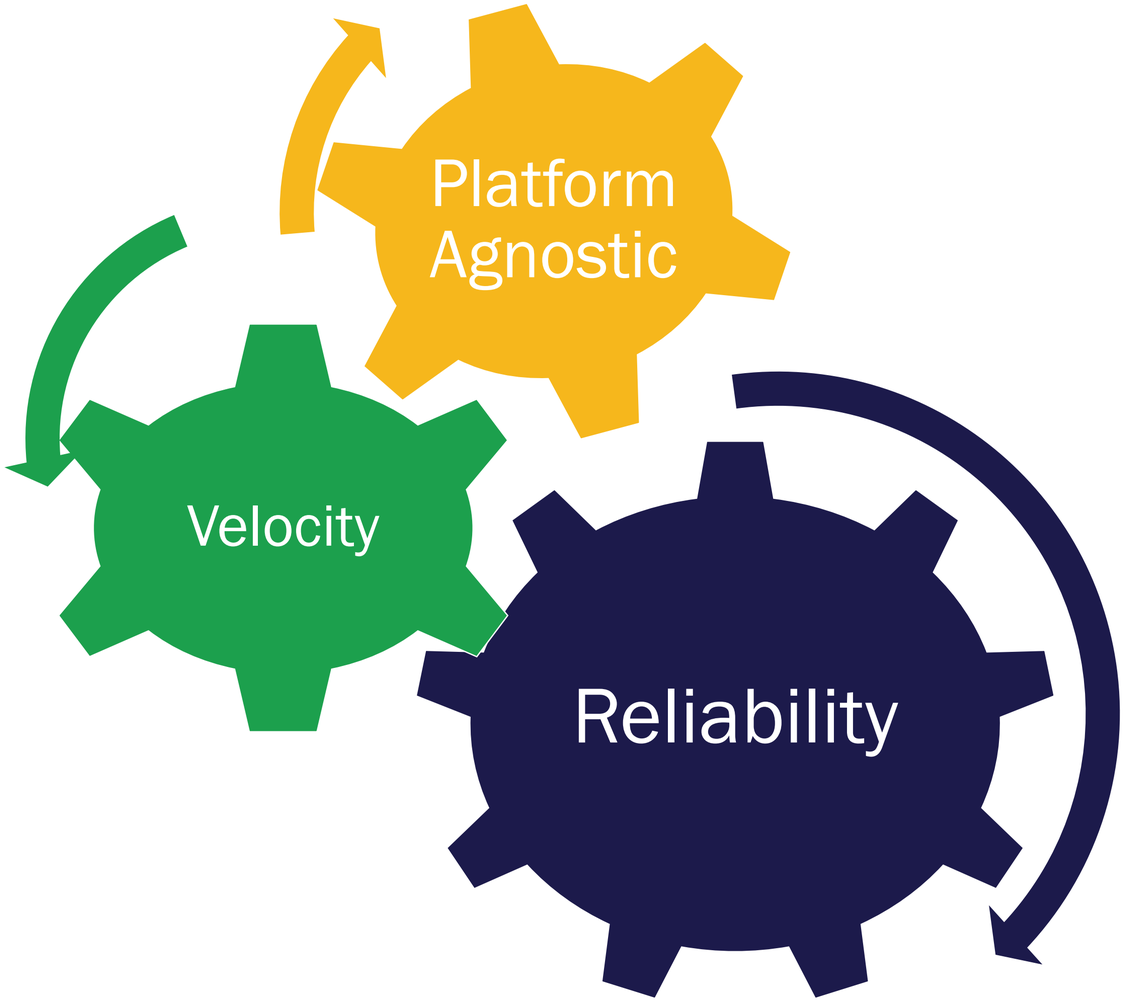

What Is New -

- SONiC supports Open Optical Monitoring (OOM)
- Richer features, thanks to
  - Alibaba: Vlan Trunk, TACACS, etc. and leading streaming telemetry work
  - LinkedIn: leading FRR integration, BGP convergence and Open 19
  - Tencent: leading VRF work
  - Mellanox: leading RDMA work
- Richer classes of devices
  - Arista: modular chassis
  - Marvell: ARM-based switch
- Richer scenarios via programmability
  - SONiC Network Virtualization

New Challenges from Broad Spectrum of Workload

Microsoft and NetApp Unleash the
Power of Data Through the Industry's
Fi

Intro
It's
the

Cray is bringing its supercomputers to Microsoft Azure

Fred

Here's
bring

Get Ready for VMware Horizon Cloud on Microsoft Azure

ement

Microsoft Runs SAP HANA Enterprise Cloud on Azure

Microsoft will run SAP HANA Enterprise Cloud on Microsoft Azure. This will allow customers to run SAP S/4HANA in a secure, managed cloud.

Related Articles

Microsoft

That's a lot of eye-crossing product names, so we asked SAP for some clarification. A spokesperson explained:

Case Study - VNET Peering in Legacy Network

Traditional Implementation

- VNET represented by VRF
- VNET1 peering with VNET2 implies copy routes from VNET1 to VNET2 and vice versa
- 1K VMs and 100 VNETs will require up to 10M routes !!!

1K VMs and 100 VNETs will require up to 10M routes !!!

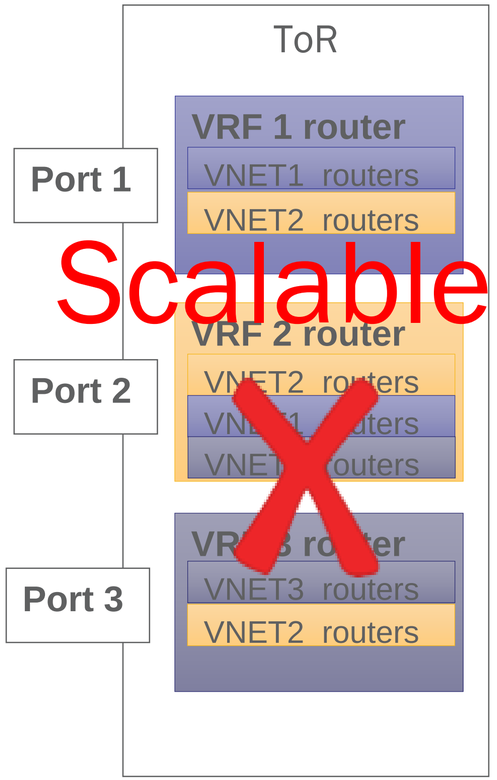

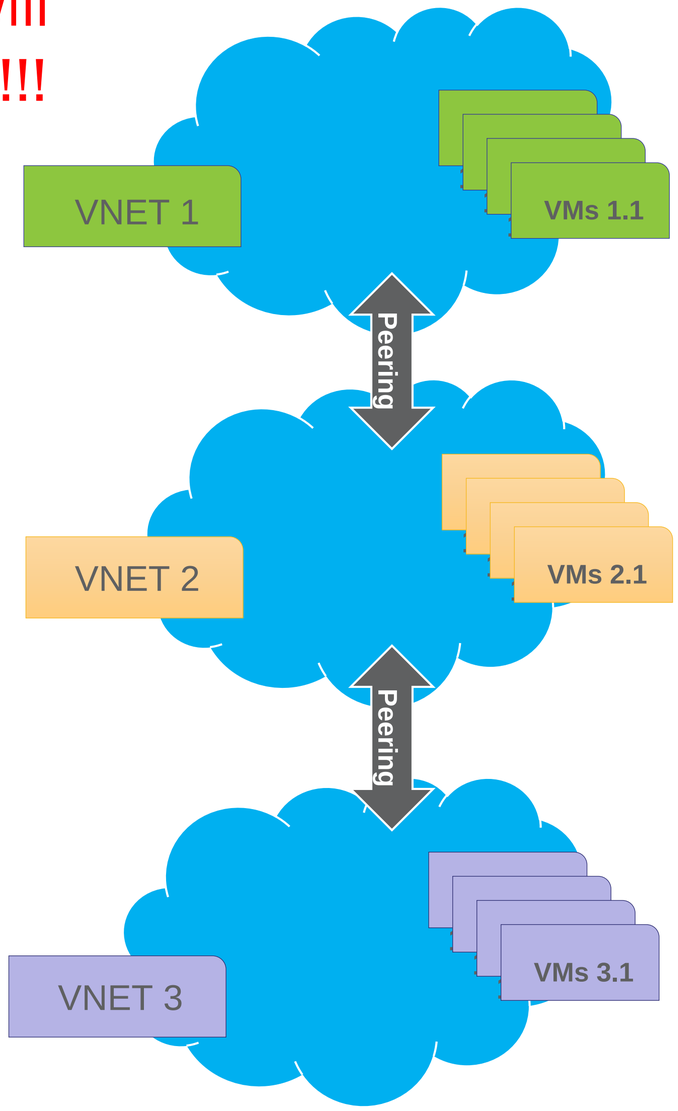

VNET Peering in Programmable Network

SONiC Implementation

- Two match action tables
- Port to VNET
  - Key: Port
  - Action Set metadata
    - Where metadata = VNET ID
- VNET routing
  - Key: metadata, prefix
    - Where metadata vector of VNET peers
  - Action: next hop
- VNET1 peering with VNET2 -> turn on VNET1 VNET ID in VNET routing metadata of all routes originated by VNET2
  - A single route per VM
  - VM update requires a single route update

1K VMs and 100 VNETs need only 100k routes

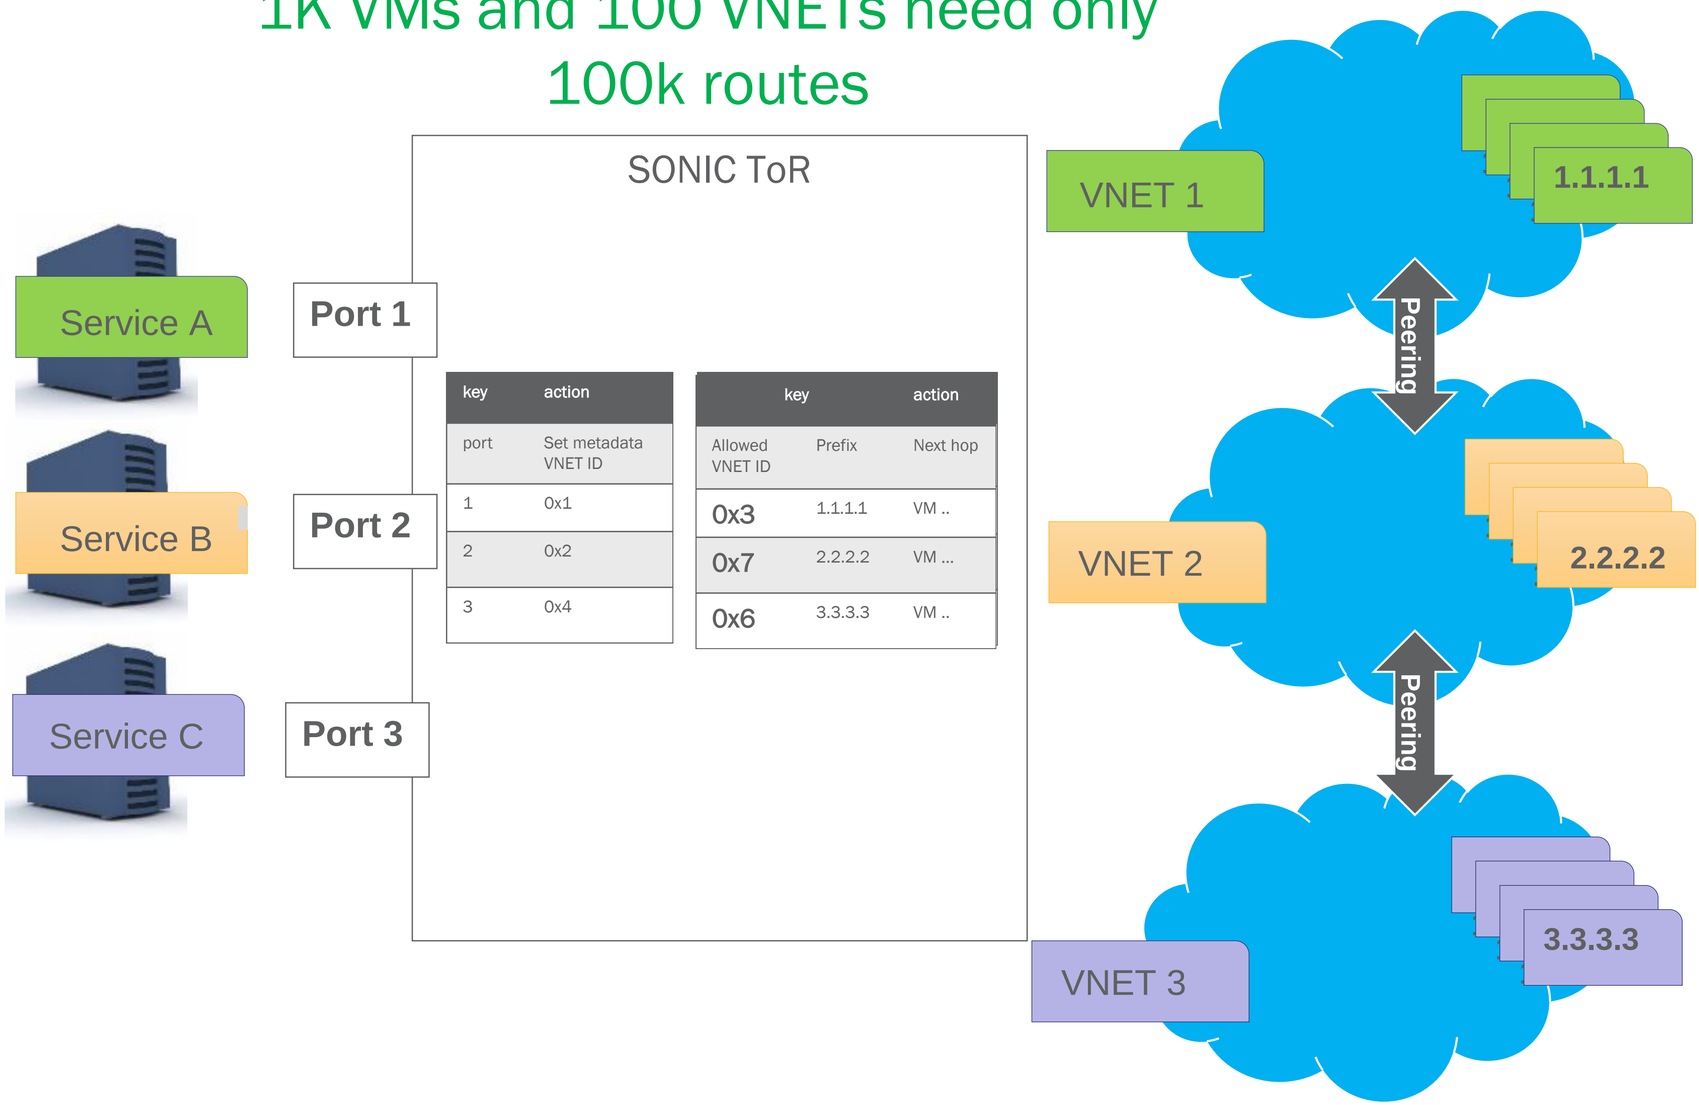

SAI Part: Programmable SAI API

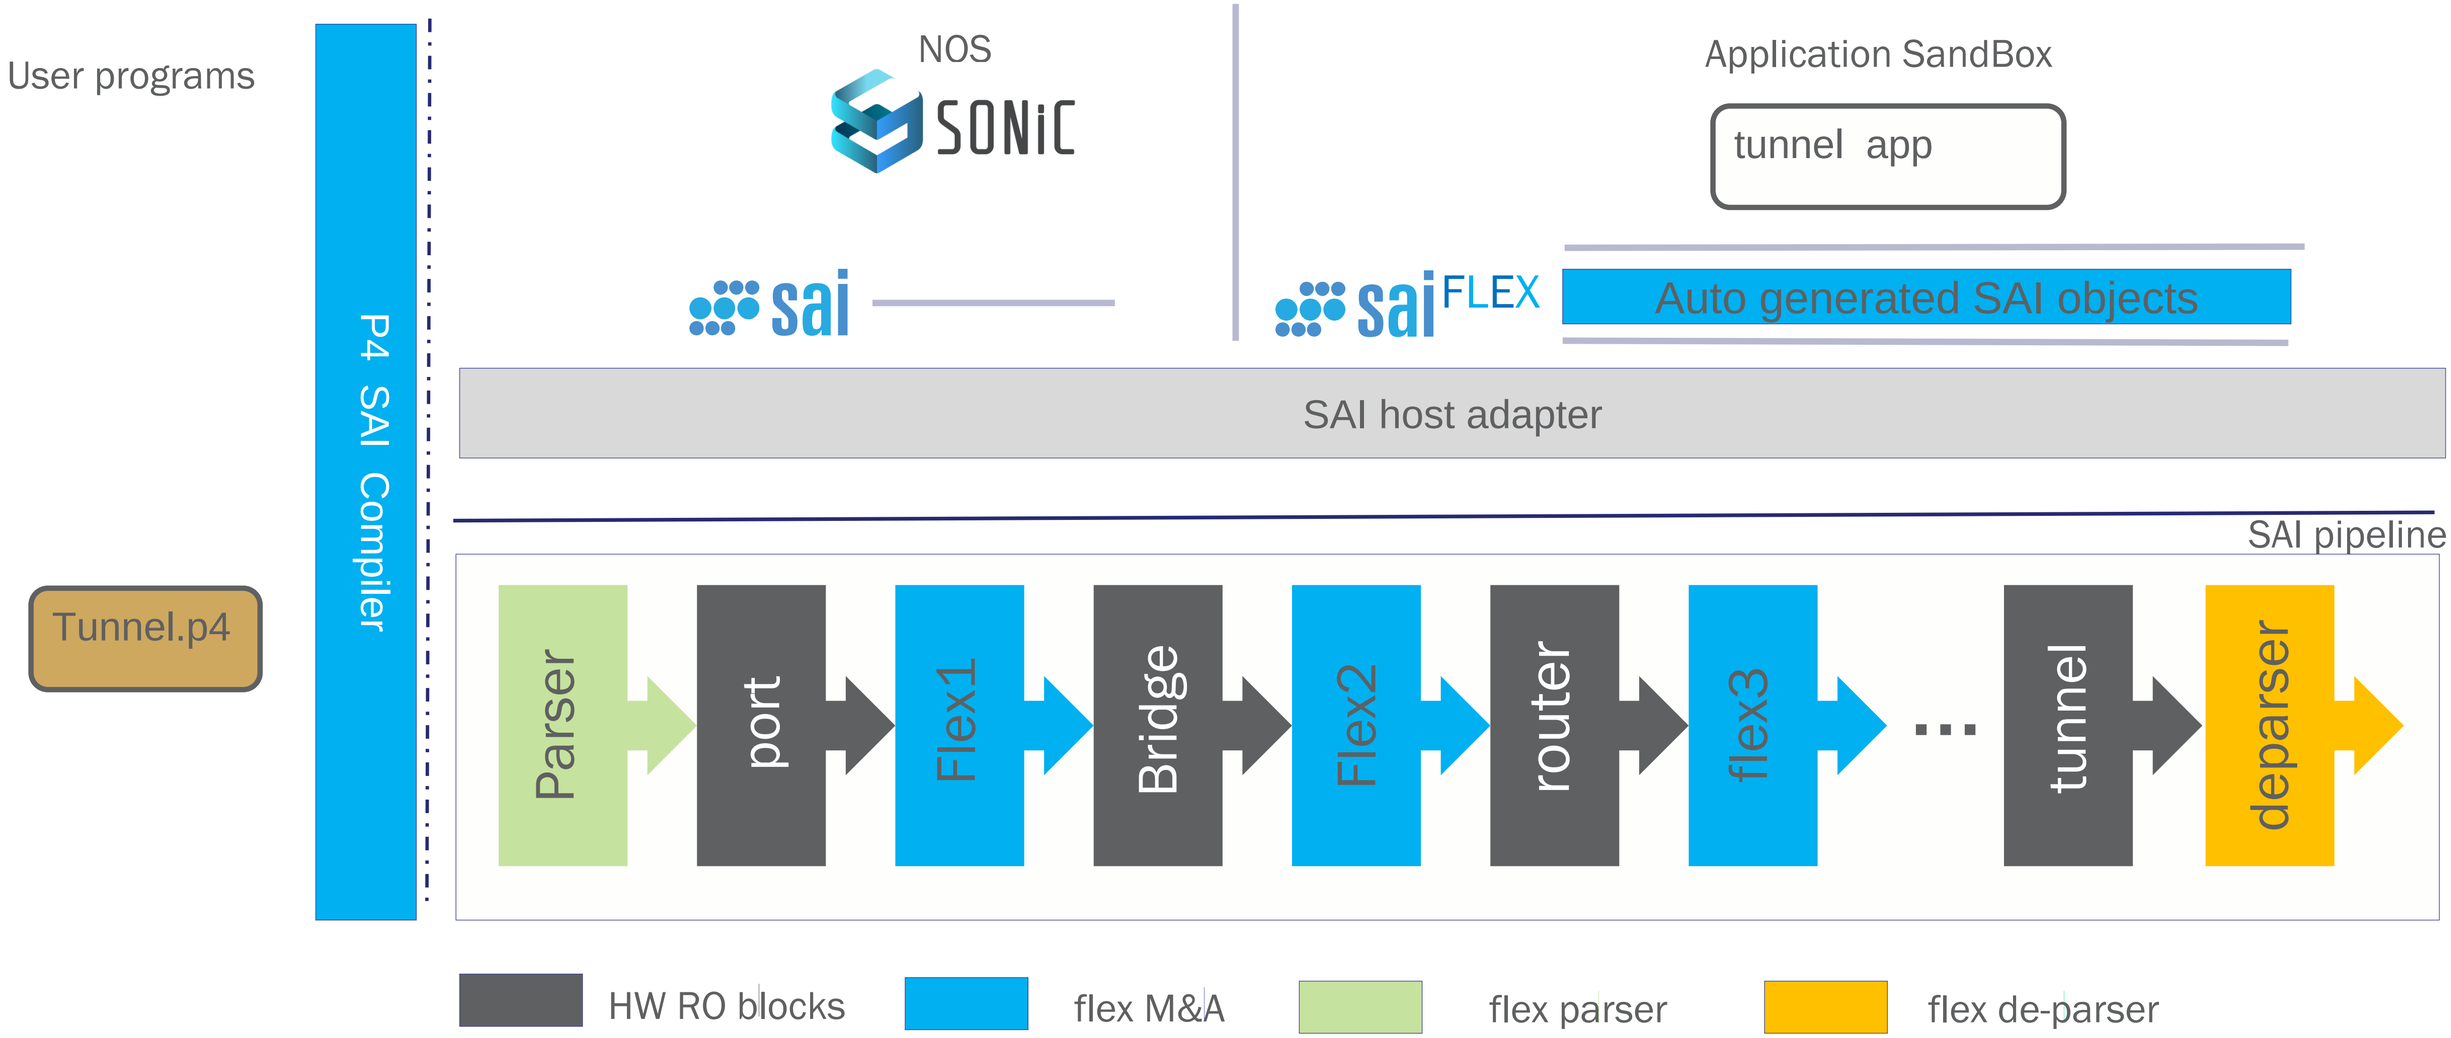

- Multiple switching SW options, develop apps
- SAIFlexAPI – uniform API for all programing language

SONIC Part: Supporting VNET

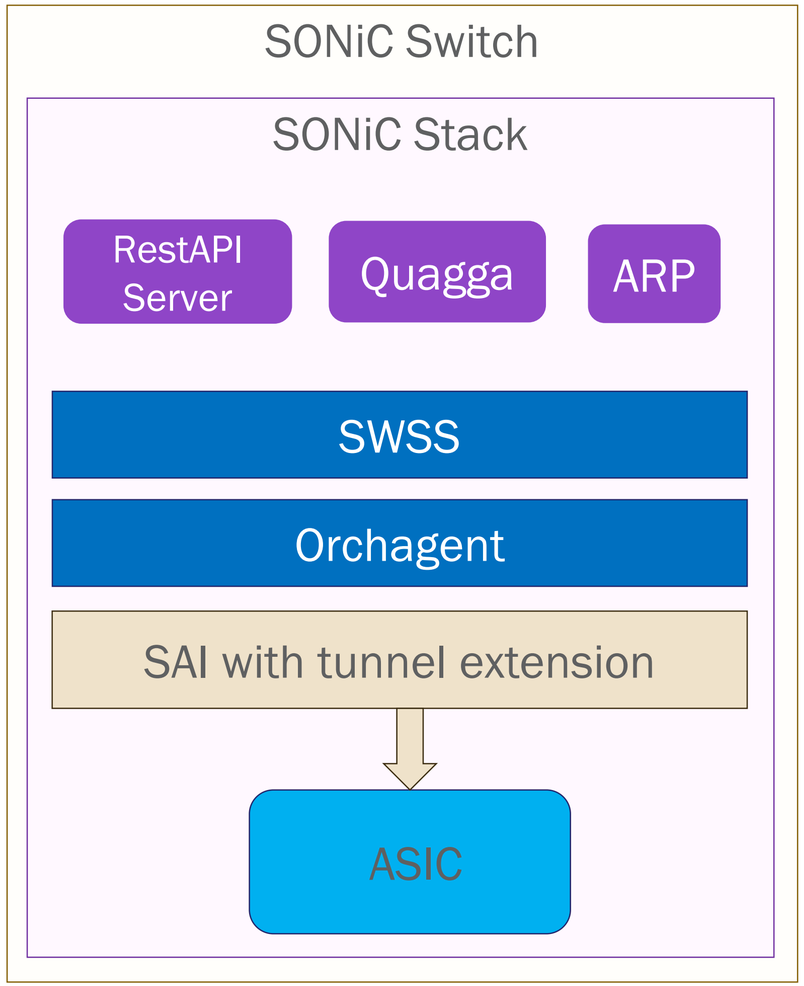

- RestAPI: Provide RestAPI for external
  - Allow external control to config the switch
  - Provides real-time data path counters and resource monitoring
  - Use OpenAPI specification (swagger)

- SWSS/Orchestration Agent
  - Use SAI with tunnel extension API
  - Provide tunnel support to upper applications

SONiC Part: Achieving Scalability via DPDK

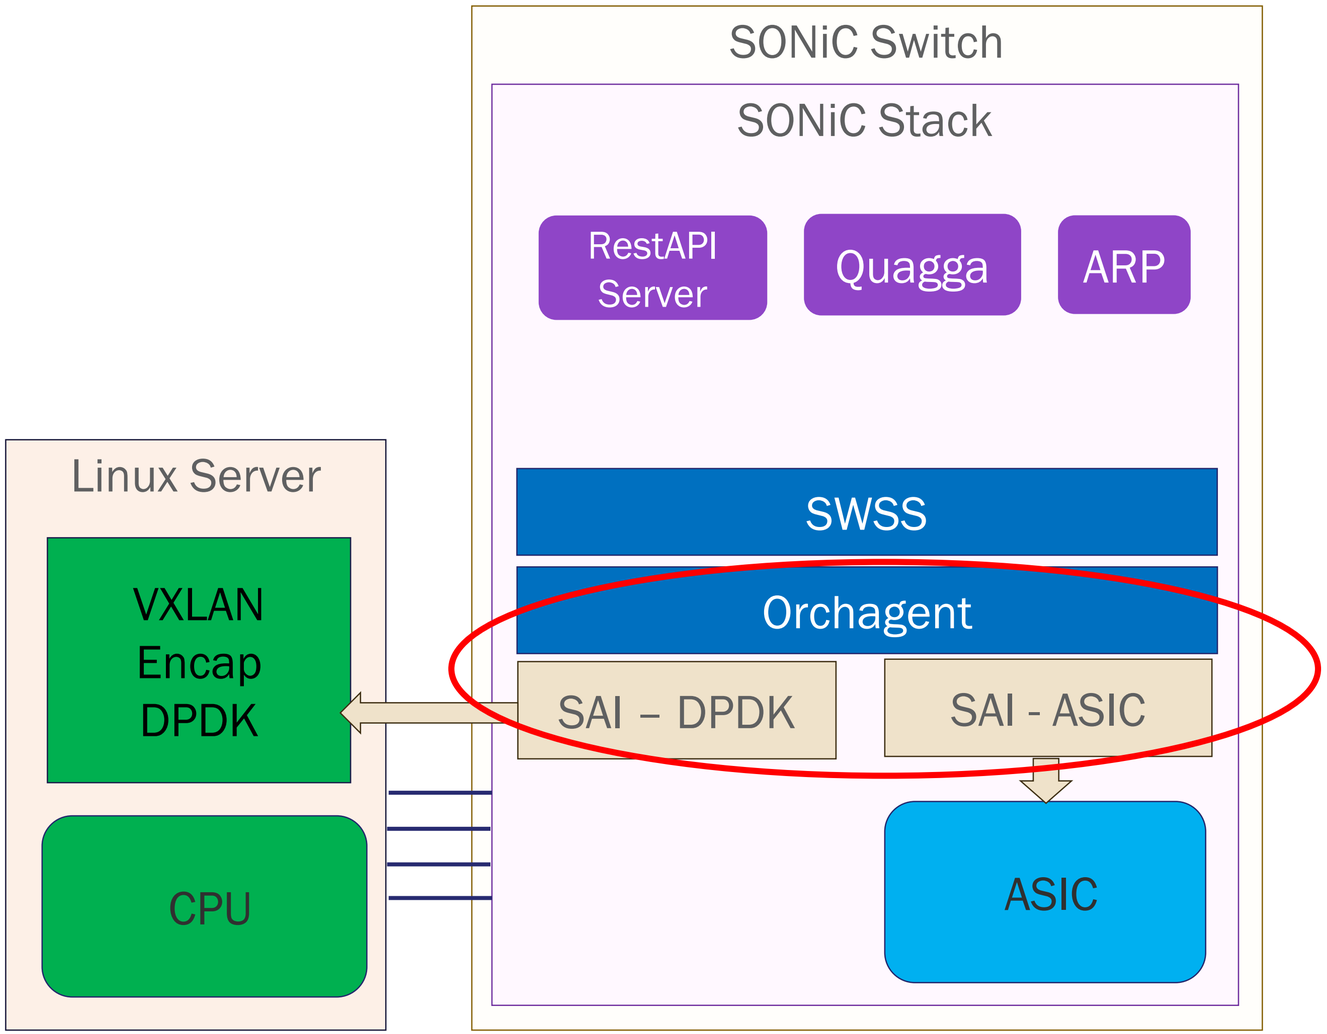

- SWSS/Orchestration agent: manage multiple SAI instances
  - Manage tunnel entry cache between DPDK and ASIC
- Server: data plane scalability and programmability
  - 16M tunnels
  - 40G/100G line rate
  - 25 ~ 30 us forwarding latency
- ASIC: High port density and rich data plane functionality
  - Tunnel entry cache
  - Underlay routing
  - Traffic policing/shaping
  - ACL
  - Mirroring

Demo: SONiC for Network Virtualization

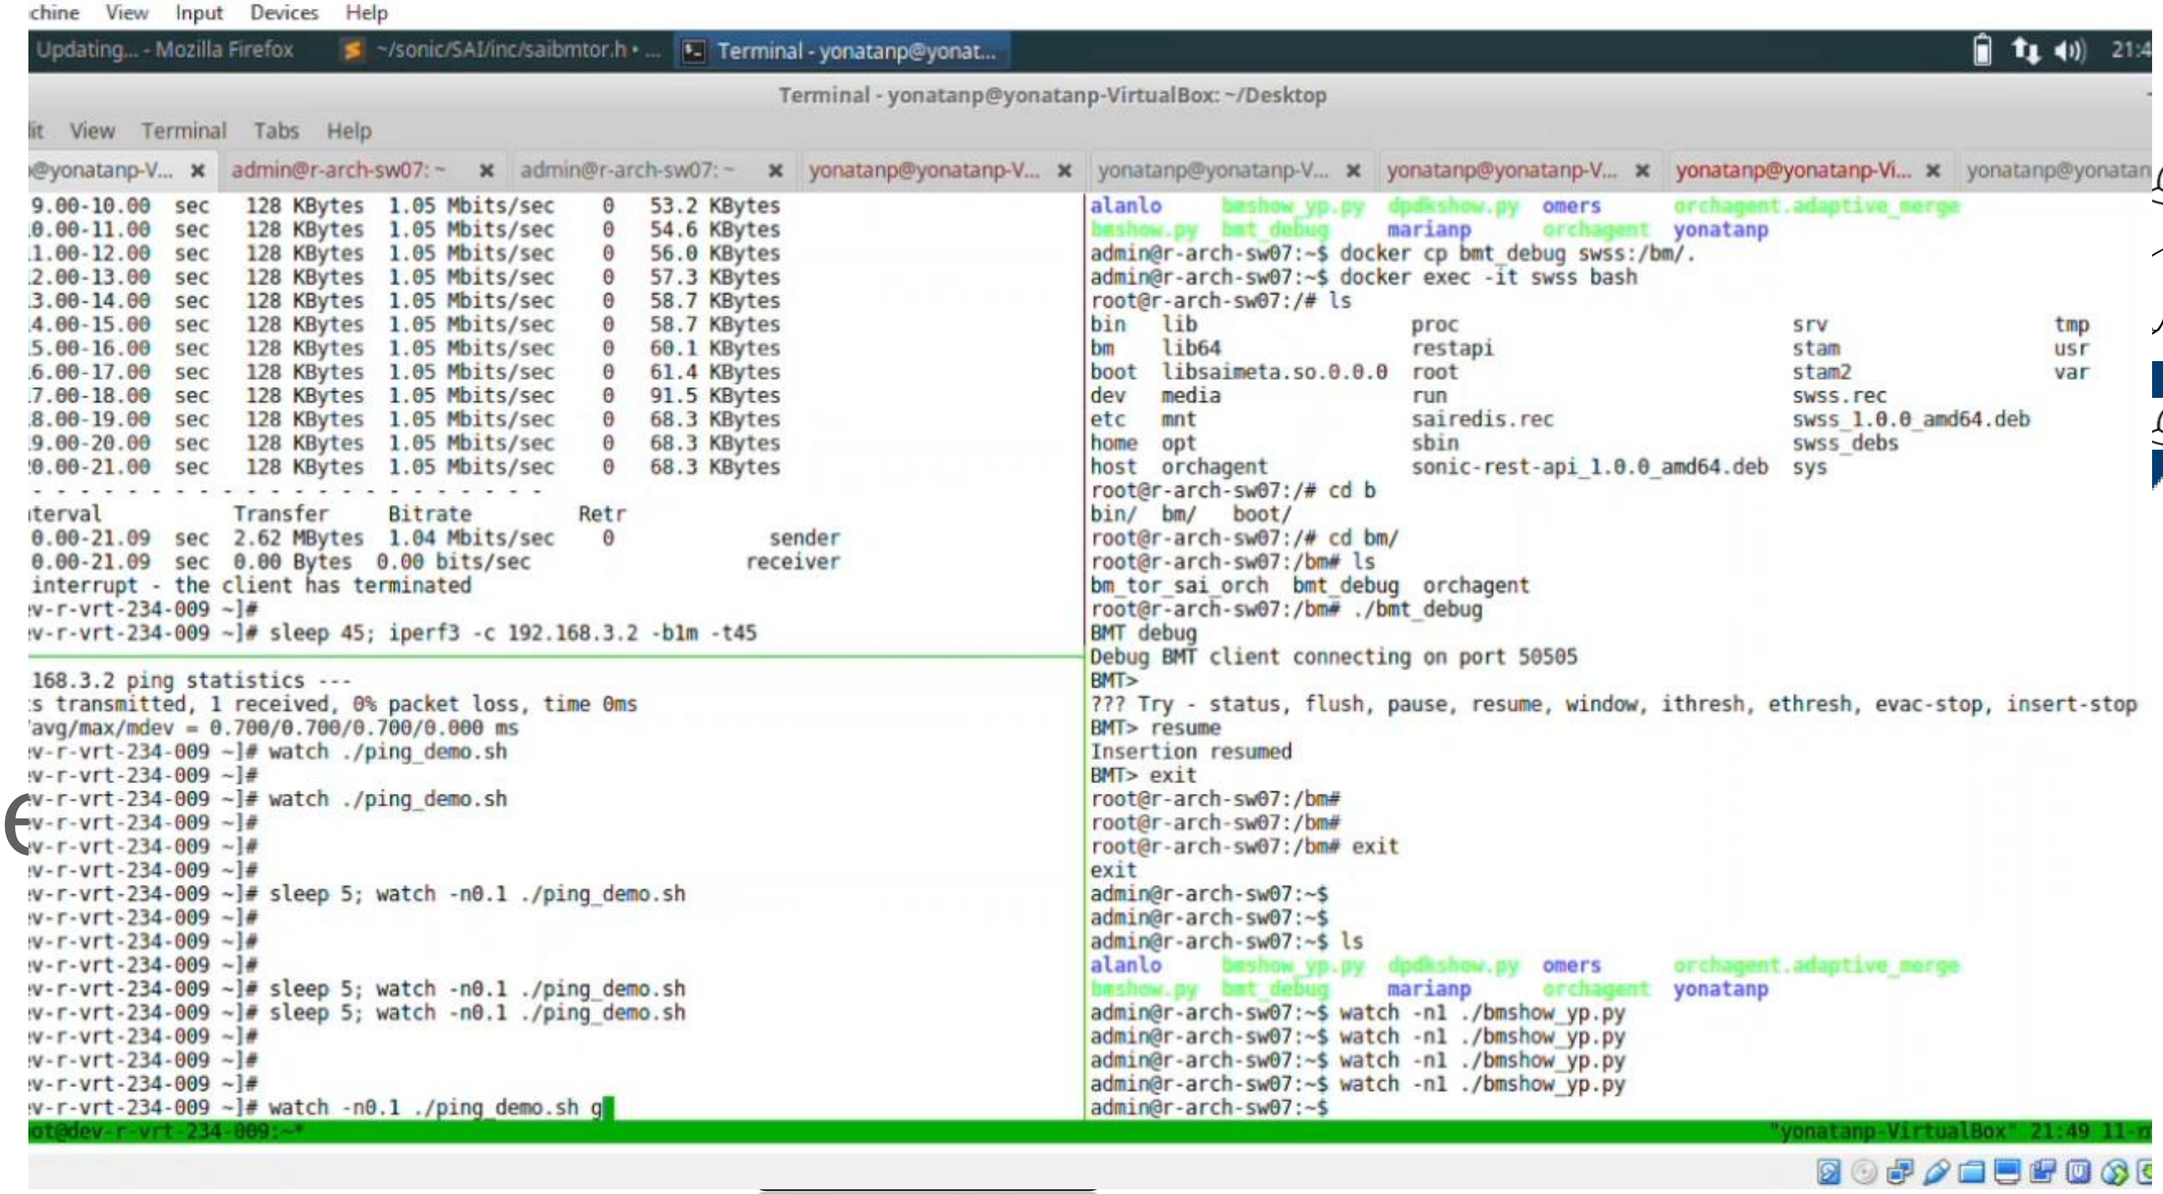

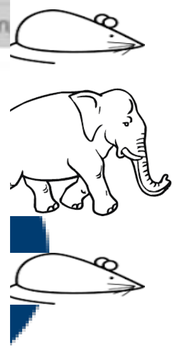

Demo: SONiC on Chassis Switch

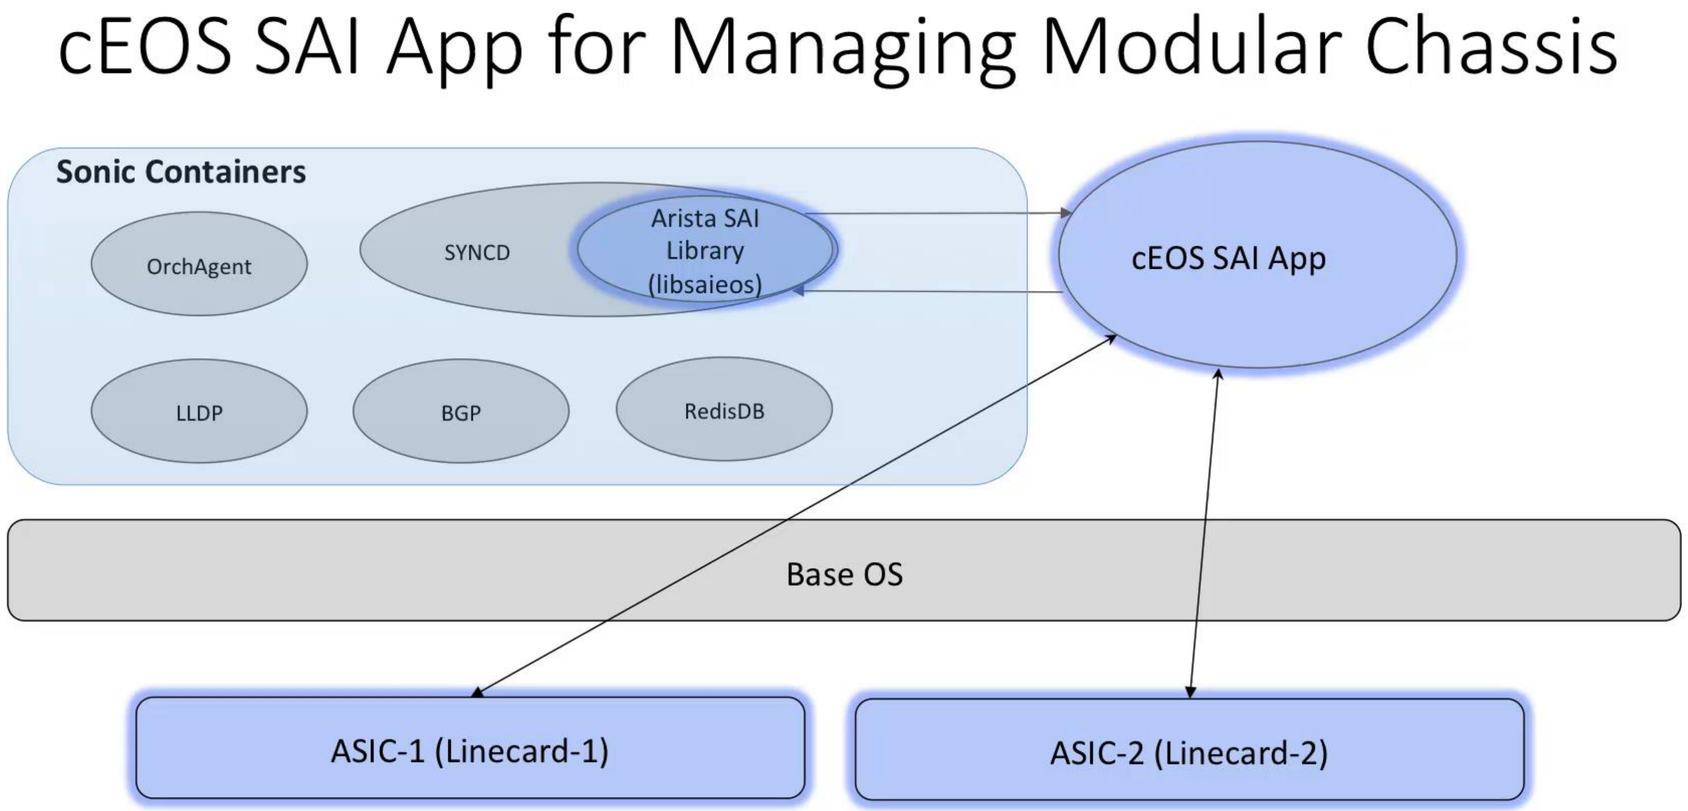

Moving Forward: Enabling WAN Scenarios

- Global network is growing exponentially, requires
  - Agility for fast Time to Market feature release and defect remediation
  - To minimize hardware dependencies
  - To scale and grow the WAN efficiently while controlling costs

- Sonic is an integral element of our cloud SDN solutions for intelligent traffic management

- Two major roles
  - Edge Peering Router
  - Backbone Router

More Demos in Microsoft (A11) and Partner Booths

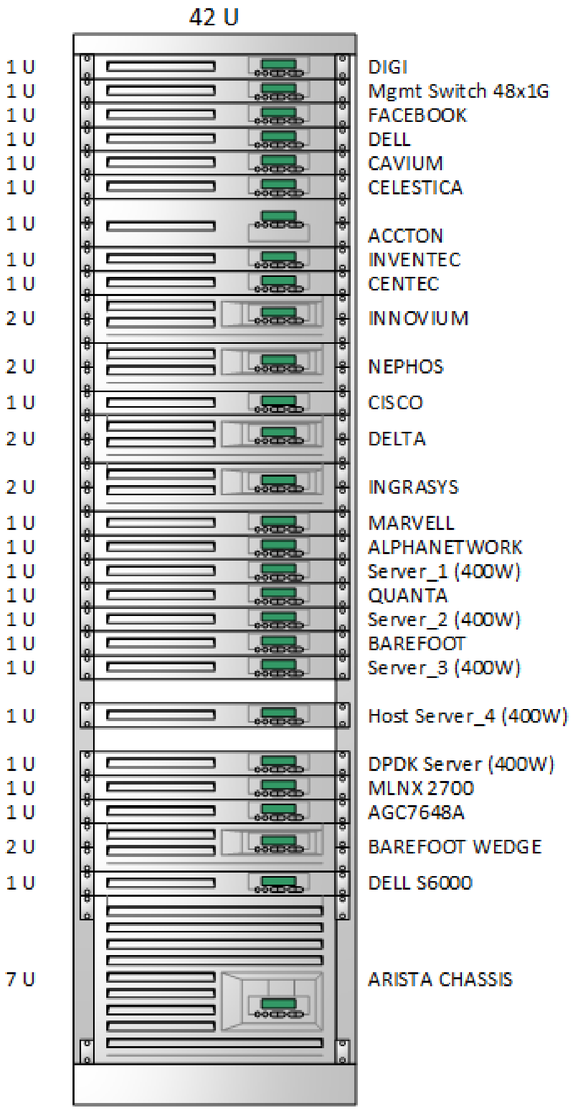

New ASIC Supported

Pizza box and Chassis

Rich Hardware Platform Supported

Virtualization

Programmability

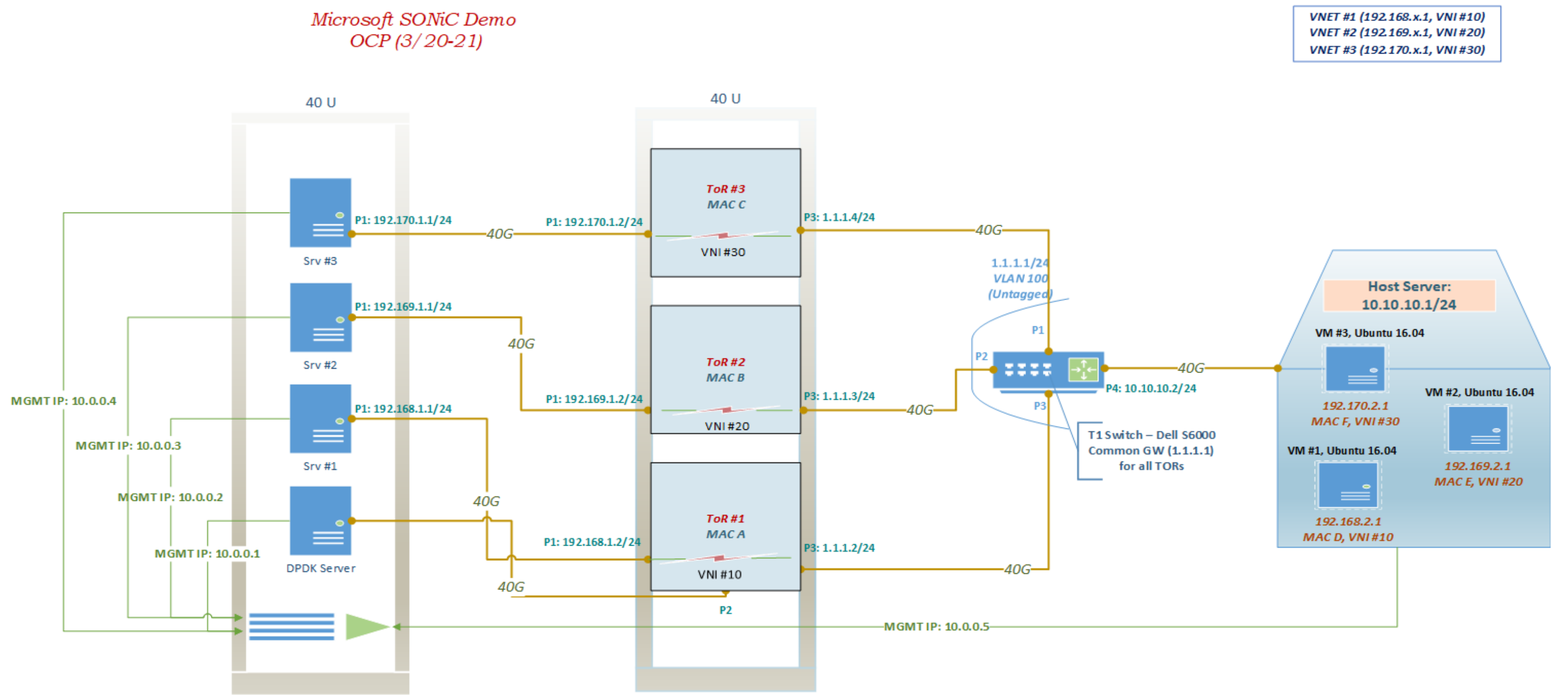

Open Invitation

- OCP SONiC/SAI workshop on 3/22

- Inviting contributions in all areas
  - SONiC/SAI
  - Hardware platform
  - New features, applications and tools
  - Download it, test it and use it!

Website:	https://azure.github.io/SONiC/
Mailing list:	sonicproject@googlegroups.com
Source code:	https://github.com/Azure/SONiC/blob/gh-pages/sourcecode.md
Wiki:	https://github.com/Azure/SONiC/wiki/

OCP SUMMIT