



# **SONIC** Innovating the Cloud Network

#### Xin Liu Principal Product Manager



#### Lihua Yuan Partner Dev Manager

Microsoft Azure Networking Team







## **SONIC** Software for Open Networking in the Cloud













centec n e t w o r k s





#### 恒为科技 EmbedWay



CISCO

碩 和





SUMMIT



















#### Welcome









## SONic Keeps Evolving

#### Telemetry

- gRPC for streaming telemetry
- Dataplane Telemetry (Dtel) extension
- Virtual Path for streaming telemetry

#### **Platform Management**

- Sensor/Transceiver monitoring
- **Dynamic Parameter Tuning**
- Platform Enhancement (PMON)

#### **RDMA**

- **PFC Watermark**
- Asymmetric PFC

#### **Data Plane**

- L3 VxLan Support
- Large Table/Deep Buffer Devices

- Juniper PTX
- Broadcom TH3, JR2 Mellanox Spectrum II Facebook Mini-pack Marvell 12.8T Falcon and ARM based switch

- Innovium Teralynx
- And more •



Reliability
-------------

- Warm Reboot
- Routing stack graceful start

#### System

- Kernel Upgrade
- Component docker upgrade
- Security patches

#### Configuration

- Incremental config
- ConfigDB

#### **Routing Stack**

- Quagga  $\rightarrow$  FRR
- cRPD from Juniper

#### **New Platforms**









#### 2019 March Workshop @Linkedin SunnyvaLe, US









## Warm Boot: A True Community Effort





Nephos

## 











имміт



#### CISCO ARISTA











#### **Data plane disruption < 30 seconds**



## Warm Boot



![](_page_9_Picture_2.jpeg)

#### **Control plane disruption < 90 seconds Data plane disruption < 1 second**

•	Warm Reboot Finishes

State Reconciliation, via SAI state-driven API

![](_page_9_Picture_7.jpeg)

![](_page_10_Figure_0.jpeg)

![](_page_10_Picture_1.jpeg)

![](_page_10_Picture_2.jpeg)

## Warm Boot Architecture

Warm boot script stores App/ASIC DB on disc Redis restores App/ASIC DB after reboot OA reads AppDB and compiles a new ASIC DB SyncD compares old/new ASIC DB, and apply diff to the ASIC Applications waking up in parallel May staged changes to App DB

OA comes in as usual, updates ASIC dB SyncD keeps syncing ASIC DB to hardware

![](_page_10_Figure_7.jpeg)

![](_page_10_Picture_8.jpeg)

![](_page_10_Picture_9.jpeg)

### warm reboot demo 0312

2019-03-13 17:11 UTC

Renorded.to Maggie Sun

Signified # Ying Xie

![](_page_11_Picture_4.jpeg)

![](_page_11_Picture_5.jpeg)

Microsoft Teams

![](_page_11_Picture_8.jpeg)

## We are not done yet – Control Plane?

![](_page_12_Figure_1.jpeg)

![](_page_12_Picture_2.jpeg)

![](_page_12_Picture_5.jpeg)

## Control Plane Assistant (Upcoming)

• ASIC  $\rightarrow$  Assistant: • ERSPAN mirror

### • Assistant $\rightarrow$ ASIC:

- Assistant encap the payload meant for neighbors
- ASIC decap and forward

![](_page_13_Picture_5.jpeg)

![](_page_13_Picture_6.jpeg)

![](_page_13_Figure_7.jpeg)

![](_page_13_Picture_9.jpeg)

## SONIC Support for Disaggregated Chassis

![](_page_14_Picture_1.jpeg)

![](_page_14_Picture_2.jpeg)

![](_page_14_Picture_4.jpeg)

## SONIC Is Powering Microsoft At Cloud Scale

![](_page_15_Figure_2.jpeg)

![](_page_15_Picture_3.jpeg)

![](_page_15_Picture_5.jpeg)

![](_page_15_Picture_6.jpeg)

## Enabling SONiC Beyond Tier 1?

![](_page_16_Figure_1.jpeg)

![](_page_16_Picture_2.jpeg)

![](_page_16_Picture_4.jpeg)

![](_page_16_Picture_5.jpeg)

## Chassis – the challenges

![](_page_17_Figure_1.jpeg)

Ethernet ports

![](_page_17_Picture_3.jpeg)

![](_page_17_Picture_4.jpeg)

#### **Power efficiency Port density** Low table scale on backend ASICs

#### ? No standard topology/connectivity **Proprietary ports/packet format Proprietary switching/load balancing**

![](_page_17_Picture_8.jpeg)

## **SONIC Support for Disaggregated Chassis**

![](_page_18_Figure_1.jpeg)

**CLOS Topology with Ethernet ports** 

#### **VXLAN-based switching**

Each front end chip is a VXLAN Tunnel End Point (VTEP)

Packets inside the chassis are encapsulated with VXLAN headers

#### **BGP-EVPN** as the internal control plane

One SONiC/BGP instance per ASIC Frontend SONiC *directly* redistribute routes using EVPN

![](_page_18_Picture_9.jpeg)

![](_page_18_Figure_10.jpeg)

![](_page_18_Picture_11.jpeg)

## SONIC Disaggregated Chassis Demo at Booth

![](_page_19_Figure_1.jpeg)

![](_page_19_Picture_2.jpeg)

![](_page_19_Picture_4.jpeg)

![](_page_20_Figure_1.jpeg)

![](_page_20_Picture_3.jpeg)

![](_page_20_Picture_5.jpeg)

# **Open Invitation**

Inviting contributions in all areas

- SONIC/SAL
- Hardware platform
- New features, applications and tools
- Download, test, deploy!

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

![](_page_21_Picture_7.jpeg)

![](_page_21_Picture_8.jpeg)

![](_page_21_Picture_10.jpeg)

# Open. Together.

![](_page_22_Picture_2.jpeg)

#### OCP Global Summit | March 14–15, 2019

![](_page_22_Picture_4.jpeg)

![](_page_22_Picture_5.jpeg)