OPEN POSSIBILITIES.

PINS – P4 Integrated Network Stack

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

PINS – P4 Integrated Network Stack

Bhagat Janarthanan, Google
Brian O'Connor, ONF
Reshma Sudarshan, Intel

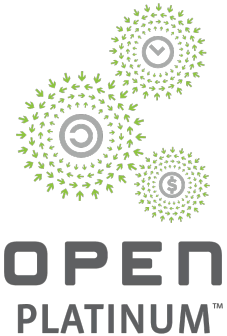

OPEN POSSIBILITIES.

SDN & SONiC

SDN

NETWORKING

- Google - a major user of SDN in Data Center [1]
  - Simpler Traffic Engineering [2]
  - Easier Debugging - Network State visible to controller [3]
  - Control Plane runs on Dedicated Fast Servers
- SDN & Remote Controller Adoption growing
- Focus: Help SDN go mainstream

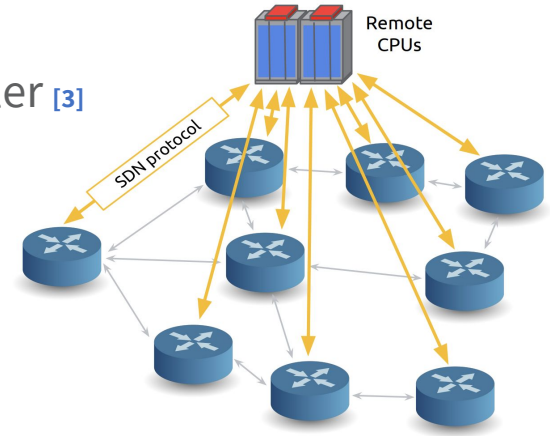

[1]: Jupiter Rising: A Decade of Clos Topologies and Centralized Control in Google's Datacenter Network
[2]: B4: Experience with a Globally-Deployed Software Defined WAN
[3]: Orion: Google's Software-Defined Networking Control Plane

SONiC

- Vendor Agnostic (thanks to SAI and common platform API)
- Decoupled Software from Hardware
- Enabling rapid innovation
- Vibrant Ecosystem
- Open Source

Can we extend SONiC to provide an incremental, opt-in path to SDN ?

- Focus on customer problems & business opportunities
- Build confidence in new infrastructure as we go

P4 Programs

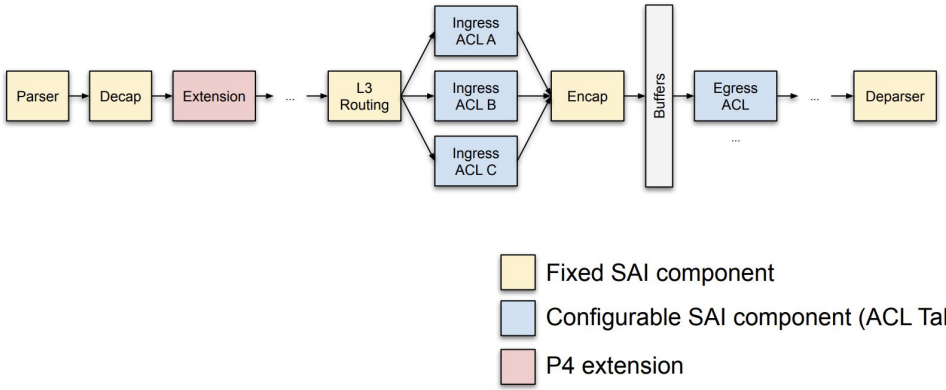

sai/routing.p4

```
...

@p4runtime_role(P4RUNTIME_ROLE_ROUTING)
 @id(ROUTING_IPV4_TABLE_ID)
 table ipv4_table {
   key = {
     // Sets vrf_id in sai_route_entry_t.
     local_metadata.vrf_id : exact @id(1) @name("vrf_id")
         @refers_to(vrf_table, vrf_id);
     // Sets destination in sai_route_entry_t to an IPv4 prefix.
     headers.ipv4.dst_addr : lpm @format(IPV4_ADDRESS) @id(2)
                                 @name("ipv4_dst");
   }
   actions = {
     @proto_id(1) drop;
     @proto_id(2) set_nexthop_id;
     @proto_id(3) set_wcmp_group_id;
   }
   const default_action = drop;
   size = ROUTING_IPV4_TABLE_MINIMUM_GUARANTEED_SIZE;
 }

...
```

P4 compiler

NETWORKING

OPEN POSSIBILITIES.

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

Solution: PINS

- P4 used to model the SAI pipeline
- SDN protocol: P4Runtime
  - Standard, open, silicon-independent
  - Enables runtime-control of data plane objects
- Management protocols: OpenConfig
  - Standard, open, widely used
  - *Already used in SONiC*

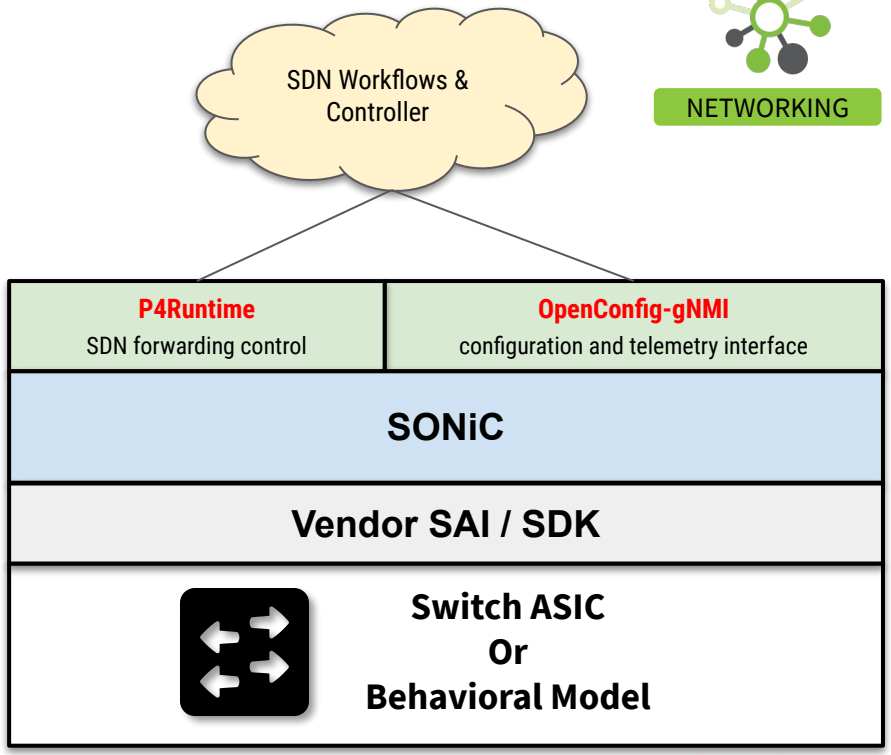

Architecture

Integration - ASIC Programming

NETWORKING

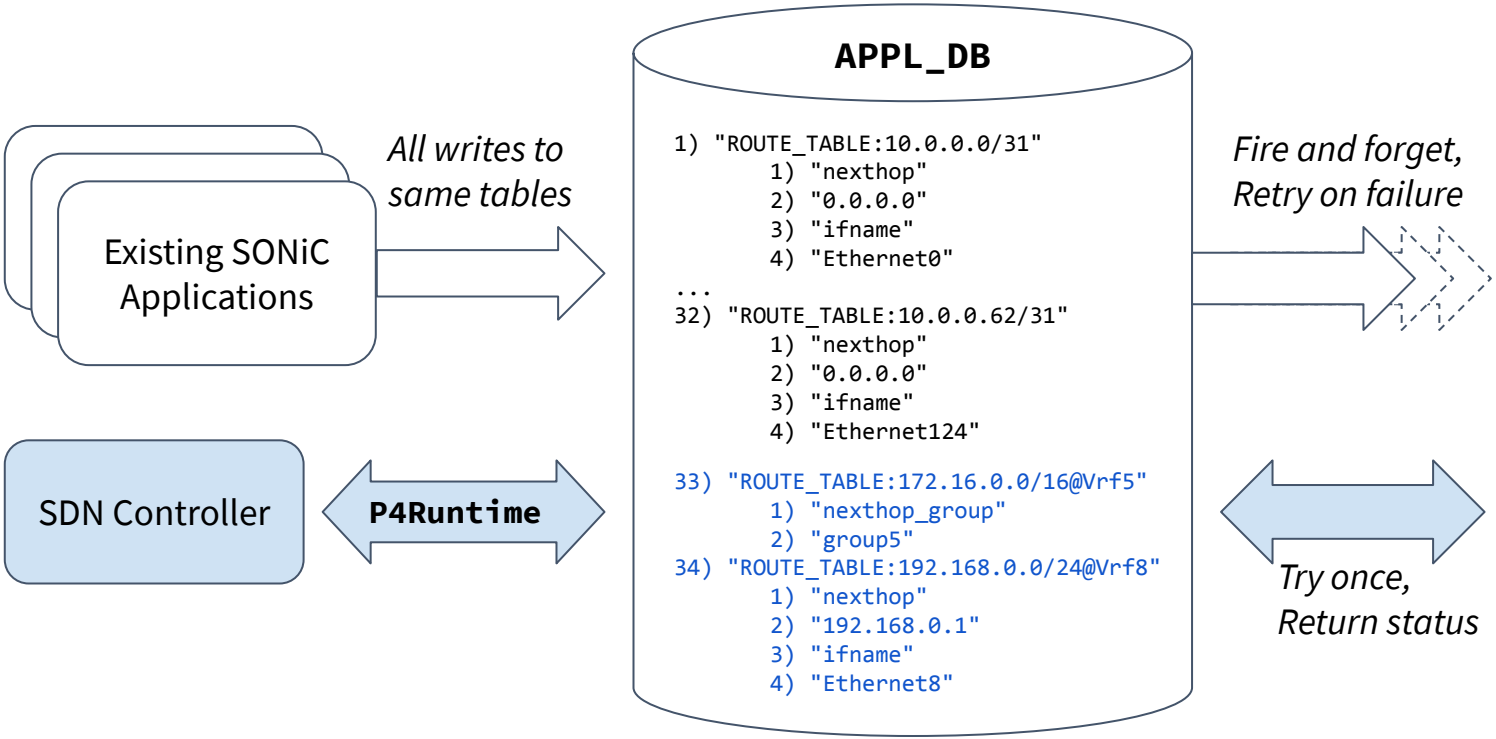

OPEN POSSIBILITIES.

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

PINS Architecture

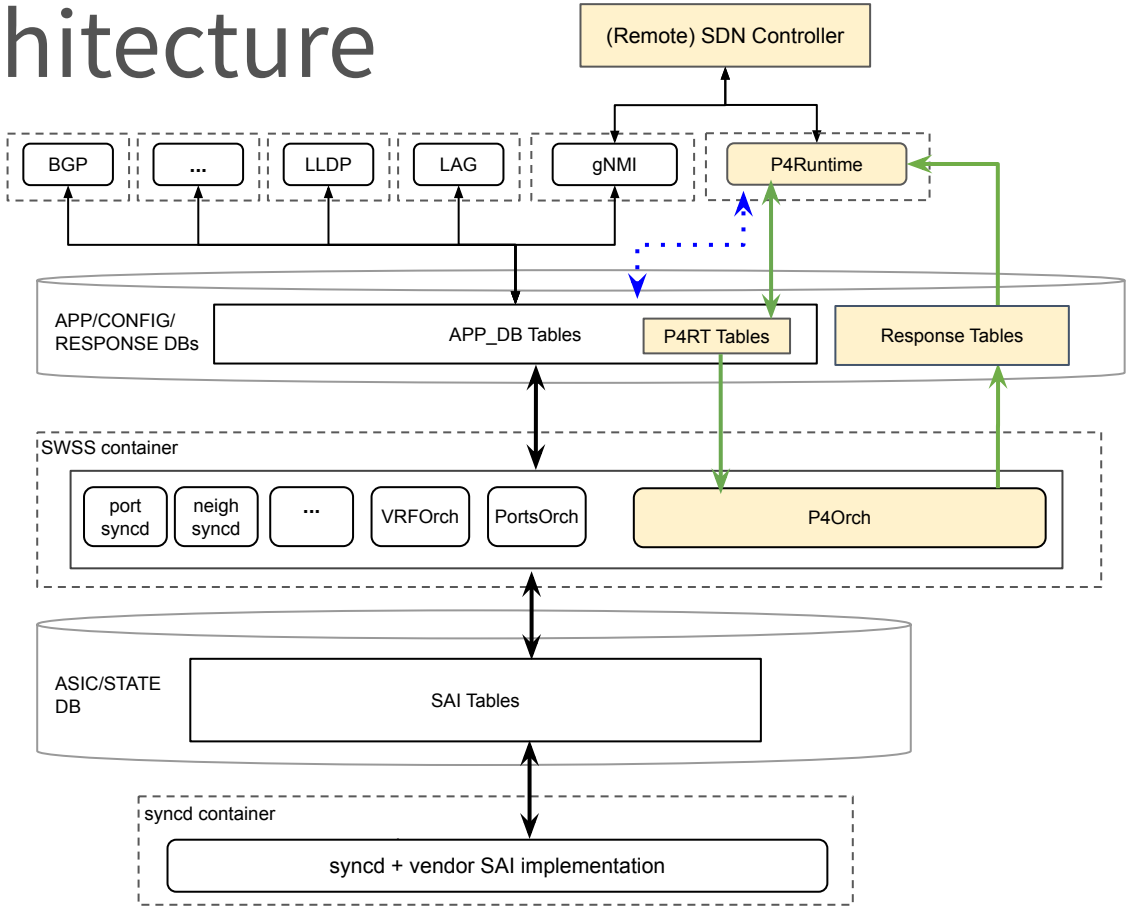

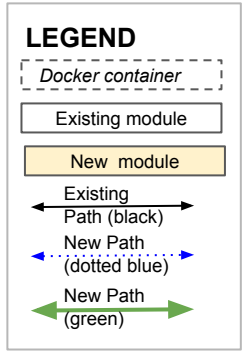

NETWORKING

OPEN POSSIBILITIES.

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

P4Runtime App

- New Docker container
- Writes controller intent to APPL_DB
- Features
  - Arbitration
  - Readback support
  - P4Info Upgrades

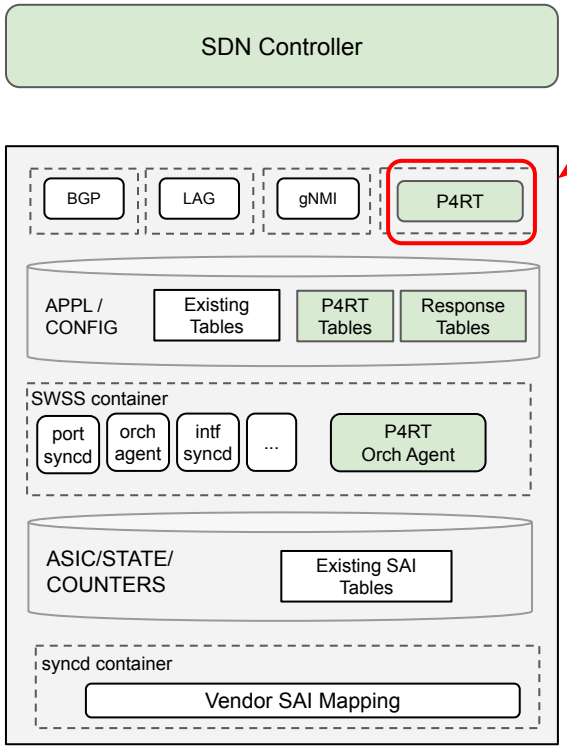

NETWORKING

OPEN POSSIBILITIES.

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

P4 Orchagent

- New orchagent in SWSS
  - Parses APPL_DB entries,
  - Maintains objects, refcounts
  - Translates intent to ASIC_DB
- Handles ordering dependencies
- Supports response path

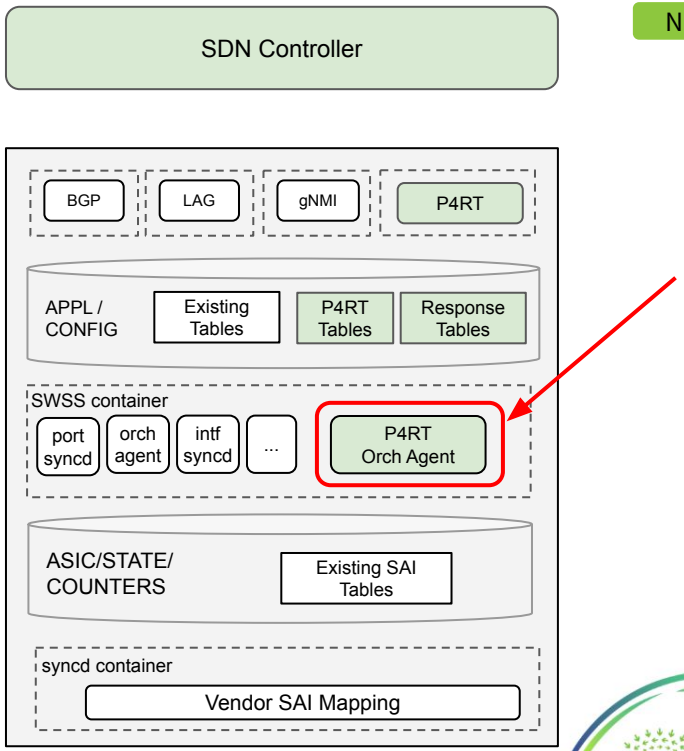

NETWORKING

Example: Installing a route

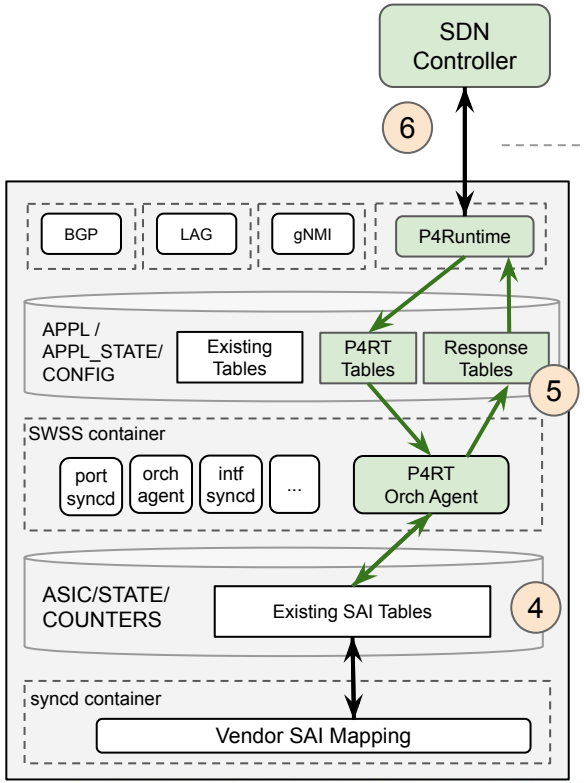

1 — P4RT RPC

```
RPC: Write
type: INSERT
ipv4_table_entry
  match
    vrf_id: "vrf-8"
    ipv4_dst: "192.168.0.0/24"
  action: set_nexthop_id
    nexthop_id: "s1"
```

2 — APPL DB

```
P4RT:FIXED_IPV4_TABLE:{"match/vrf_id":"vrf-8",
                       "match/ipv4_dst":"192.168.0.0/24"}
  "action" = "set_nexthop_id"
  "param/nexthop_id" = "s1"
```

3 — ASIC DB

```
"ASIC_STATE:SAI_OBJECT_TYPE_ROUTE_ENTRY:{"dest":"192.168.0.0/24",
      "switch_id":"oid:0x21000000000000","vr":"oid:0x3000000000008"}
  "SAI_ROUTE_ENTRY_ATTR_NEXT_HOP_ID" = "oid:0x500000000097b"
```

NETWORKING

OPEN POSSIBILITIES.

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

Packet I/O

Integration - Packet I/O

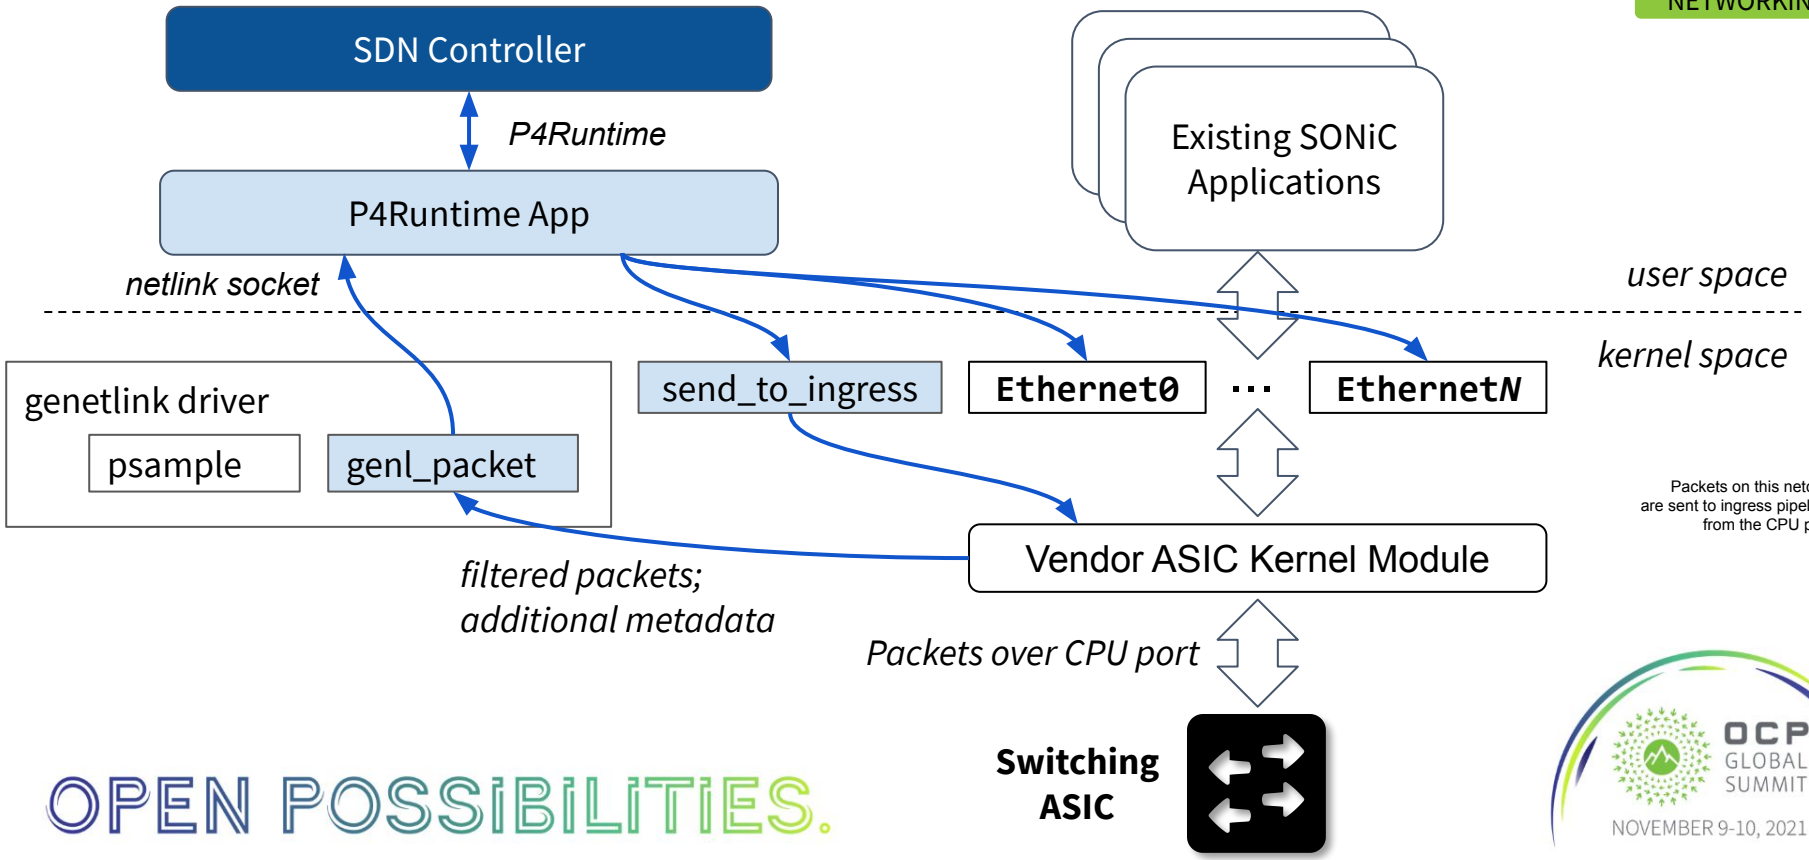

NETWORKING

OPEN POSSIBILITIES.

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

Demo

Use Case: L3 Routing & WCMP

NETWORKING

Use a converged view of the fabric topology to:

- Compute the global routing table
- Assign optimal weights for balanced traffic

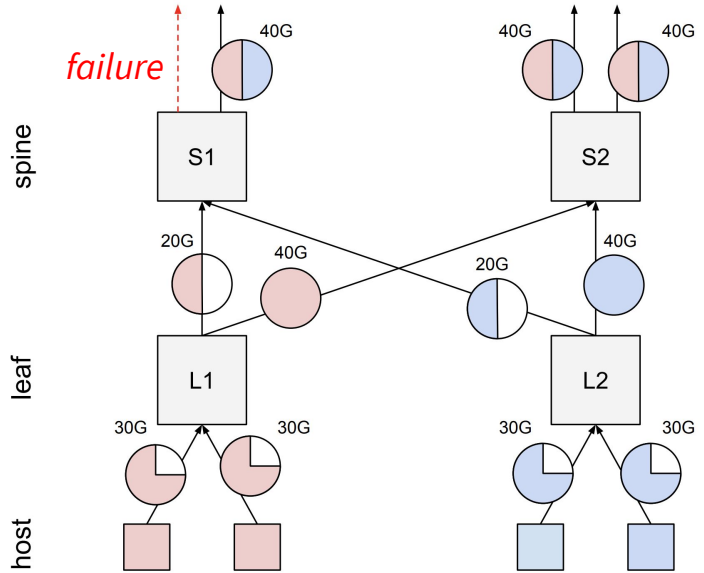

Use case details: [PINS: P4 Integrated Network Stack, 2021 P4 Workshop]

Demo: L3 Routing & WCMP

NETWORKING

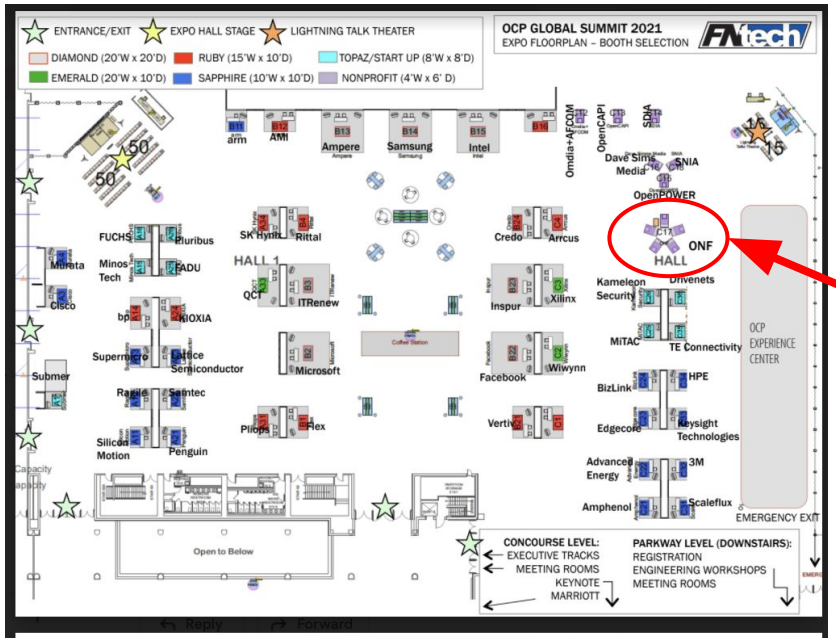

ONF Booth Location

ONOS® Open Network Operating System

TRELLIS

SDN Controller

P4Runtime
gNMI

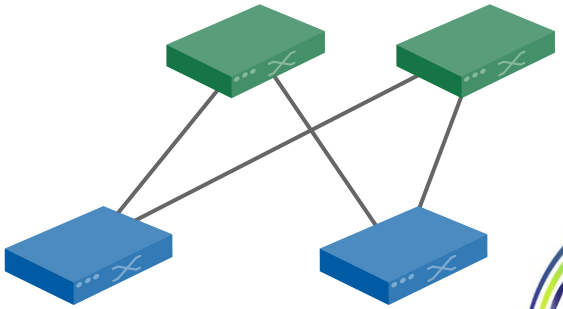

PINS

OPEN POSSIBILITIES.

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

Roadmap

Roadmap

SONiC Release	202111 - MVP	202206	Future
Release Date	Nov. 30, 2021	June 30, 2022	2022+

Features	• P4Runtime Application
	• SAI P4 Program
	  ○ L3 Admit / VRF
	  ○ IPv4 / IPv6 Routing
	  ○ WCMP Next Hop Groups
	  ○ Configurable ACLs
	• Packet I/O
	• P4Orch with Response Path
	• Integration with ONOS
	• Internal System Testing
	• SAI P4 Extension Path (delayed)
		• SAI P4 Program
		  ○ Hash configuration
		  ○ L2 FDB / VLAN
		  ○ SVI
		• Critical State
		• Open System Test Framework
		  ○ Test Cases
			• Response Path in other Orchs
			• SAI P4 Program
			  ○ VxLAN
			  ○ Other tunnels
			• gNOI support

Community – PINS Working Group

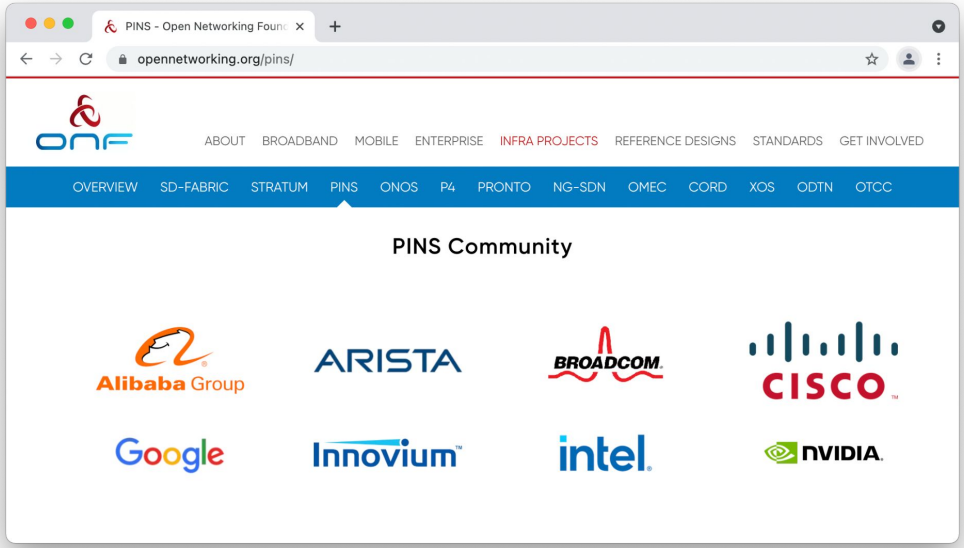

- Working group formed in 2019
- Multiple Use Cases (Data center, 5G, …)

PINS "MVP" is part of the SONiC.202111

OPEN POSSIBILITIES.

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

Call to Action

- Try out PINS in the SONiC.202111 release!
- Join the SONiC PINS Subgroup
- Comment on and contribute to PINS high-level designs (HLDs)
- Help us build the next set of features for upcoming releases
  - PINS “MVP” Pull Requests
  - SONiC repos: sonic-pins, sonic-swss, sonic-swss-common, sonic-buildimage
  - Active development in the ONF PINS Working Group repos

Thank you!

OCP GLOBAL SUMMIT

NOVEMBER 9-10, 2021