

An abstract graphic on the left side of the image, composed of numerous thin, wavy green lines that swirl and overlap to form a complex, organic shape. The lines are a vibrant green color against the dark blue background.

Open. Together.



OCP
SUMMIT



NETWORKING

SONiC Deployments Powered by Programmable Data Plane

Arkadiy Shapiro

Product Line Manager, Core Software & Technical Partnerships
Barefoot Networks



OPEN
COMMUNITY®



@ArkadiyShapiro



Open. Together.

SONiC Background

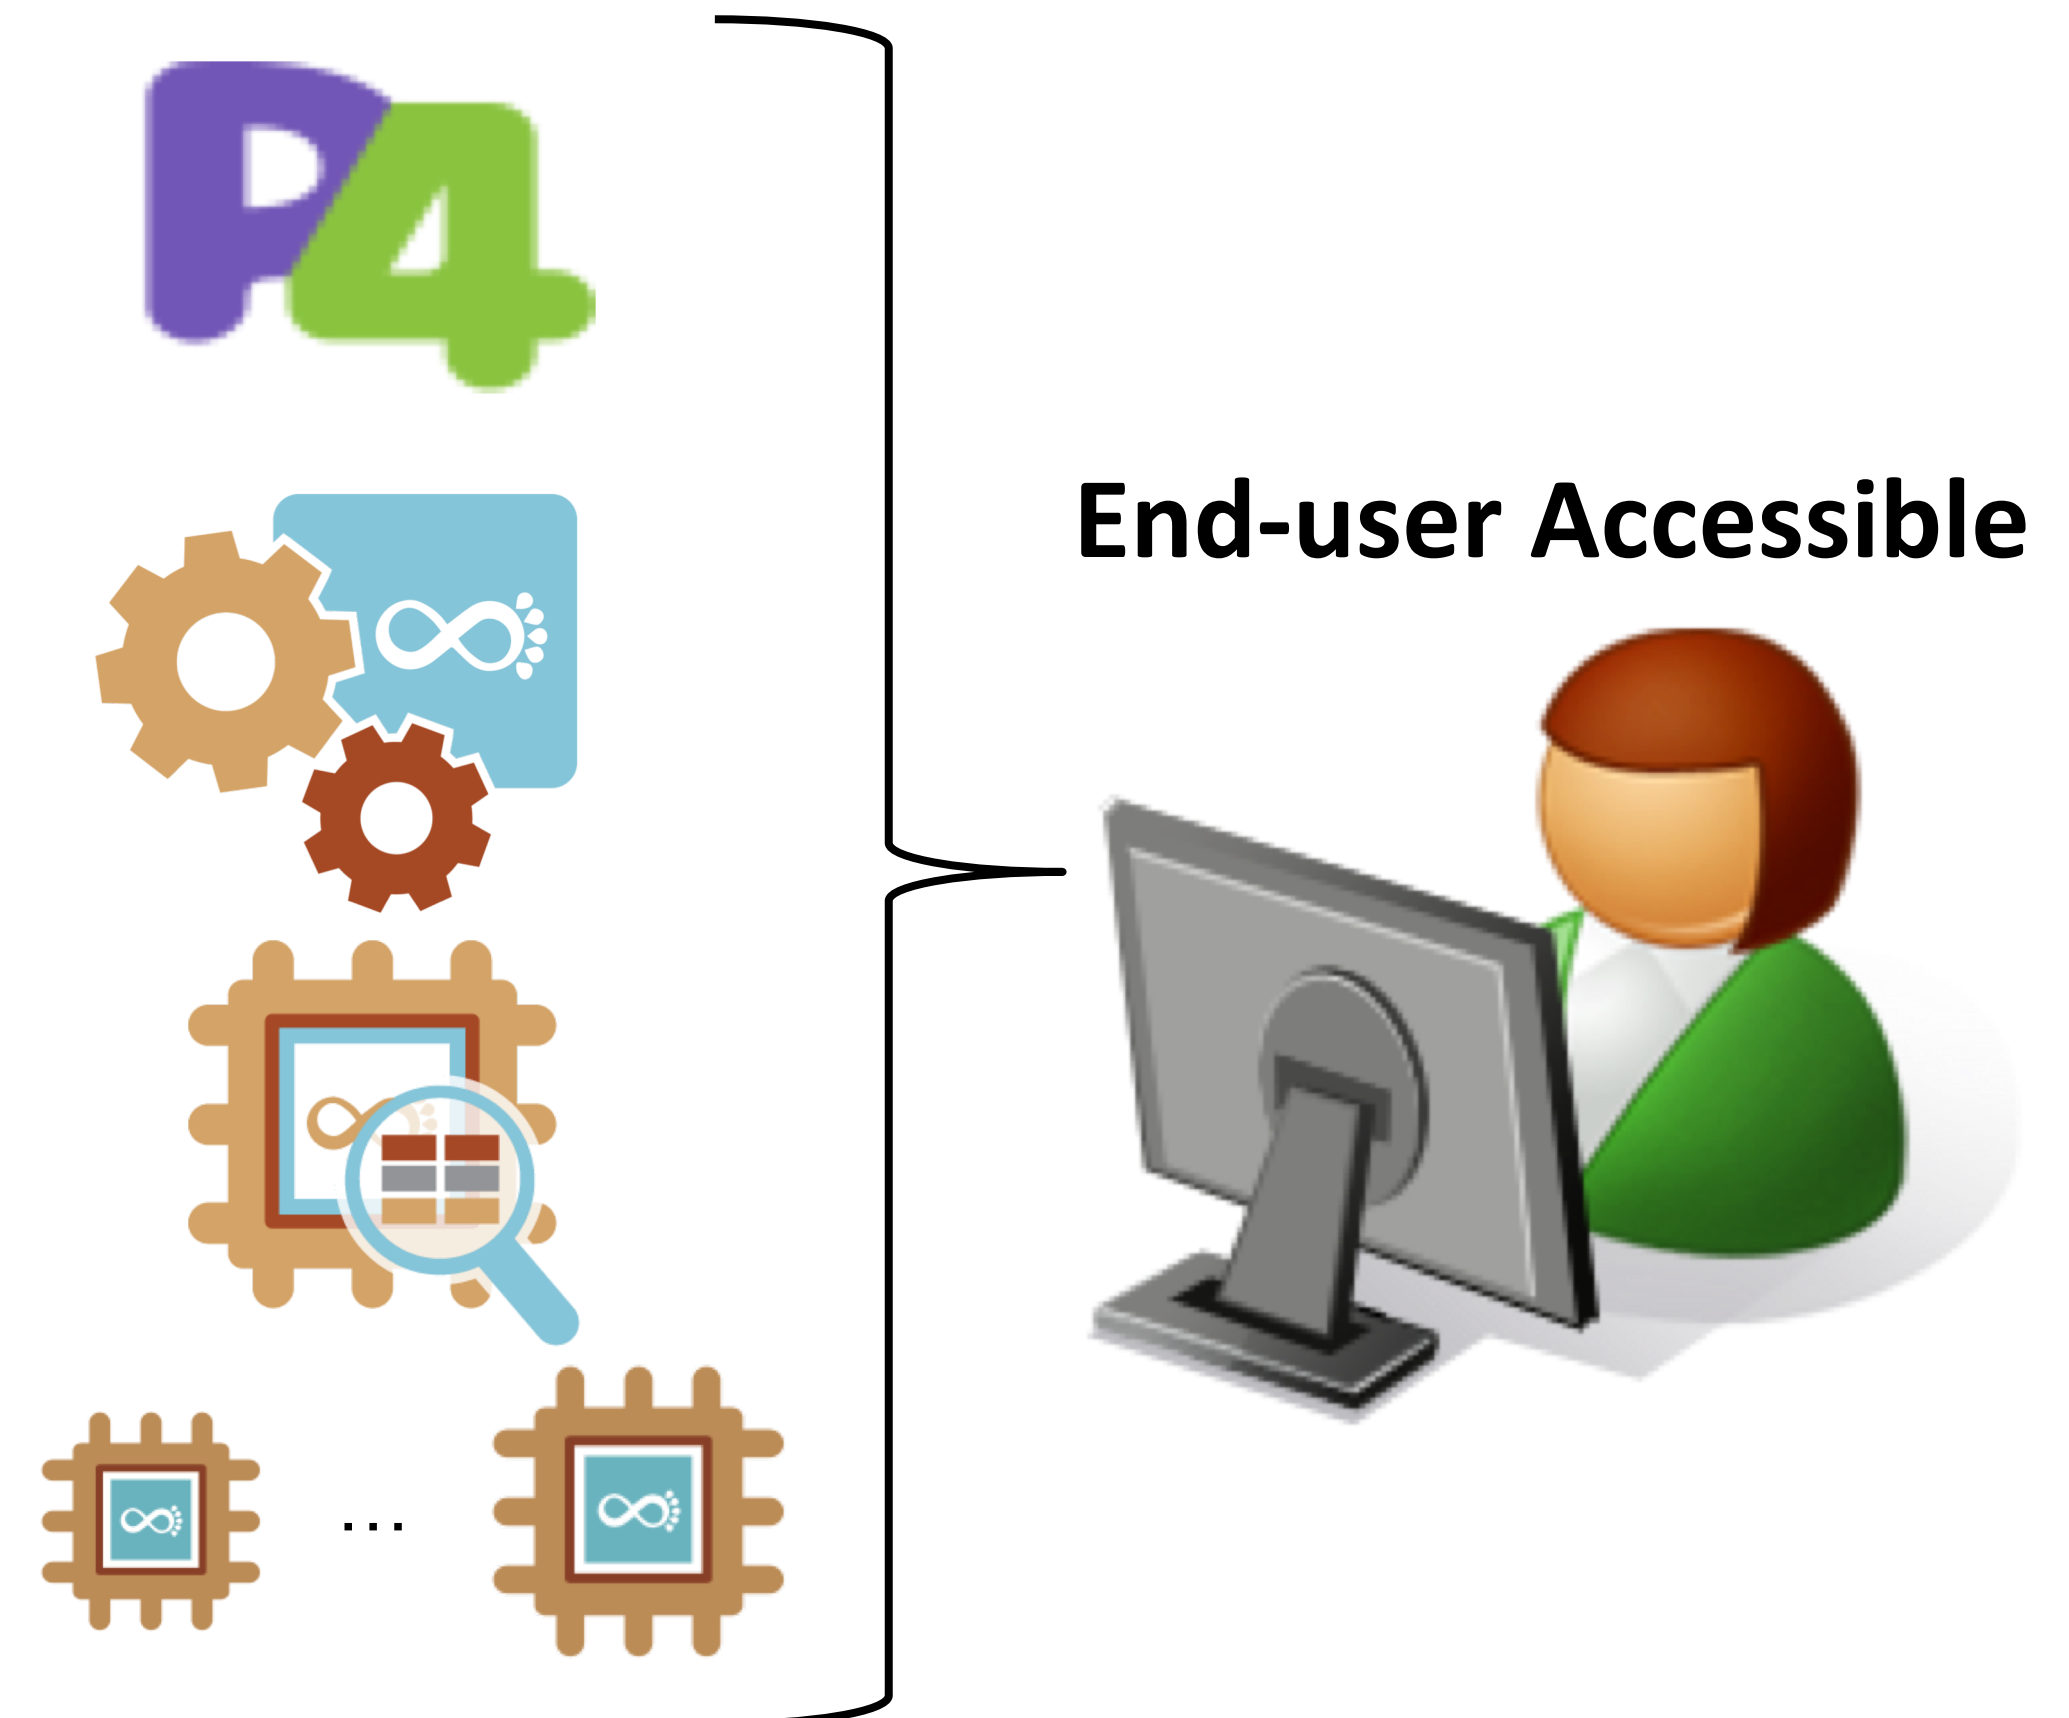
- Announced @ March 2016 OCP Summit
- Growing open-source network OS project with multiple contributors (MSDC, ASIC, ODM, OEM)
- **Originally focused on fixed function switches, so how can we leverage programmable data plane?**



What makes you programmable?

AVOIDING “PROGRAMMABILITY” WASHING...

- 1 Open programming language
- 2 Compiler & Chip Simulation Model
- 3 Visualization Tools
- 4 Fully Programmable Silicon

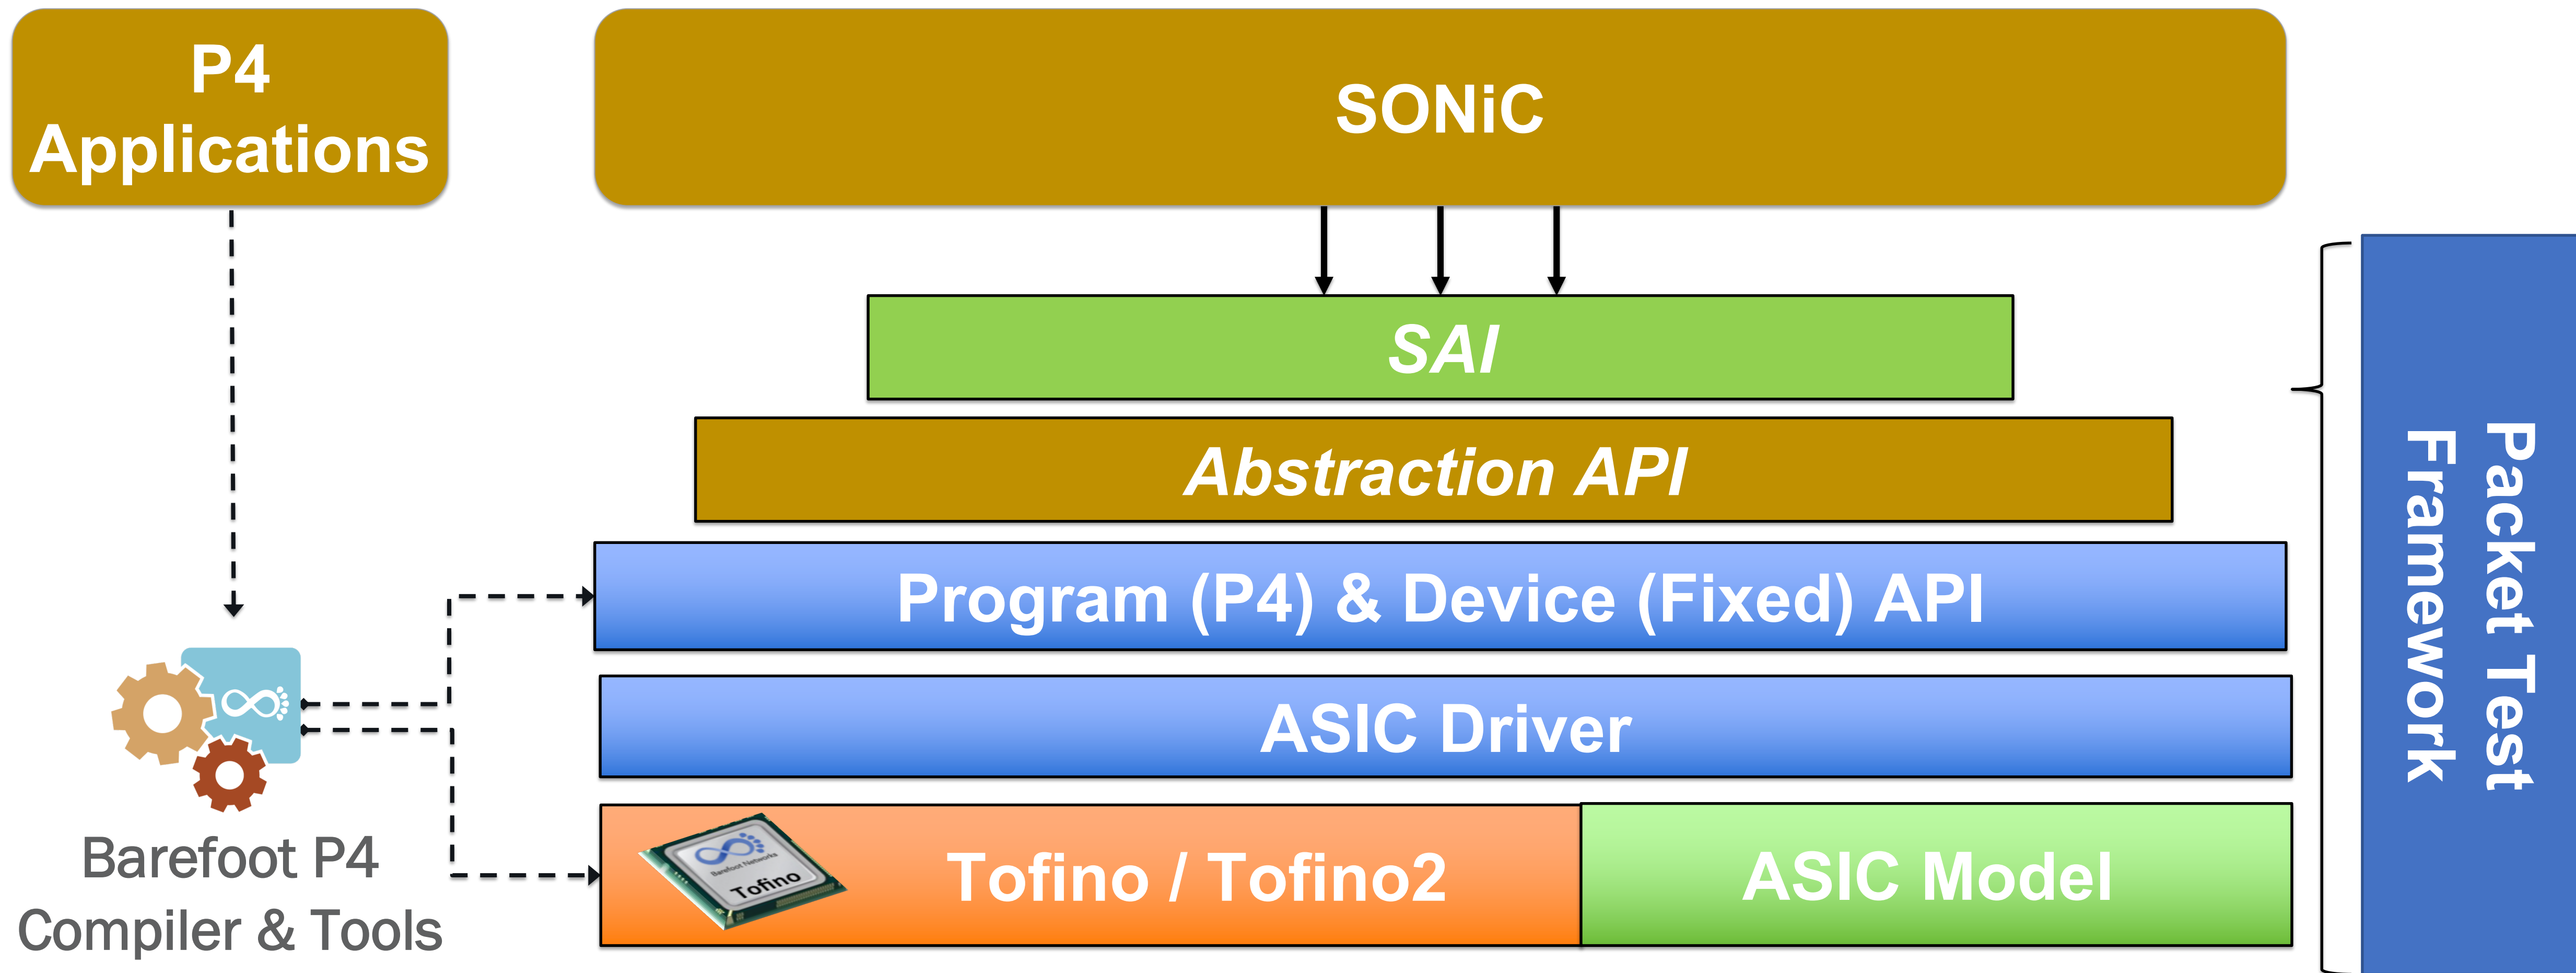


Barefoot SONiC Update

SINCE OCP U.S. SUMMIT 2018

- Continuous **upstreaming of changes** to support latest SAI / SONiC releases
- **Data plane telemetry** feature contribution and participation in SAI TAM discussions
- Expand set of **supported platforms**
- Leading platform to validate **VxLAN** support on SONiC with **cloud scale**
- Support latest SONiC **high availability** features like warm reboot with **zero service disruption**
- Enable **multiple use-cases** powered by P4 inside one SONiC image

Barefoot Software Support for SONiC



Barefoot Tofino Platform Support for SONiC



STORDIS®



Inventec



Sample Switch.P4 Features for SONiC

- **Ethernet switching**
 - VLAN Flooding
 - MAC Learning & Aging
 - STP state
 - VLAN Translation
- **IPv4 and IPv6 routing**
 - Unicast Routing
 - Routed Ports & SVI
 - VRF
 - Unicast RPF - Strict and Loose
 - Multicast - PIM-SM/DM & PIM-Bidir
- **QOS**
 - QoS Classification & marking
 - Drop profiles/WRED -
 - RoCE v2 / PFC
 - CoPP (Control plane policing)
 - WRED-based ECN marking
- **Tunneling: VxLAN (v4/v6), IP-in-IP, GRE**
- **ACL**
 - Ingress MAC ACL, IPv4/v6 ACL, RACL
 - Egress MAC ACL, IPv4/v6 ACL, RACL
 - QoS ACL, System ACL, PBR
 - Port Range lookups in ACLs
- **Security Features**
 - Storm Control,
 - IP Source Guard
- sFlow
- PTP
- **Counters**
 - Route Table Entry Counters
 - VLAN/Bridge Domain Counters
 - **Port/Interface Counters**
 - ACL stats
- **Barefoot Dataplane Telemetry**

SONiC with Programmable Silicon

- P4 program defines a **use-case** (ToR, Spine, Gateway):
 - **Features** enabled / disabled (L2, L3, MPLS, VxLAN, Telemetry etc.)
 - **Scale** for each table (MACs, routes, ACLs etc.)
- Delivering different data planes with operating system:
 - Option 1: Compile different images
 - **Option 2: Single image includes multiple P4 programs**



Switching Between Different Data Planes

AVAILABLE FROM 201807 SONIC RELEASE IN A SINGLE BAREFOOT DEBIAN PACKAGE

1. In SONiC, edit
/etc/sonic/config_db.json to
include the p4_profile attribute:

```
...  
"DEVICE_METADATA": { "localhost": {  
...  
  "p4_profile": "<P4 program name"&#34; } }
```

2. Load the updated
config_db.json:

```
sudo config load -y
```
3. Reboot the switch for the new data
plane to take effect


Embedded P4 programs:

- **MSDC_PROFILE**
 - Default for baseline SONiC features
- **MSDC_IPV4_PROFILE**
 - Focused on high scale tunneling
- **MSDC_LEAF_DTEL_INT_PROFILE**
 - Dataplane telemetry leaf
- **MSDC_SPINE_DTEL_INT_PROFILE**
 - Dataplane telemetry spine

SONiC Delivery

HOW TO GET STARTED

Option	Scenario
Binary file on SONiC community page	Quick start with SONiC supported features
Binary files provided by Barefoot support	Quick start with features not upstreamed
Compile from P4 Studio SDE	Data plane / SAI implementation change



Use-case 1: SONiC and Table Scale

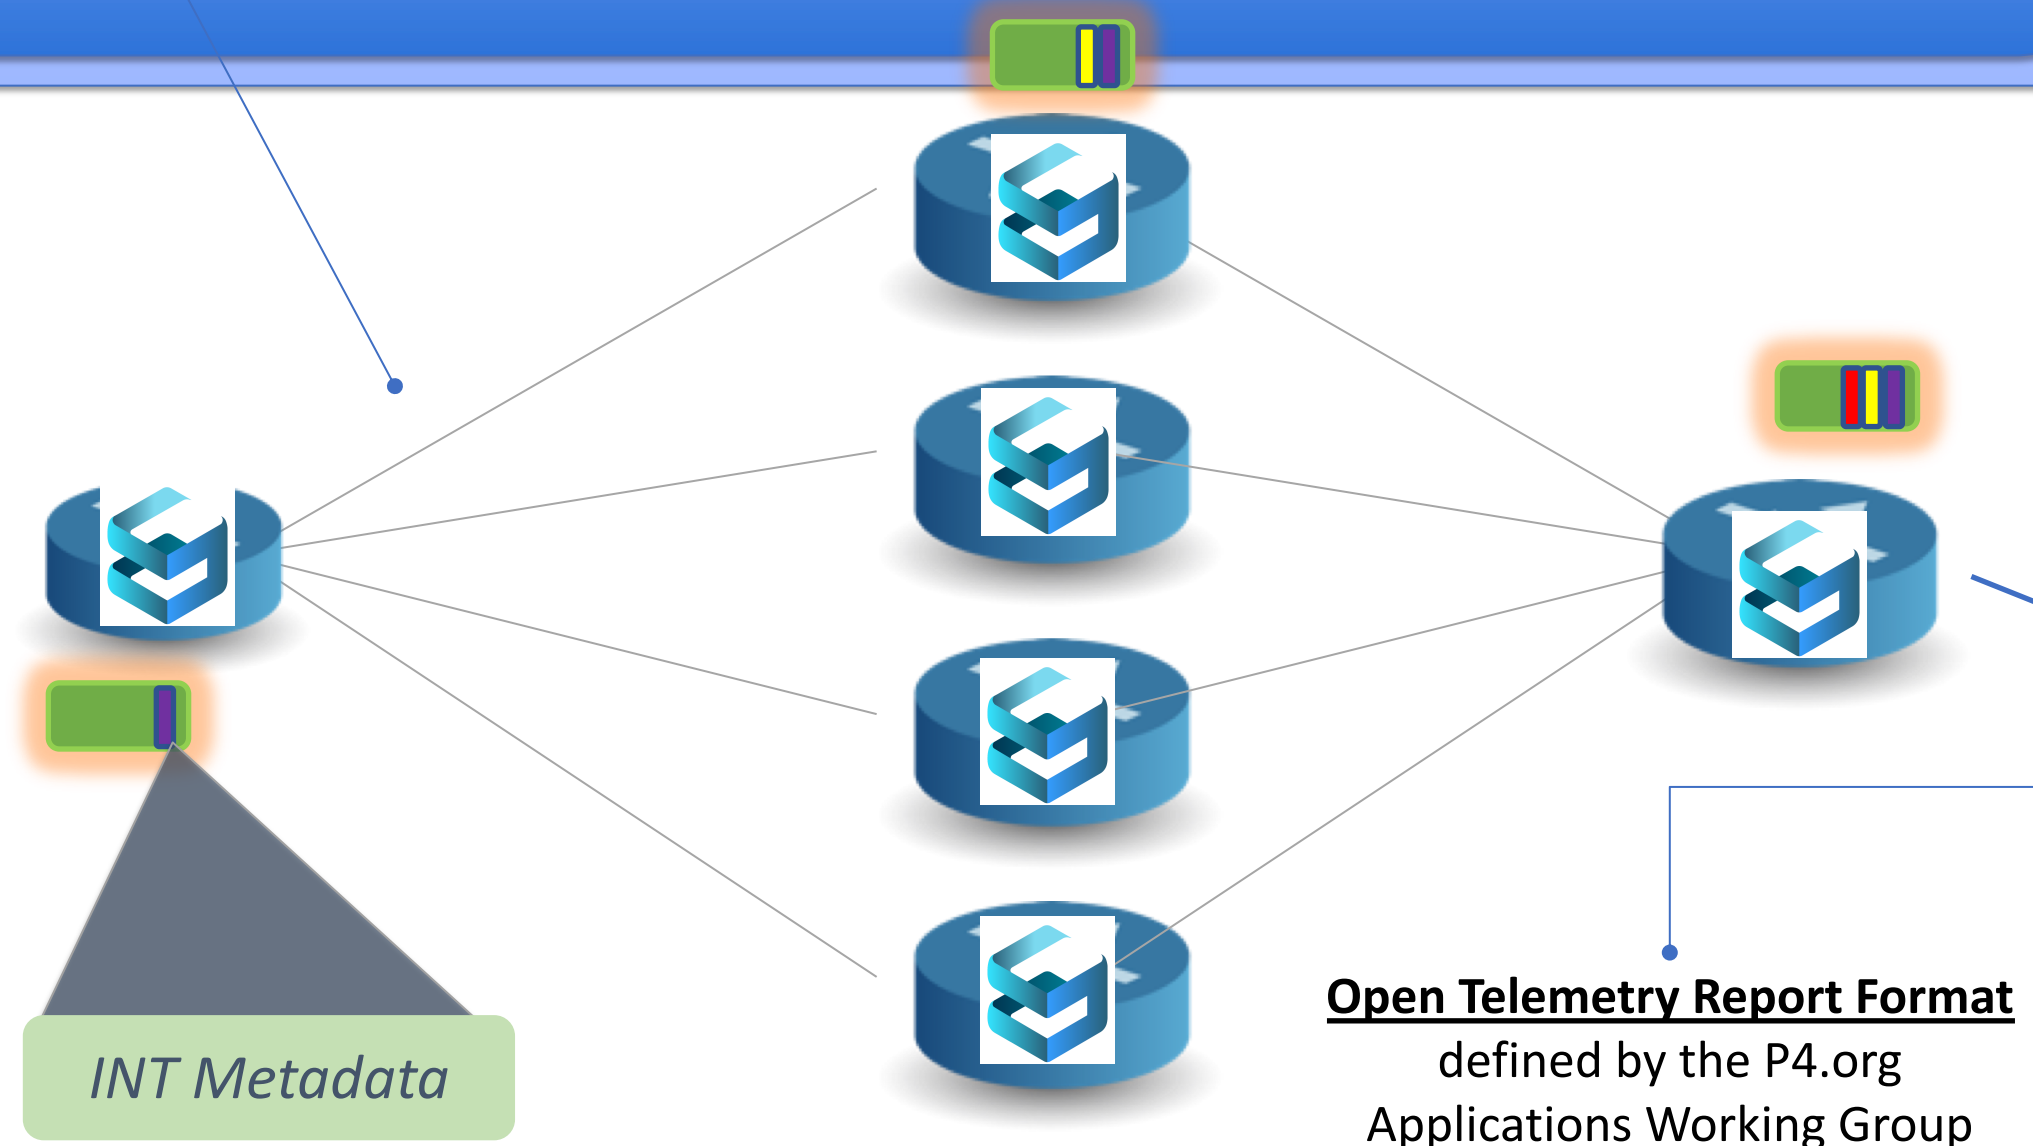
- Different table sizes for leaf and spine
- Different table sizes for different deployments
- Example: IPv4 vs IPv6 heavy fabric

Parameter	Scenario 1	Scenario 2
IPv4 Host Local	128K	32K
IPv4 LPM	16K	4K
IPv6 Host	4K	4K
IPv6 LPM	10K	60K


Use-case 2: SONiC & Data-Plane Telemetry

Barefoot Data-Plane Telemetry

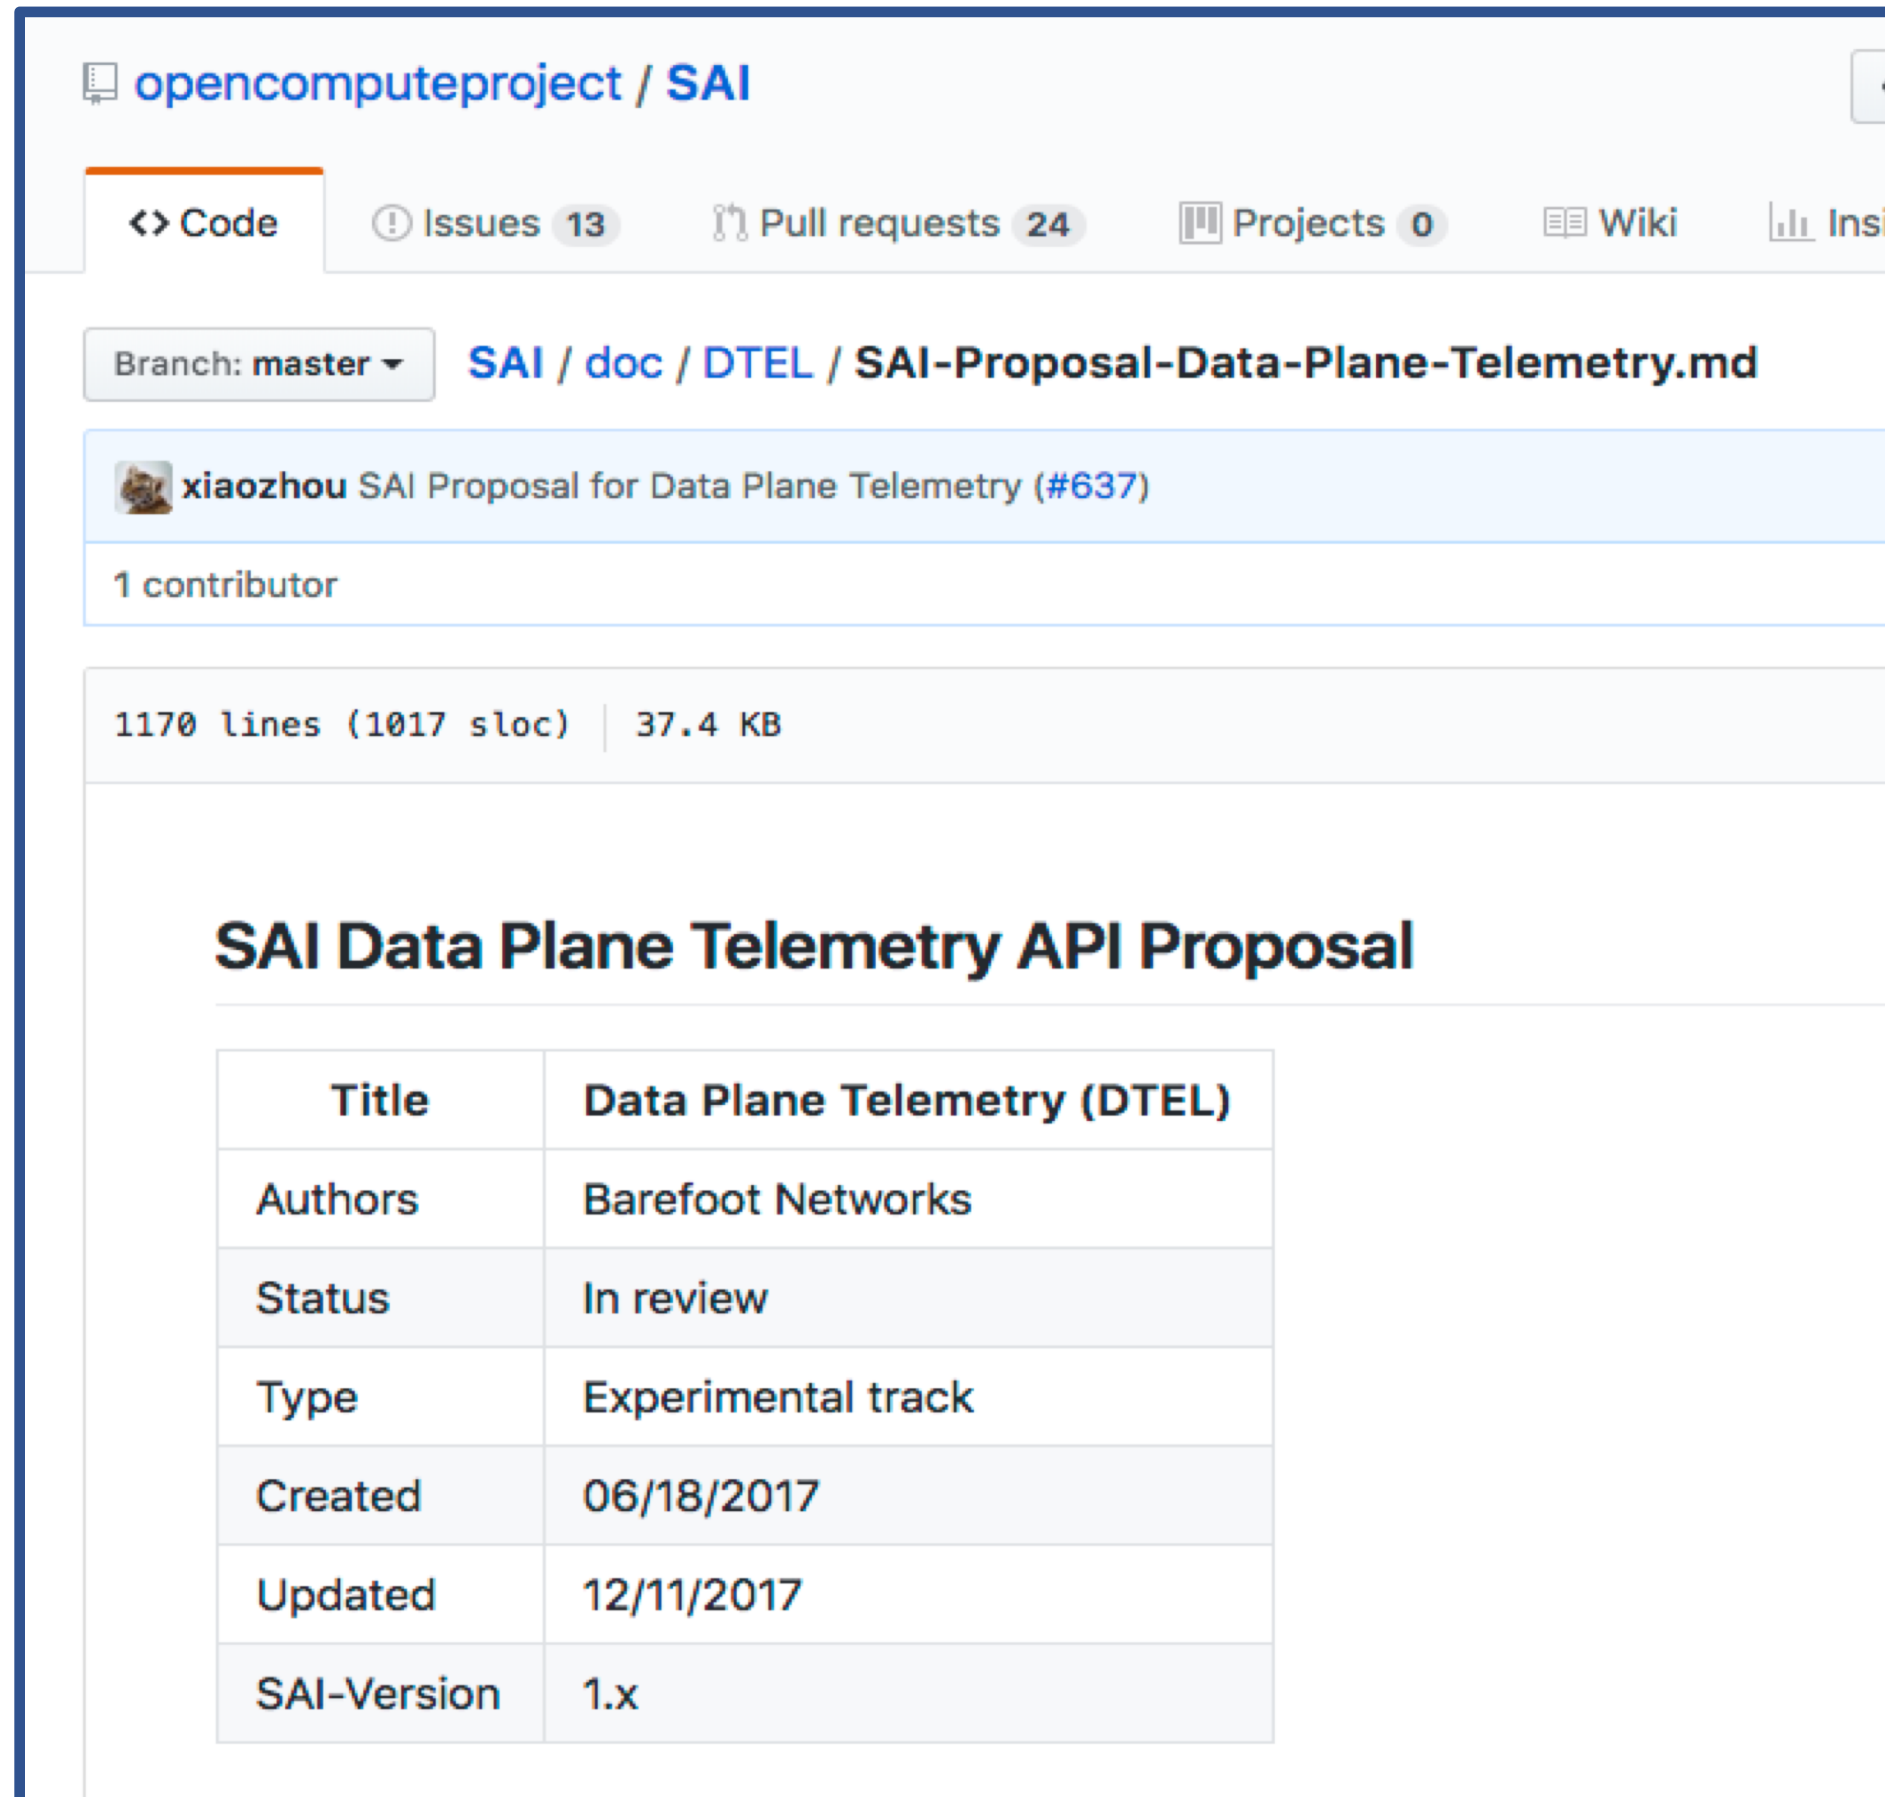
- In-Band Network Telemetry (INT)
- Intelligent Deduplication and Triggers
- Path & Latency Change Detection
- Microburst detection
- Report on Drop



3rd Party Network Management Solutions



SAI Data Plane Telemetry APIs



The screenshot shows a GitHub pull request interface for the 'opencomputeproject / SAI' repository. The pull request is titled 'SAI Proposal for Data Plane Telemetry (#637)' and is authored by 'xiaozhou'. It shows 1 contributor, 1170 lines of code (1017 sloc), and a file size of 37.4 KB. The pull request is for the 'master' branch and is for the file 'SAI / doc / DTEL / SAI-Proposal-Data-Plane-Telemetry.md'. Below the pull request details, there is a table titled 'SAI Data Plane Telemetry API Proposal'.

Title	Data Plane Telemetry (DTEL)
Authors	Barefoot Networks
Status	In review
Type	Experimental track
Created	06/18/2017
Updated	12/11/2017
SAI-Version	1.x

- Part of SAI 1.3
- Defines entire spectrum of data plane telemetry configuration
 - Flow watchlists
 - Switch ID
 - Report destination

SONiC and Data Plane Telemetry

- SONiC Telemetry Feature
- New tables in several SONiC databases
- Configuration script
- Released in SONiC 201807 release!

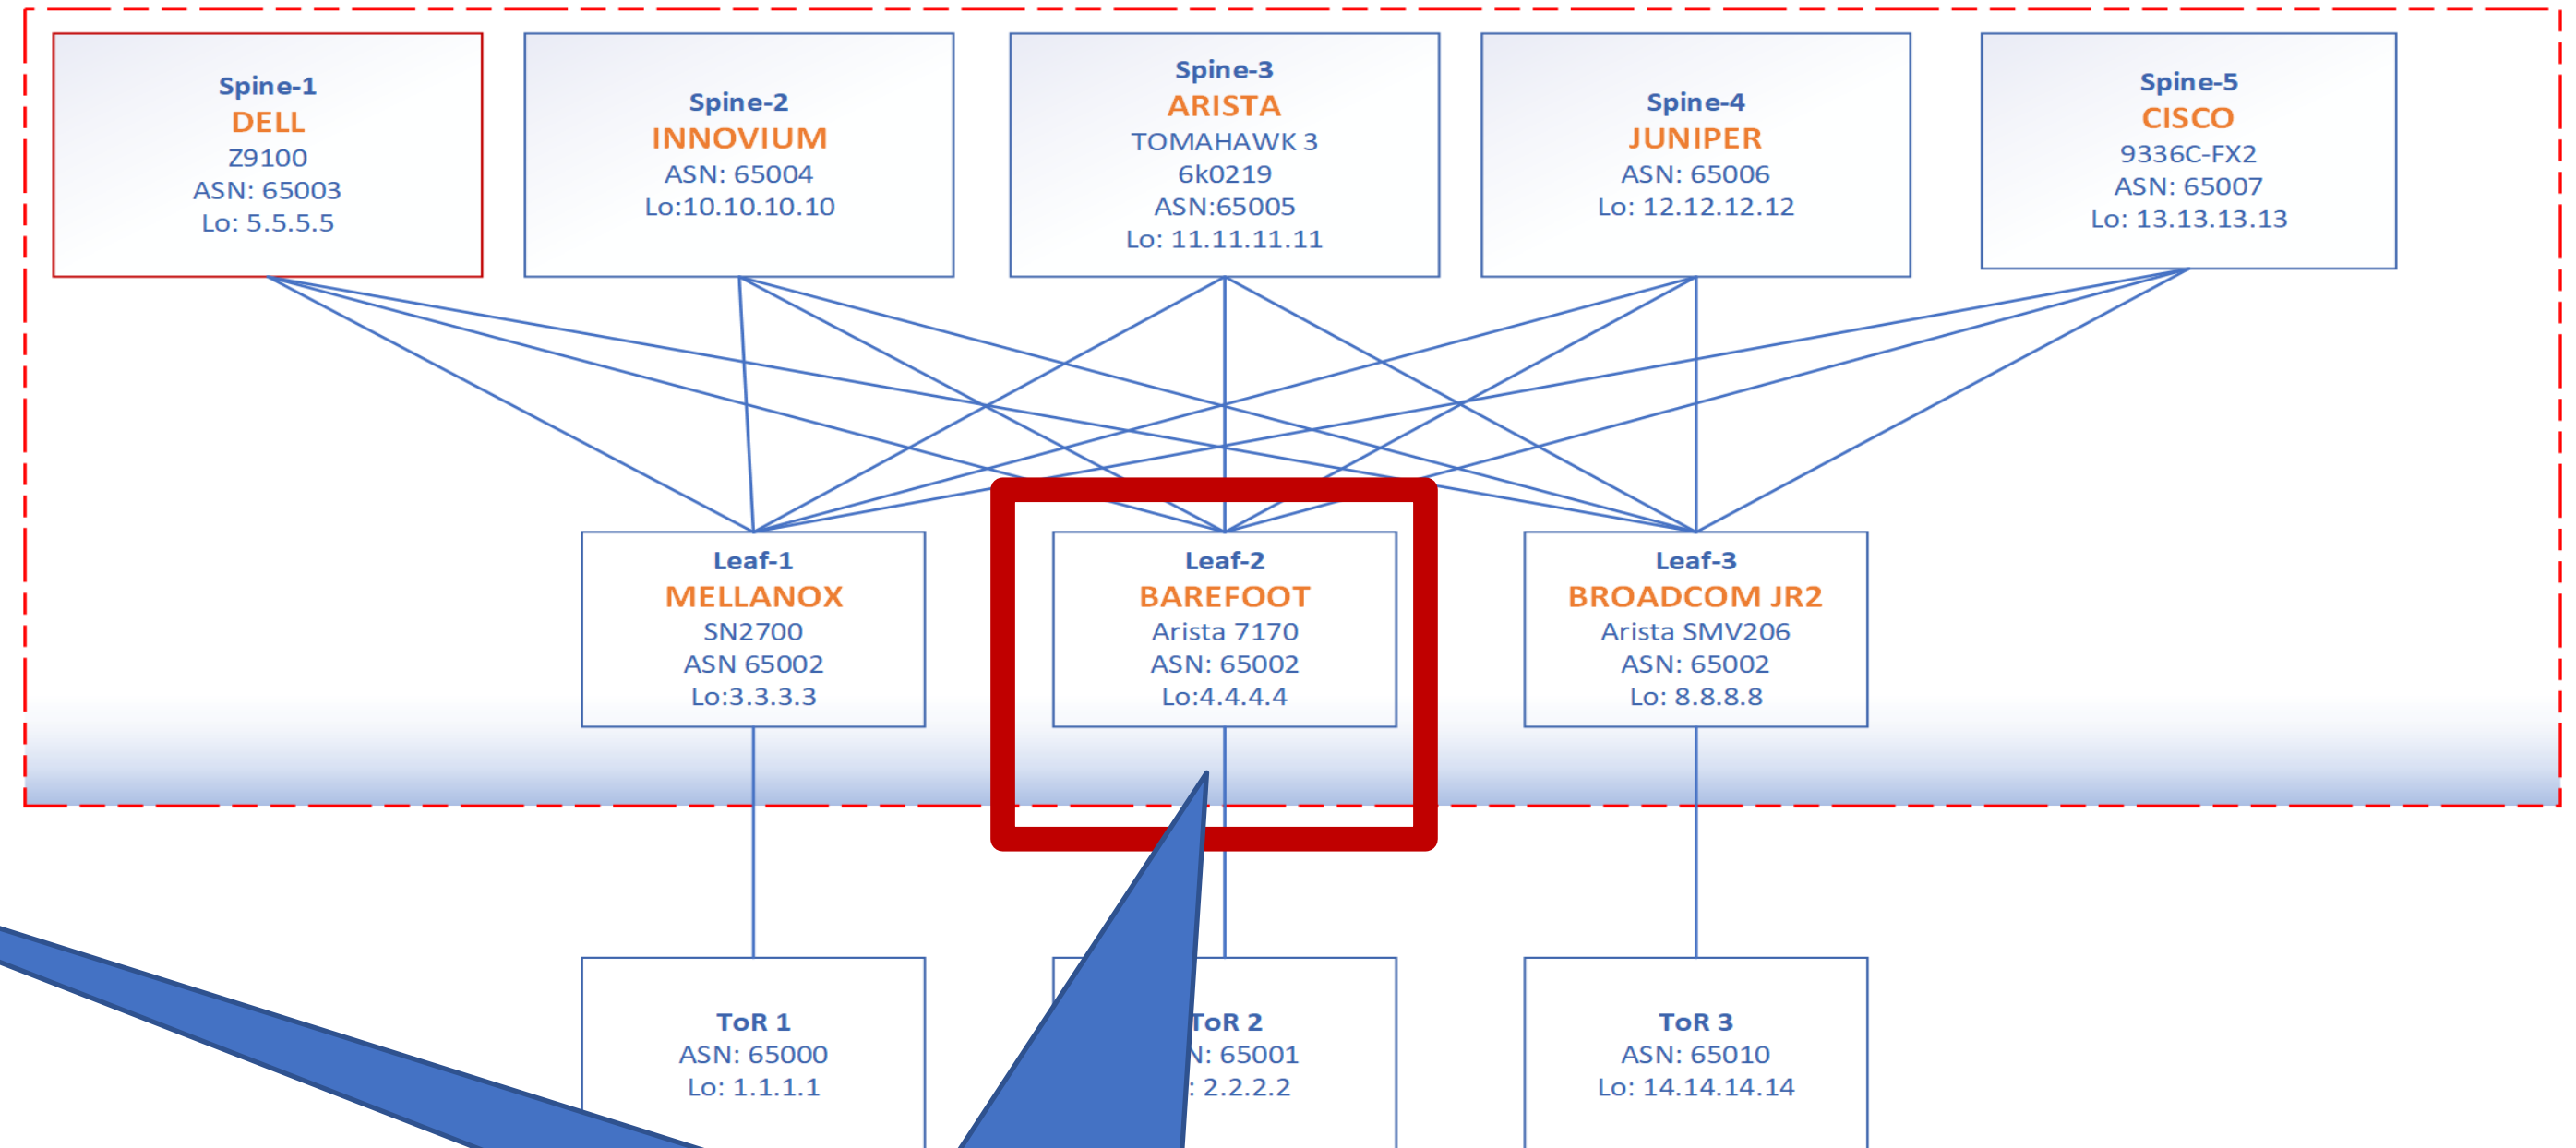
```
# Instantiate a switch:
my_switch = sonic_switch.SONiCSwitch(dtel_switch_id='123',
                                     management_ip='10.10.10.10',
                                     dtel_monitoring_type='int_endpoint')

# Create a report session:
rs = my_switch.create_dtel_report_session('192.168.0.1')

# Create a watchlist:
wl = my_switch.create_dtel_watchlist('flow')

# Add entries to the watchlist:
wl.create_entry(priority=10,
               src_ip='10.131.0.0',
               src_ip_mask=11,
               dst_ip='10.131.0.0',
               dst_ip_mask=11,
               dtel_sample_percent=100,
               dtel_report_all=True)
```


Baremetal Server Hosting / Disaggregated Chassis



Powered by Arista 7170 and
Barefoot Tofino

SONiC Deployment Challenges

IMPROVEMENT OPPORTUNITIES TO DRIVE BROADER SONIC ACCEPTANCE

- Stable released branch / image vs master
- SAI incompatibility with prior versions
- Community tests to become more flexible
 - Not tied to specific fan-out switch
 - Increase use-case coverage

Call to Action

- See Barefoot related presentations and demos:
<https://www.barefootnetworks.com/blog/bringing-ocp-vision-reality-programmable-dataplanes-progress-report/>
- Use latest SONiC binary for Barefoot to try out some of the use-cases
<https://github.com/Azure/SONiC/wiki/Supported-Devices-and-Platforms>
- Use Barefoot P4 Studio and P4 to modify data plane to address new use cases



Open. Together.

OCP Global Summit | March 14–15, 2019

