



**OCP**  
SUMMIT

March 20-21  
**2018**  
San Jose, CA

**OPEN. FOR BUSINESS.**



# SAI Updates and Looking Forward

Guohan Lu, Principal Dev Manager

Xin Liu, Principal Product Manager

Microsoft Azure Networking

**OPEN. FOR BUSINESS.**




# Switch Abstraction Interface (SAI)

Network Applications


Hello

Switch Abstraction Interface


частный




你好



नमस्ते



Bonjour





SONiC

OPX

OS10

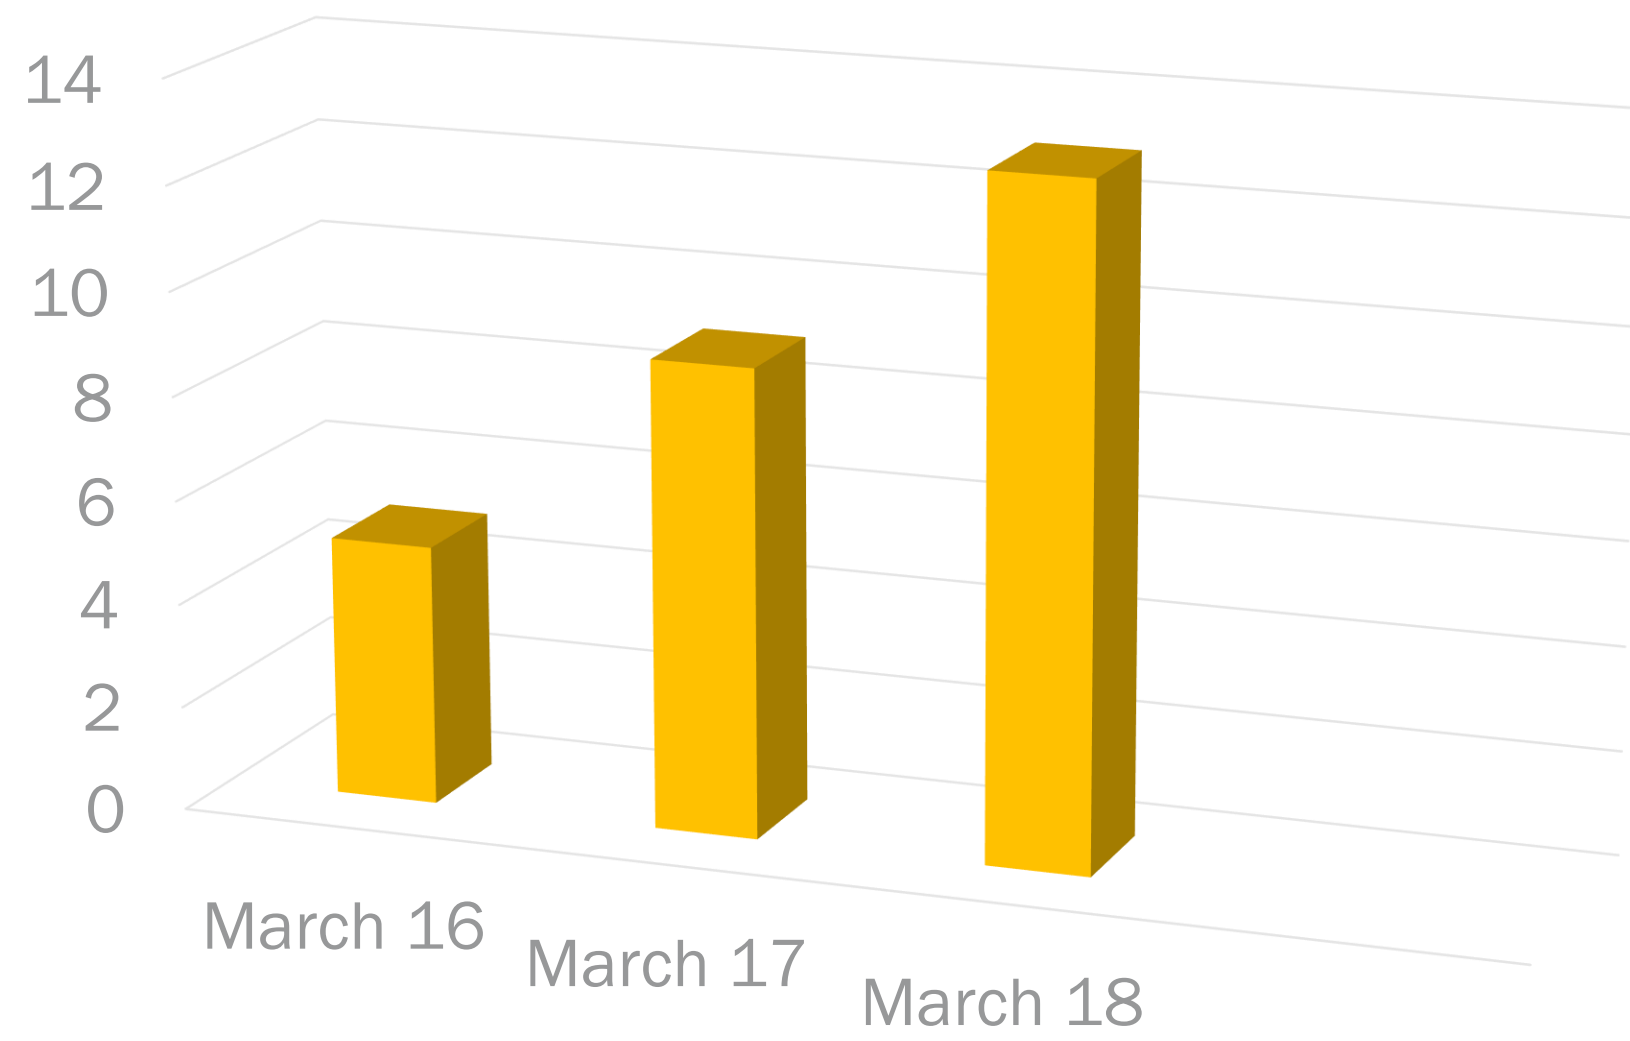
Metaswitch

CNOS

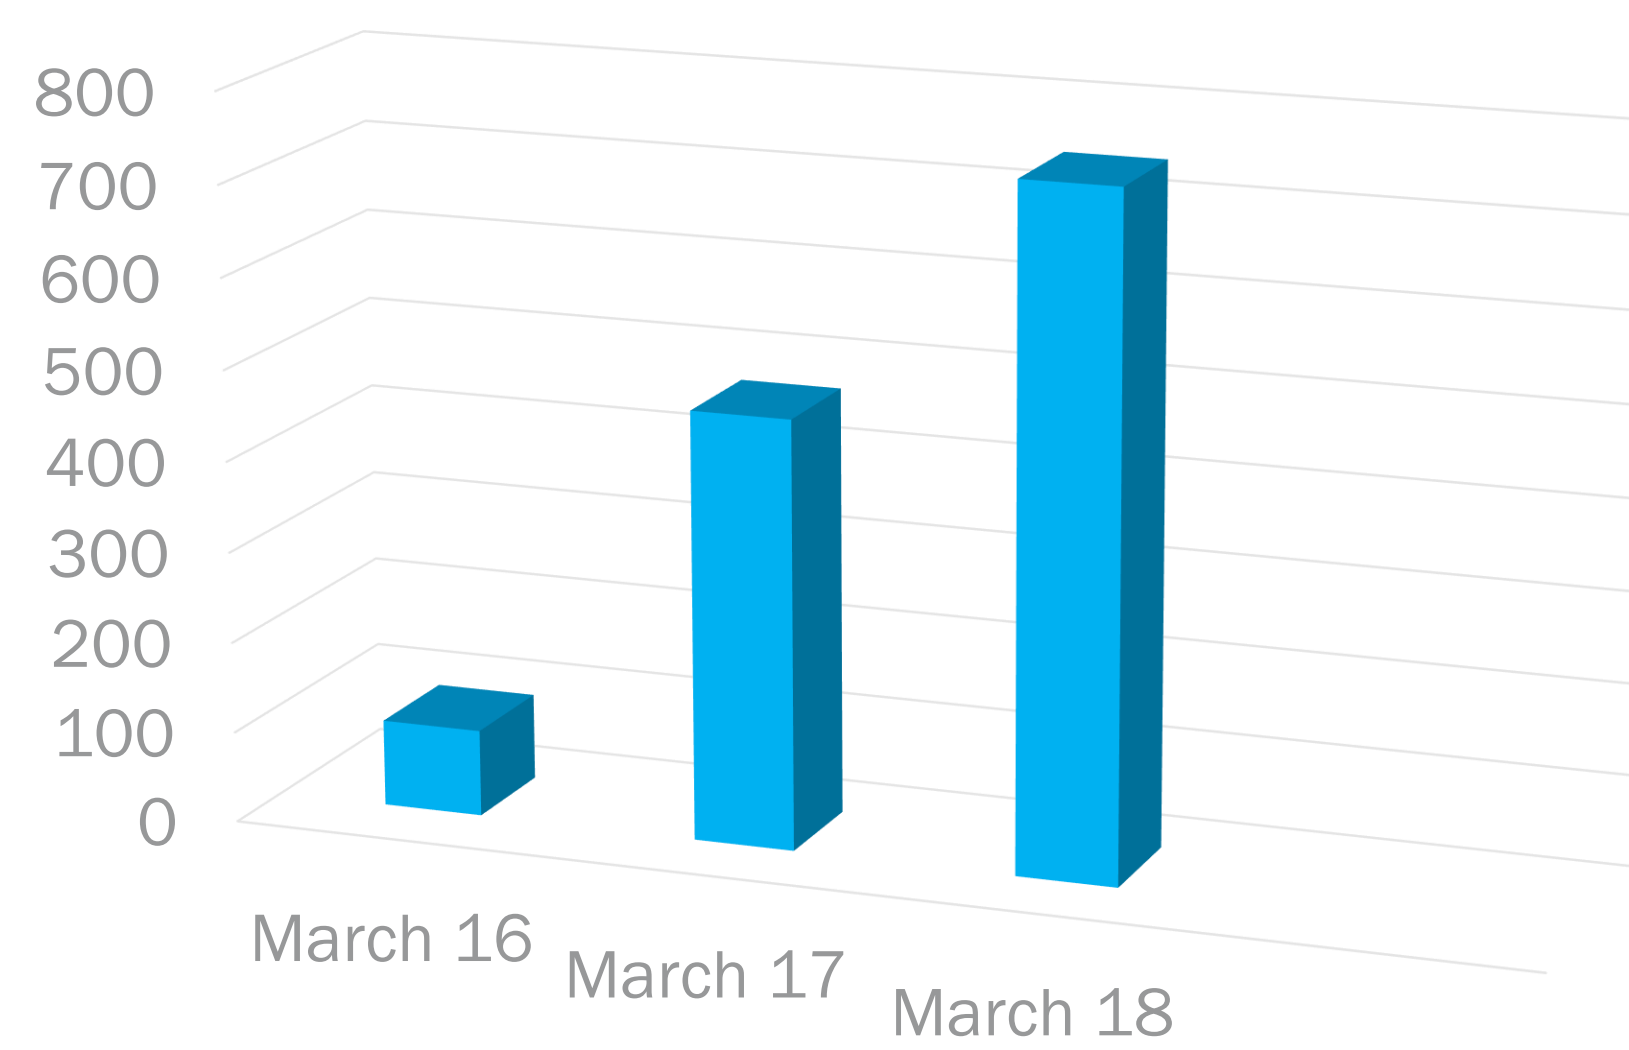
dNOS





# Organizations



# Commits



# Proposals



# SAI Achievements In 2017

## Monitoring

TAM [Broadcom]

Microburst  
[Marvell]

Critical Resource  
Monitoring [MSFT]

Data Plane Telemetry  
[Barefoot]

## Protocols

MPLS [Mellanox]

BFD [Dell]

Segment Routing  
[Cavium]

VxLan-Ext [MSFT]

## Reliability/QoS

L3 Fast Reroute  
[Metaswitch]

ECN [Dell]

## Infrastructure

SAI P4 Model  
[Mellanox]

Multi-NPU [Dell]

Capability Query  
[Mellanox]

SAI Ext API [Dell]

# SAI Challenges

- SAI capabilities
  - Different ASICs has different capabilities
  - SAI now supported by 8 ASIC vendors
  - Applications need to accommodate different SAI capabilities
- SAI for specialized forwarding functions
  - Networking gears in data center and WANs are still using traditional packet forwarding functions, L2 forwarding, L3 routing
  - However, certain SDN features need customized forwarding functions
  - Allow user to customize SAI pipeline and expose programmable pipeline

# Applications Adapt to ASIC Capability Through SAI

- Applications query the SAI capability
- Applications adapt its behavior based on the SAI capability
- **Example: ACL port range** (`permit tcp any any 1024 1048`)
  - When ASIC has port range checker, apps creates SAI port range check
  - When ASIC does not have port range checker, application can use multiple ACL rules to emulate port range checker



# Generic API for SAI capability

- SAI consists of objects and attributes
  - Query(object\_type, attribute)
  - Query(object\_type, attribute, valid enum range)
  - ACL Range
  - sai\_query\_attribute\_capability(SAI\_OBJECT\_TYPE\_ACL\_RANGE, SAI\_ACL\_RANGE\_ATTR\_TYPE)
  - sai\_query\_attribute\_enum\_values\_capability(SAI\_OBJECT\_TYPE\_ACL\_RANGE, SAI\_ACL\_RANGE\_ATTR\_TYPE, enum\_values\_capability)
  - enum\_values\_capability: SRC\_PORT\_RANGE, DST\_PORT\_RANGE, ...

# Infrastructure to Expose SAI Capability


- SAI meta data infrastructure
- Allow vendor to extend meta to include their own information

```
/* @type sai_object_id_t
 * @flags CREATE_AND_SET
 * @objects SAI_OBJECT_TYPE_NEXT_HOP,
            SAI_OBJECT_TYPE_NEXT_HOP_GROUP,
            SAI_OBJECT_TYPE_ROUTER_INTERFACE,
            SAI_OBJECT_TYPE_PORT
 * @allownull true
 * @default SAI_NULL_OBJECT_ID
 */
SAI_ROUTE_ENTRY_ATTR_NEXT_HOP_ID,
```

```
const sai_attr_metadata_t sai_metadata_attr_SAI_ROUTE_ENTRY_ATTR_NEXT_HOP_ID = {
    .objecttype           = SAI_OBJECT_TYPE_ROUTE_ENTRY,
    .attrid               = SAI_ROUTE_ENTRY_ATTR_NEXT_HOP_ID,
    .attridname           = "SAI_ROUTE_ENTRY_ATTR_NEXT_HOP_ID",
    .brief                 = "Next hop or next hop group id for the packet, or
a router interface in case of directly reachable route, or the CPU port in case of
IP2ME route.",
    .attrvaluetype        = SAI_ATTR_VALUE_TYPE_OBJECT_ID,
    .flags                 =
(sai_attr_flags_t)(SAI_ATTR_FLAGS_CREATE_AND_SET),
    .allowedobjecttypes   =
sai_metadata_SAI_ROUTE_ENTRY_ATTR_NEXT_HOP_ID_allowed_objects,
    .allowedobjecttypeslength = 4,
    .allowrepetitiononlist = false,
    .allowmixedobjecttypes = false,
    .allowemptylist       = false,
    .allownullobjectid    = true,
    .isoidattribute        = (4 > 0),
    .defaultvaluetype     = SAI_DEFAULT_VALUE_TYPE_CONST,
    .defaultvalue         =
&sai_metadata_SAI_ROUTE_ENTRY_ATTR_NEXT_HOP_ID_default_value,
    .defaultvalueobjecttype = SAI_OBJECT_TYPE_NULL,
    .defaultvalueattrid    = SAI_INVALID_ATTRIBUTE_ID,
    .storedefaultvalue    = false,
    .isenum               = false,
    .isenumlist           = false,
    .enummetadata         = NULL,
    .conditiontype        = SAI_ATTR_CONDITION_TYPE_NONE,
    .conditions           = NULL,
    .conditionslength     = 0,
    .isconditional        = (0 != 0),
    .validonlytype        = SAI_ATTR_CONDITION_TYPE_NONE,
    .validonly            = NULL,
    .validonlylength      = 0,
    .isvalidonly          = (0 != 0),
    .getsave              = false,
    .isvlan               = false,
    .isaclfield           = false,
    .isaclaction          = false,
    .ismandatoryoncreate = false,
    .iscreateonly         = false,
    .iscreateandset       = true,
    .isreadonly           = false,
    .iskey                 = false,
    .isprimitive          = true,
};
```

# SAI for Common Functionalities

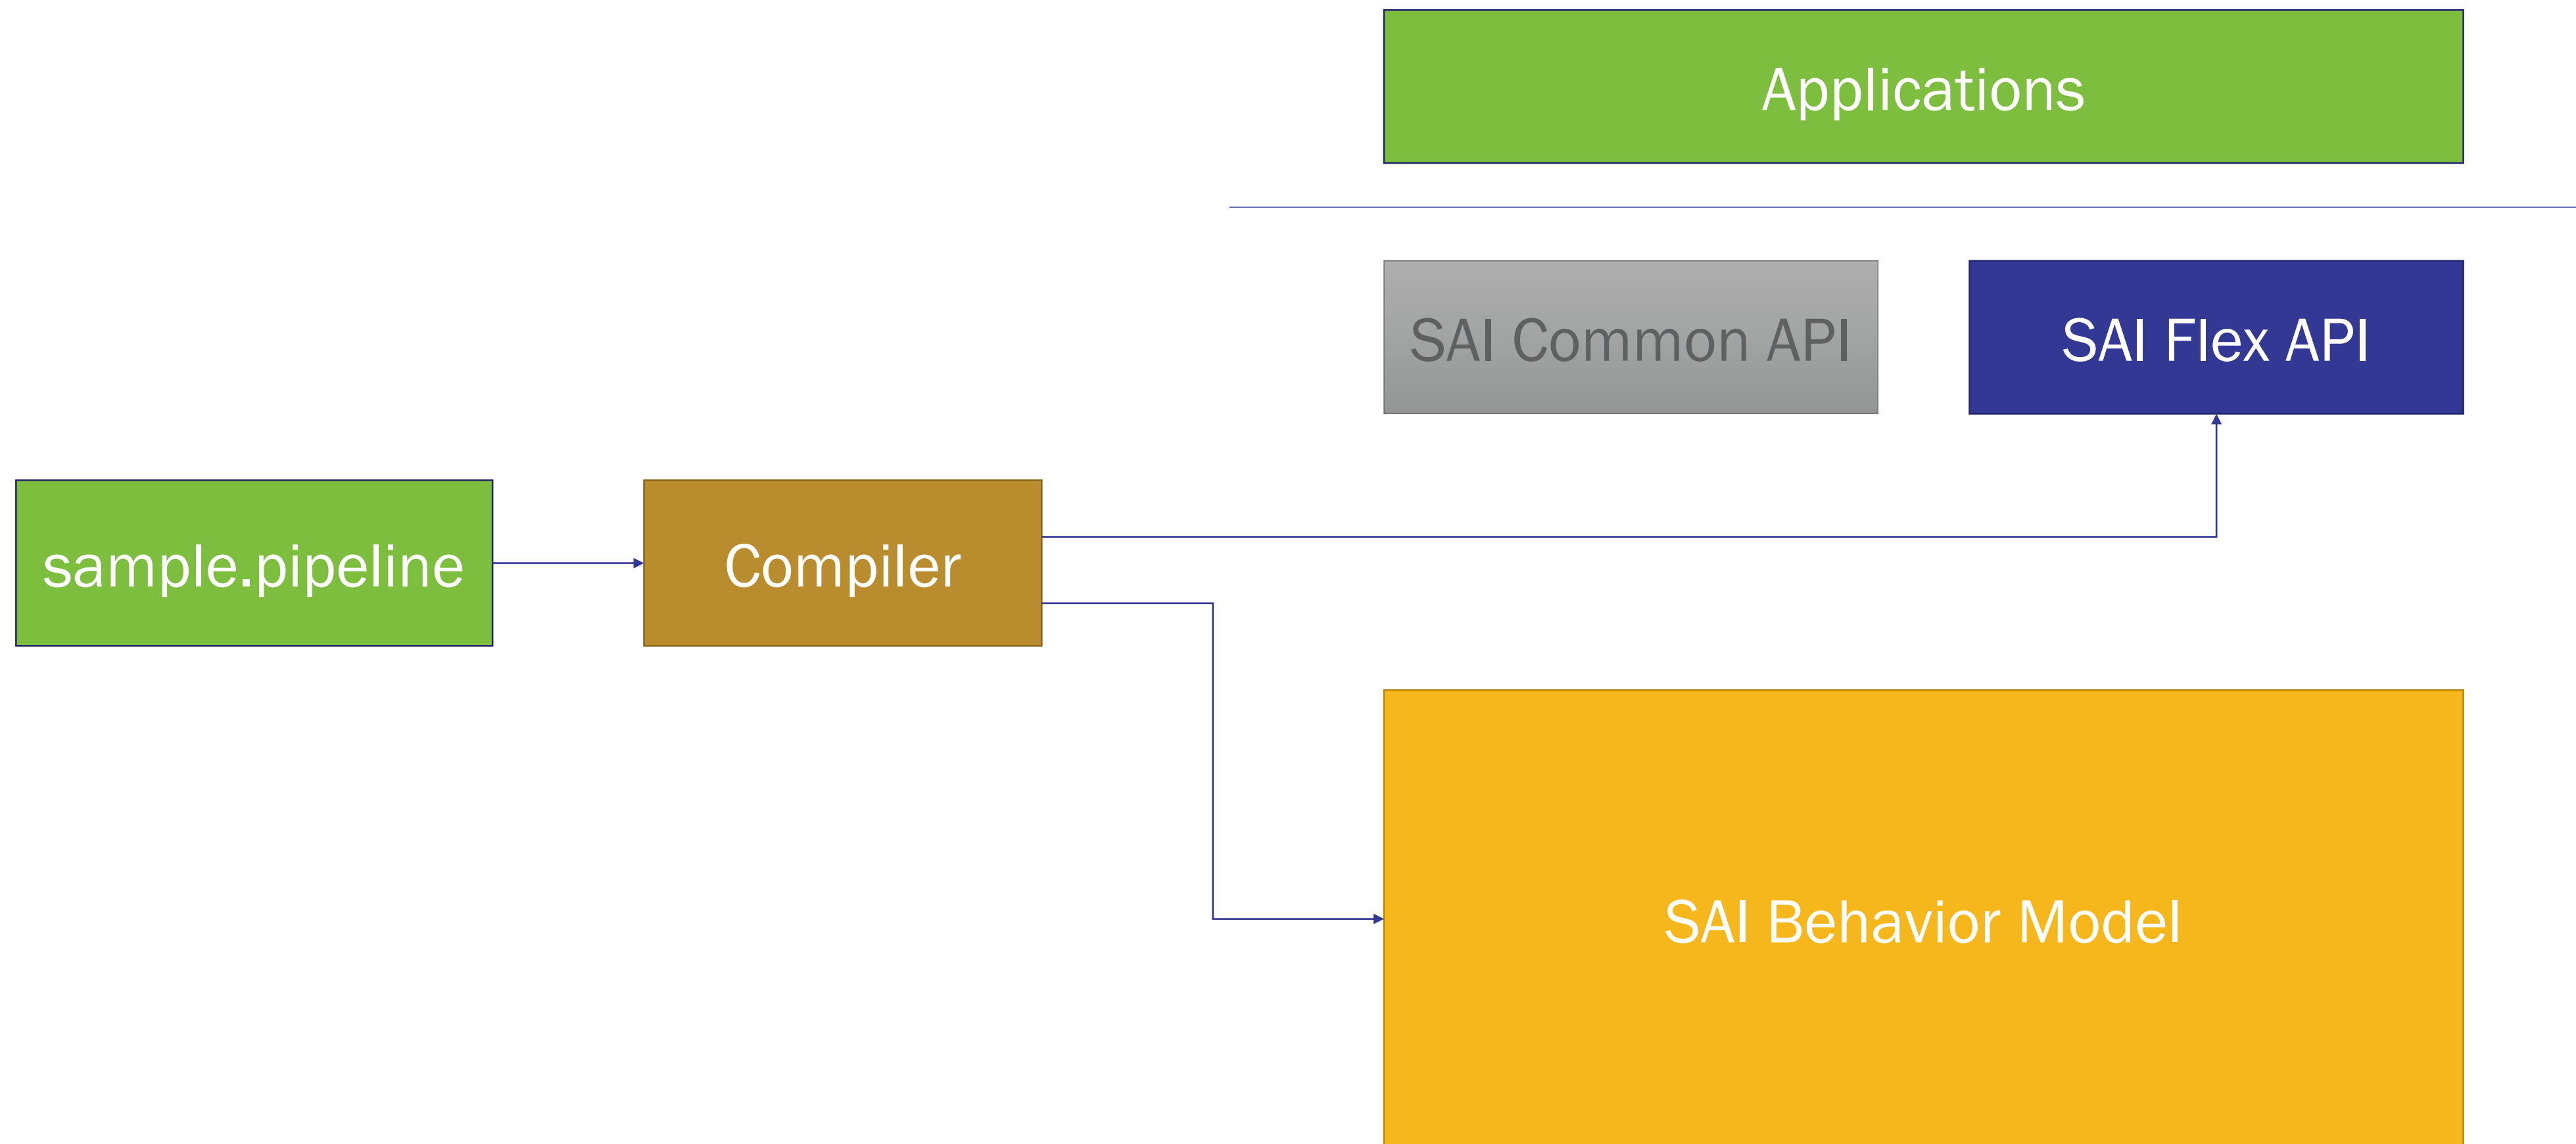
- Goal: provide common API for different ASICs
- SAI works very well for common packet forwarding functionalities
  - Define common APIs
  - Define SAI packet pipeline model (SAI behavior model)
  - Each ASIC maps the SAI API to their ASIC SDK based on SAI behavior model
- Example, L3 packet forwarding model



# SAI for Specialized Functionalities

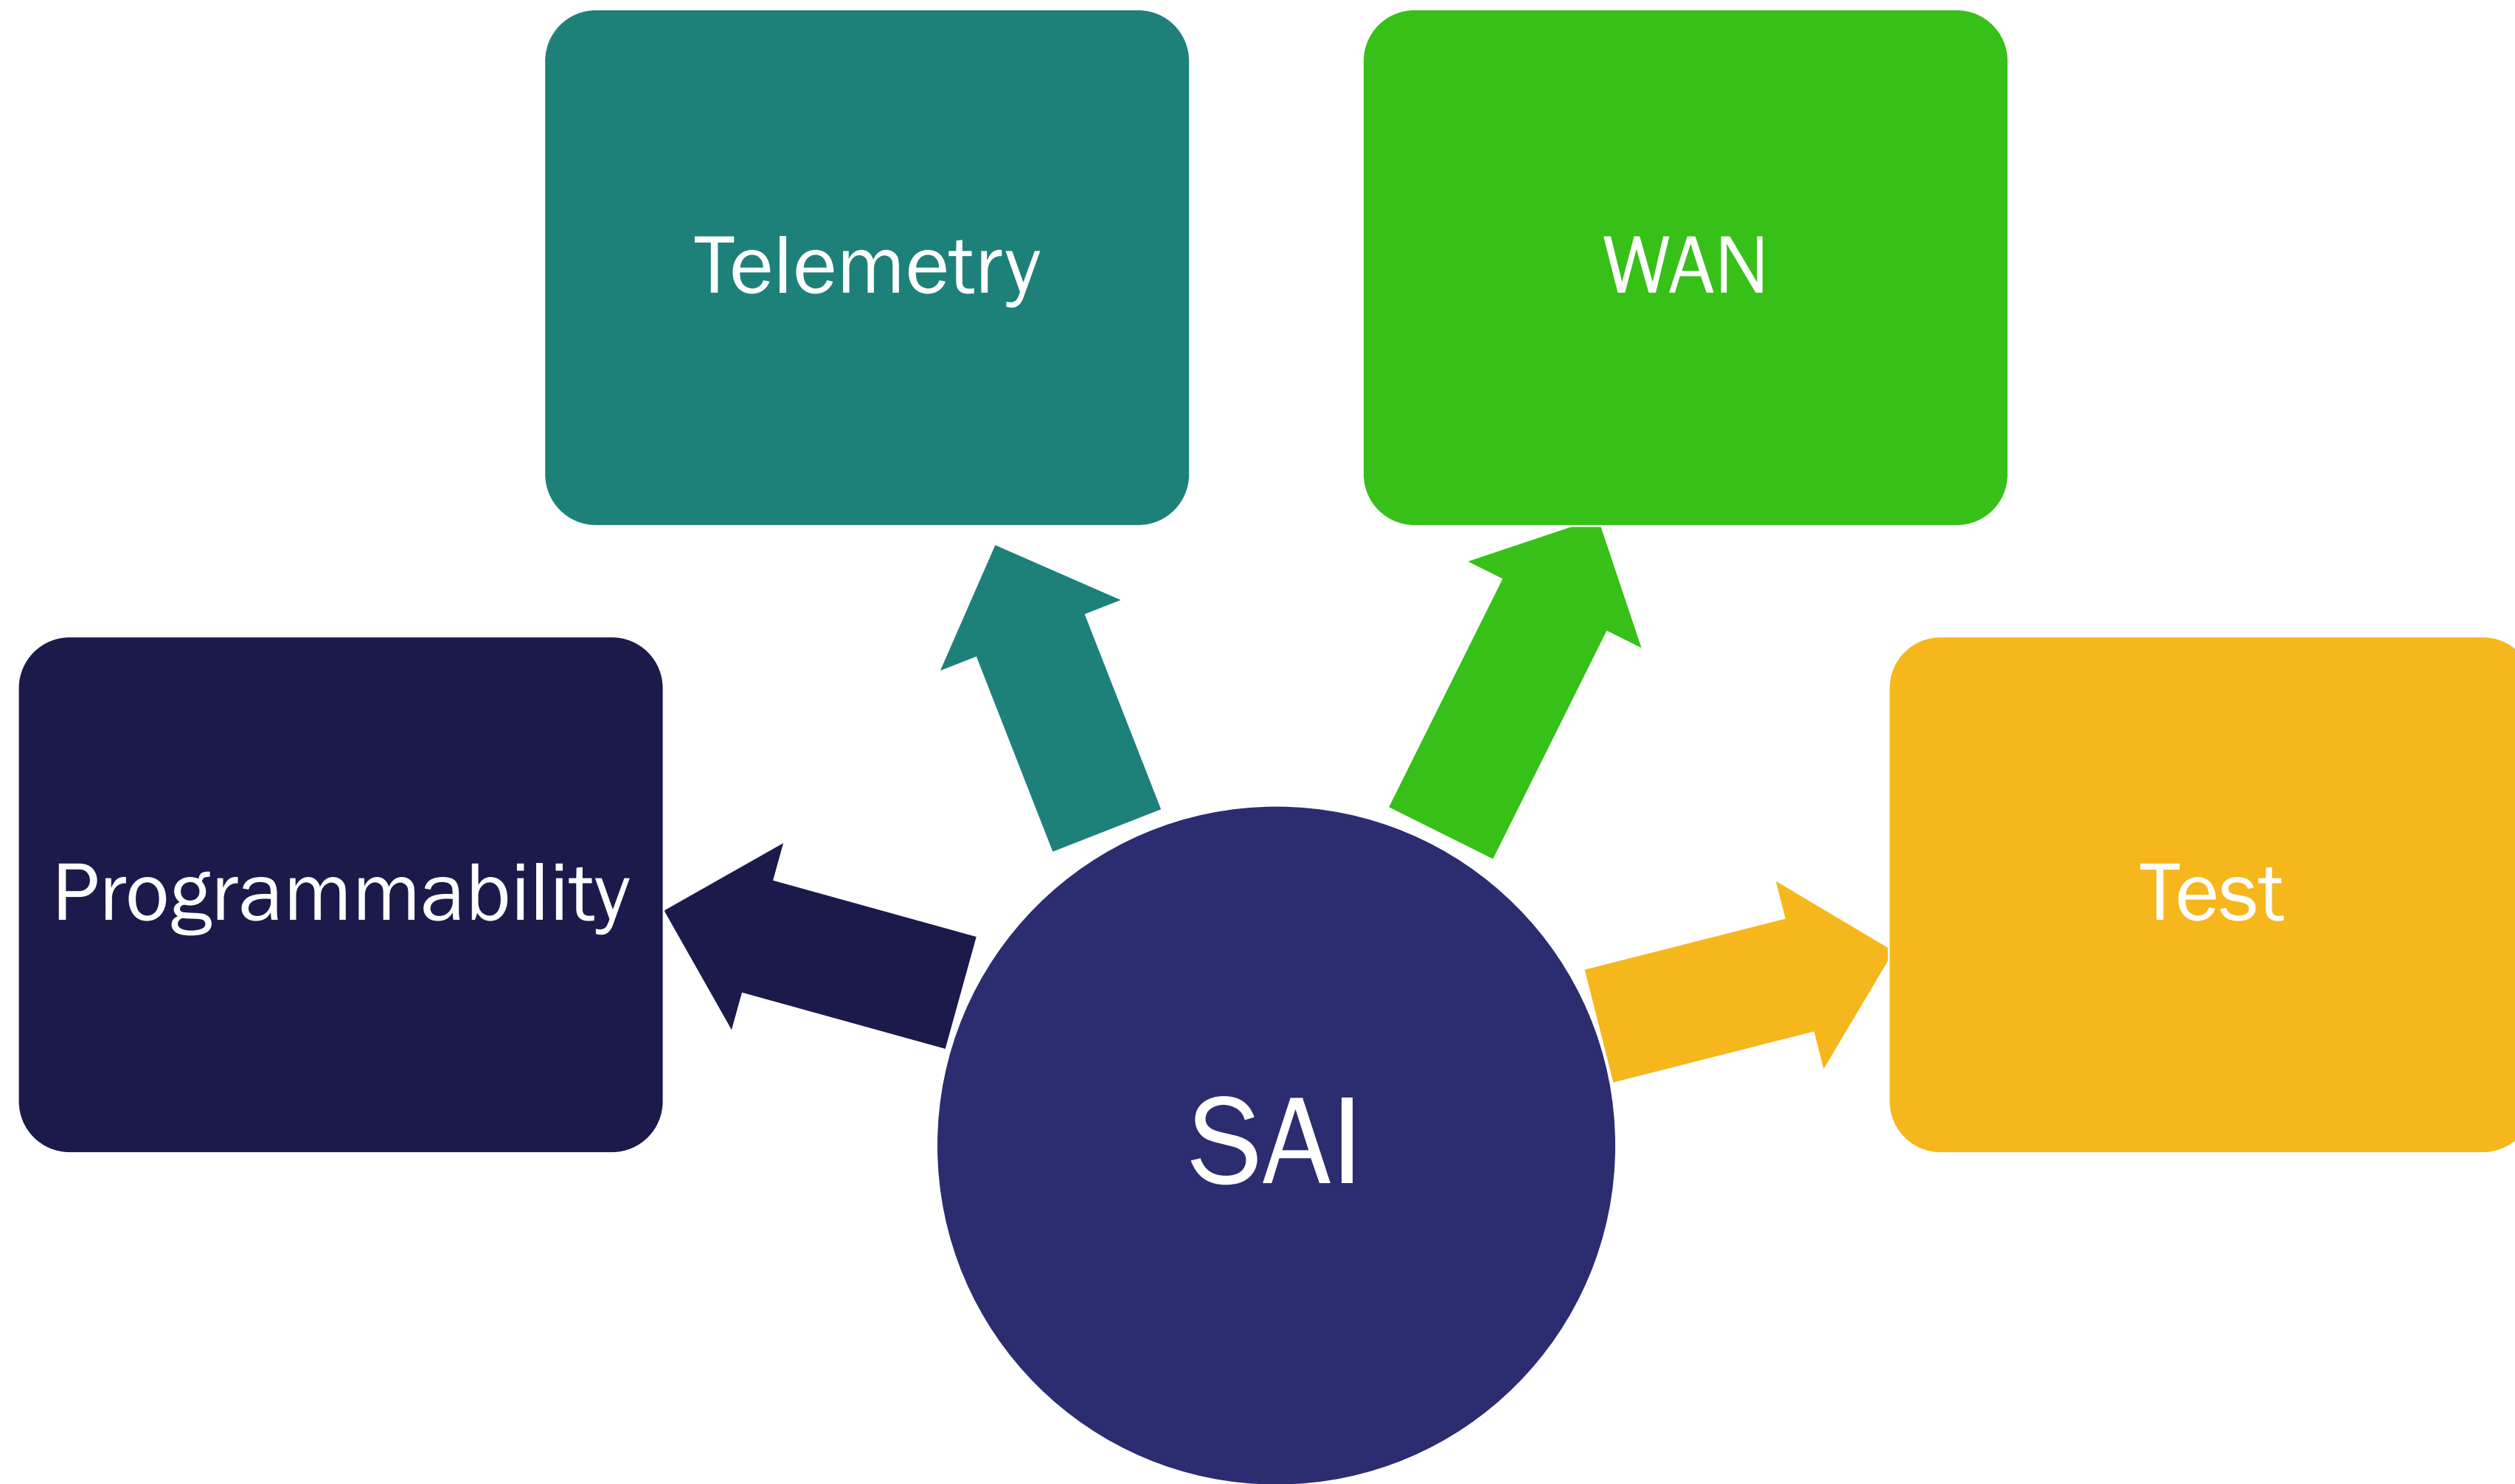
- Goal: Providing flexibility on top of existing SAI pipeline
- Approach
  - Provide SAI flex API generated by pipeline description language on top of existing SAI behavior model
- Benefits
  - Existing SAI APIs are still available
  - Users can customize the pipeline and use the generated SAI API

# SAI Flex API






# SAI Road Map 2018 and Beyond



# Open Invitation I – OCP Workshop on 3/22



MAR  
22

**OCP SONiC/SAI  
Engineering Workshop -  
hosted by LinkedIn**

by Open Compute Project Foundation

Free

[↑](#) [🔖](#) Sold Out [DETAILS](#)

#### DESCRIPTION

- March 22nd 8:00am - 5:00pm
- Dinner to Follow - when registering, please let us know if you can join us for dinner. Location TBD
- Breakfast/Lunch will be provided.

Dear SONiC, SAI community and Open Networking enthusiasts :  
We would love to invite you to join us on 3/22 to discuss the latest SONiC and SAI developments, roadmaps, and meet face-to-face with your fellow community members.

#### DATE AND TIME

Thu, March 22, 2018  
8:00 AM – 5:00 PM PDT  
[Add to Calendar](#)

#### LOCATION

ZGC Capital Corporation  
4500 Great America Parkway  
Room 1025~1027  
Santa Clara, CA 95054

# Open Invitation II

- Inviting contributions in all areas:
  - Bring up new proposals
  - Test and contribute test cases
  - Use it and report bugs
- Github <https://github.com/opencomputeproject/SAI>
- Mailing list [opencompute-sai@lists.opencompute.org](mailto:opencompute-sai@lists.opencompute.org)
- Meeting <https://attendee.gotowebinar.com/register/6659543304101781761>



# OCP SUMMIT