



SONiC: Software for Open Networking in the Cloud

Lihua Yuan
Microsoft Azure Network Team
for
the SONiC Community

SONiC Community @ OCP Mar 2018

Application & Management



CANONICAL



ARISTA



Tencent 腾讯

SONiC [Software For Open Networking in the Cloud]

Switch Platform



Inventec



Quanta Computer



Celestica™



SAI [Switch Abstraction Interface]

Silicon/ASIC



nephos

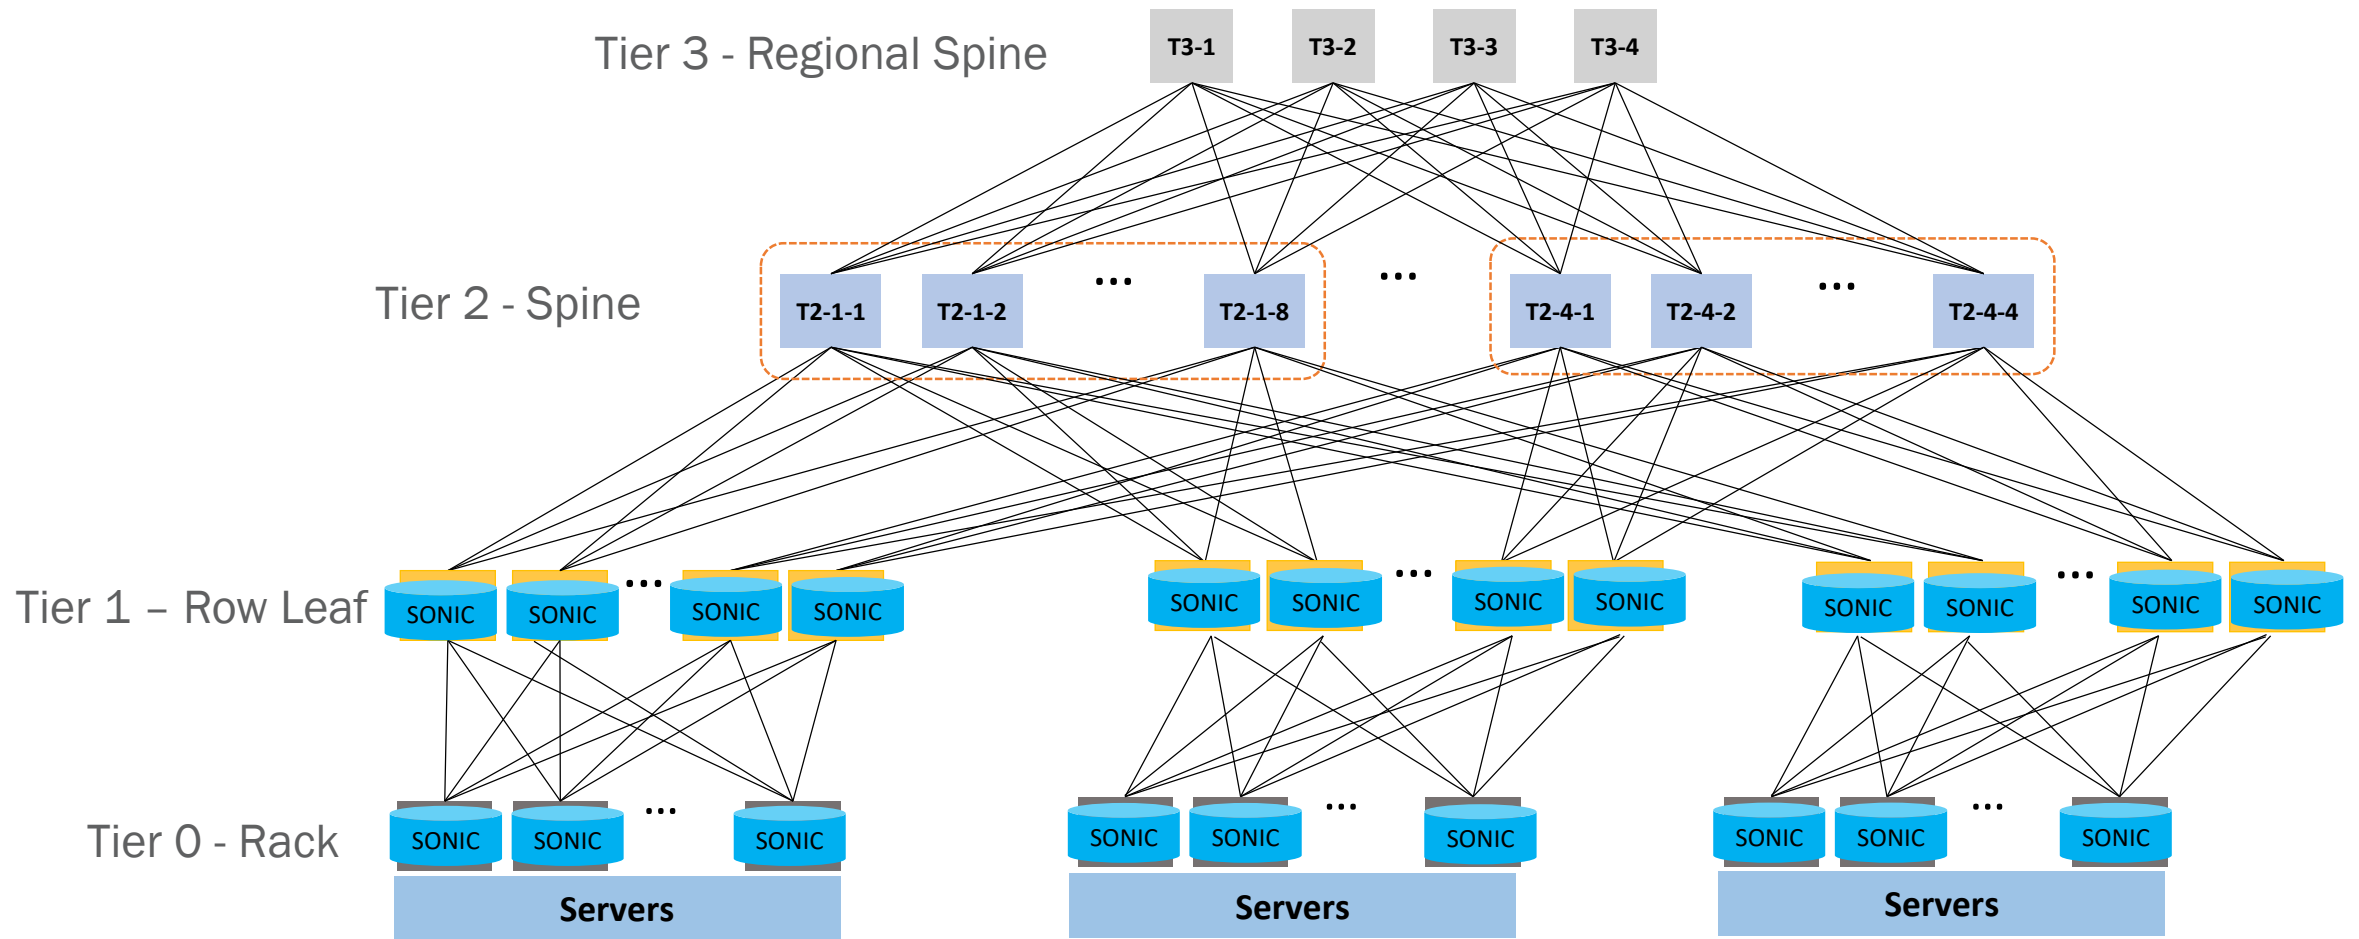


Innovium™

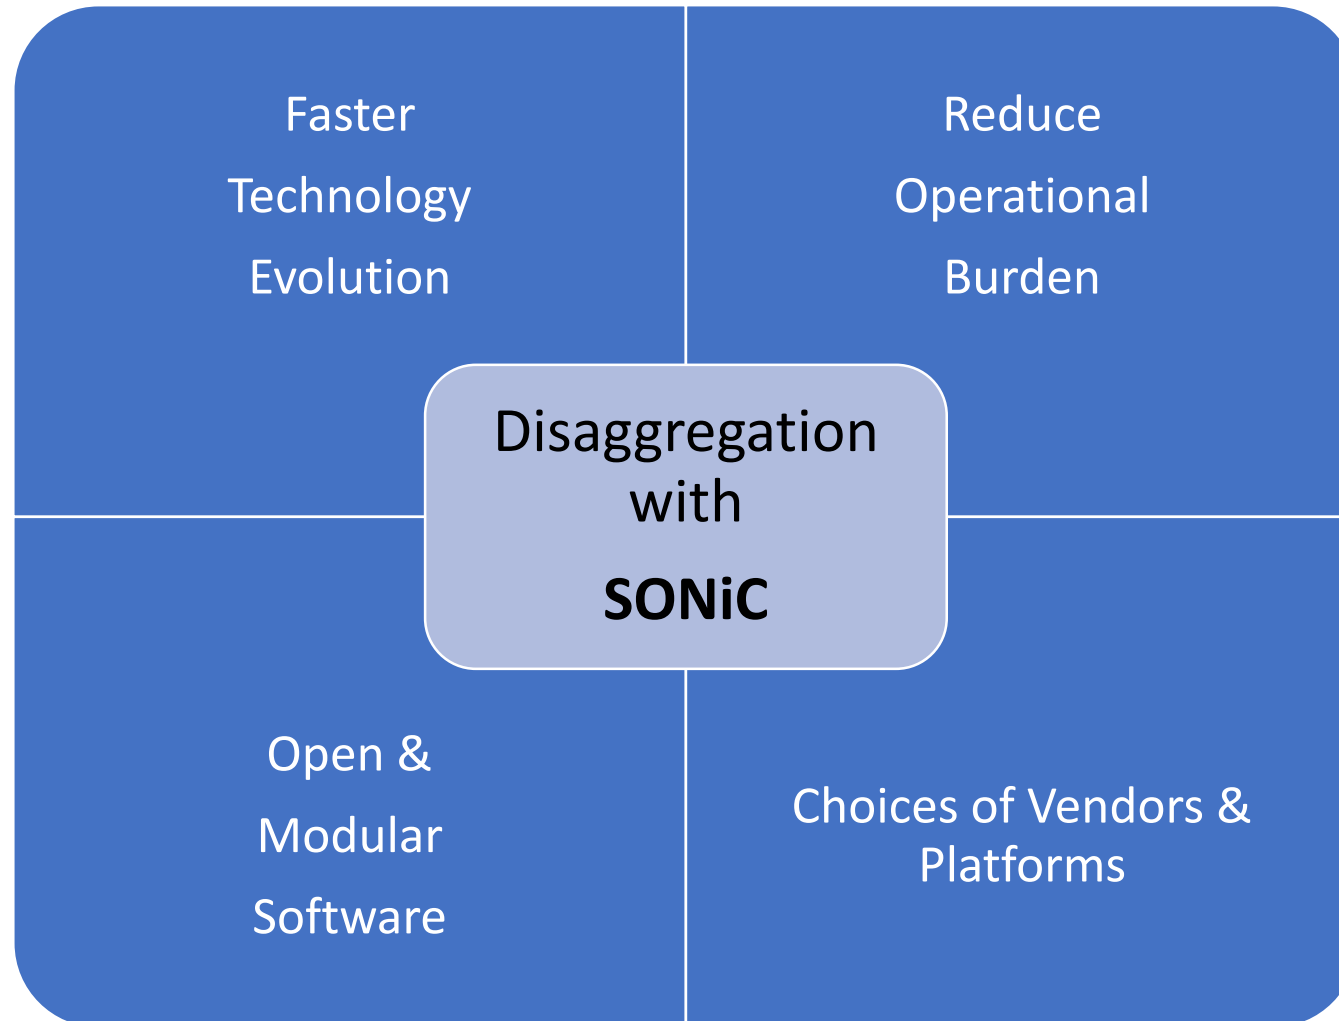
Microsoft Cloud Network - 15 Billion Dollar Bet



SONiC Is Powering Microsoft Cloud At Scale



Goals of SONiC



SONiC: Software for Open Networking in the Cloud

- Switch Abstraction Interface (SAI)
 - Cross-ASIC portability
- Modular Design with Switch State Service (SwSS)
 - Decoupling software components
 - Consistent application development model
- Containerization of SONiC
 - Serviceability
 - Cross-platform portability
- SONiC Operational Scenarios
 - Hitless upgrade
 - Network emulation with CrystalNet

Switch Abstraction Interface (SAI)

Network Applications

Hello

Switch Abstraction Interface

Simple, consistent, and stable network application stack

Helps consume the underlying complex, heterogeneous hardware easily and faster

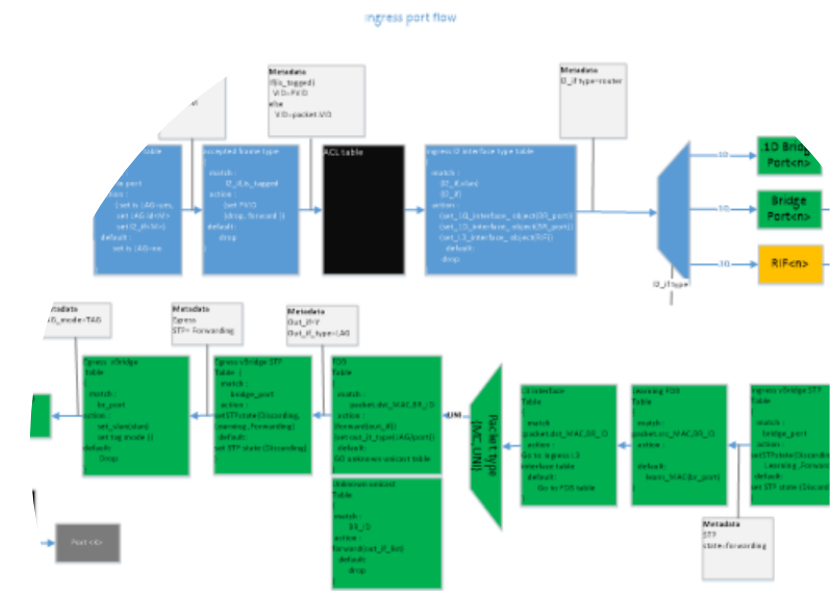


Switch Abstraction Interface

- CRUD operation over extensible Entity/Attribute/Value data model
- Reference data-plane behavior model supports various devices
- Significant feature/partner growth since announcement in 2015
- <https://github.com/opencomputeproject/SAI>

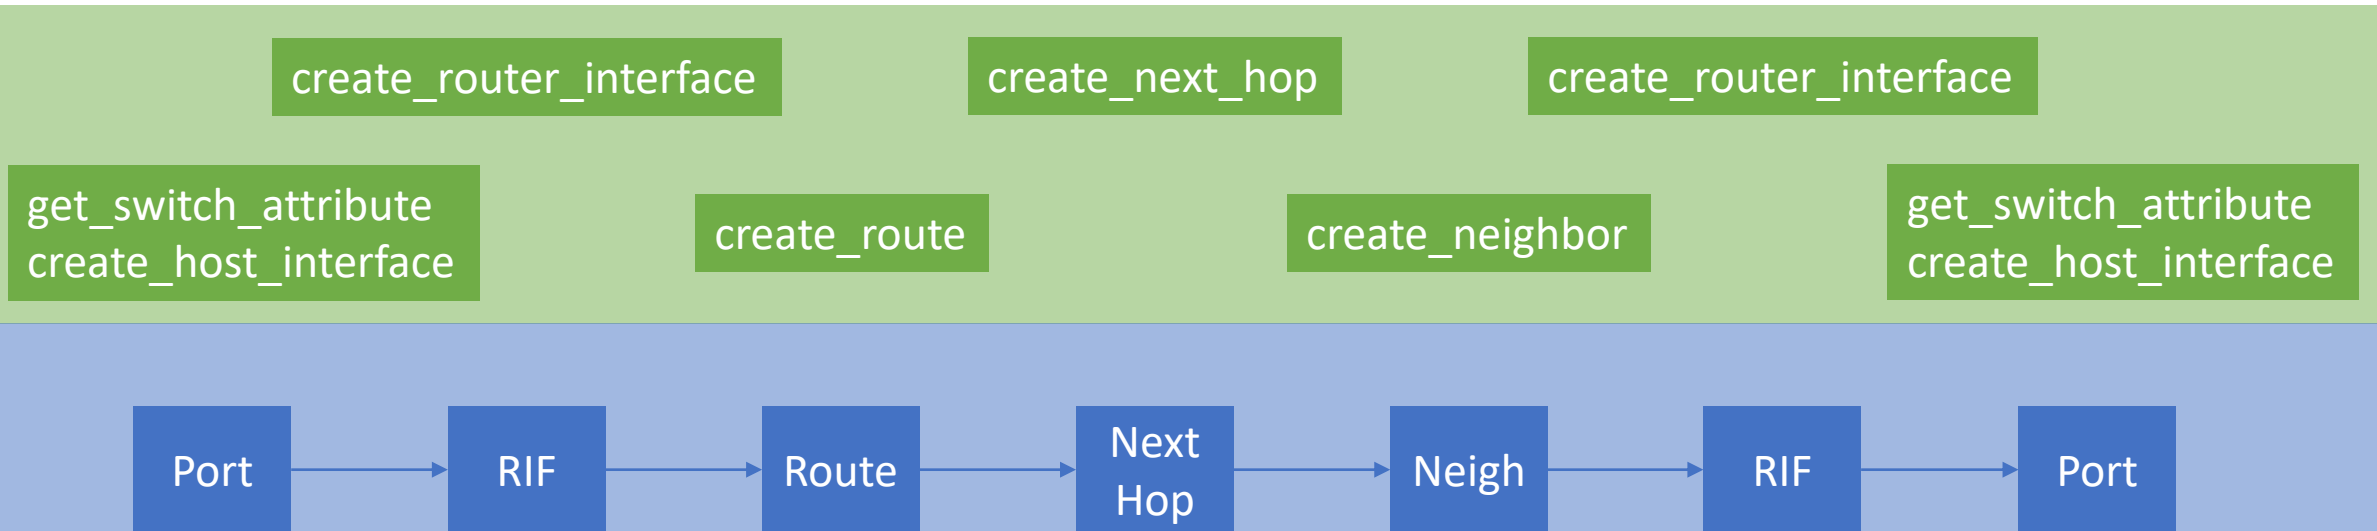
```
* @param[out] bridge_port_id
* @param[in] attr_count number of attributes
* @param[in] attr_list array of attributes
*
* @return SAI_STATUS_SUCCESS on success
*         Failure status code on error
*/
typedef sai_status_t (*sai_create_bridge_fn)(
    _Out_ sai_object_id_t* bridge_id,
    _In_ uint32_t attr_count,
    _In_ const sai_attribute_t *attr_list
);

remove bridge
```



SAI Function Calls & Pipelines in SONiC

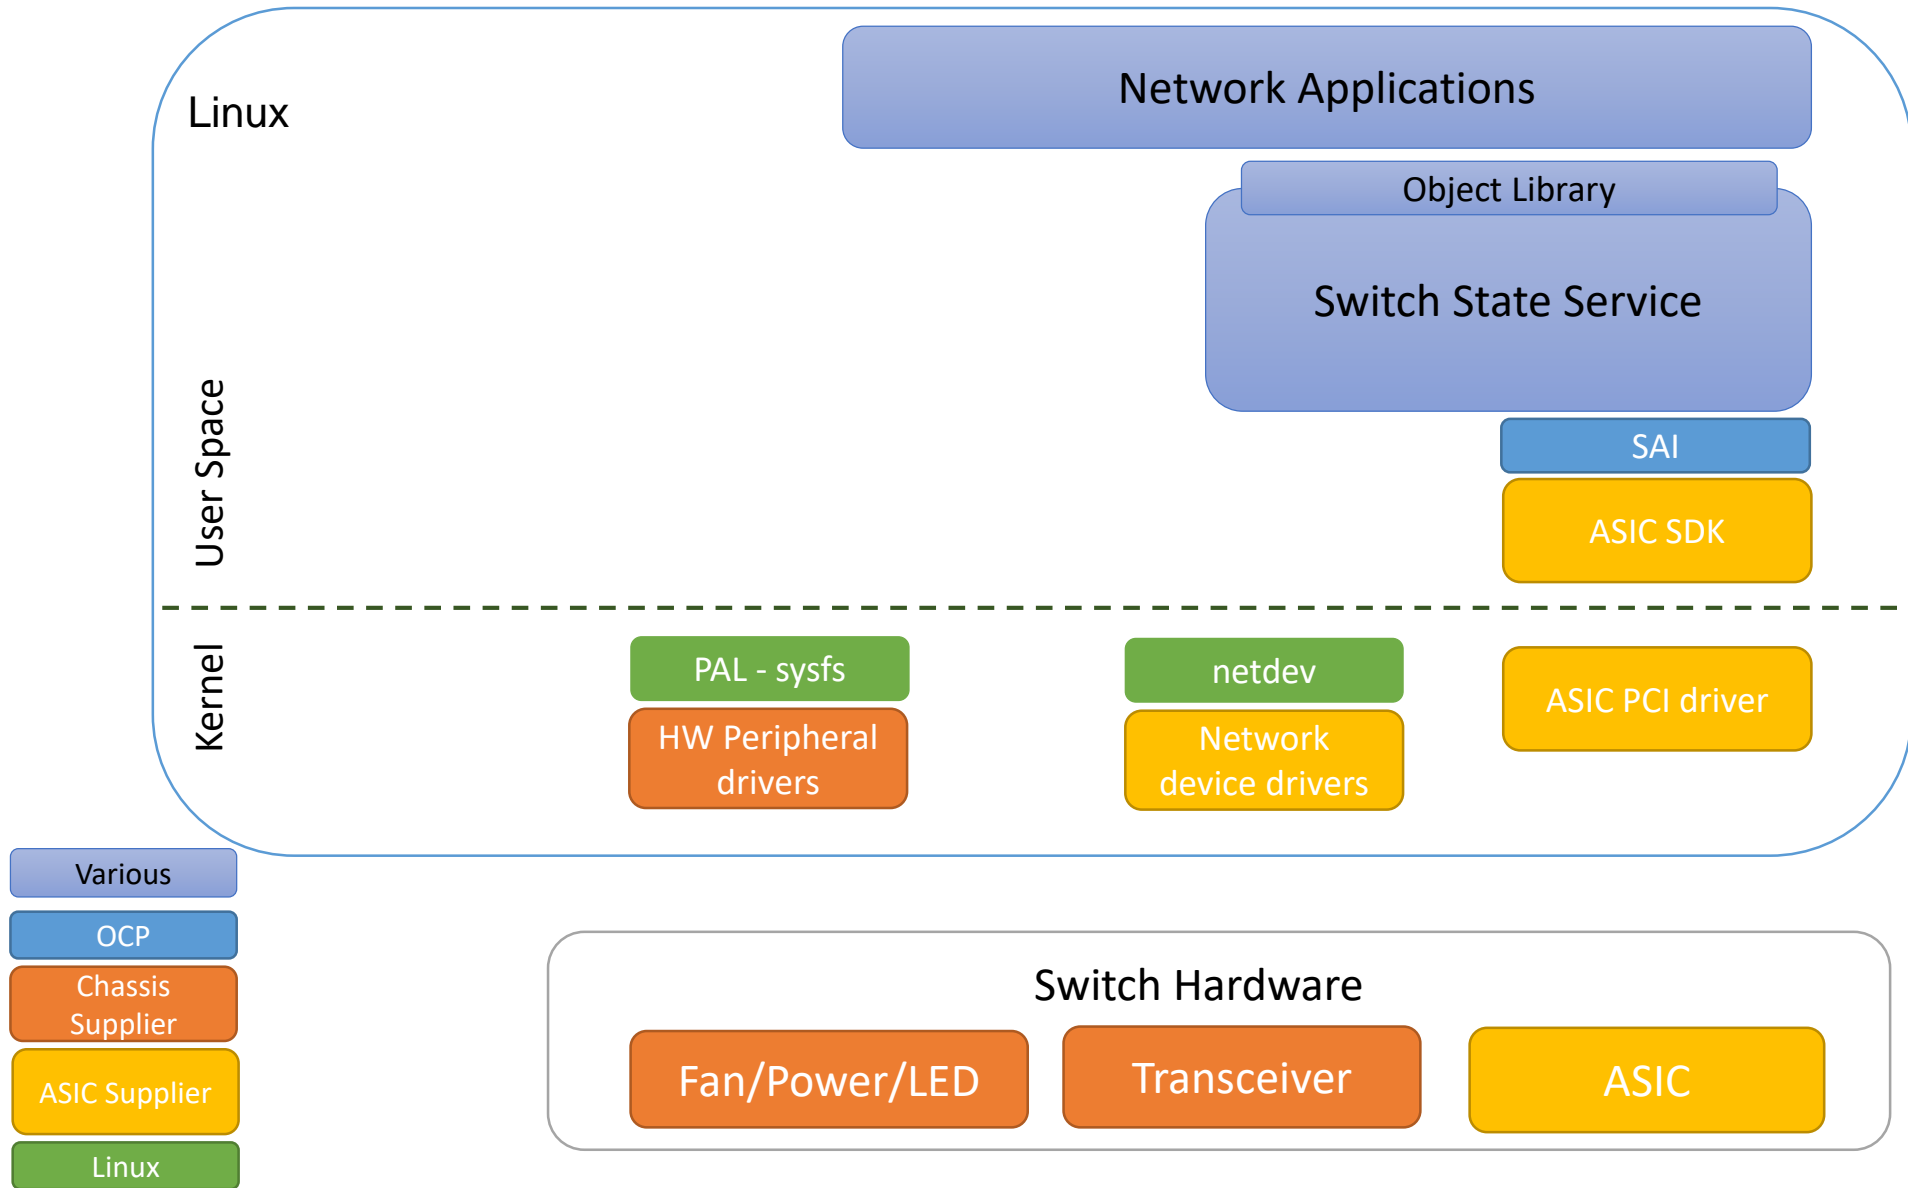
- Register SAI APIs to achieve specific functions and control logging separately
- Construct the right order to execute SAI function calls
- Refer to SAI Pipelines to create the right packet action flow



SONiC: Software for Open Networking in the Cloud

- Switch Abstraction Interface (SAI)
 - Cross-ASIC portability
- Modular Design with Switch State Service (SwSS)
 - Decoupling software components
 - Consistent application development model
- Containerization of SONiC
 - Serviceability
 - Cross-platform portability
- SONiC Operational Scenarios
 - Hitless upgrade
 - Network emulation with CrystalNet

SONiC High Level Architecture



Switch State Service (SwSS)

SAI DB: persist SAI objects

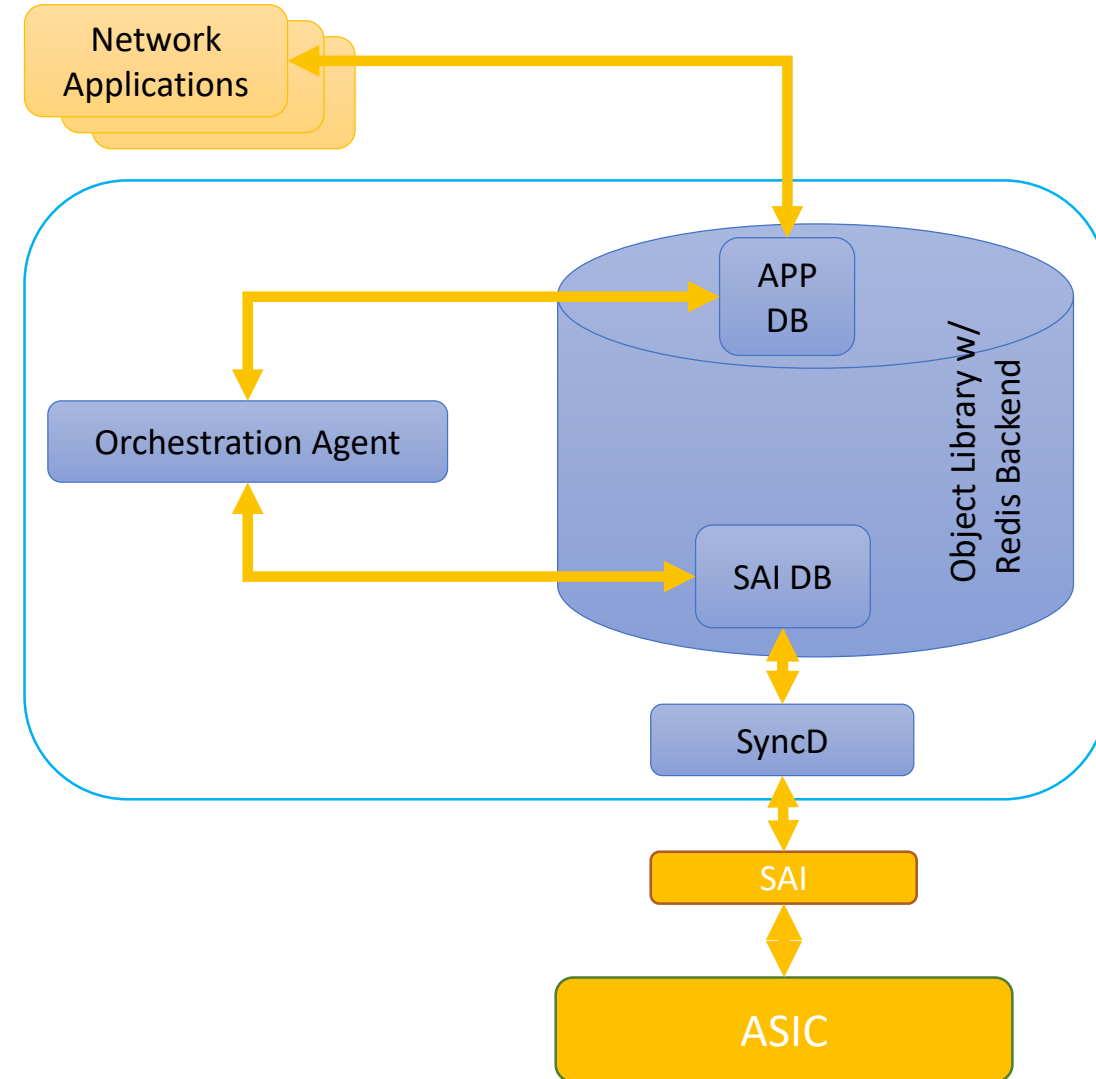
App DB: persist App objects

DB backend: redis with object library

SyncD: sync SAI objects between software and hardware

Orchestration Agent: translation between apps and SAI objects, resolution of dependency and conflict

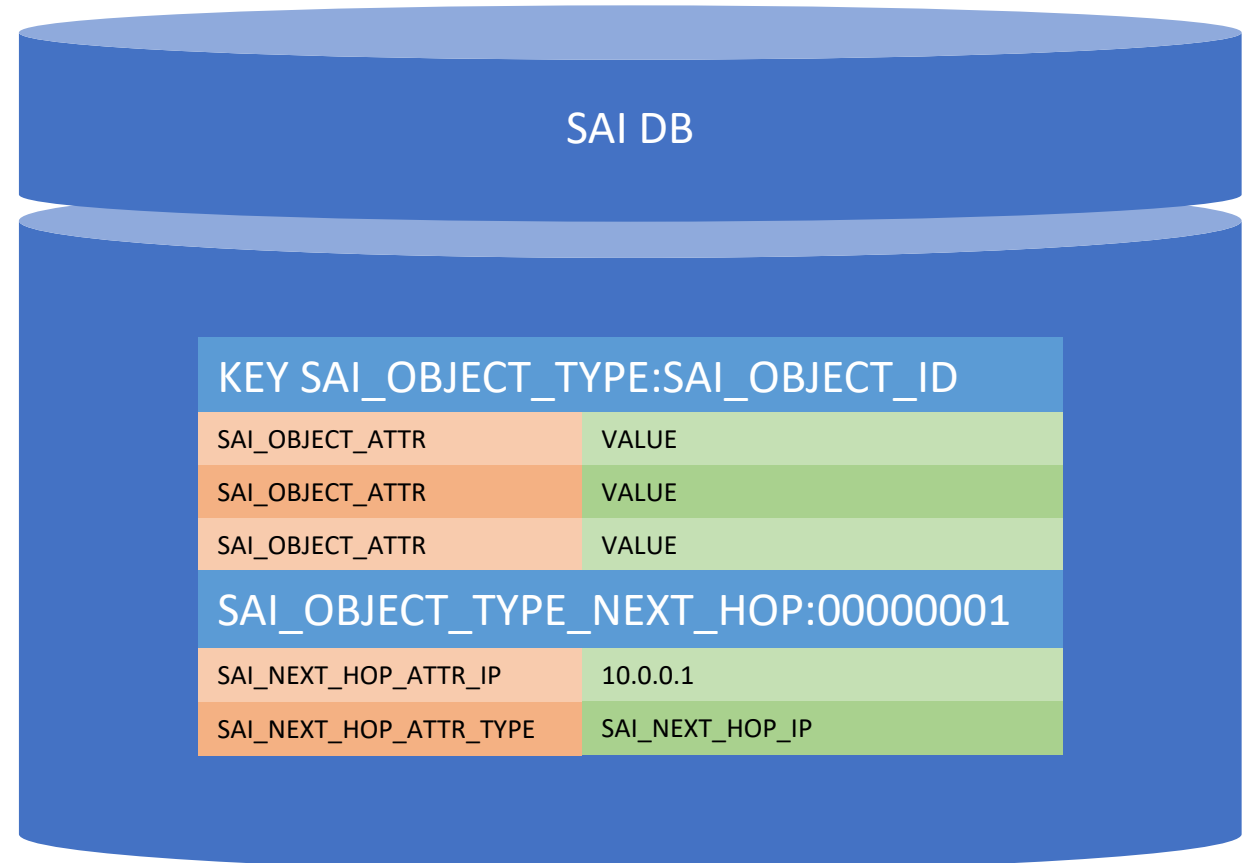
Key Goal: Evolve components independently



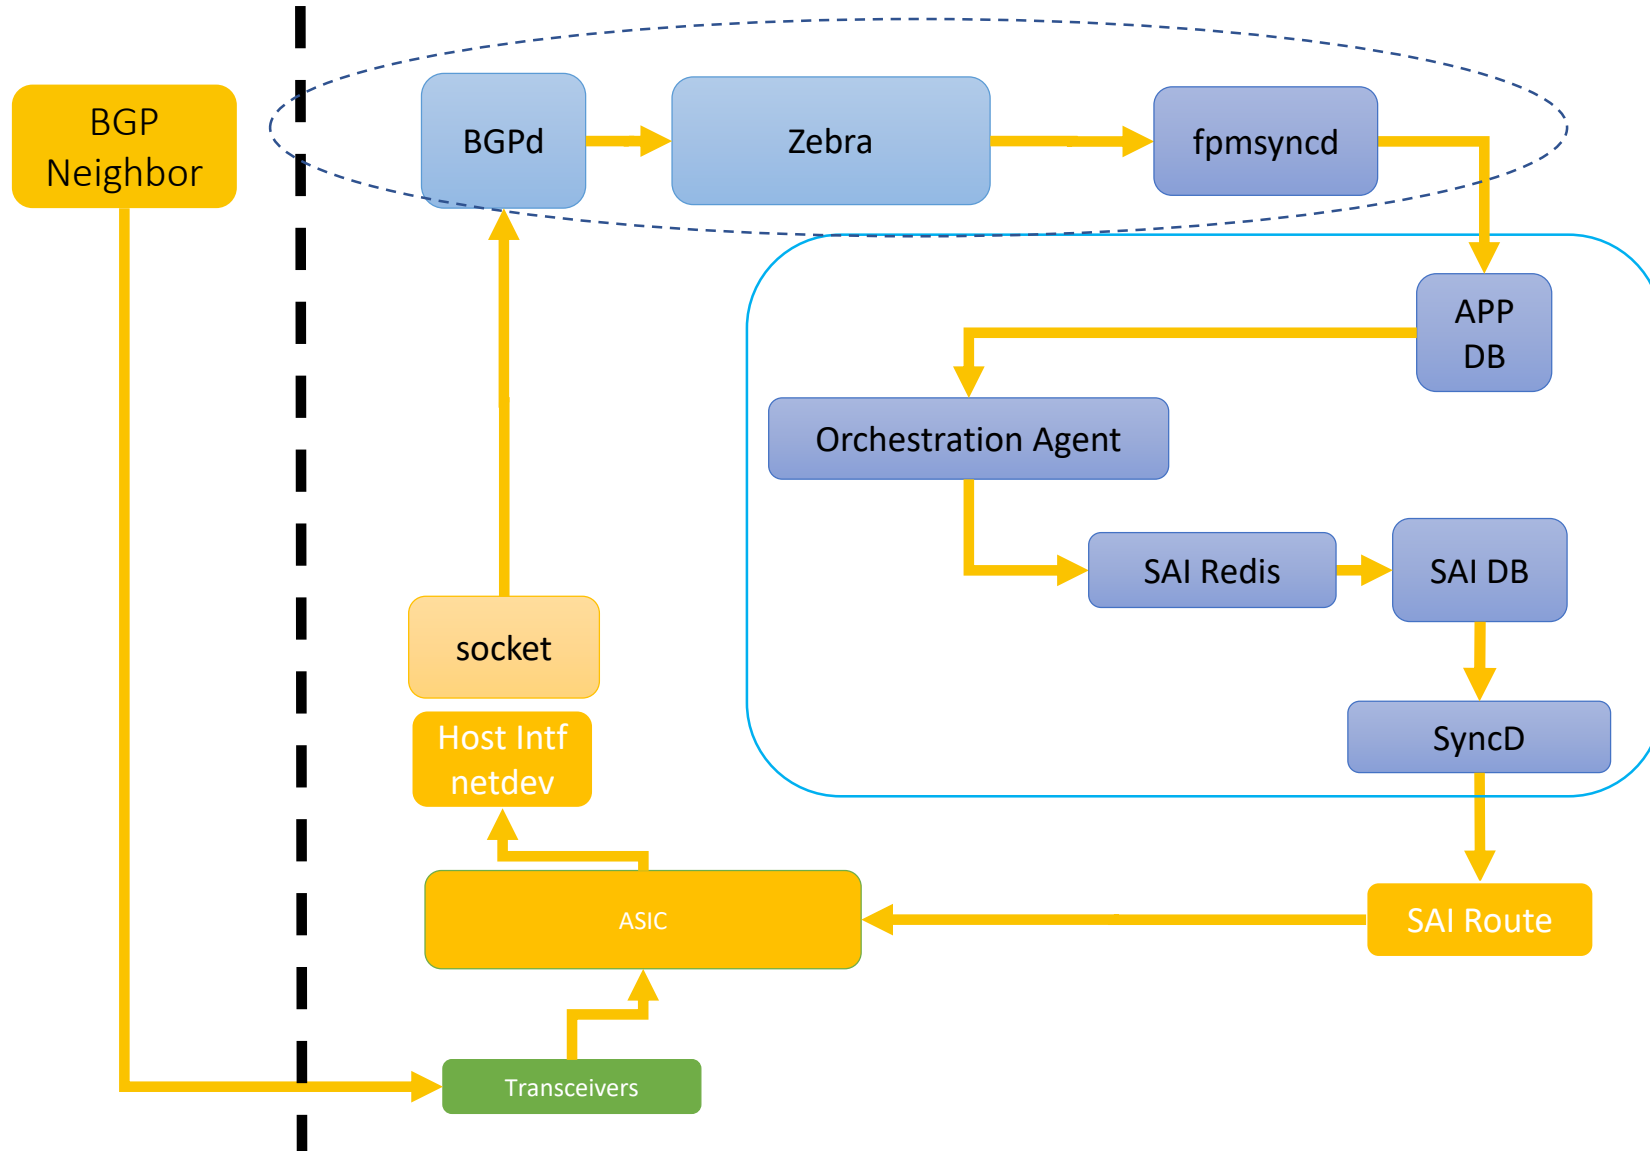
How SAI Objects are Stored in SAI Database

- APP DB:
 - Application oriented
 - Human readable
- SAI DB:
 - SAI oriented
 - Query-able and machine parse-able
 - Snapshot of the current SAI state
- APIs:
 - sai_create_api
 - sai_set_api
 - sai_get_api*
 - sai_remove_api

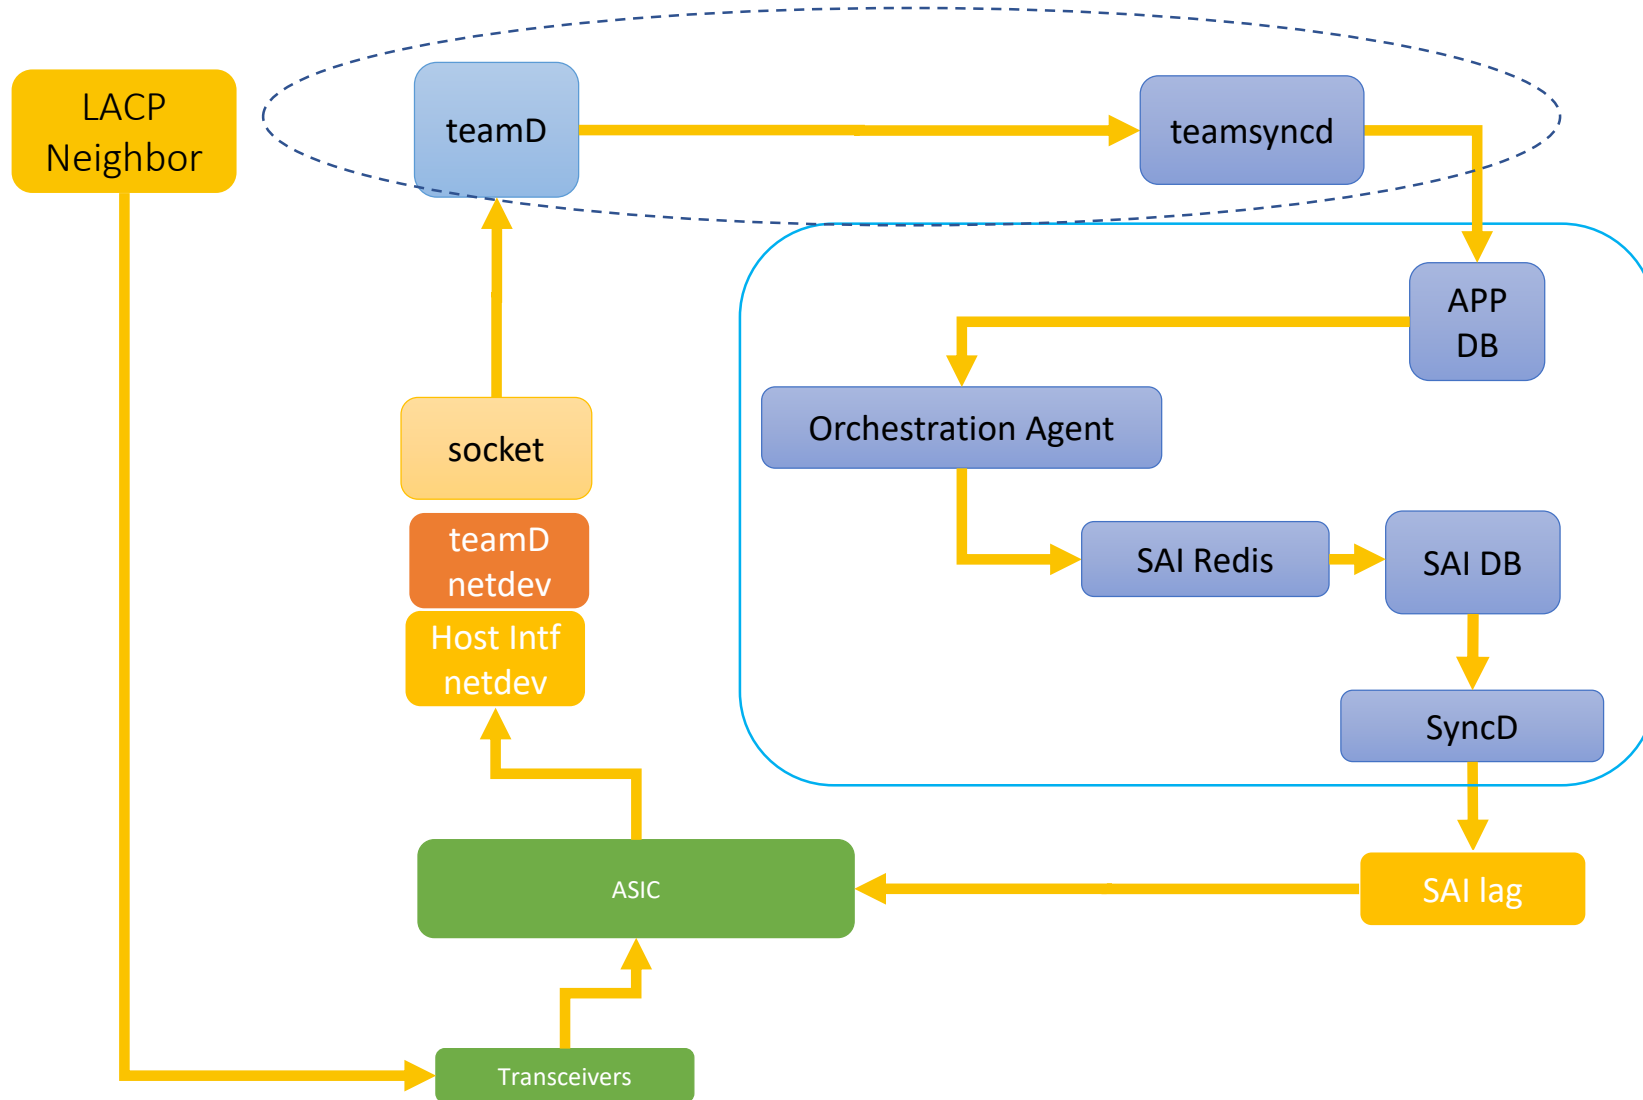
*Used only when querying ASIC settings and configurations



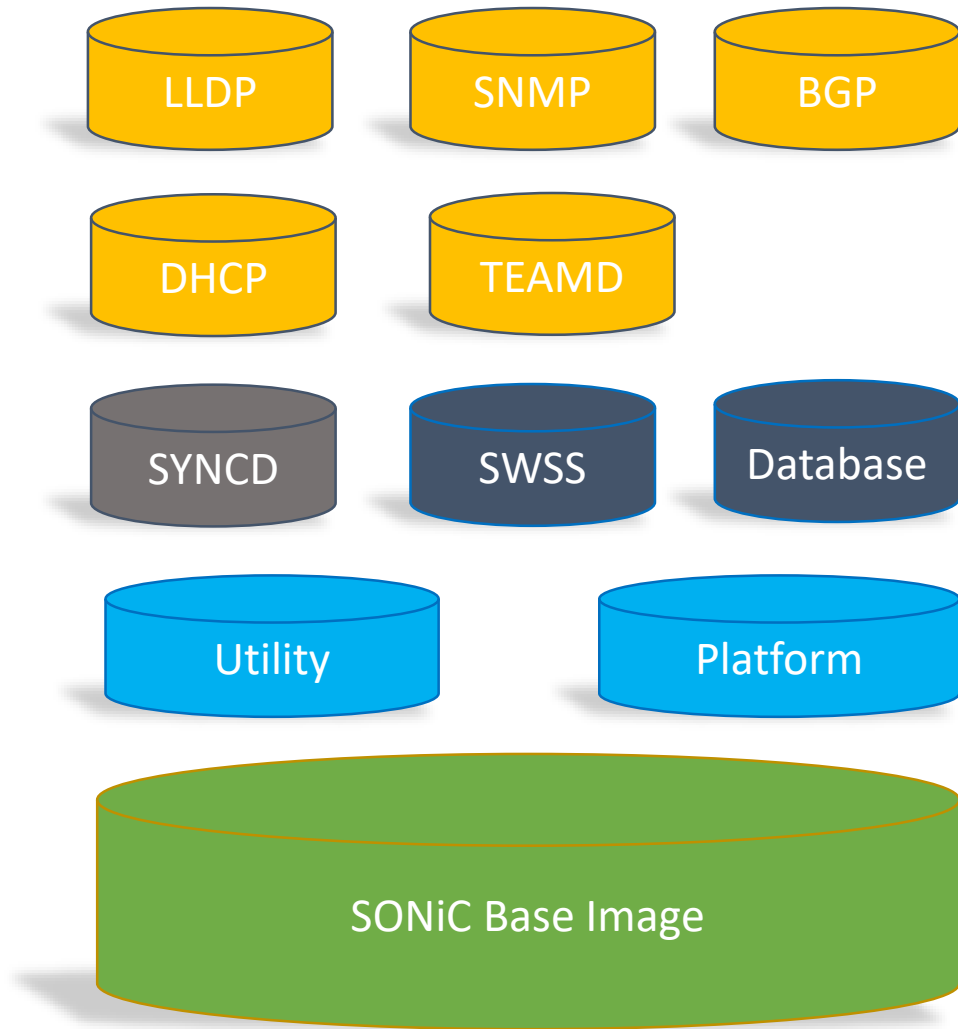
How Routing Works in SONiC



How LAG Works in SONiC



SONiC Containerization



- [TEAMD](#)
- [LLDP](#)
- BGP: Quagga or FRR
- SNMP: Net-SNMP + SNMP subagent
- DHCP Relay: isc dhcp
- Platform: sensors
- SWSS: switch state service
- DB: Redis
- Syncd: sairedis + syncd agent

SONiC Containerization

Container Strength

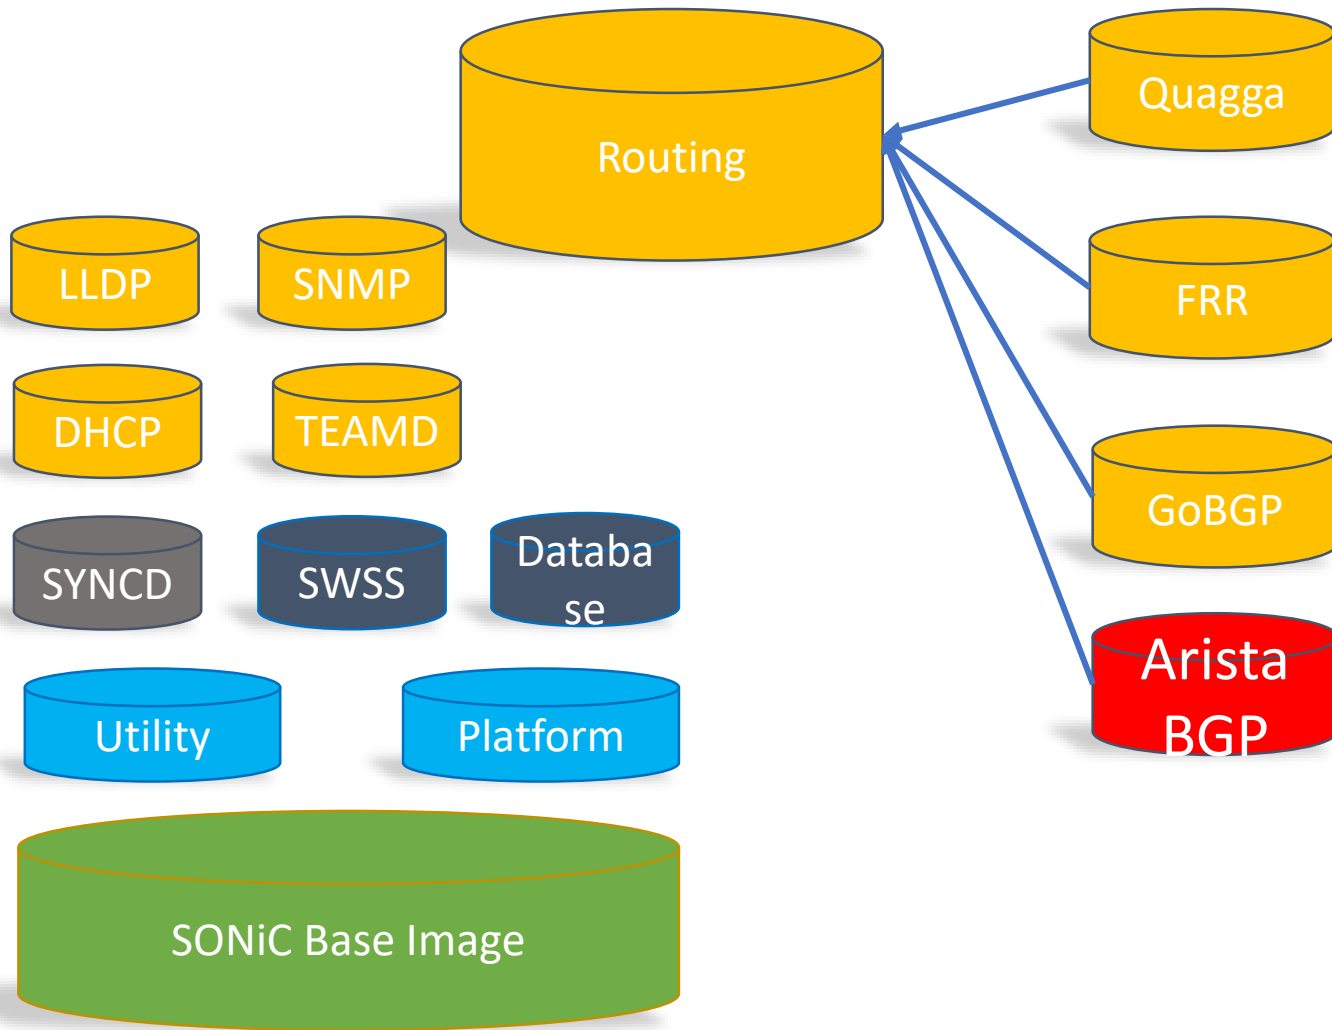
- Clean isolation
- Easy deployment
- Transactional
- Run universal



SONiC Benefits

- Serviceability
- Extensibility
- Development agility
- Cross-platform

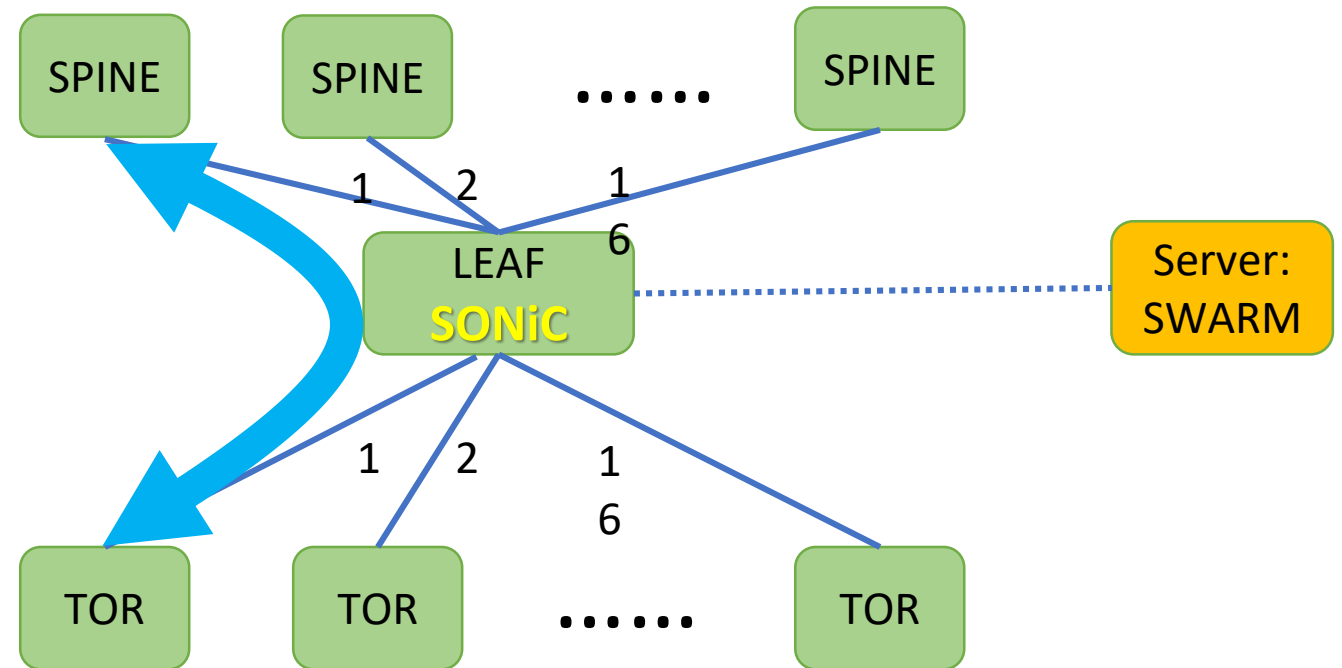
SONiC Containerization



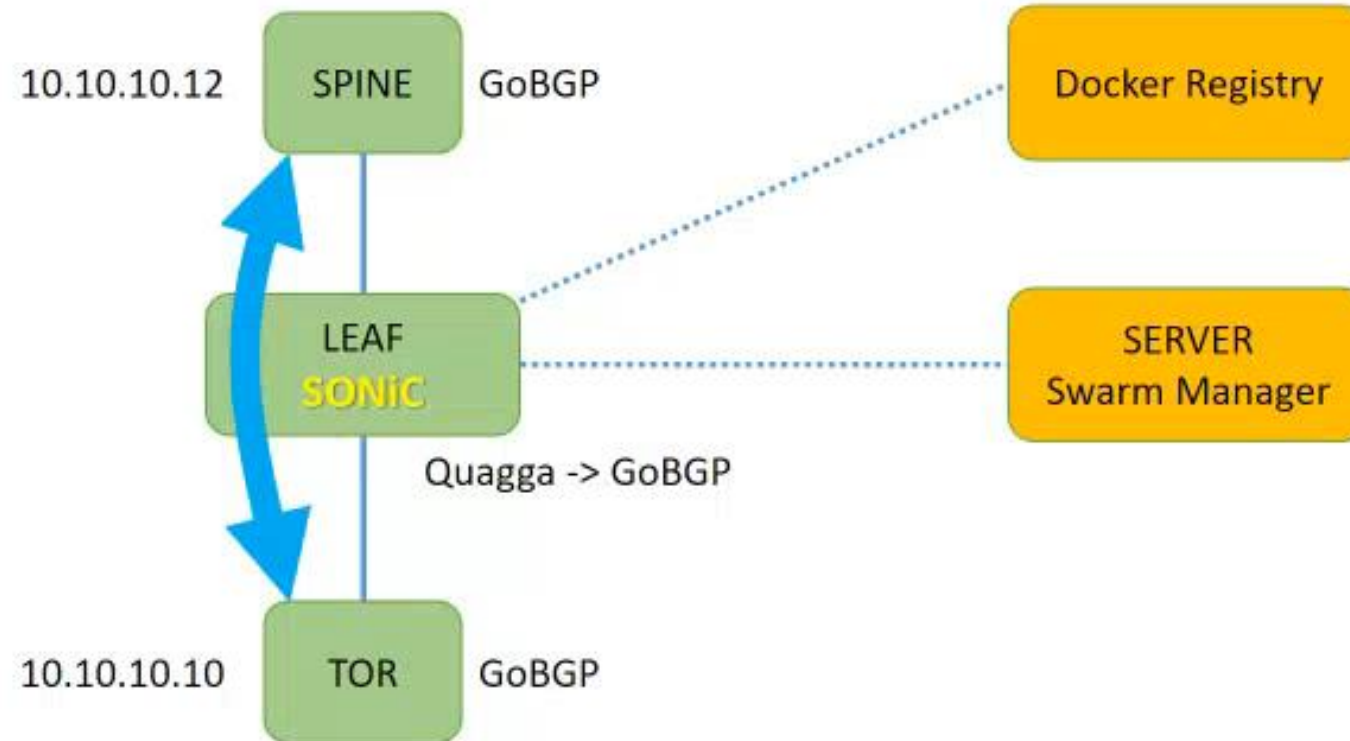
- Components developed in different environments
- Source code may not be available
- Enables choices on a per-component basis

Hitless Upgrade Through Docker

- Cloud users demand zero-down time, but
- Security patches, bug fixes need to be rolled out in hours
- New features need to be rolled out in days





Hitless BGP Upgrade Through Docker Swarm





Before they fly the space shuttle

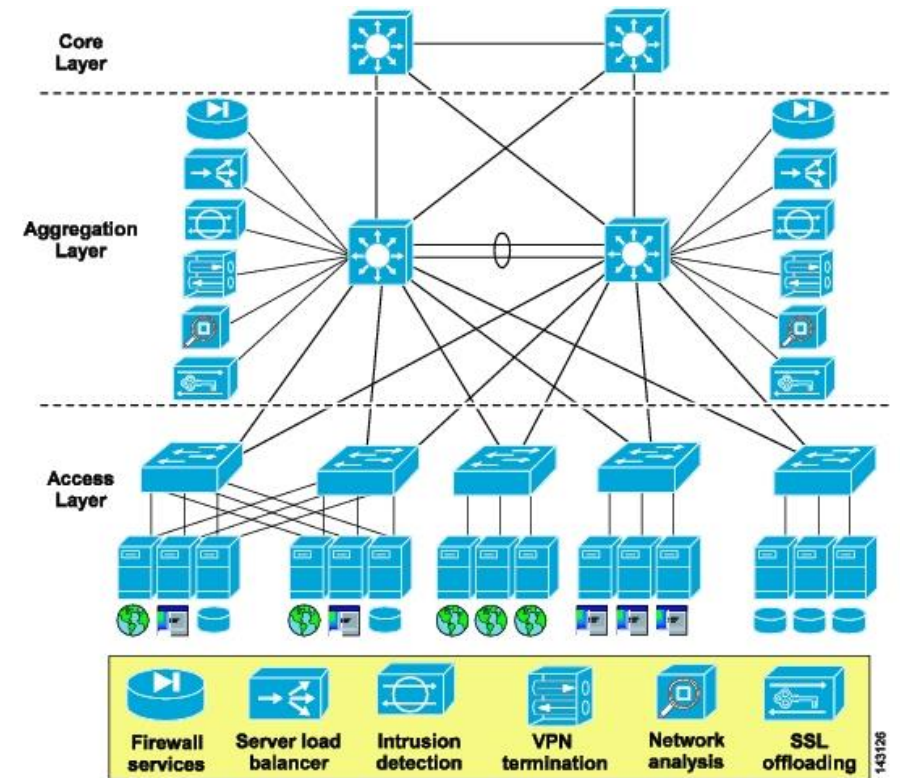



They practice all activities and test all equipment extensively in a high-fidelity simulators

CrystalNet: A high-fidelity, cloud-scale network emulator



SONiC Container
Soft Switch



Open Invitation

- Inviting contributions in all areas
 - **New ideas on white/open network devices**
 - New features, applications and tools
 - Download it, Test, Deploy!

Website: <https://azure.github.io/SONiC/>

Mailing list: sonicproject@googlegroups.com

GitHub: <https://github.com/Azure/SONiC>

Wiki: <https://github.com/Azure/SONiC/wiki/>

