OCTOBER 3

P4

WORKSHOP 2024

Centralized Telemetry and Security Enforcement Using SONiC and P4

Shekher Bulusu , Ravi Kumar GV Subrahmanya,
James Choi , Pawan Ravi, Don Newton, Ofer Gill

Agenda:

1. GEICO Network Components
2. GEICO Security Enforcement
3. GEICO Smart Network Remediation
4. SONIC-PINS Components
5. PINS P4Orch & SAI Dependency
6. PINS Driver Dependency
7. GEICO Next-step

GEICO Network Components

- **SONiC switch fabric:** The Data Center network infrastructure built using community SONiC software.
- Geico Software Defined Networking(SDN) is a platform which includes closed-loop network operations - Network Switch inventory database, Observability Assurance, Operations and Automation.
- The Network Switch inventory is a centralized database for all the intended network switches and cabling. It establishes a baseline for comparison against the observed state.
- Observability Assurance collects the observed state(operational and configuration) from the SONiC switch fabric. This validates the delta between intended and observed states, ensuring that the SONiC Switch fabric operates as intended.
- Operations module executes workflows that modify and update the intended states to the Switch inventory database.
- Automation modules pushes the desired state to the SONiC switch fabric using P4 Runtime(P4RT)/gNMI.

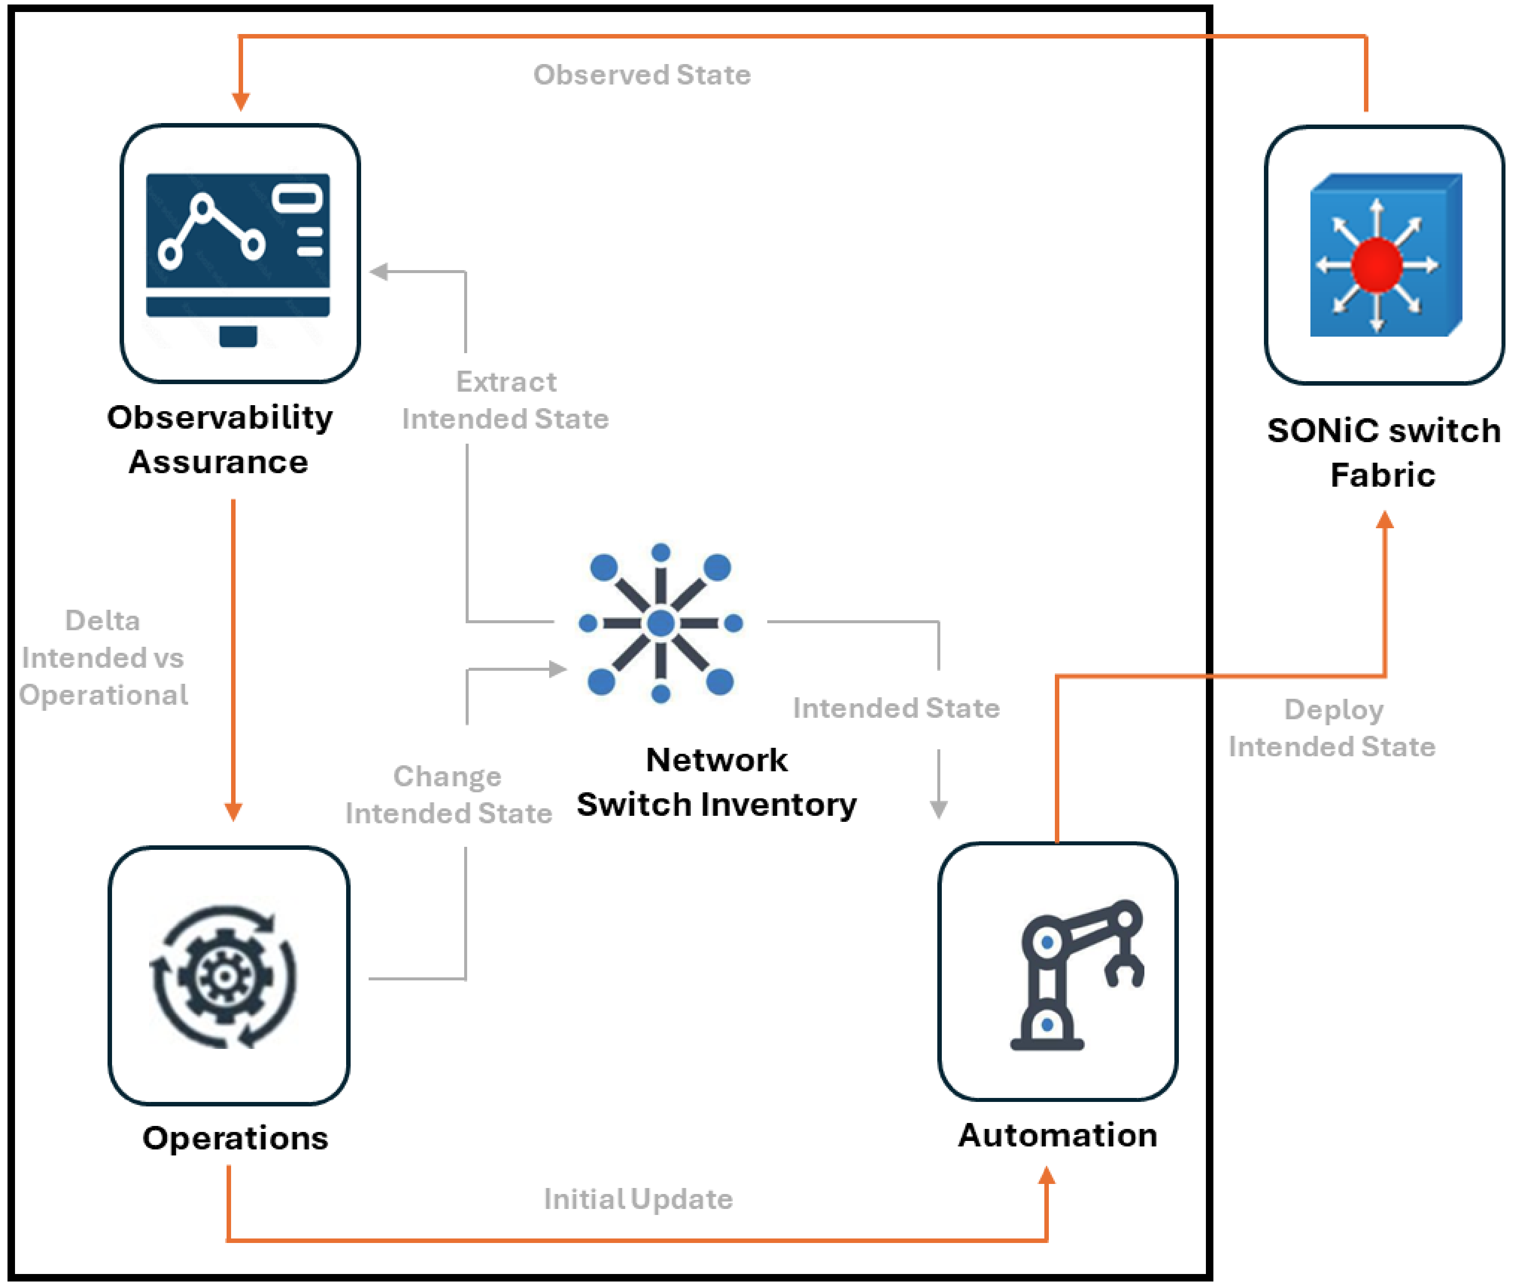

GEICO Software Defined Networking (SDN)

Sensitivity: Company-Internal

GEICO Security Enforcement

- Observability platform monitors network traffic using IPFIX to detect security violations.
- Example:
  - TCP SYN flood is a denial-of-service(DDoS) attack that sends massive numbers of SYN requests to a server to overwhelm it with open connections.
  - Observability detects TCP SYN flood by the presence of large number of half open connections.
  - Automation workflow remediates the anomaly by programming ACL rules via P4RT to match the source IP of the client and set DROP action.

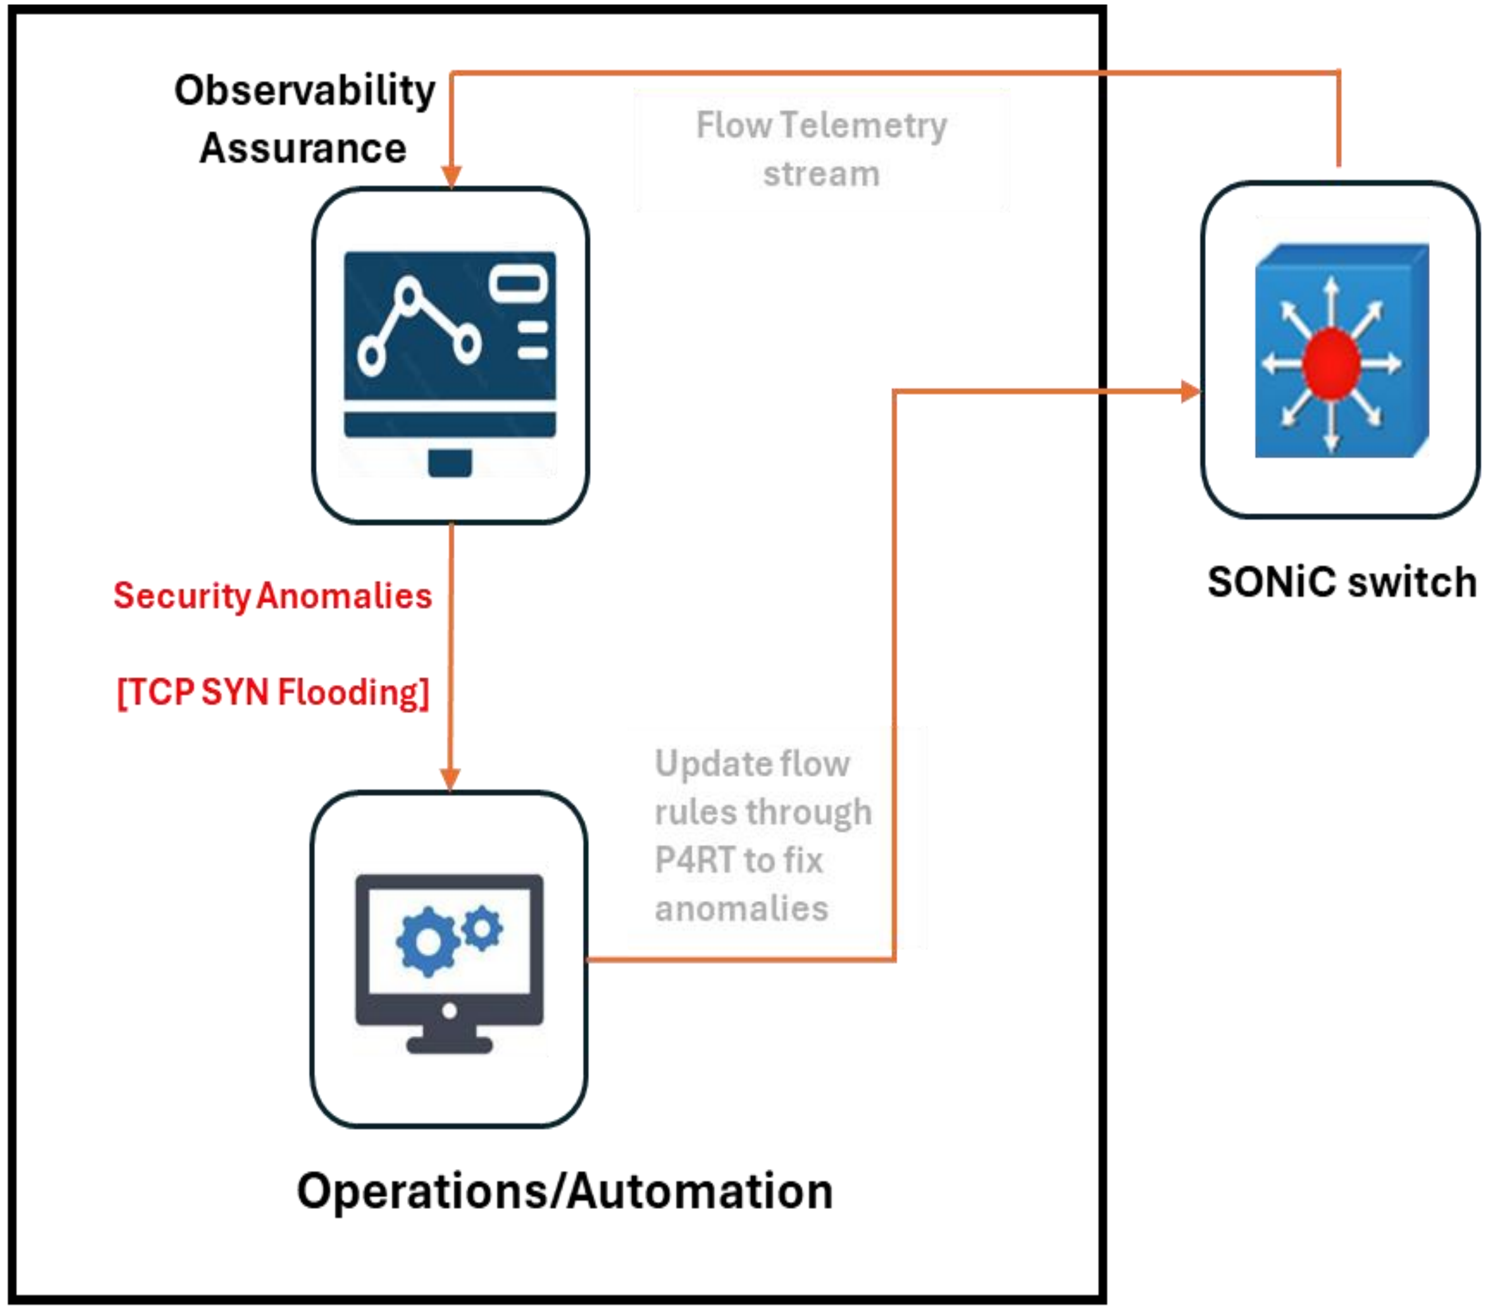

GEICO Software Defined Networking (SDN)

GEICO Smart Network Remediation

- Observability continuously monitors device health stats, counters, Route table and switch state using gNMI from SONiC switch.

- Example:

    - Missing route(s) during a link flap, is a scenario where the route tables see a loss of routes during link flapping.

    - Observability detects missing routes from the continuous updates of route table through gNMI.

    - Automation workflow remediates the anomaly by programming L3 flow rules via P4RT to program the missing routes during link flap.

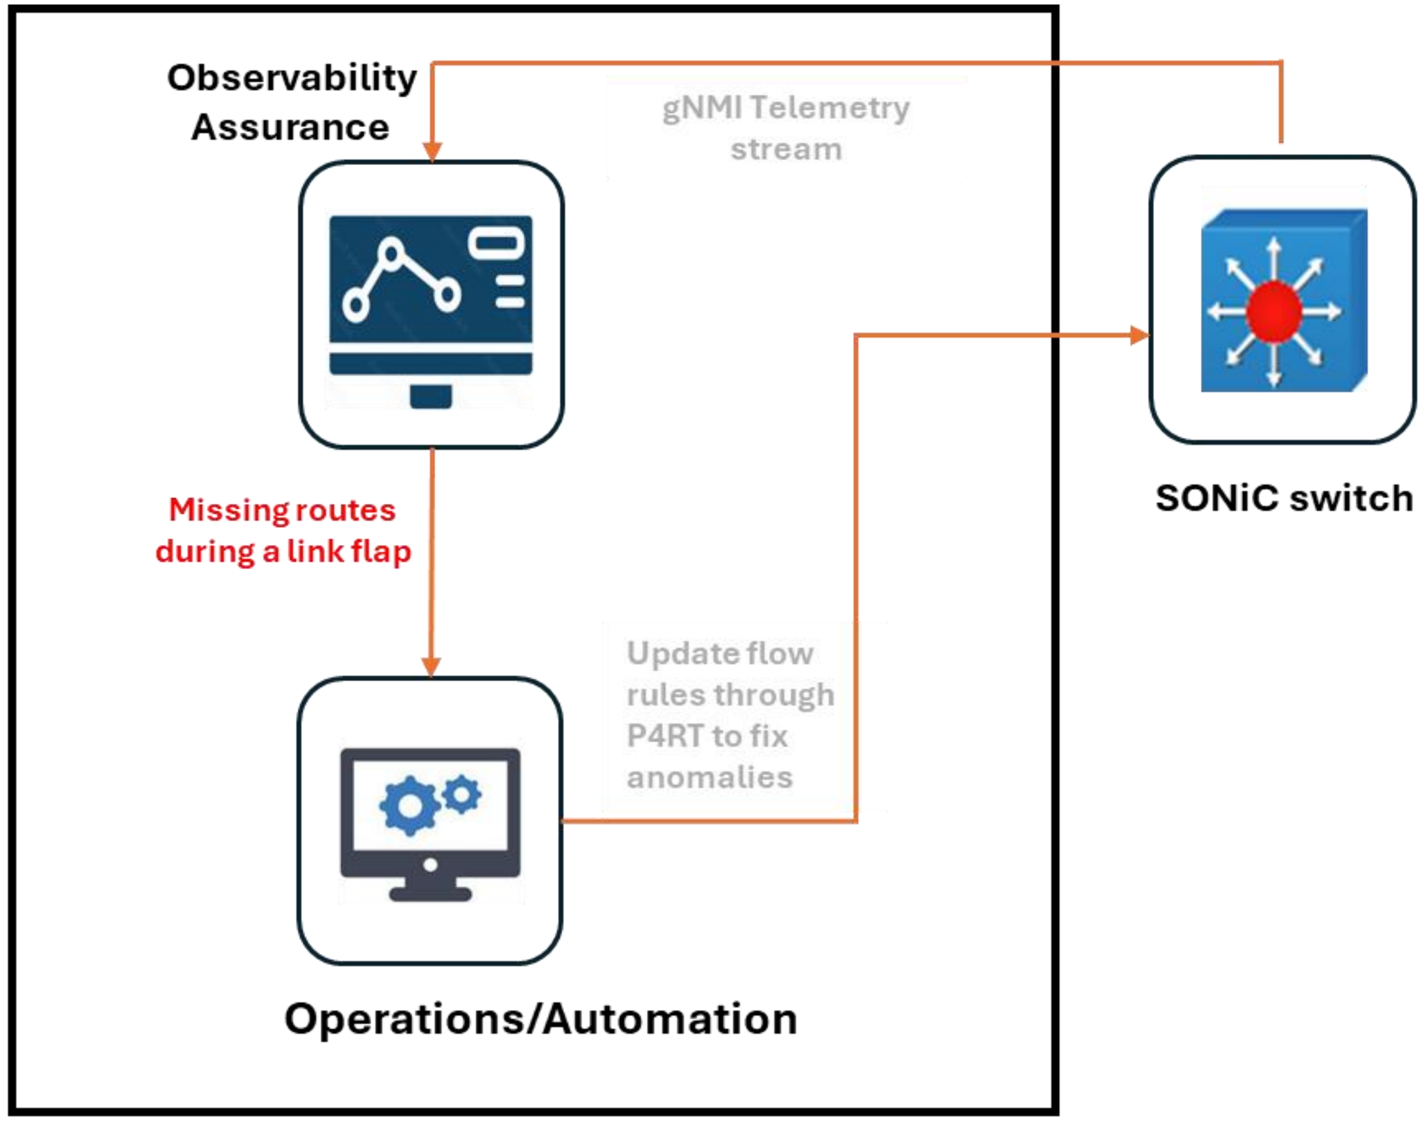
Observability Assurance
gNMI Telemetry stream
SONiC switch
Missing routes during a link flap
Update flow rules through P4RT to fix anomalies
Operations/Automation

GEICO Software Defined Networking (SDN)

Sensitivity: Company-Internal

GEICO SONiC Components

GEICO SONiC Model

- Open Community SONiC code
- GEICO features and fixes
- Vendor SDK: SAI + Platform Packages

SONiC PINS Components

- P4Runtime Server
  - Database Adapter: P4Runtime -> APP_DB
  - GENL-PACKET receiver, packet tx sockets
- APP_DB & APP_STATE_DB
  - P4 pipeline objects
- P4OrchAgent
  - Map P4 APP_DB objects to SAI objects
- ASIC_DB
  - Holds SAI objects
- SyncD (Vendor SDK)
  - SAI -> Vendor SDK API
  - Vendor SDK API -> Driver & HW API
- Kernel drivers & modules (Vendor SDK)
  - Vendor driver and modules

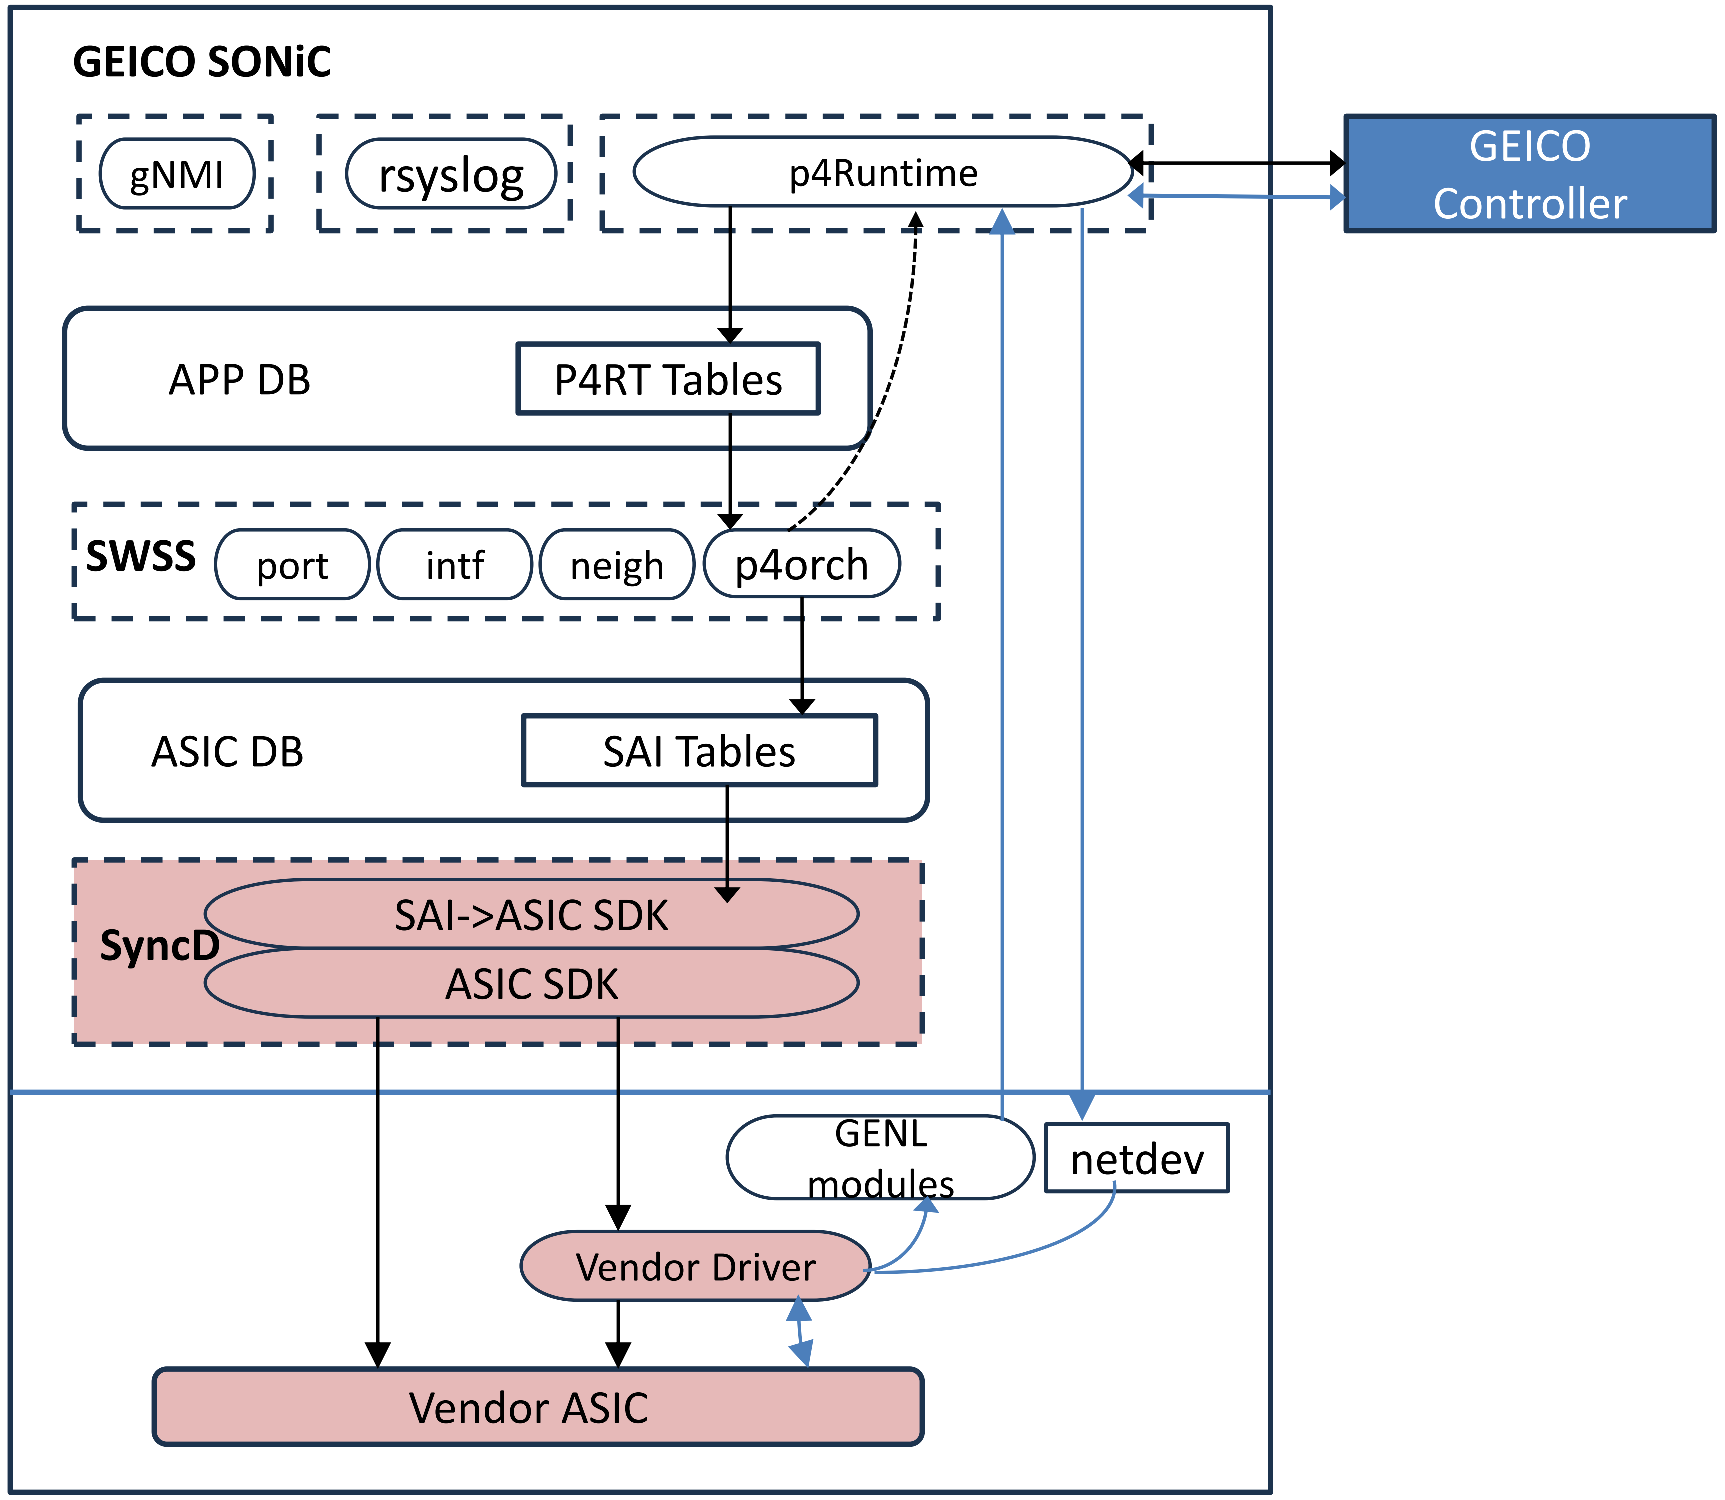

PINS P4Orch & SAI Dependency

Action:

- Create Ingress ACL to trap LLDP & ARP packets

Issues:

- P4Orch tries to create pre-ingress ACL when ingress ACL entry is being created.

Recommendation:

- SAI object dependency should be reflected in P4 Program and P4Orch code
- Useful to make explicit the SAI object dependencies in P4Orch code.

```
 // Copy the packet to the CPU. The original packet is dropped.
action acl_trap(@sai_action_param(QOS_QUEUE) @id(1) qos_queue_t qos_queue) {
  acl_copy(qos_queue);
  mark_to_drop(standard_metadata);
 }

 table acl_ingress_table {
  key = {
    ......
  }                              const default_action = NoAction;
  actions = {                      meters = acl_ingress_meter;
  @proto_id(2) acl_trap();         counters = acl_ingress_counter;
  @defaultonly NoAction;           size = ACL_INGRESS_TABLE_MINIMUM_GUARANTEED_SIZE;
  }
}
```

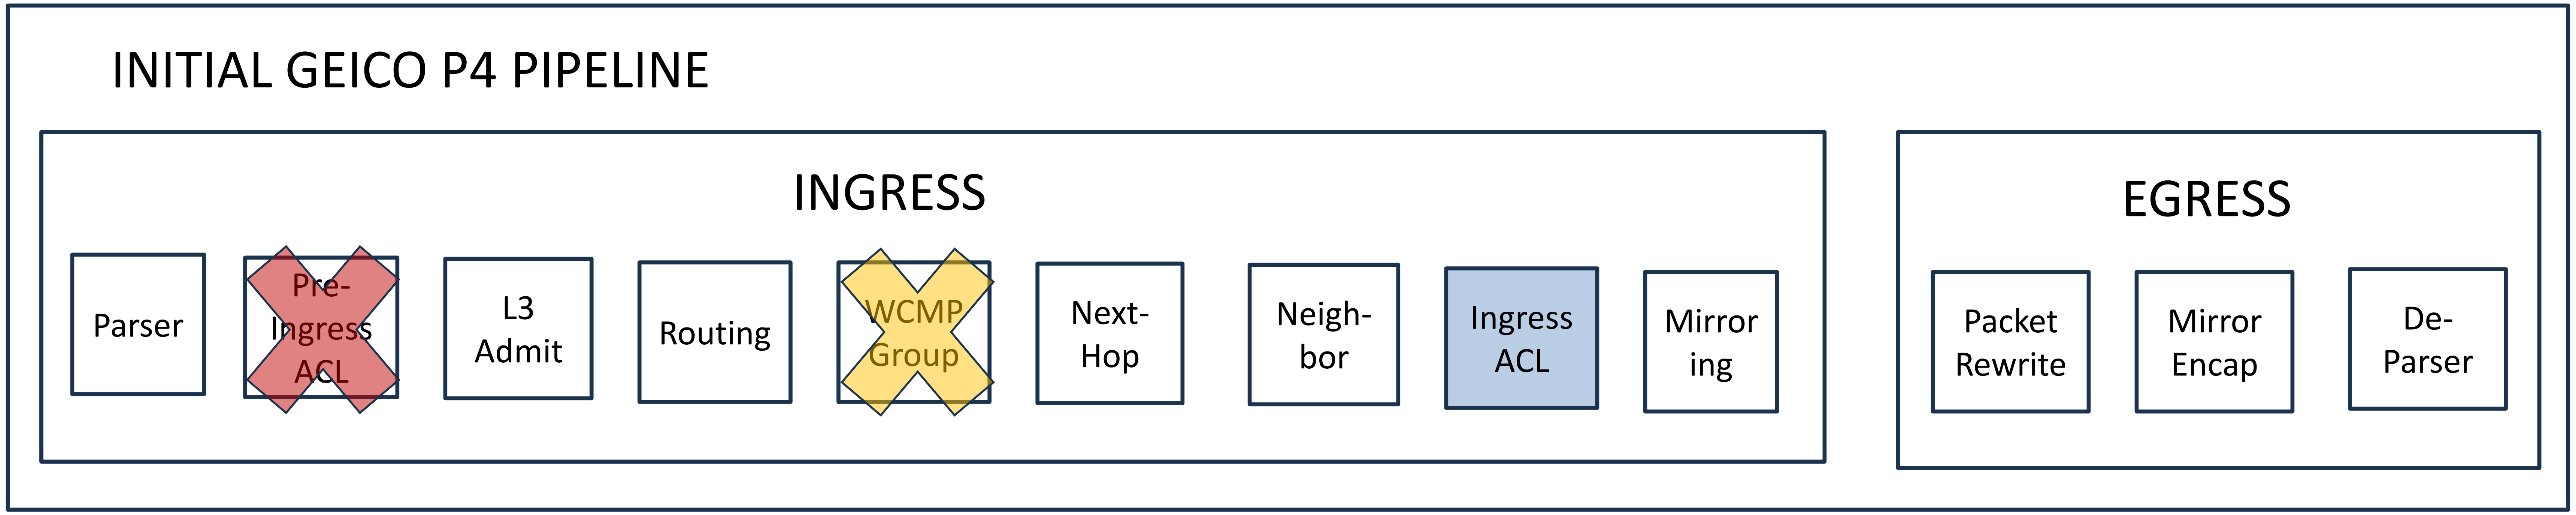

PINS Driver Dependency

Action:
- Send genl-packet msg for LLDP / ARP punted by Ingress ACL

Issues:
- Without ASIC driver support or source code, not able to send on genl-packet family

Solution:
- Used netfilter PRE_ROUTING hook to send genl-packet msg, instead of sending the msg from vendor driver

Recommendation:
- Recommend having vendor independent framework for generating genl msg that includes P4 defined user metadata.

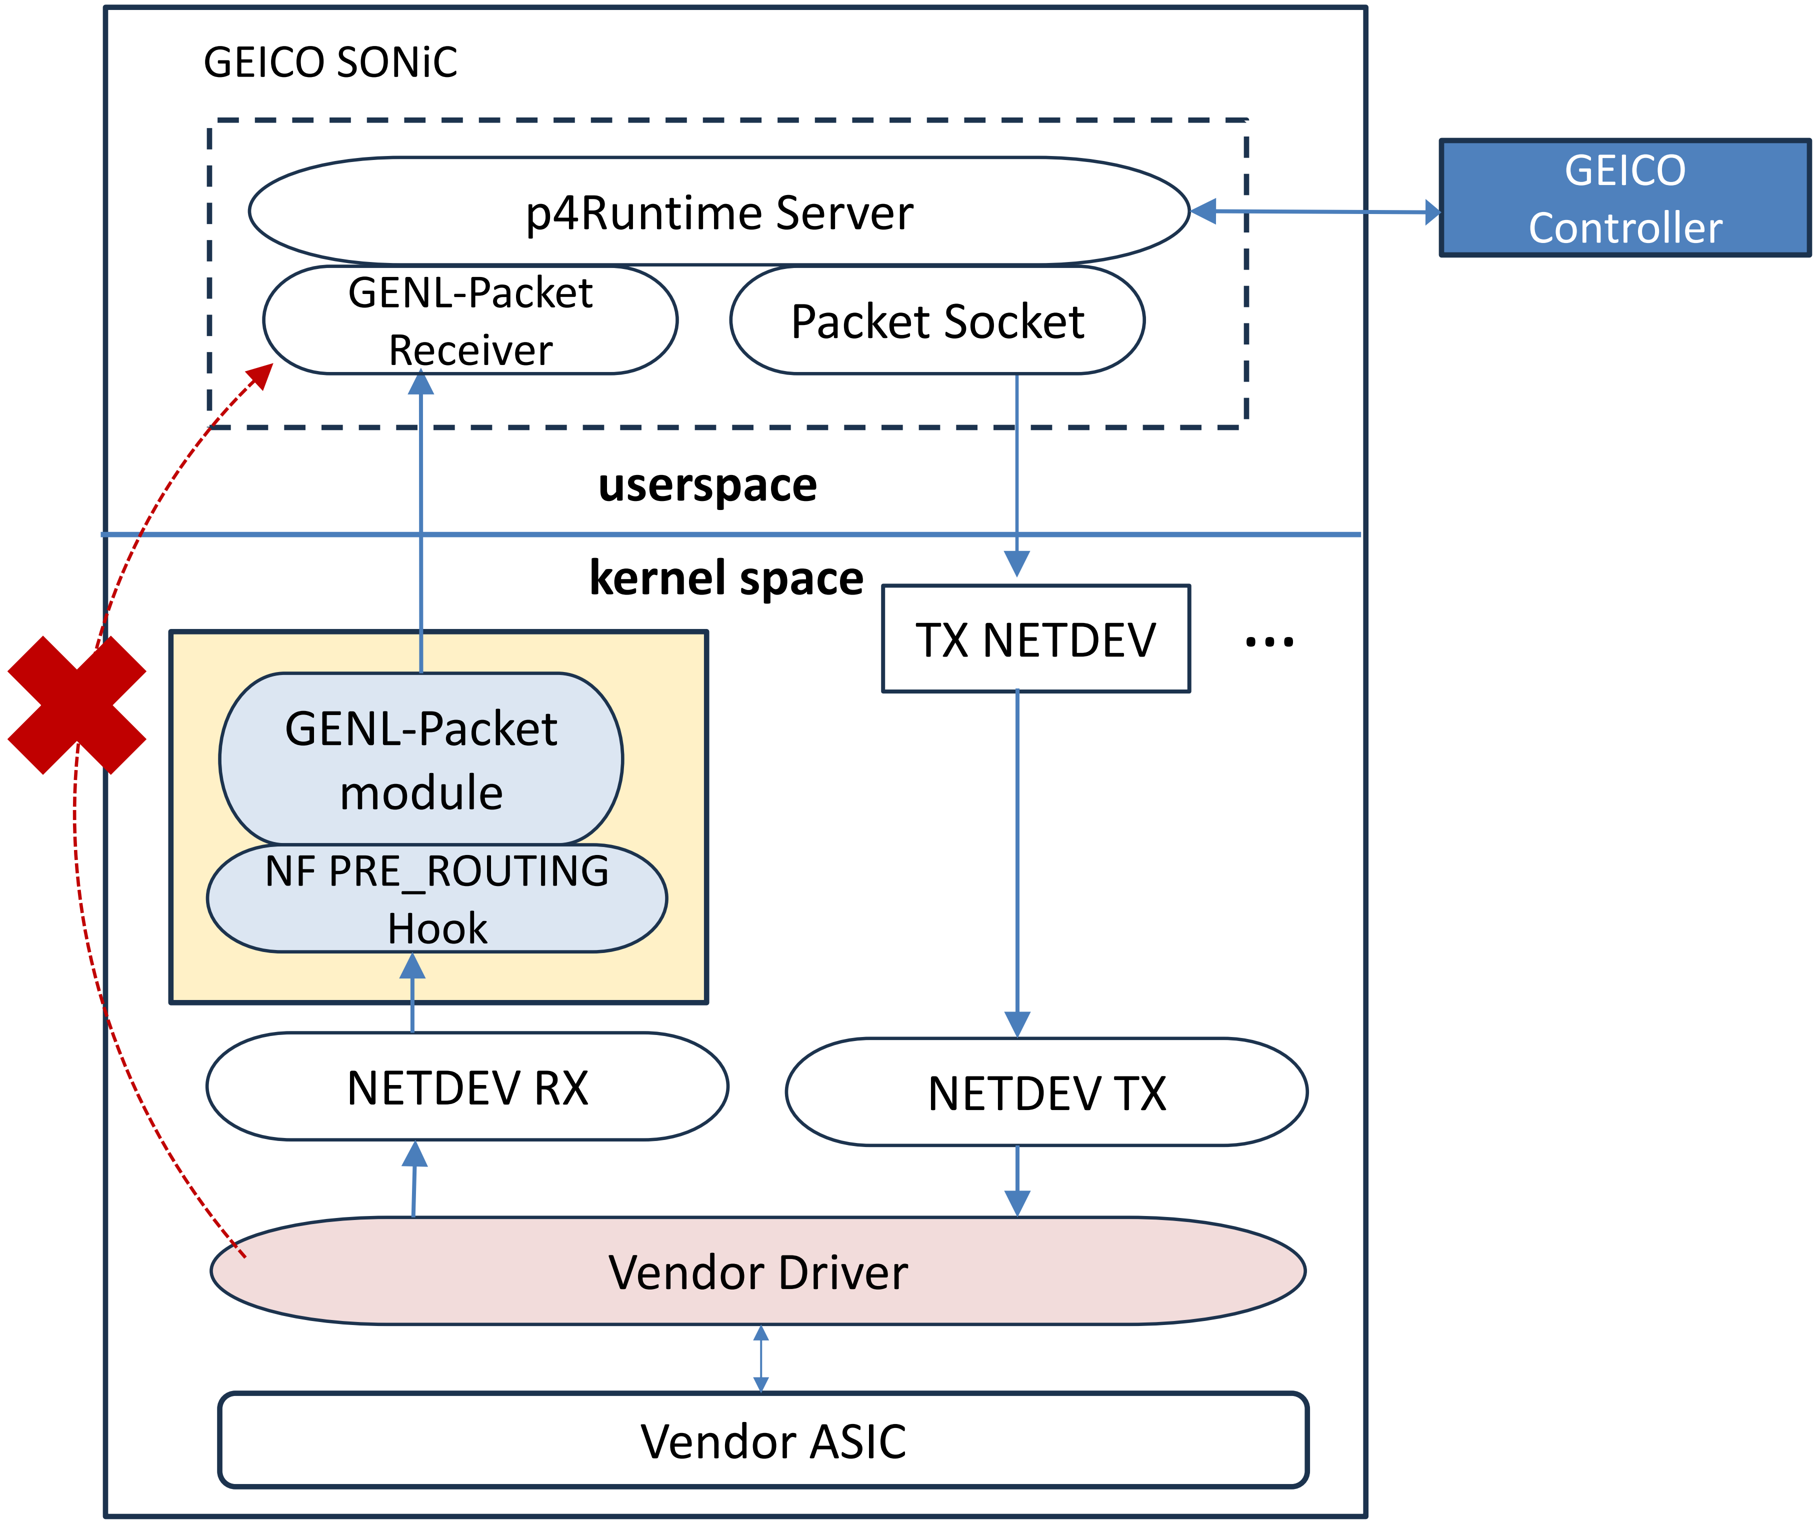

Sensitivity: Company-Internal

GEICO Next Step

Action:
- Move to the latest version of SONIC-PINS
- Implement VxLAN Tunneling

Issues:
- Hard to find accompanying P4Orch and other SONIC component code changes that go with latest SONIC-PINS changes

Recommendation:
- Make P4Orch changes available early to community

Call To Action:
- Join GEICO team in the SONiC Community to collaborate on SONIC-PINS development!

Thank You