BGP Flow Specification for Source Address Validation

draft-geng-idr-flowspec-sav-03

Nan Geng (Huawei), Dan Li (Tsinghua University)

July 2024

Use Case: BGP FS for SAV

- Use Case: BGP FS for SAV
  - Enhance source address validation when routers have not been upgraded to support SAVNET mechanisms
  - Unlike existing SAV, FS supports flexible traffic handling actions.
- SAV rule
  - <src, incoming-interface, …>
- How to generate SAV rules
  - Run SAVNET mechanism in controller or security center
  - IRR/RPKI/BGP data can be used for SAV rule computation

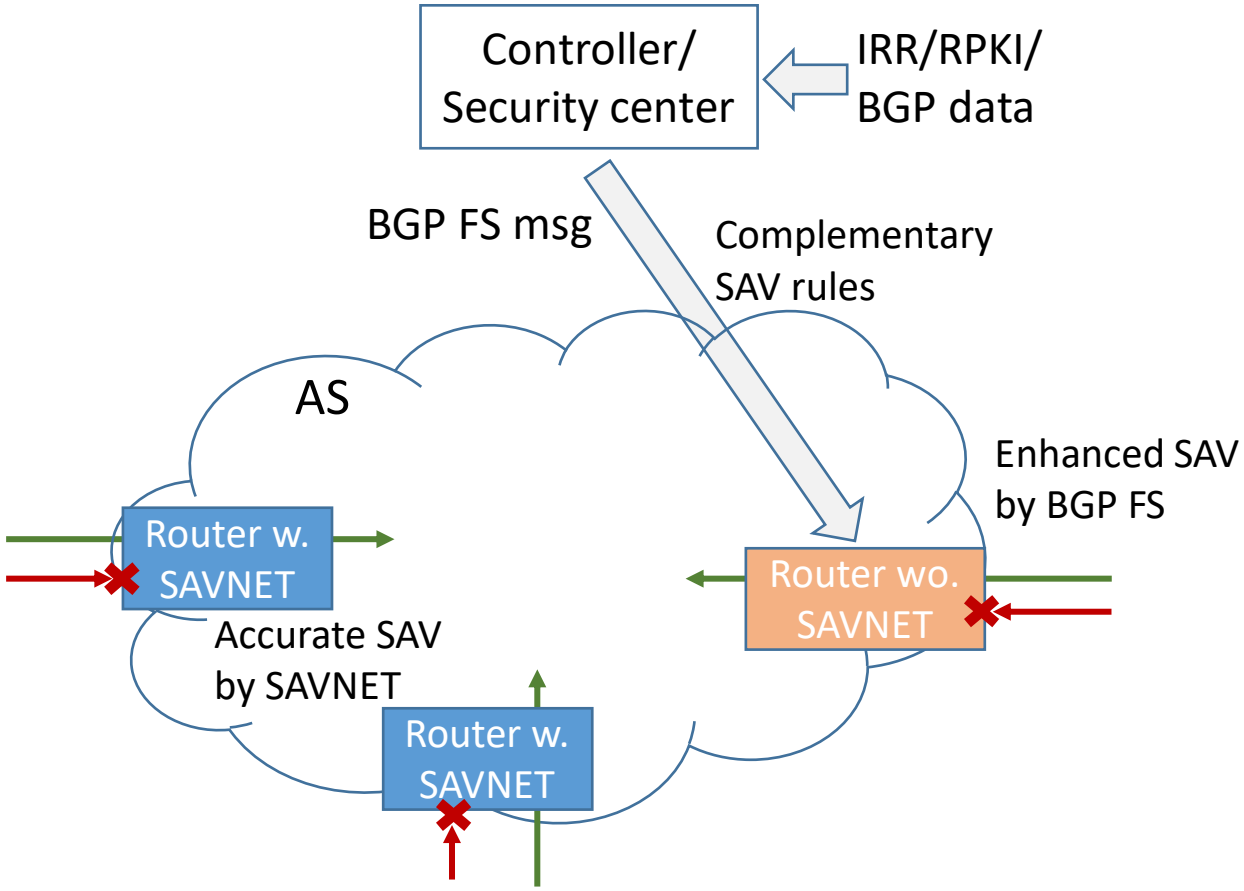

Similar use case in draft-tong-savnet-sav-enhanced-by-controller-00

Some discussion points

- How to make good scalability?
- Where to put Interface-set, Extended community or NLRI or others?
- How about grouping interfaces by using neighboring ASN which may be in NLRI or community?
- Whether a flag can be used to indicate that the FS route is specific for SAV and will be installed in a SAV table?

How to make good scalability?

- Group source prefixes. Different granularities:
  - a specific source prefix
  - source prefix group identifier
  - router-id
  - origin AS number
  - peer or peer AS number (all routes received from that peer)
  - etc.
- Group interfaces. Different granularities:
  - a specific interface
  - interface group identifier
  - neighboring AS number
  - etc.
- It is not new to group them in FS: draft-wang-idr-flowspec-dip-origin-as-filter, draft-ietf-idr-flowspec-interfaceset, etc.

Interface-set in Extended community or NLRI?

- Interface-set in Extended community
  - Pro: It can be re-written in the inter-domain flowspec case, while NLRI cannot.
  - Con: Need to maintain interface-set id mappings in the inter-domain flowspec case. (right?)
- Interface-set in NLRI
  - Pro: In terms of meaning, Interface-set is more like a filter instead of an action. Filter decides which flow to take actions on, while action decides which operation will be taken on the matched flow.
    - Interface-set is one of the factors deciding which flow to take actions on. Unlike action, Interface-set itself is applied to routers not flows.
    - OpenFlow also considers in_port as a match field instead of an instruction. (https://opennetworking.org/wp-content/uploads/2014/10/openflow-spec-v1.3.0.pdf)
  - Con: Interface-set (Group Identifier) has very local meaning. Same value may map to different interface sets among ASes. May not suitable to inter-domain flowspec case.
- Q: can interface-set has non-local meaning?

Can interface-set have non-local meaning?

- Use a Flag bit to indicate whether Group Identifier has a local or a non-local meaning.
- For example, a "non-local meaning" Group Identifier field carries AS number, which represents all the interfaces connected to the neighboring AS with the AS number.
- When BGP messages with "non-local meaning" Group Identifier is propagated across ASes, the groupId can still be recognized by other ASes.

- Q: How to verify the routes with only src+interface in the inter-domain flowspec case?
- A: Only accept the FS routes with source prefixes originated from the same AS as the FS route, which is similar to draft-geng-idr-bgp-savnet-03.
  - Use case: A customer advertises SAV rules to its provider. The provider installs SAV rules for the source prefixes of the customer and helps block the forged source prefixes transiting the provider AS.

A flag to indicate that the FS route is specific for SAV and will be installed in a SAV table?

- BGP FS rules are mostly installed in ACL table/firewall table. Not dedicatedly designed for source prefix filtering. (draft-huang-savnet-sav-table-06)
- Suppose BGP FS is used to distribute SAV rules. Where to install SAV rules can be implicitly or explicitly indicated.
- Order: SAV rules are behind other BGP FS rules, which is what routers do now (e.g., ACL first, then uRPF, then FIB).

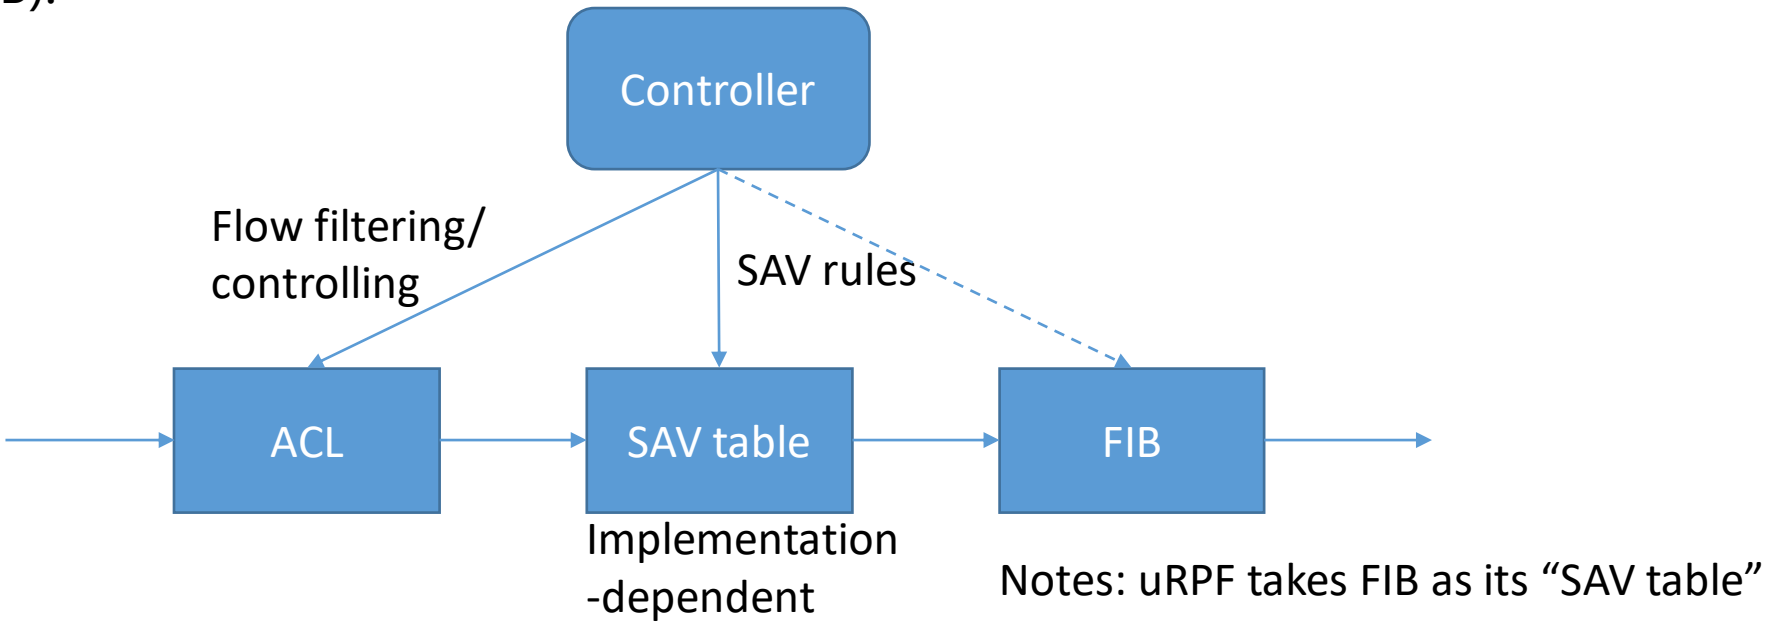

Conclusion

- The draft-geng-idr-flowspec-sav-03 does not include all the discussion points. Will update the draft according to the feedback received during IETF 120.

- Thanks for the comments from Jeff Haas, Sue Hares, Nat Kao, Randy Bush, etc.

- Welcome to leave your comments or questions.

Thanks!