# Flow Specification v2

draft-hares-idr-flowspec-v2

Hares and Eastlake

# FSv2

- FSv2 requires 2 new SAFIs (IP, IPVPN) + 1 Open Capability (FSv2)
- FSv2 components = FSv1 components + new FSv2 components
- FSv2 actions – ordered by number [redefine the FSv1 actions]
- Match rules + actions need to be ordered (FSv1 or FSv2)
- FSv2 orders rules by:
  1. User defined ordered (UD-Order)
  2. If UD-order, then order by FSv2 components
  3. If UD-order + FSv2 components same, then values
  - FSv1 in DB follows FSv2 (allows for easy deployment [Keyur])
- Actions chain is ordered in FSv2 by:
  - User defined action order
  - If action's UD-order is same, then by Action type
  - If UD-order + action type is the same, then by value (per action)

# FS data base for FSv2 Node

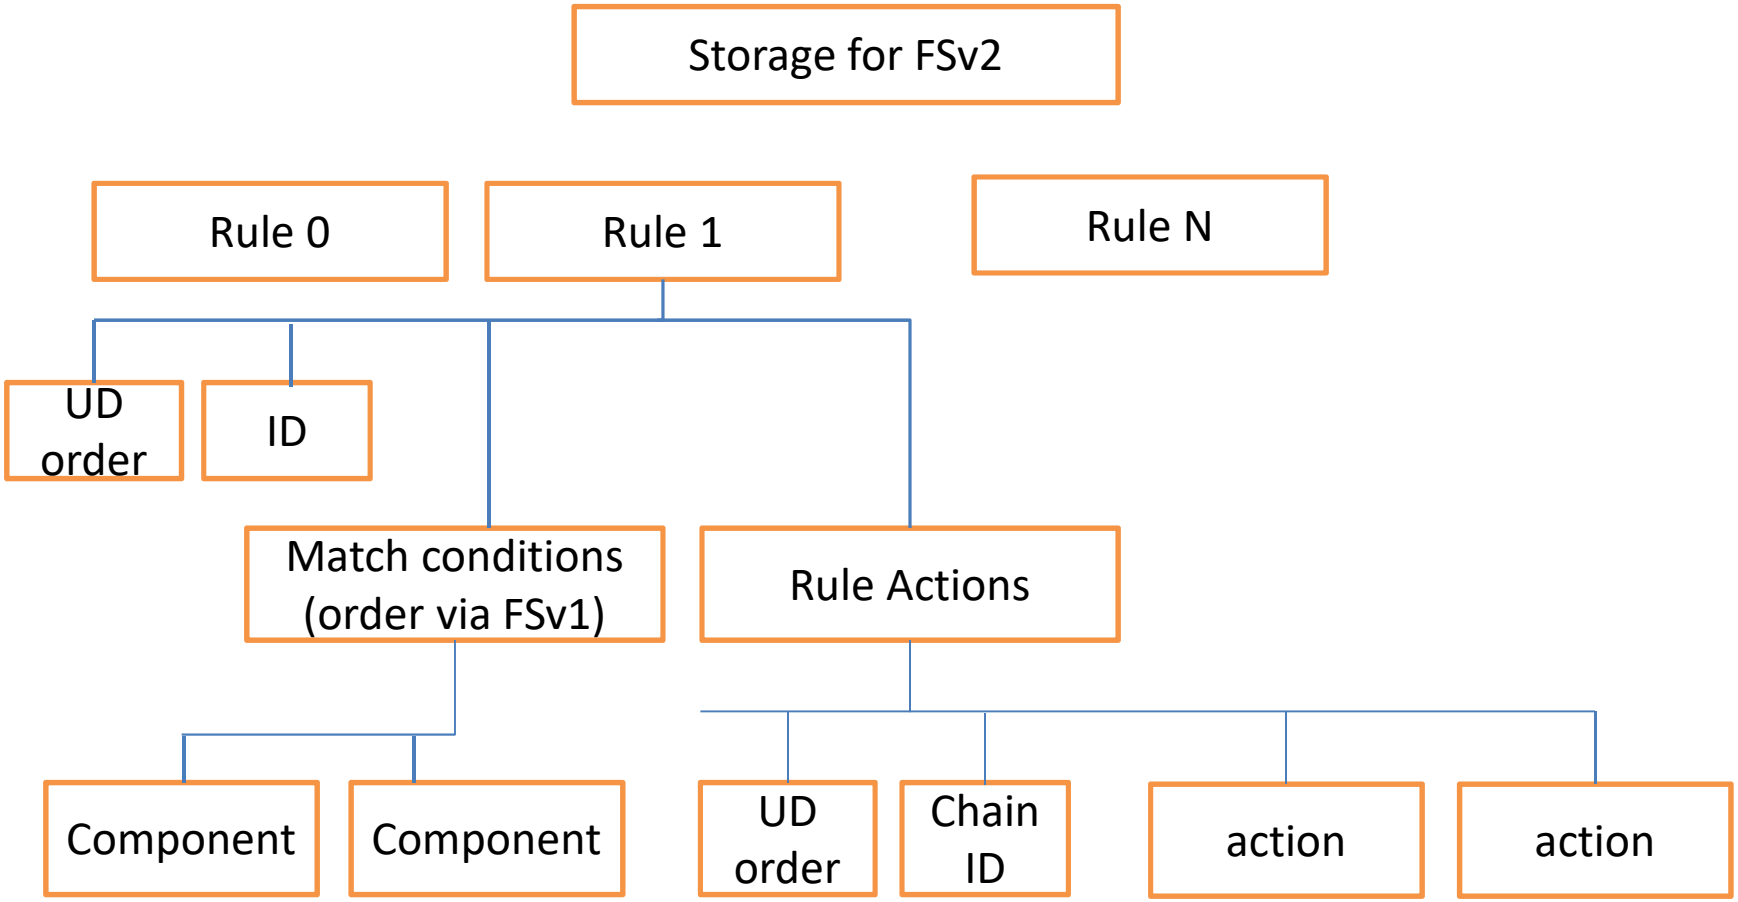

# Basic principles of Filter DB

- Filter DB simple for FSv1 only or FSv2 only
  - Rule-zero – 0/0 with permit all
  - Rule 1-N – FS filters
- BGP peers with FSv1 + FSv2 need to
  - Rule-zero – 0/0 with permit all
  - Rule 1 to  N – FSv2 filters
  - Rule N to end – FSv1 filters

# FSv2 + FSv1 SIN

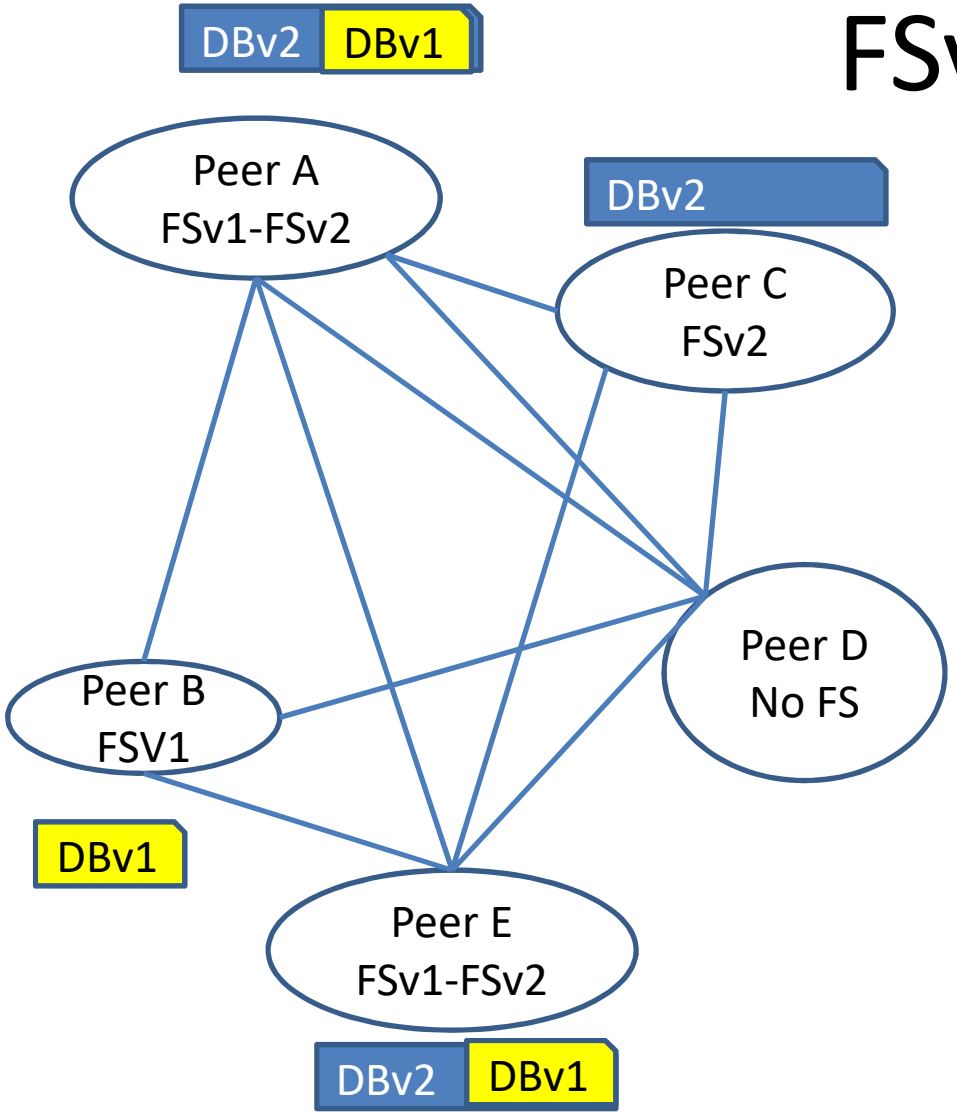


- SIN – Ships in Night NLRIs
  - No BGP importing of FSv1 into FSv2
- 5 BGP Peers in under administrative domain
  - Complete mesh of Peers (not all links shown)
  - Peer-A-B – pass FSv1
  - Peer A-C – pass FSv2
  - Peer A-D – no FS
  - Peer A-E – FSv1 and FSv2
- Peers
  - Establish by capabilities
  - Pass DBs

# Action Chains with User-define Ordering

- Deterministic User defined Ordering (UDO)
  - Action zero (default) defined as “permit all”
  - User-defined Order value
  - Ordering within the same user-defined order:
    - Action type, then Action Value
    - Actions must define value comparison
  - What happens when Actions fail to complete

- Issues
  - Operational issues with NLRI associated action vs Denial of Service “Die-Die-Die Internet Worm”
  - Action Chain Ordering – Default and changes
  - Some action chains will need conditional branch points

# Action Chain Operation

- FSv2 Actions must plan for failure
  - Default – stop upon failure
  - Other options:
    1. Continue on failure
    2. Do all or nothing
    3. Conditional continue

# DDoS Response Requested

- BGP
  - Delivers NLRIs + Communities – with good scaling properties
  - does not have action-response function
- NETCONF/RESTCONF
  - Action/response built into monitoring capabilities
  - Push/Pull – with filtering have been worked out
- Best scaling
  - Use BGP to deliver NLRI
  - Set YANG Model with monitoring action that sends information when filters are installed
  - Filters already capable to tune response flow for massive bursts

# Drafts Considerd

Components
- draft-li-flowspec-srv6-07.txt
- draft-ietf-idr-flowspec-l2vpn.txt
- RFC9015

New AFI/SAFI
- draft-ietf-idr-flowspec-nv03.txt

Actions
- draft-ietf-idr-flowspec-path-redirect
- draft-ietf-idr-flowspec-interface-set
- draft-ietf-idr-flowspec-ip-02.txt
- RFC9015 (SFC flow specification action)
- draft-ietf-idr-flowspec-l2vpn

- IDs in process
  - draft-dong-idr-flowspec-scalable-prefix-steering-01
  - draft-ietf0-idr-srv6-flowspec-path-redirect-06
  - draft-wang-idr-flowspec-dip-origin-as-filter-04
  - draft-jiang-idr-ts-flowspec-srv6-policy-04
  - draft-xiong-idr-detnet-flow-mapping-00

# Issues to Discuss

- FSv2 Actions
  - What happens if an Action fails in a chain?
  - How are new actions deployed? What happens if failure?
- FSv2 NLRI filters
  - How should nodes handle unknown filter components?
  - How can we incrementally deploy new filter components
- FSv2 + FSv1 nodes
  - BGP FSv2 and FSv1 routes are Ship-in-Night
  - Filters for FSv2-FSv1 need deterministic order
- FS nodes (v1 or v2) versus “no FS” nodes
  - What Operational issues have we left out?
- Error handling
  - Another embedded NLRI error handling case
  - Look at Validation + Error handling section

# Discussion Topology

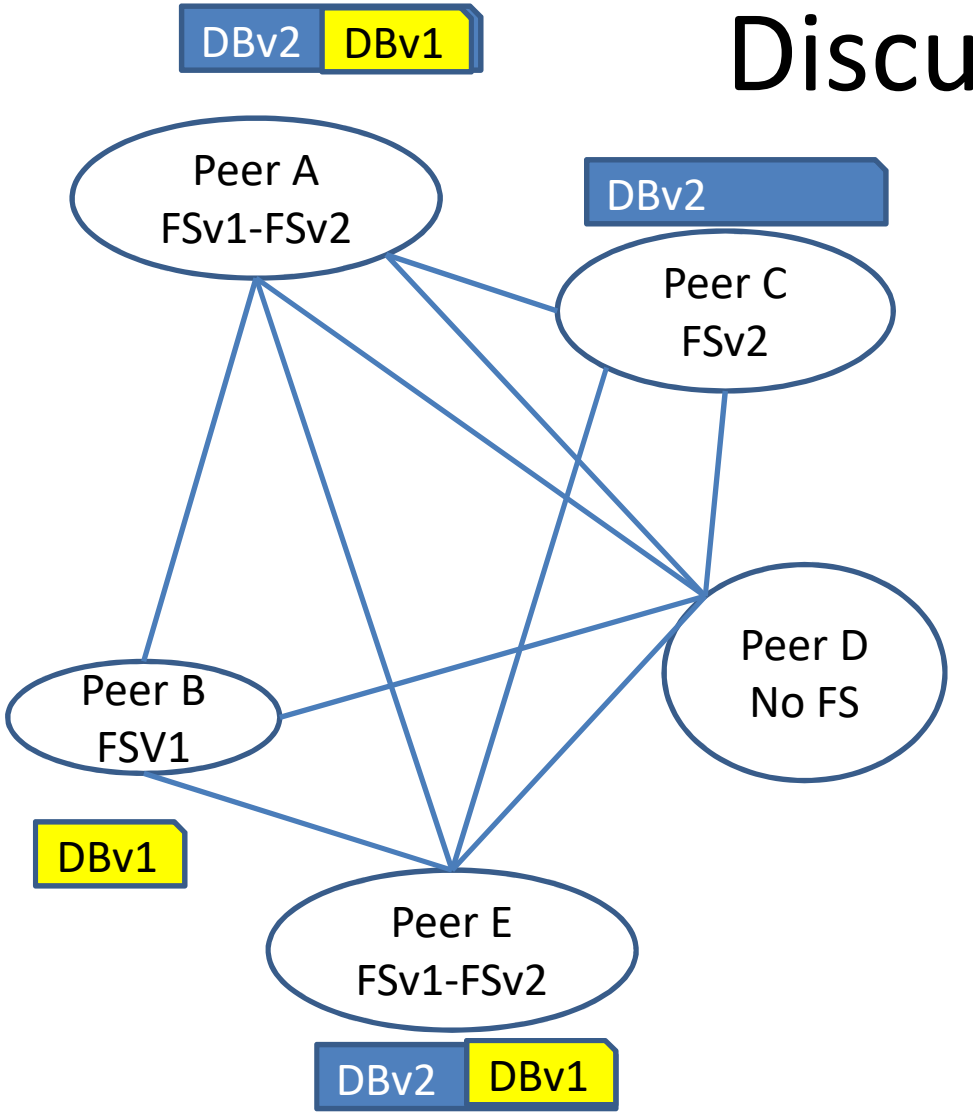


- SIN – Ships in Night NLRIs
  - No BGP importing of FSv1 into FSv2
- 5 BGP Peers in under administrative domain
  - Complete mesh of Peers (not all links shown)
  - Peer-A-B – pass FSv1
  - Peer A-C – pass FSv2
  - Peer A-D – no FS
  - Peer A-E – FSv1 and FSv2
- Peers
  - Establish by capabilities
  - Pass DBs

# Discussion RR topology

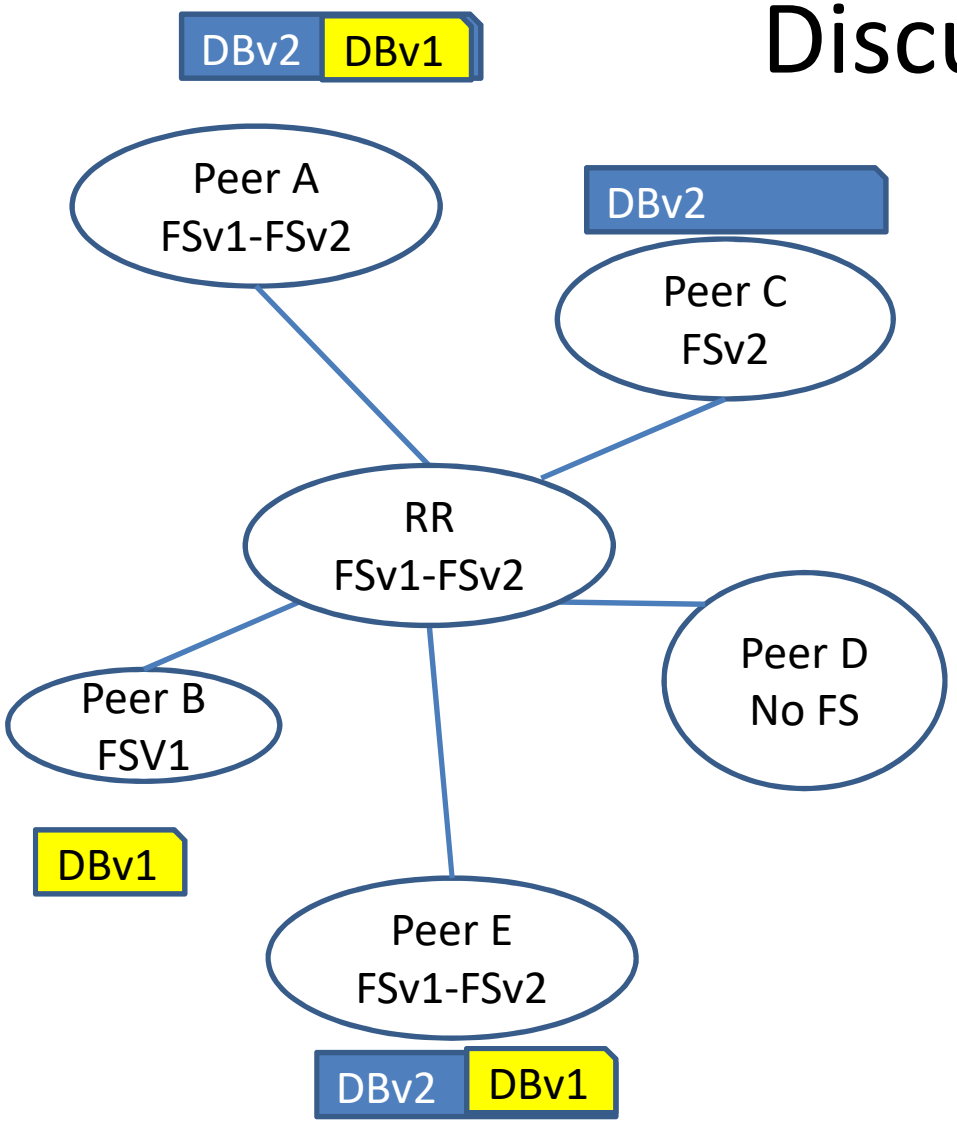


- SIN – Ships in Night NLRIs
- RR clients – take different NLRIs

- 5 BGP Peers in under administrative domain
  - Complete mesh of Peers (not all links shown)
  - Peer-A-B – pass FSv1
  - Peer A-C – pass FSv2
  - Peer A-D – no FS
  - Peer A-E – FSv1 and FSv2
- Peers
  - Establish by capabilities
  - Pass DBs

# Blank Slide for topology drawings

# Questions or Thoughts