

FILTERING OVERLAPPING ROUTES

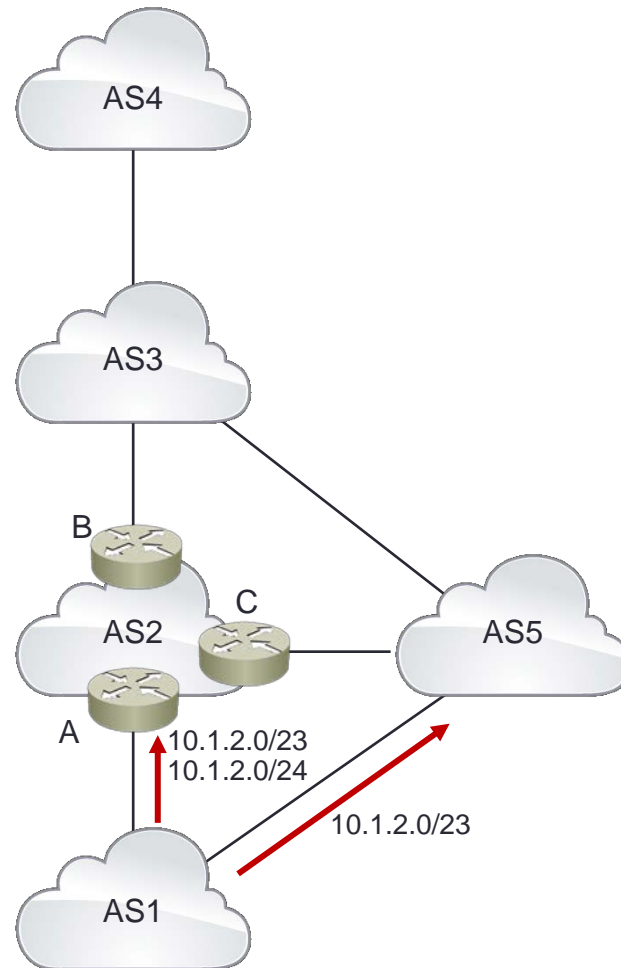
draft-white-grow-overlapping-routes

Filtering Overlapping Routes

- Inbound
 - When a router receives two overlapping prefixes...
 - Mark the longer one with the BOUNDED community
 - Locally defined community
 - Set a cost community on the longer prefix so it is preferred within the AS
 - This is essentially a tie breaker, leaving current usage of LOCAL_PREF intact
- Within the AS
 - Normal BGP decision process
 - (Optional) Don't install the route in the local RIB
- Outbound
 - Normal BGP decision process
 - Don't advertise routes to eBGP peers if they're marked with BOUNDED


Filtering Overlapping Routes

- AS1 advertises...
 - 10.1.2.0/24 to AS2
 - 10.1.2.0/23 to AS2
 - 10.1.2.0/23 to AS5
- Normally...
 - AS3 and AS4 will receive both 10.1.2.0/23 and 10.1.2.0/24
 - This increases their routing table sizes, but doesn't provide much additional usable information



Filtering Overlapping Routes

- New behavior:
 - Router A
 - marks 10.1.2.0/24 as BOUNDED
 - sets the cost community so 10.1.2.0/24 is preferred AS2
 - Router B filters 10.1.2.0/24 towards AS3
- Table size reduction:
 - AS2 may not install 10.1.2.0/24 in the local RIB/FIB
 - AS3 and AS4 now only receive the shorter prefix route



Filtering Overlapping Routes

- The key is to remove routing information when it's no longer needed
 - If the route isn't impacting traffic flow, take it out of the system
 - How do we know it's not impacting traffic flow? Because it is overlapped by a shorter prefix route that leads to the same destination
- This is not aggregation
 - Aggregation increases stretch in the network
 - Aggregation requires ownership of the shorter prefix

Next Steps

- Comments/Questions on List
- Make this a working group doc