

Multisession -04

John Scudder jgs@juniper.net

IDR, July 30, 2009

Multisession Refresher

- Facilitates establishment of multiple BGP sessions between a pair of speakers
 - Based on AFI/SAFI or other criteria
- Why
 - Finer-grained fault management and isolation
 - Management of control plane resource contention between services
- Mechanism
 - Revise rules for exchanging OPEN messages
 - Backward-compatible with 4271

Problems With Version -03

- draft-ietf-idr-bgp-multisession-03 has no means of telling peer an alternate port number to use
 - This severely hampers ability to use multiple BGP daemons since handing off an open TCP session is not straightforward

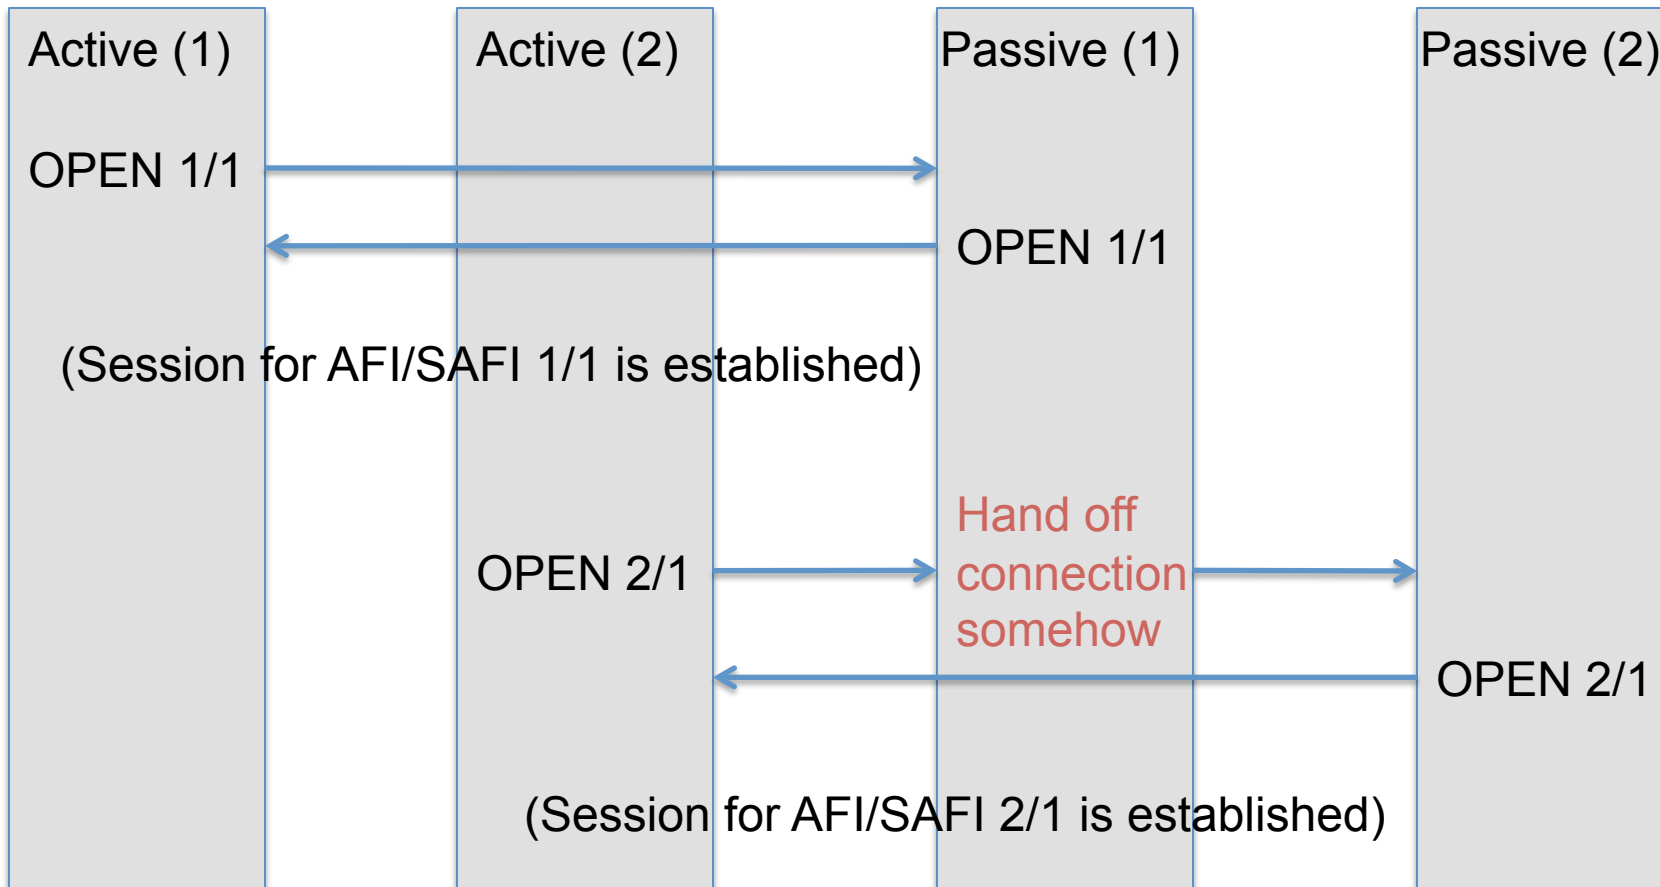
Proposed Solution

- Active speaker has no problem — the appropriate BGP process initiates the connection
- Allow passive speaker to request a redirect to a new port number

Operation in -03

- BGP process on active speaker sends an OPEN including Multisession capability and other capabilities such as MP-BGP
- Passive speaker analyzes OPEN, replies with its own
 - Any handoff between BGP processes on passive speaker must be done “under the covers”

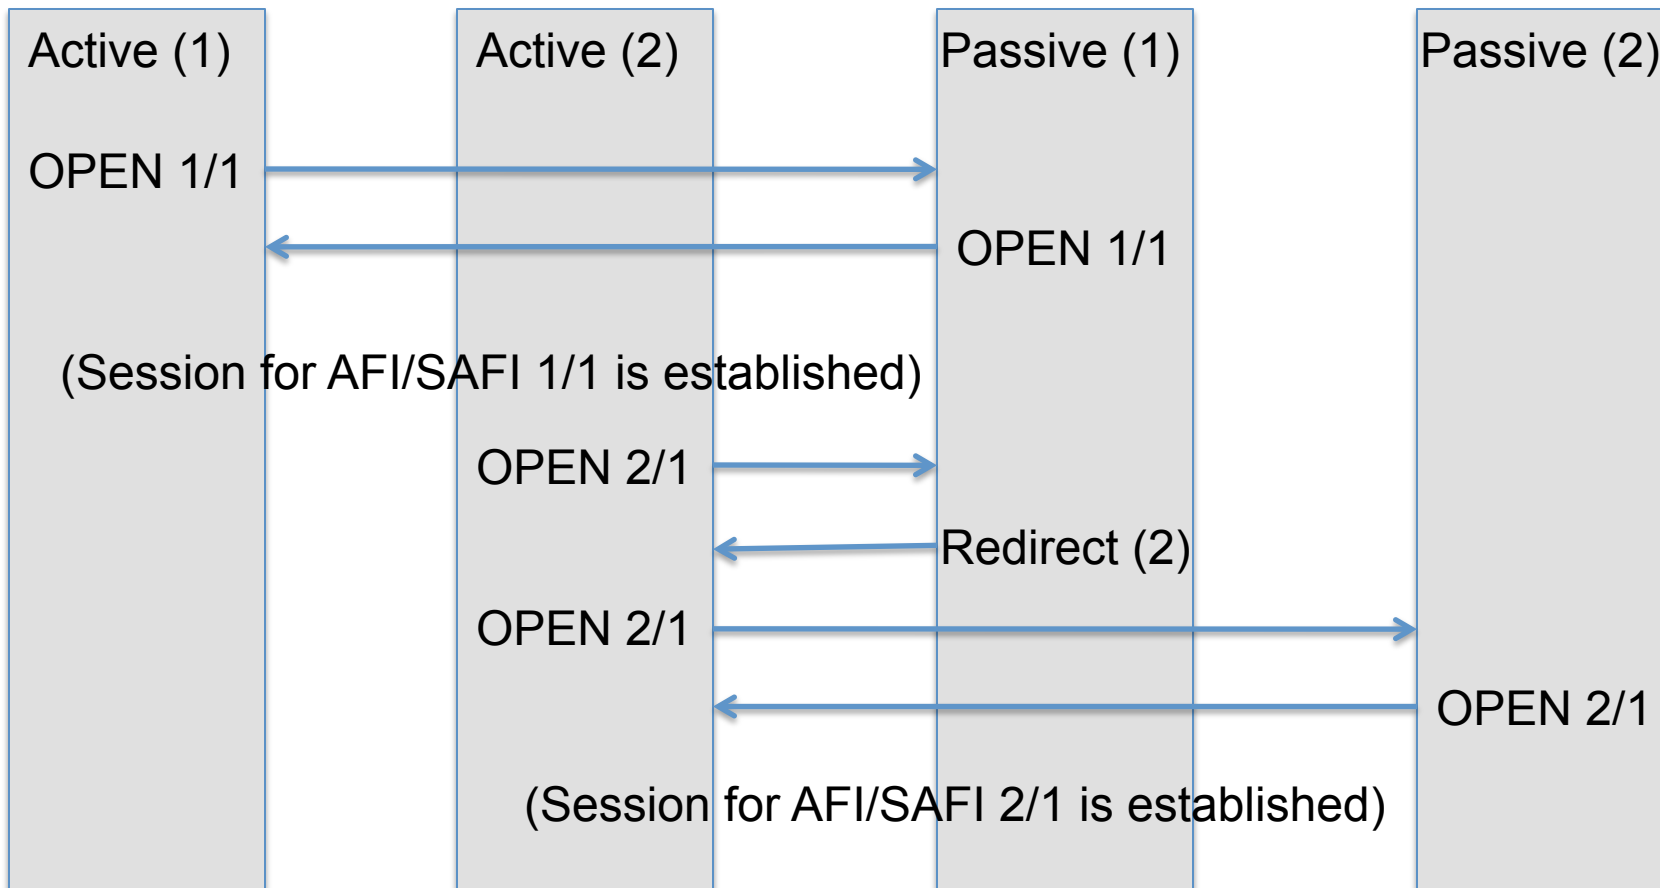
Example of -03 Operation



Changes

- Update Multisession Capability to add the “(R)edirect” bit
 - If R bit is set, a port number follows
- Only passive speaker may send R bit
 - Because active speaker has no problem
- Active speaker the drops the connection, reconnects to the specified port

Example of -04 Operation



Other Changes

- Two new error codes
 - “Redirecting Now” for active speaker to orderly close session before opening the new one
 - “Redirect Required” for passive speaker to complain if active speaker failed to redirect
- Plus editorial improvements

Backward Compatibility

- Redirection is optional — encoding is not changed if it's not used
- Current deployment is restricted anyway
 - Likely due in part to issue mentioned in problem statement

Comparison to Transport Instance BGP (TI-BGP) Proposal

- Both: different sessions over different port numbers between different BGP daemons
- TI-BGP divides BGP into two classes
 - Or some a priori fixed number
- Multisession allows any number of classes
- In computing, the answer is usually 0, 1, or N. Usually not 2.

Next Steps

- Implementation experience with -04
- Then iterate or advance spec