



FlexEthernet (FlexE) Use Cases

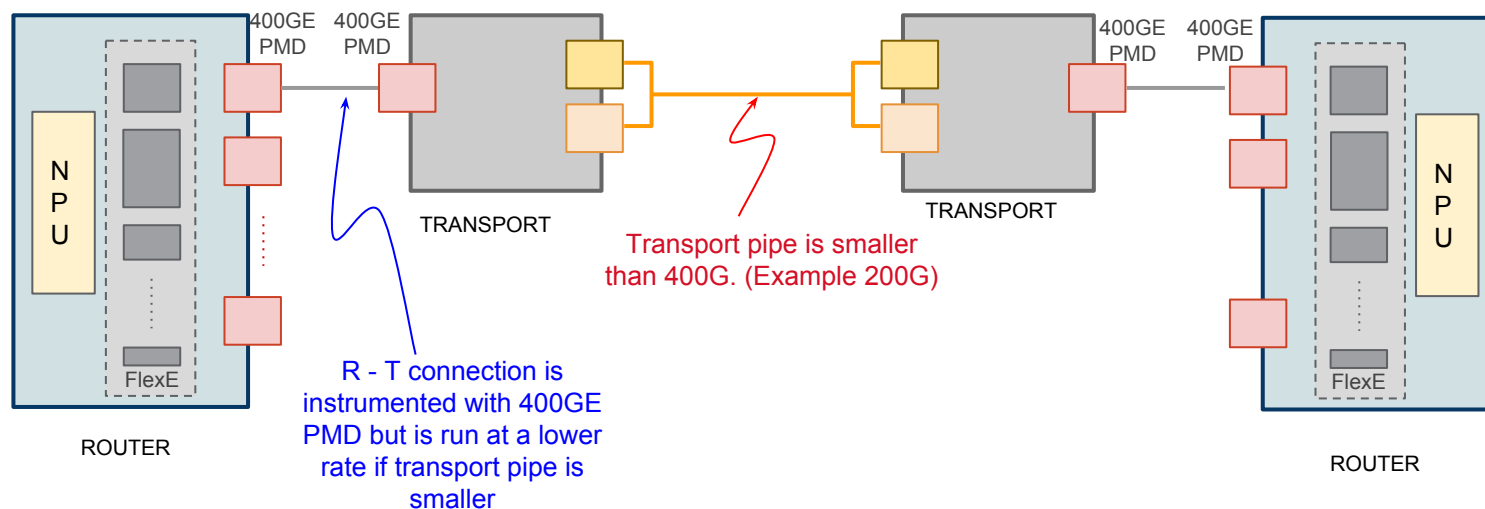
Vijay Vusirikala

TEF : The Rate Debate

Goals of FlexEthernet

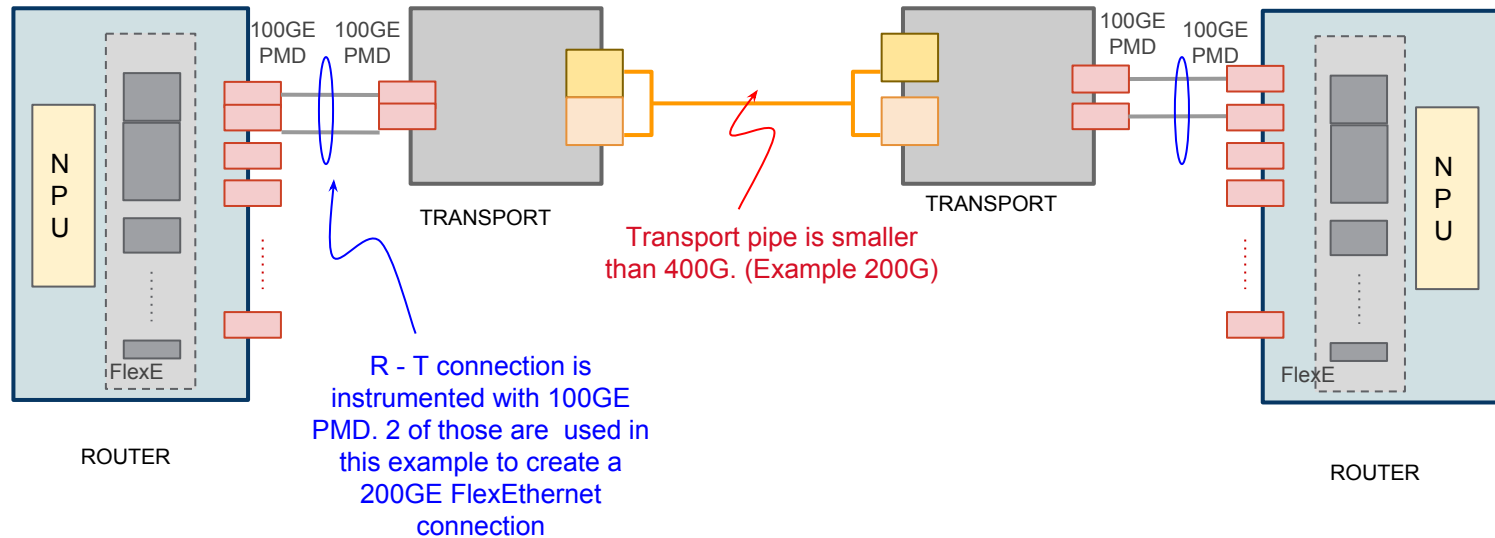
- Flexible size Ethernet instead of fixed 100G / 400G now and 1T (future)
 - Functionality of LAG without the hashing inefficiency and network state explosion
- Ability to use 100GE/400GE PMDs and any future PMDs without requiring IEEE standard process for a new Ethernet rate
 - i.e., decouple PMD development from other layers
- Utilize full capacity of NPU / Packet Forwarding Engines without waiting for future, higher speed, fixed Ethernet rates to be standardized
- Better match Router port rates with Long Haul (LH) transmission rates which have a finer granularity than 100/400/1000G because of physical layer constraints and Shannon's law

Use Case I - Transport Link Optimization w/ 400GE PMD



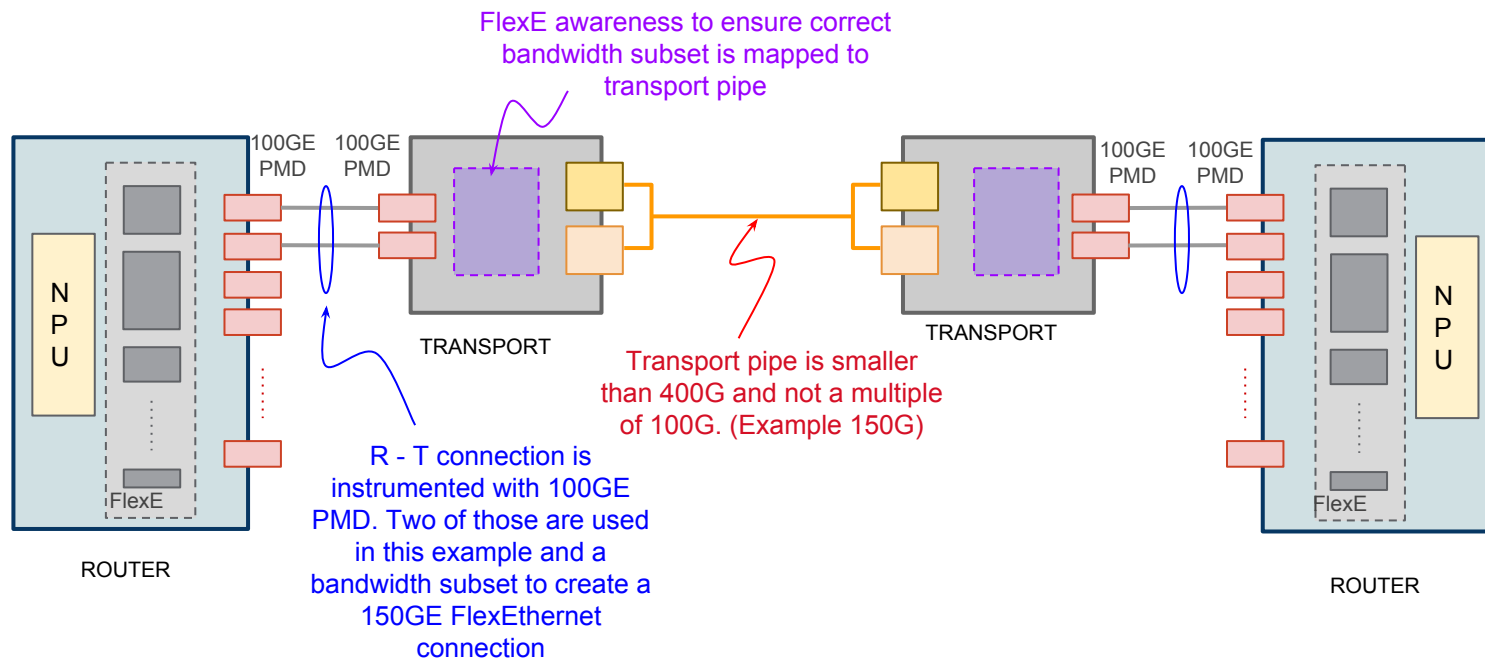
- *Value* : Cases where LH transport is substantially more expensive than client connections. Does not force LH pipe to be either 100G or 400G
- An additional benefit of this use case is that the Router port rate can be configured to match a change in transport pipe size dynamically w/out changing R-T connectivity
- *Down side* : Suboptimal utilization (stranding) of 400GE PMD and associated NPU capacity.

Use Case II - Transport Link Optimization with 100GE PMD



- In this use case, transport pipe is a multiple of 100G and the granularity of the FlexEthernet is 100G.
- No changes to Transport box (no FlexE awareness needed).
 - Current transponders can be used. (only requirement is to route all wavelengths associated with a specific FlexE connection to be co-routed to stay within FlexE skew limit)
 - Reduces the link state from $nx100G$ LAG to $1x n00G$ FlexE
 - In this use case, there is no NPU capacity stranding and no router port capacity stranding
- This use case is not limited by 400GE PMD and can scale up to the maximum capacity of the NPU.

Use Case III - Optimization with Fine Granularity



- This use case illustrates sub-100G granularity
- Value of sub-100G granularity is debatable - especially looking 2-3 years out
 - While 150G waves with 8QAM will have applicability, it is likely that 2 waves with 8QAM will co-propagate to have a 300G channel or 8QAM will be supported at 200G with higher baud rate
- The main complication of this use case is the modification to transport gear to incorporate FlexE awareness. Incorporating FlexEthernet at router endpoints is not sufficient.

Additional Use Cases & Notes

- In-service increase or decrease of FlexEthernet pipe size
 - Nice to have but not mandatory
 - For in-service pipe adjustment to be practically useful, we need to deploy highest capacity client (R-T) interfaces (i.e., Use Case I) or an aggregation or circuit switch between R and T (adds unnecessary complexity)
- In-service resizing (even if implemented) does not need to be hitless
 - It is OK to drop a few packets during the reconfiguration
 - Main benefit of in-service for us is remote (possibly SDN driven) capacity adjustment without manual reconfiguration