

Ethernet Flexes Its Ports - FlexEthernet

Tom Issenhuth
OIF Board Member
Microsoft



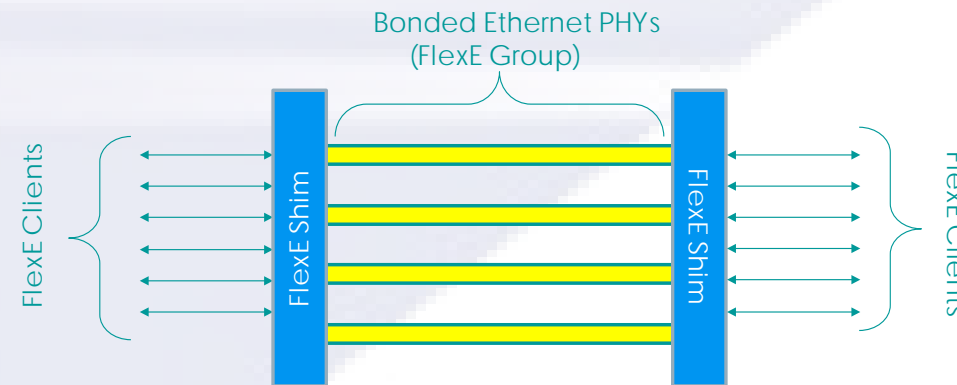
FlexEthernet History

□ History

- **Effort started by end users in 2013**
 - Google, Microsoft, Comcast, Telus, Time Warner and Verizon
- **Series of meetings with end users, system vendors and component vendors during 2013 & 2014**
- **Initial FlexMAC technical proposal by Cisco**
 - Proposal evolved throughout 2014 by group of interested parties
 - Name morphed to FlexEthernet (FlexE)
- **OIF chosen as best place to standardize**
 - Standardization process kicked off in January 2015 OIF meeting
 - Version 1.0 of the Implementation Agreement (IA) finalized earlier this year

FlexE Overview

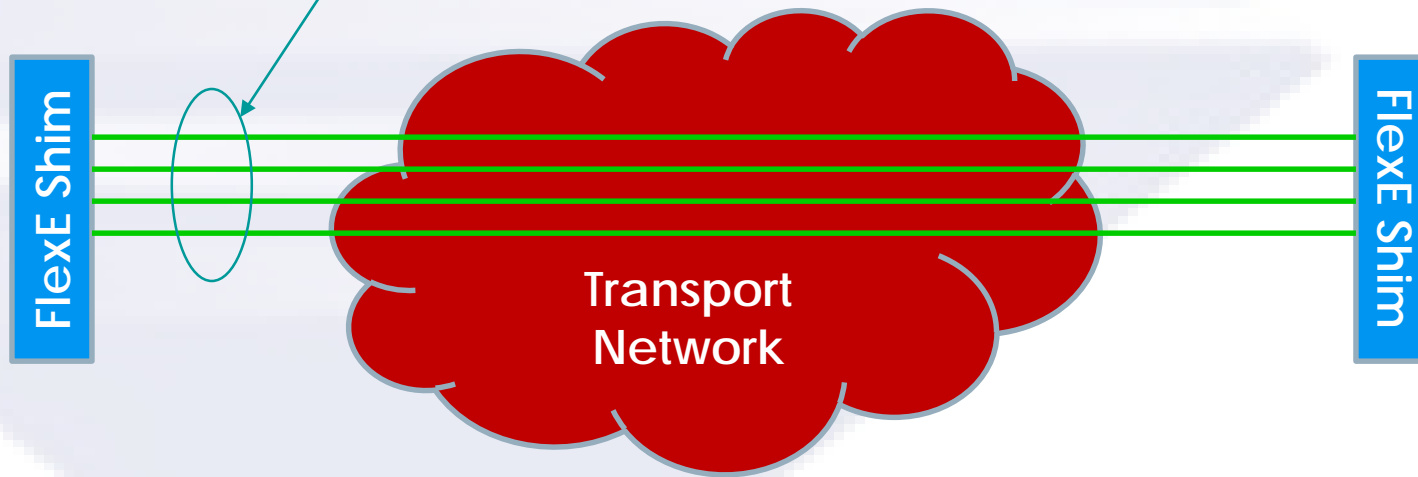
- **FlexE 1.0**
 - Provides a mechanism to map one or more FlexE clients over a group of 100 Gb/s Ethernet PHYs
 - The client rates supported are 10G, 40G, and $n \times 25G$
- **General Structure of FlexE**



FlexE 1.0 Overview

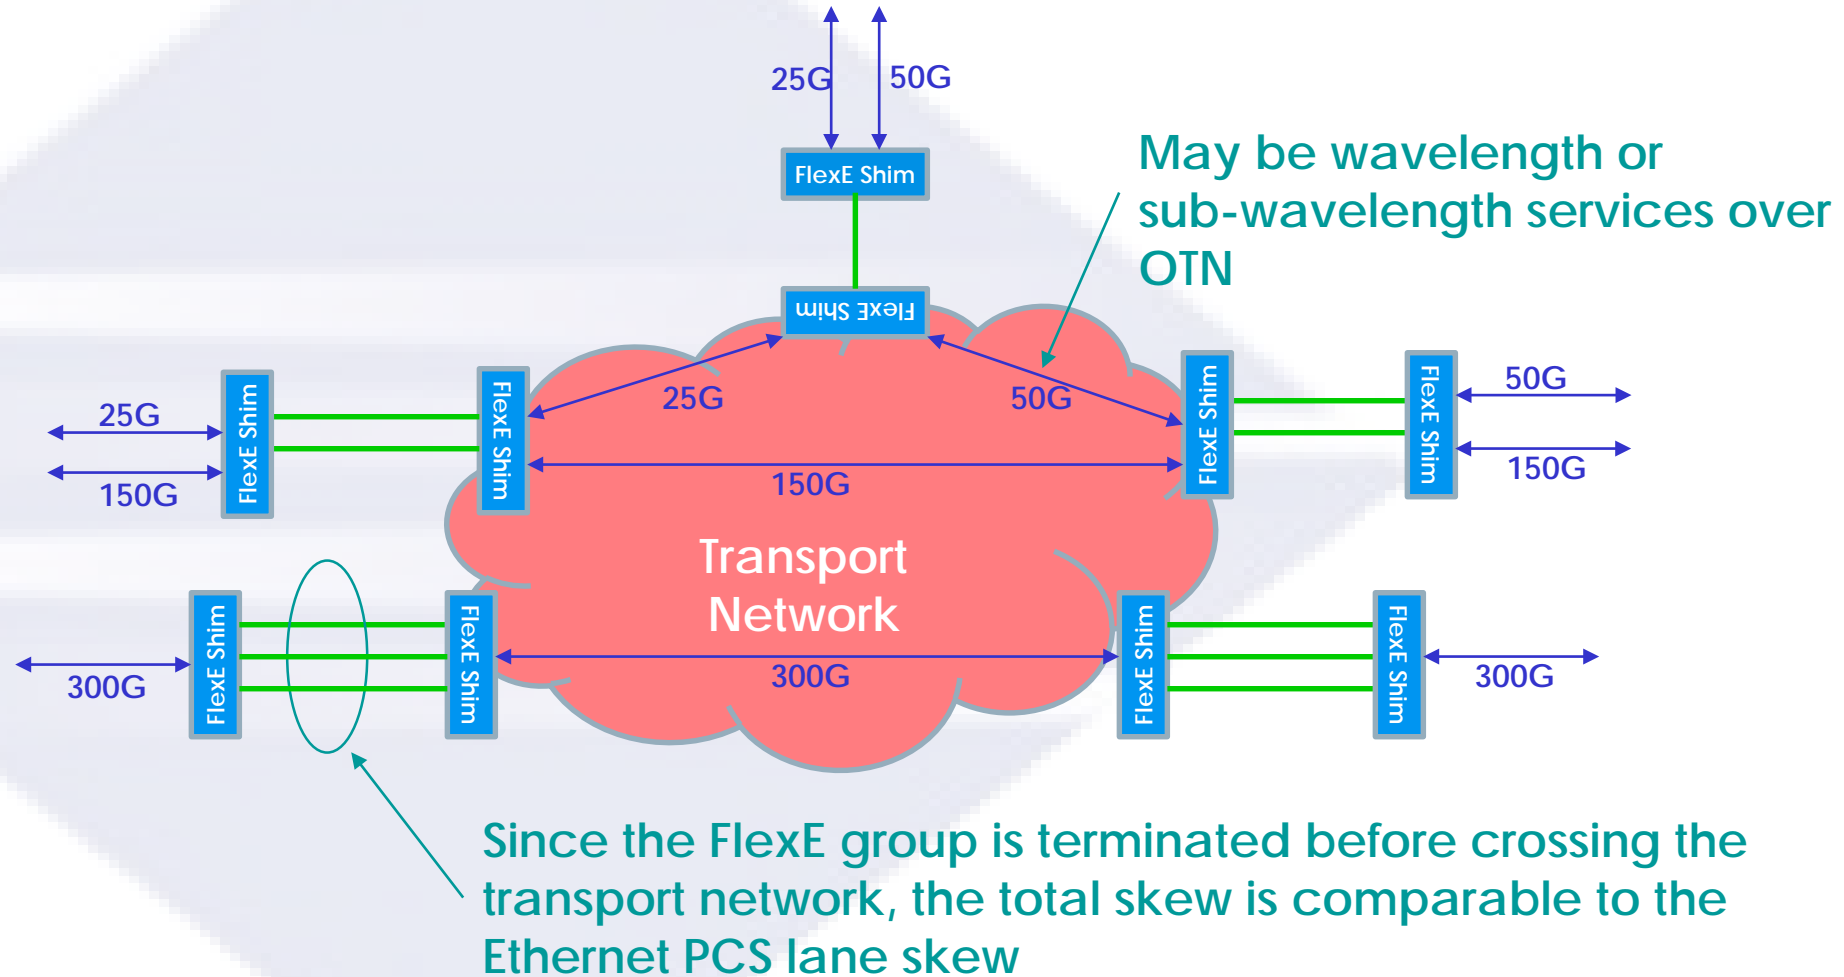
FlexE unaware transport use case

All PHYs of the FlexE group are carried independently, but over the same fiber route, over the transport network. Deskew across the transport network is performed in the FlexE shim



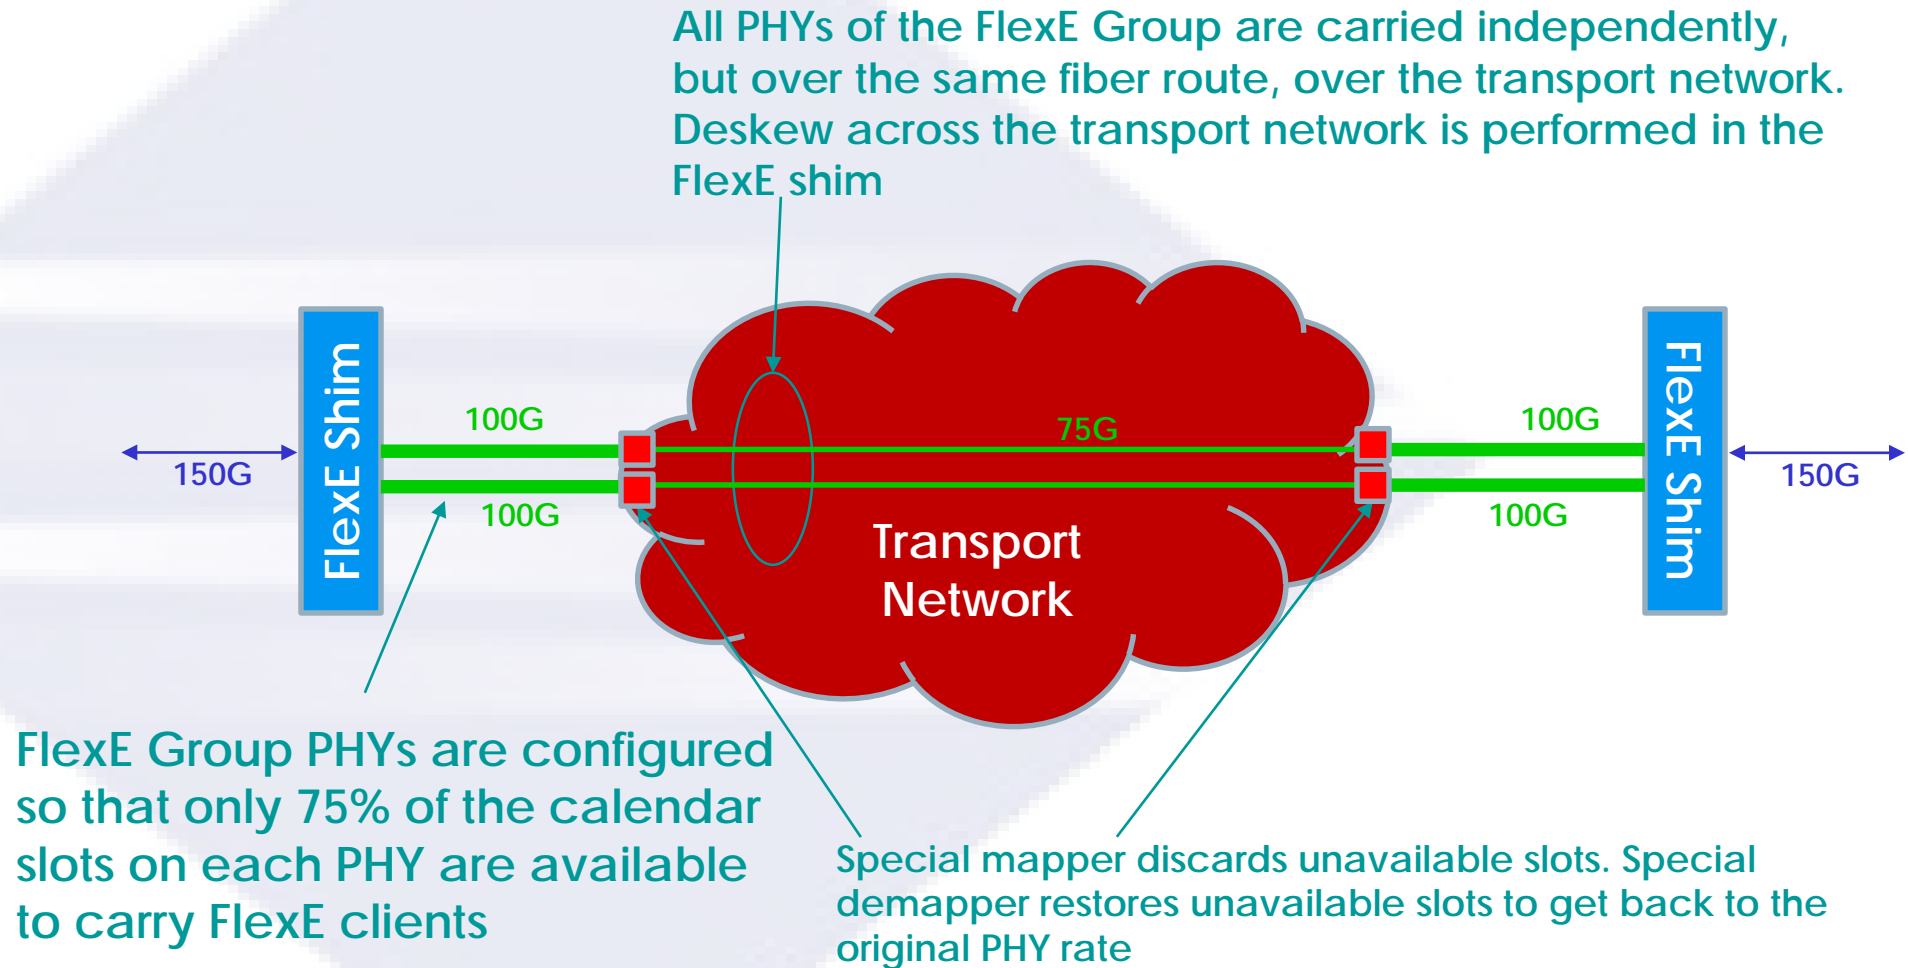
FlexE 1.0 Overview

FlexE terminating transport use case



FlexE 1.0 Overview

FlexE aware transport use case



FlexE Next Steps

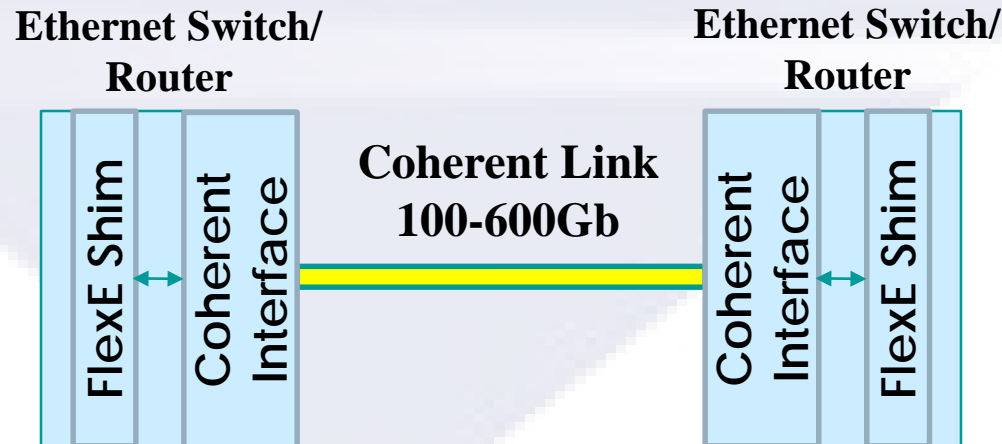
□ FlexE 2.0

- Project start proposal expected at Q4 meeting
 - Adds support for nx200Gb/s and nx400Gb/s FlexE groups
 - More detail on use of FlexE management channels
 - Consider coarser calendar granularity to reduce gate count for high bandwidth devices
 - Management of skew for specific applications
 - Transport of frequency or time by the FlexE group

Microsoft's Perspective

Additional Considerations

- Two different use cases: outside and inside the data center
- Outside the data center
 - Next generation coherent interfaces will support a range of baud rates and constellation sizes to maximize capacity with available OSNR. This will result in many different supported data rates with the same hardware.
 - Need to “squeeze” every possible bit out of these links and the Ethernet rates need to match the available capacity
 - FlexE granularity needs to match switching silicon SERDES granularity ie 50Gb starting in 2018
 - A 350Gb coherent link would need a matching 350Gb MAC rate



CAL
RNETWORKING
UM

Additional Considerations

□ Inside the data center

- Desire Ethernet links larger than available client rates
- Could be supported with LAG but not ideal
- Granularity can be limited to a subset of Nx values of the SERDES rates of available switching silicon
 - 50Gb x1 (50Gb), x2 (100Gb), x4 (200Gb), x8 (400Gb), x16 (800Gb) and x32 (1,600Gb)

