

High-speed data options from warm board to RCEs


Matt Graham
April 21, 2016

Intro/overview/opinions/disclaimers



- I'll talk about the data path from (after) the cold FPGA/COLDATA to the RCE
- Two options on table: multiplex on the flange or not
- The RCE system is flexible enough that, from that narrow viewpoint, the details of the warm interface don't matter much
- That said, from a system view, I (Matt Graham) am not sure that it makes a lot of sense (some other people agree)
- I'll show what the two options look like from a data path view and touch on what the warm interface will need to do if it touches the data

The (RCE) preferred scheme...

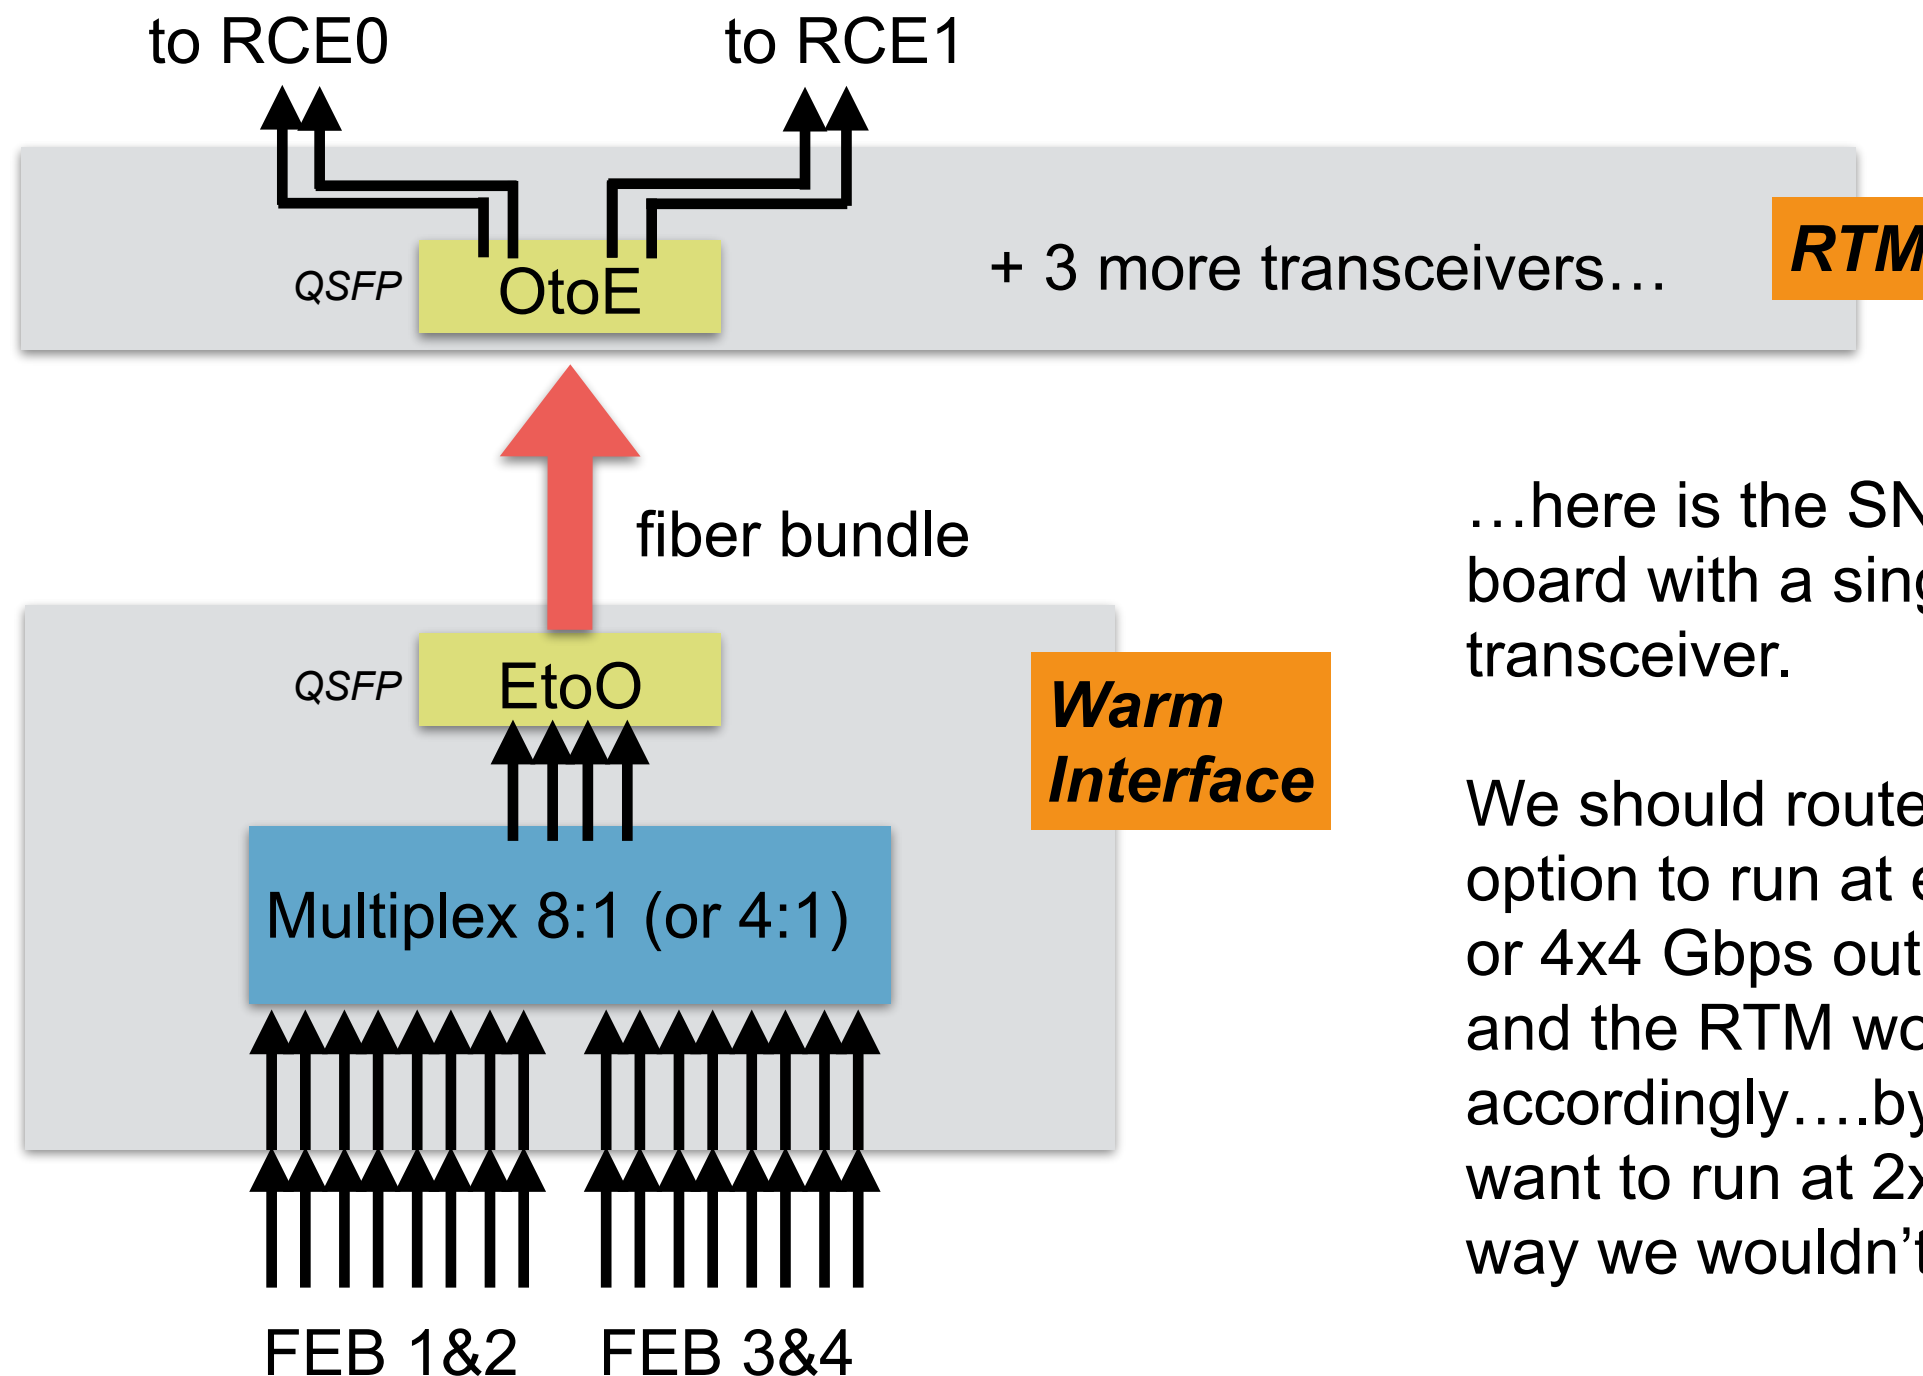


This follows the spirit of 35ton, where the RCE FPGA is the first “smart” element to see the data...

Fewest interfaces/firmware/chances for failures to creep in.

Also gives a lot of flexibility in links/RCE...we can reduce if we need more bandwidth/computing (reduces risk) or increase if we want to test potential DUNE configuration (value engineering). (Manager speak!)

Multiplexed-to-RCE scheme




...here is the SNBD-like warm board with a single QSFP transceiver.

We should route it so there is an option to run at either 2x8 Gbps or 4x4 Gbps out of the WIB FPGA and the RTM would be routed accordingly....by default, we'd want to run at 2x8 Gbps. That way we wouldn't have to re-align.

What the concentrator has to do when it receives raw data...


First time we see the data out of the cold, we need to align it and (in the process) check for missing beginning & end of frames...

Alignment



Error detection and "correction"


the quality of the high-speed links depend a lot on the quality of the clock (we learned all to well in 35 ton)



Data formats from a multiplexed system to RCE

If we the data does get multiplexed and serialized again through the warm board, we need to be careful about the format so that the de-multiplexing doesn't take many cycles. The minimal thing is to just wrap up the 8 (or 4) aligned streams and send them on, keeping the lane order fixed.

At any rate, the RCE team would like ownership of the post-alignment firmware on the warm



Summary

- We can make either scheme work...the no-touching scheme is more flexible for numerology of the RCEs (a bit), safer (a bit), less maintenance (a bit), and easier to maintain (Matt G's opinions, again)
- If we decide on the warm FPGA, the data fiddling (after receiving) should be the DAQ groups' responsibly