Chip-to-Chip Communication (Interlaken) for Enterprise and Cloud

Sundeep Gupta, *Sr Director, SoC IP Engineering*
Ketan Mehta, *Director, SoC IP Product Marketing*

Agenda

- Market Applications
- Need for Chip-to-Chip Connectivity
- Interlaken as Protocol of Choice
- Case Study: Interlaken In Data Center Servers
- Case Study: High Performance Compute Clusters
- Introduction to Interlaken IP
- Key Features
- Configuration
- TX & RX Path Data Flow
- Implementation

Interlaken-LL: Market Applications & Case Study

HPC, Enterprise Networking and Cloud Applications

Data Intensive Applications

Networking

AI/ML

Data Center

HPC/Cloud

- Pushing the limits on data traffic going in and out of an SoC
- More data => IOs require high speed interfaces
- Disaggregation is the new trend : Multi-Die, Multi-Chip, Multi-CPU

"It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected"

- Gordon E. Moore, Intel

Need for Chip-to-Chip Connectivity

- **HPC and Data Center applications**
  - SoCs are reaching maximum reticle size limiting the scalability
  - Splitting chips due to either area, power, process yield or cost limitations

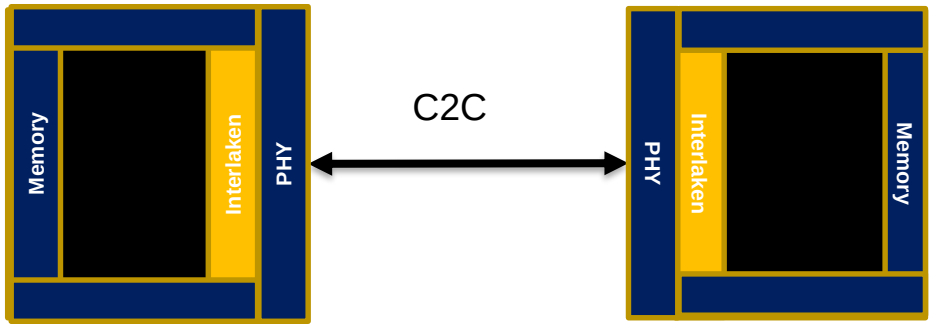

- **AI/ML and HPC applications**
  - Driving Cluster/Matrix architecture with distributed processing
  - Higher Bandwidth requirements

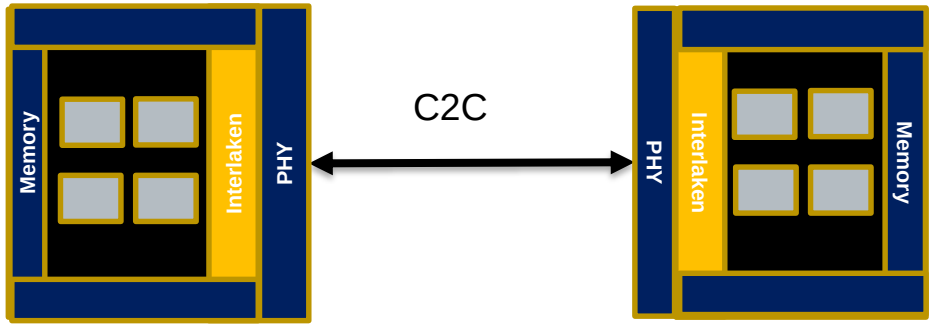

"Instead of the large devices and clusters, we disaggregated the network into small identical units – server pods – and created uniform, scalable, high-performance connectivity between all pods in the data center" – Facebook https://engineering.fb.com

Interlaken as Protocol of Choice

- **Narrow Interface**
  - Few IOs per each physical high speed lane
- **High Speed and Throughput**
  - PHY agnostic interface (For e.g. can support up to 112G/lane)
- **Channelized Packet Interface**
  - Allows up to 64K virtual channels with QoS
- **Highly Scalable**
  - Supports from 1 lane up to 48 lanes
- **Low Latency**
  - Low protocol overhead and modular design allowing customized latency

Interlaken with Low Latency

- For NIC accelerator, HPC and NAND Flash controller applications, latency is very critical and customers are requesting latency of 20 to 40 cycles end to end
- SiFive Low Latency Interlaken provides ~10-20%* latency improvements in both directions
  - Further customization is possible in various modules inside Tx/Rx blocks
  - For more information please visit: https://www.sifive.com/soc-ip/interlaken

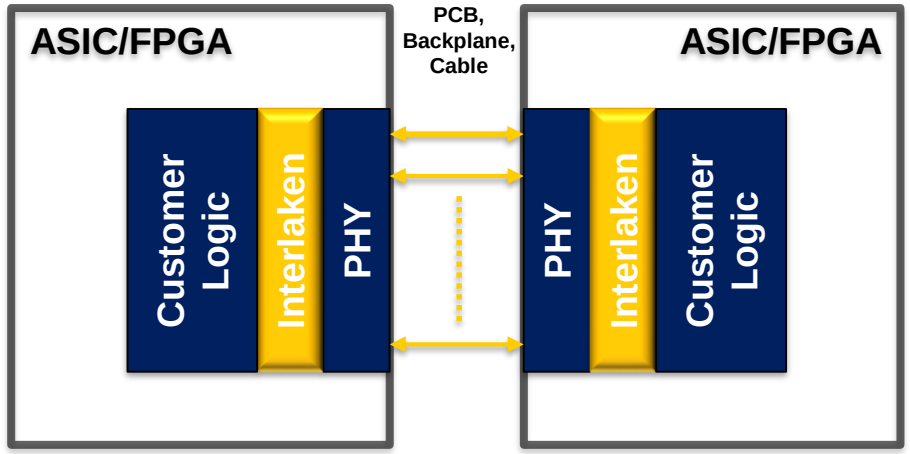

High Speed & High Bandwidth
Chip to Chip Communication Interface

*Certain configurations in cut-through mode w.r.t. legacy Interlaken IP

Interlaken In Data Center Servers

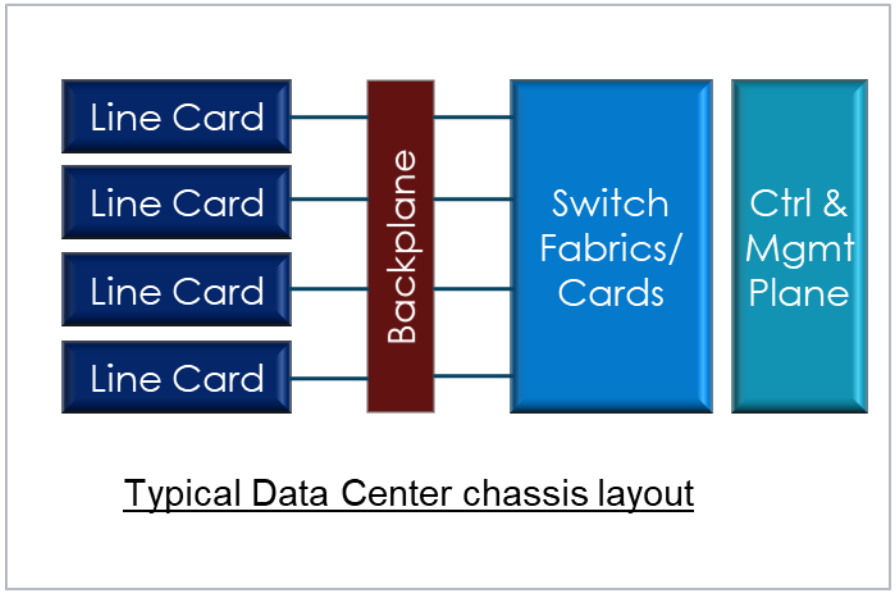

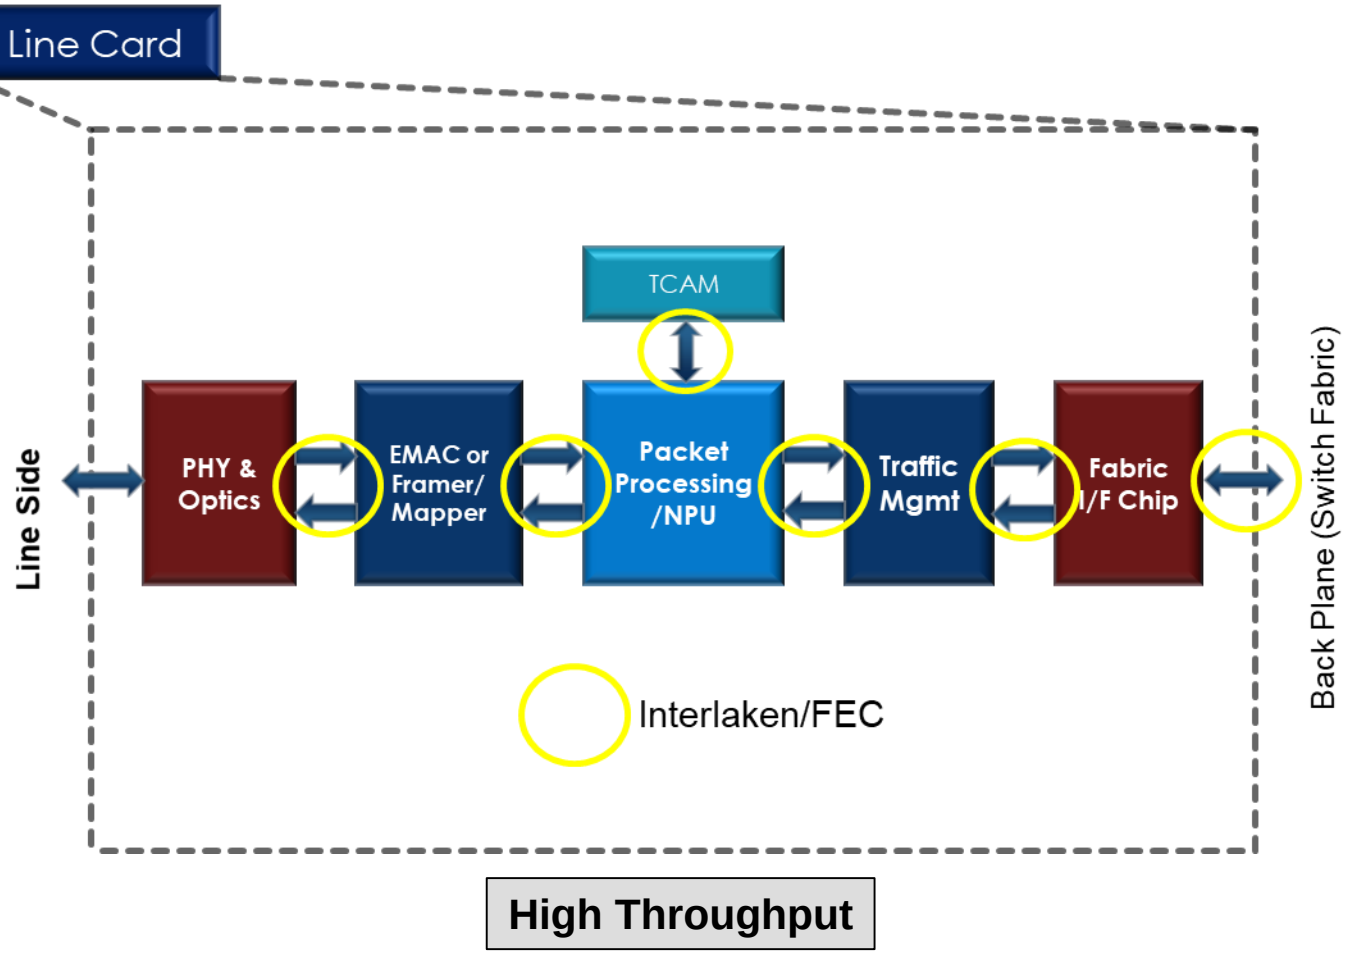

High Throughput

Other Applications

- Packet Processing/NPU
- Traffic Management
- Switch Fabric Interface
- Framer/Mapper
- Serial Memory (INLK-LA)

Interlaken in High Performance Compute Clusters

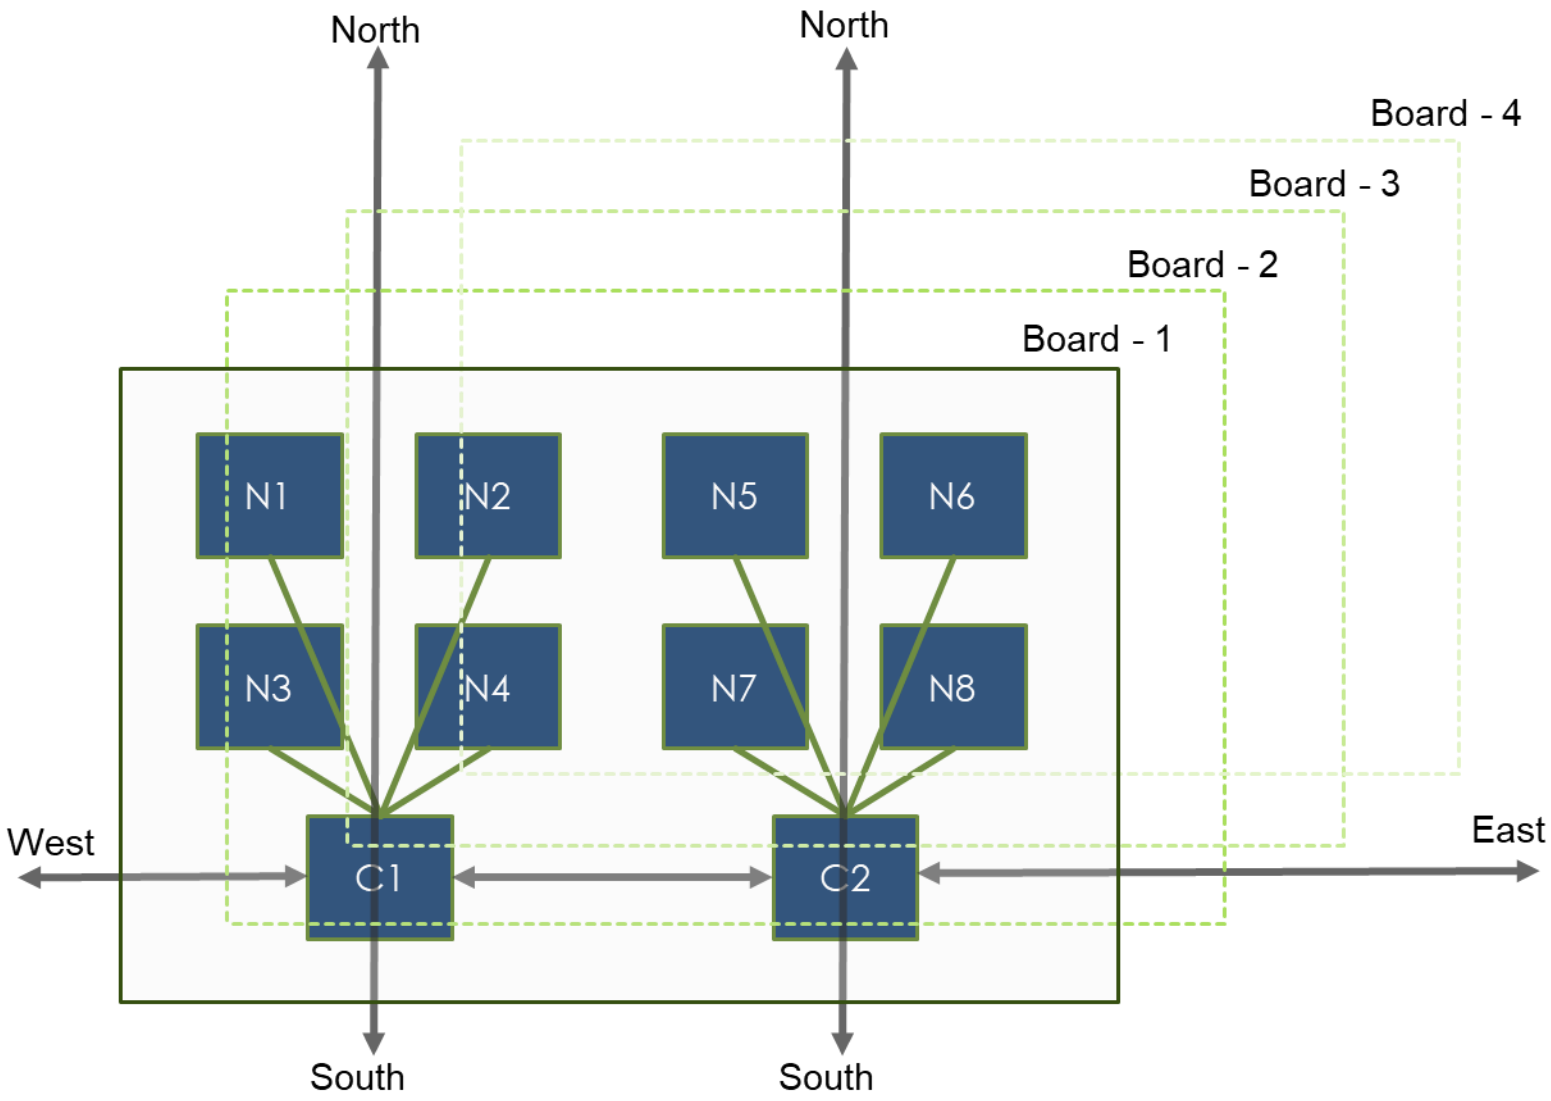

Scalable Architecture

Interlaken Product Portfolio

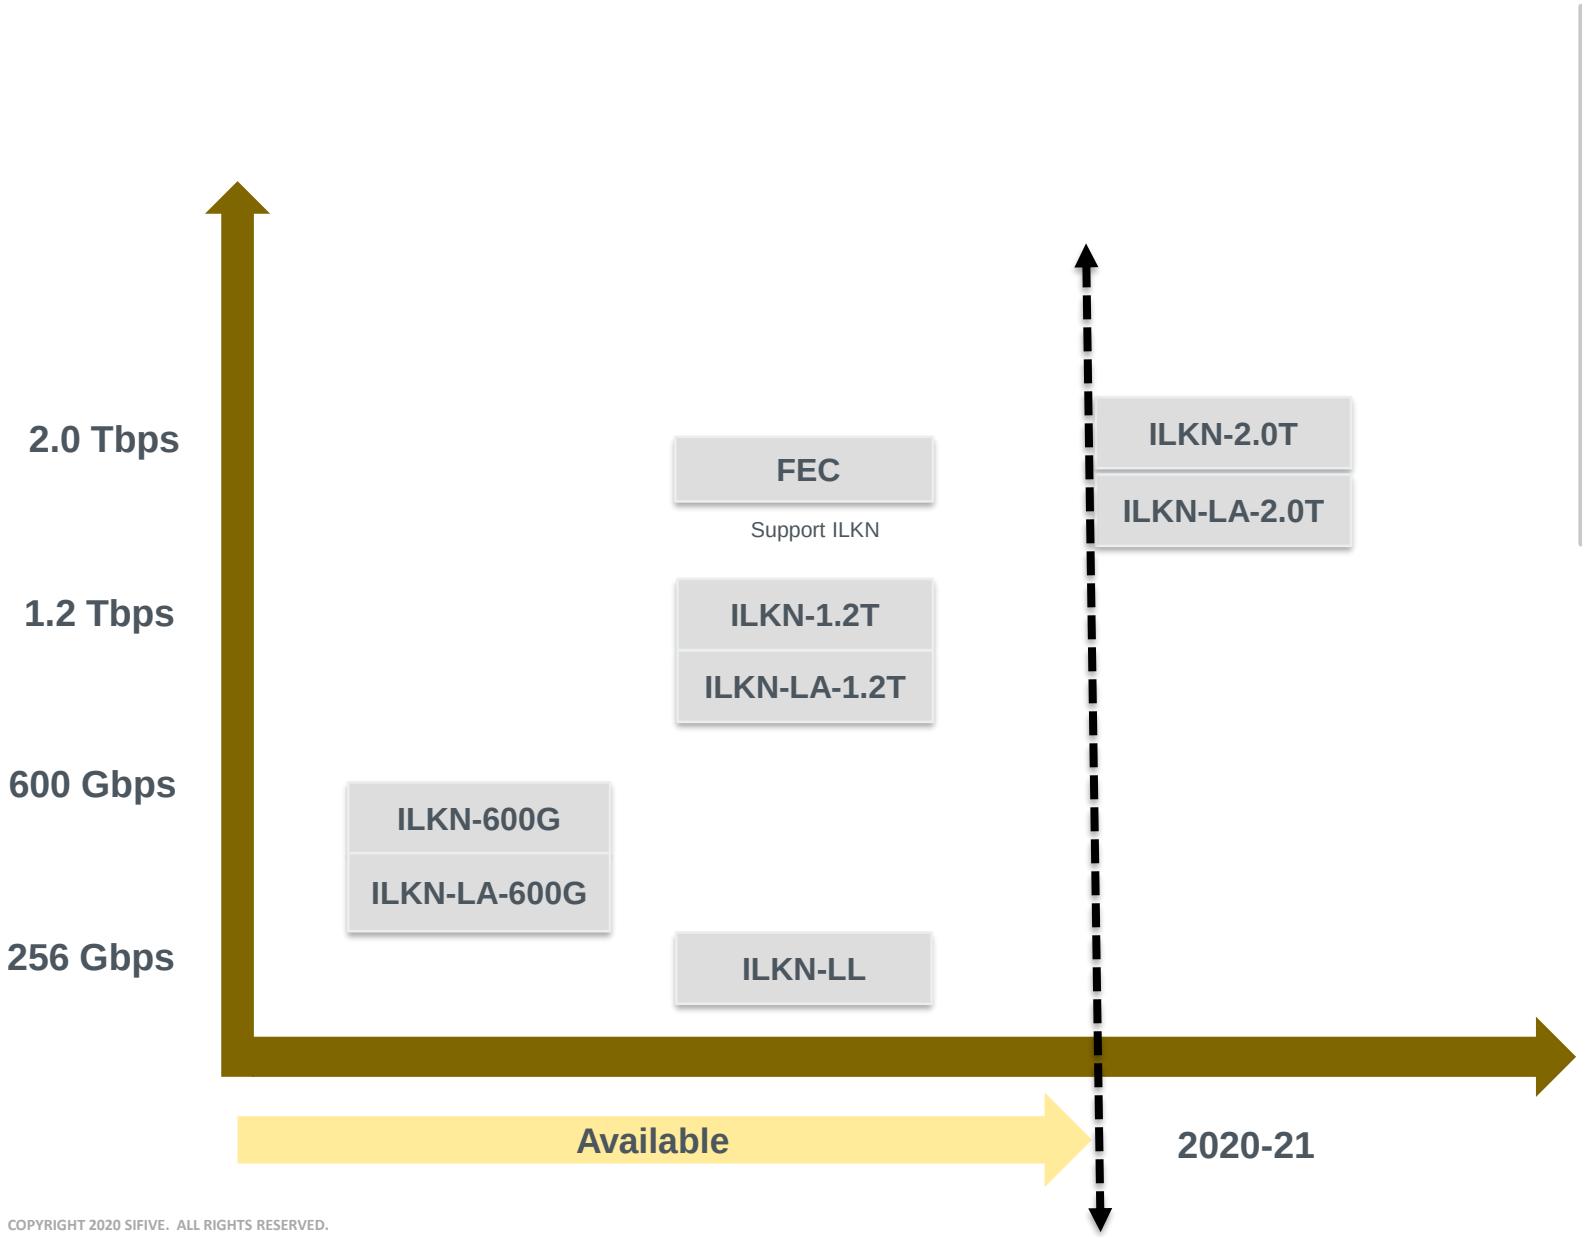

FEC	Forward Error Correction
ILKN	Interlaken
ILKN-LA	Interlaken Look Aside
ILKN-LL	Interlaken Low Latency

SiFive Interlaken IP

- Specifications supported: www.interlakenalliance.com
- The Interlaken IP supports the following Interlaken Alliance specifications:
  - Interlaken Protocol Definition, v1.2
  - Interlaken Look-Aside Protocol Definition, v1.1
  - Interlaken Retransmit Extension, v1.2
  - Interlaken Dual Calendar Extension v1.0
  - Interlaken Interop Recommendations, v1.7
- SiFive Interlaken IP is available in 2 broad flavors:
  - High-Bandwidth (up to 1.2Tbps) Interlaken Core
  - Low-Latency (↓10-30%) Interlaken Core.

SiFive High-Bandwidth Interlaken IP – Key Features

- 1.2Tbps high-bandwidth performance
- Up to 48 SerDes lanes with custom width support
  - Scalable up to 1.2Tbps with SerDes rates from 3.125Gbps to 56G PAM4.
- KP4 FEC support.
- Flexible User Interface options
  - 128b: 1x128b, 2x128, or 4x128b
  - 256b: 1x256b, 2x256b, 4x256 or, 8x256b
- Support for 256 (extendible up to 64K) logical channels
- Lane resiliency support (independent Lane enable/disable)
- Lane bifurcation support (1:2 and 1:4)
- Interlaken Retransmit Extension support
- Interlaken-LA protocol support
- Simultaneous In-band and Out-of-Band flow control with programmable dual calendar
- Built-in error detection/injection for testability

SiFive Interlaken IP Core – Configurability - One Segment User Interface

User Interface bandwidths supported

- 1x256 User I/F for up to 150 Gbps
- 1x128 User I/F for up to 100 Gbps

1 x 128 bit
1 x 256 bit
2 x 128 bit
2 x 256 bit
4 x 128 bit
4 x 256 bit
8 x 256 bit

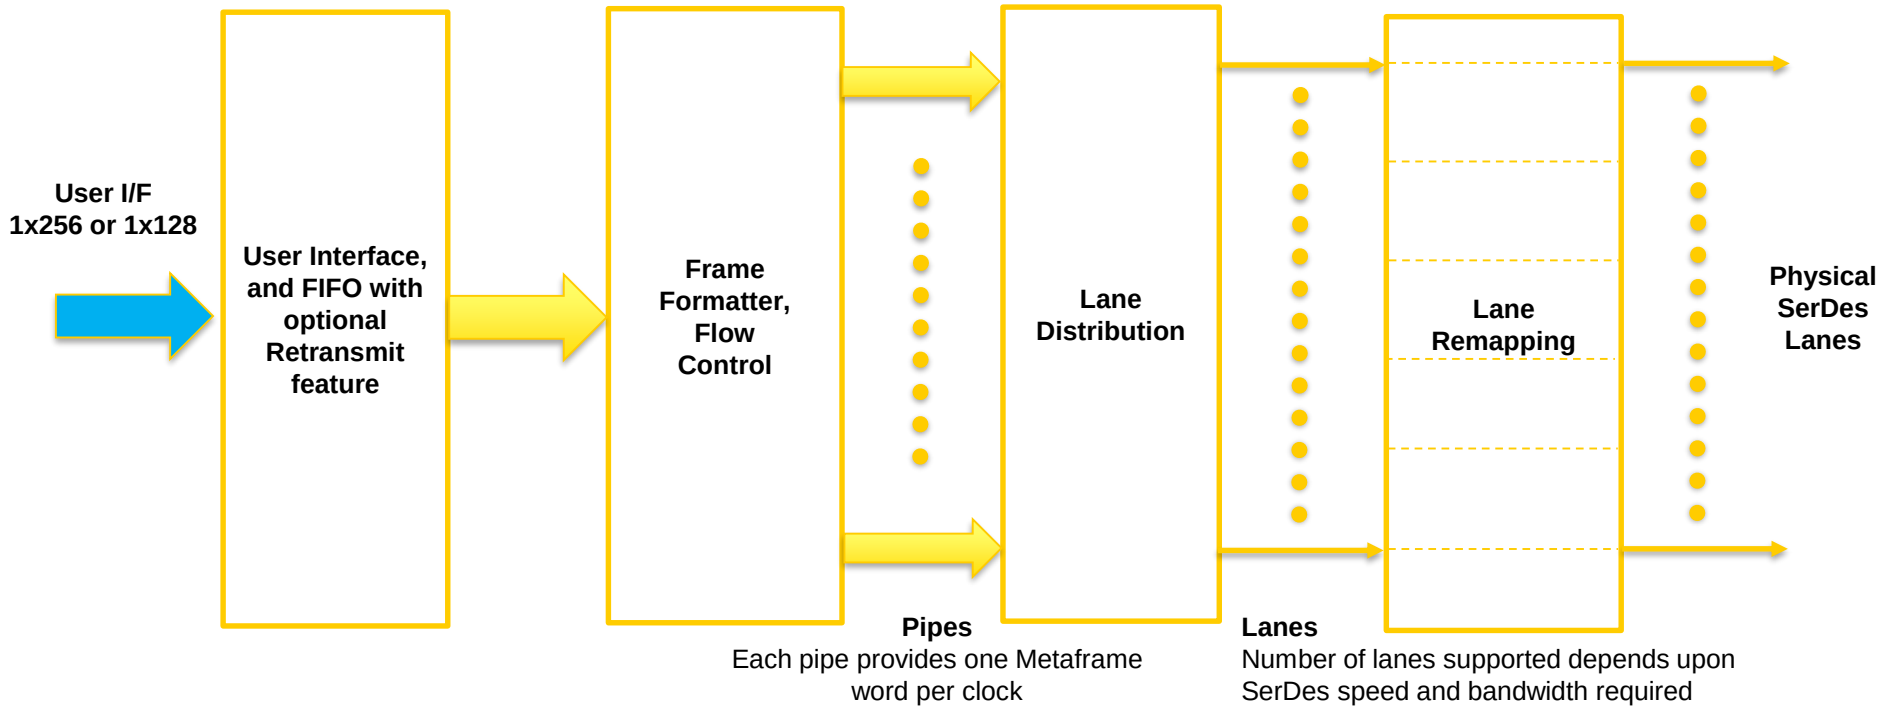

SiFive Interlaken IP Core – Configurability - Two Segment User Interface

User Interface and bandwidths supported
- 2x256 User I/F for up to 300 Gbps
- 2x128 User I/F for up to 200 Gbps

1 x 128 bit
1 x 256 bit
2 x 128 bit
2 x 256 bit
4 x 128 bit
4 x 256 bit
8 x 256 bit

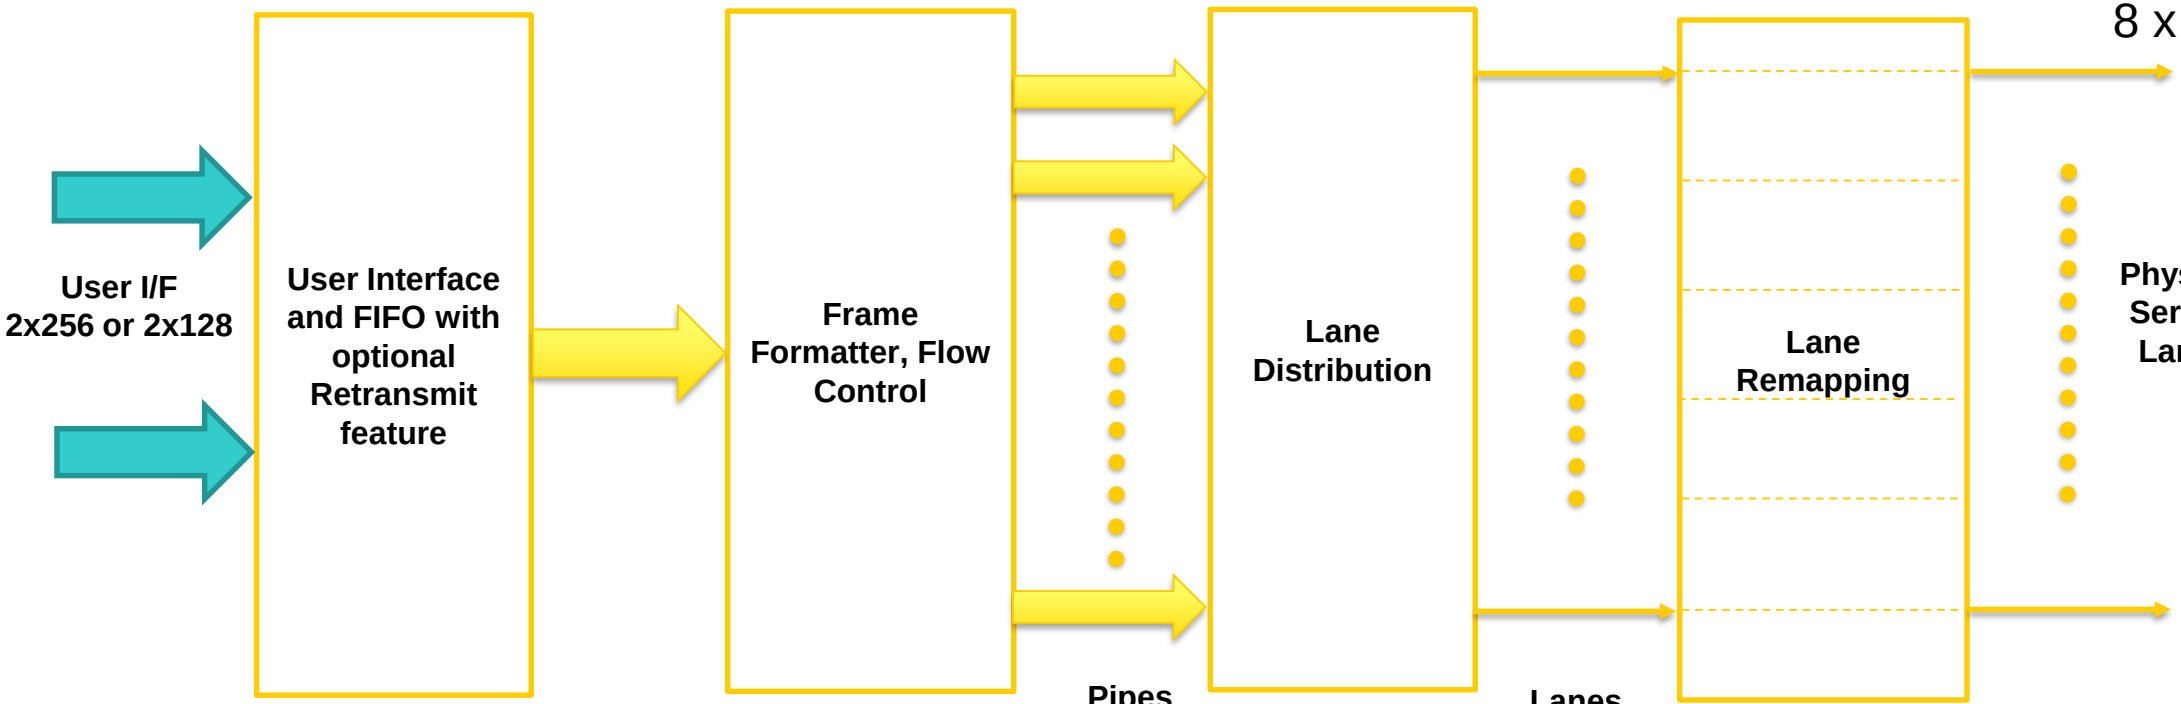

Pipes
Each pipe provides one Metaframe word per clock

Lanes
Number of lanes supported depends upon SerDes speed and bandwidth required

SiFive Interlaken IP Core – Configurability - Remapping Lanes

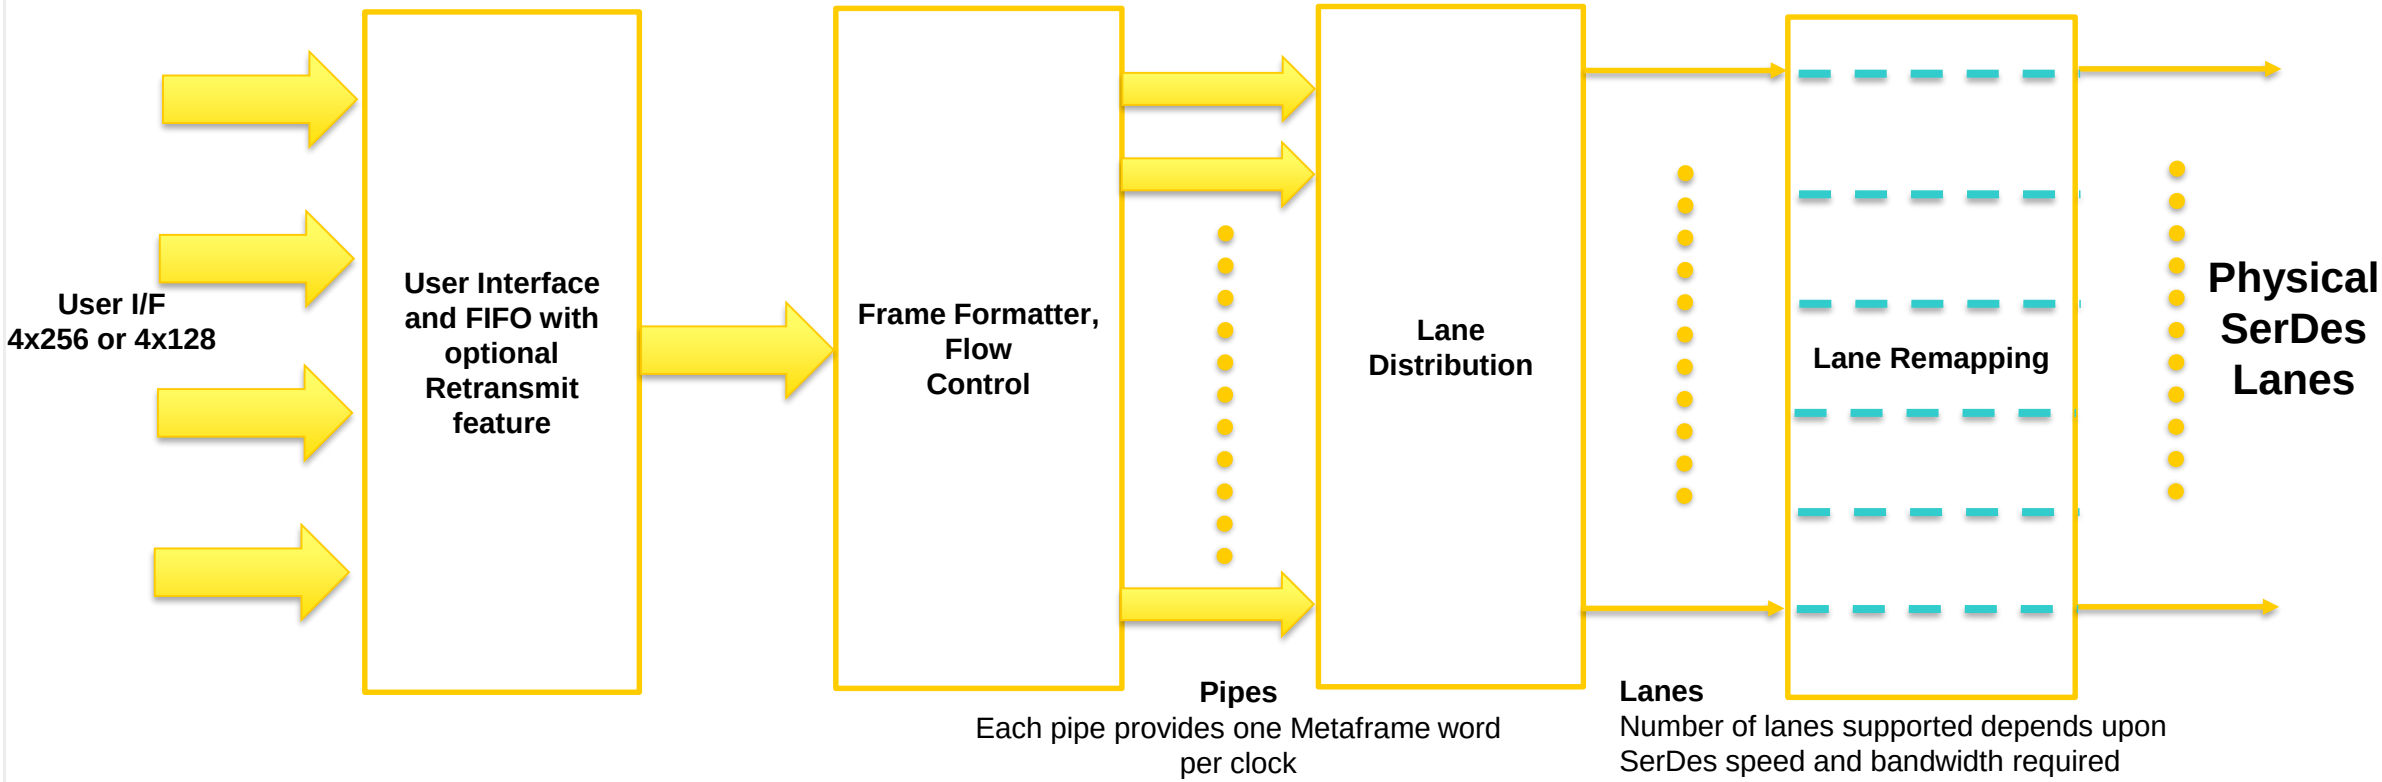

SiFive Interlaken IP Core – Configurability - Remapping Lanes

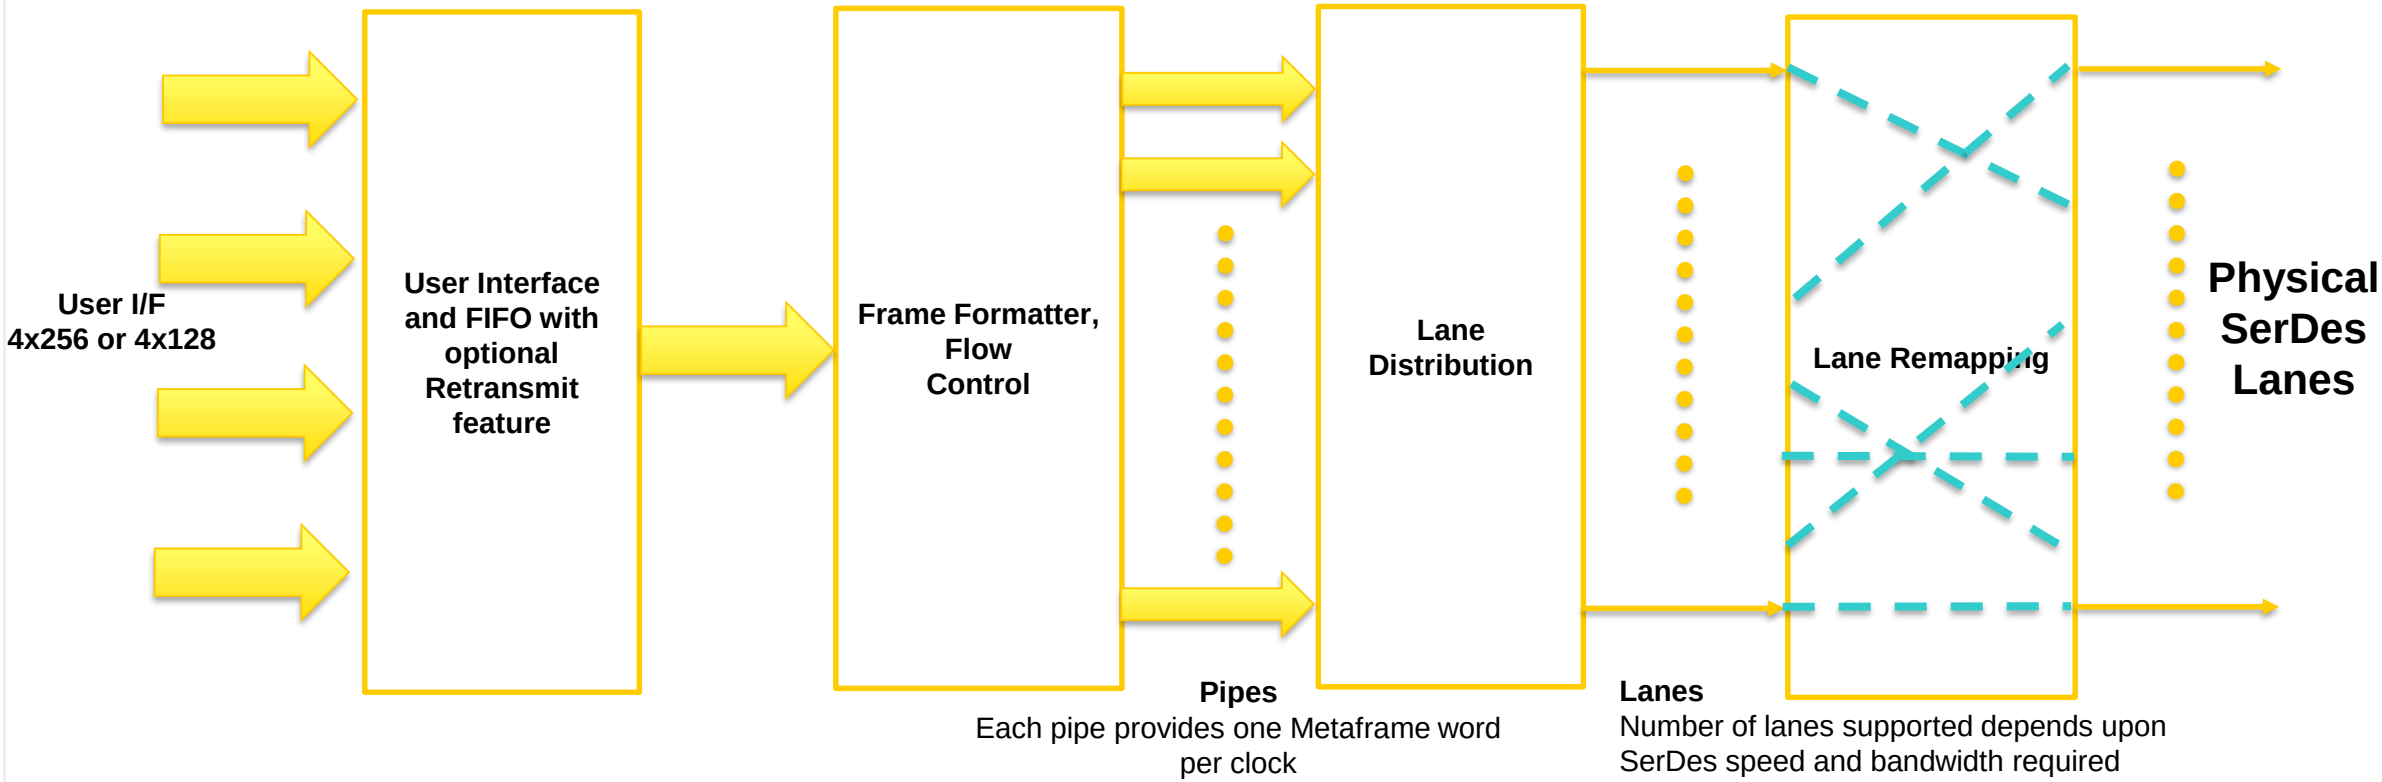

SiFive Interlaken IP Core – Configurability - Multi Core Configuration – Single Core

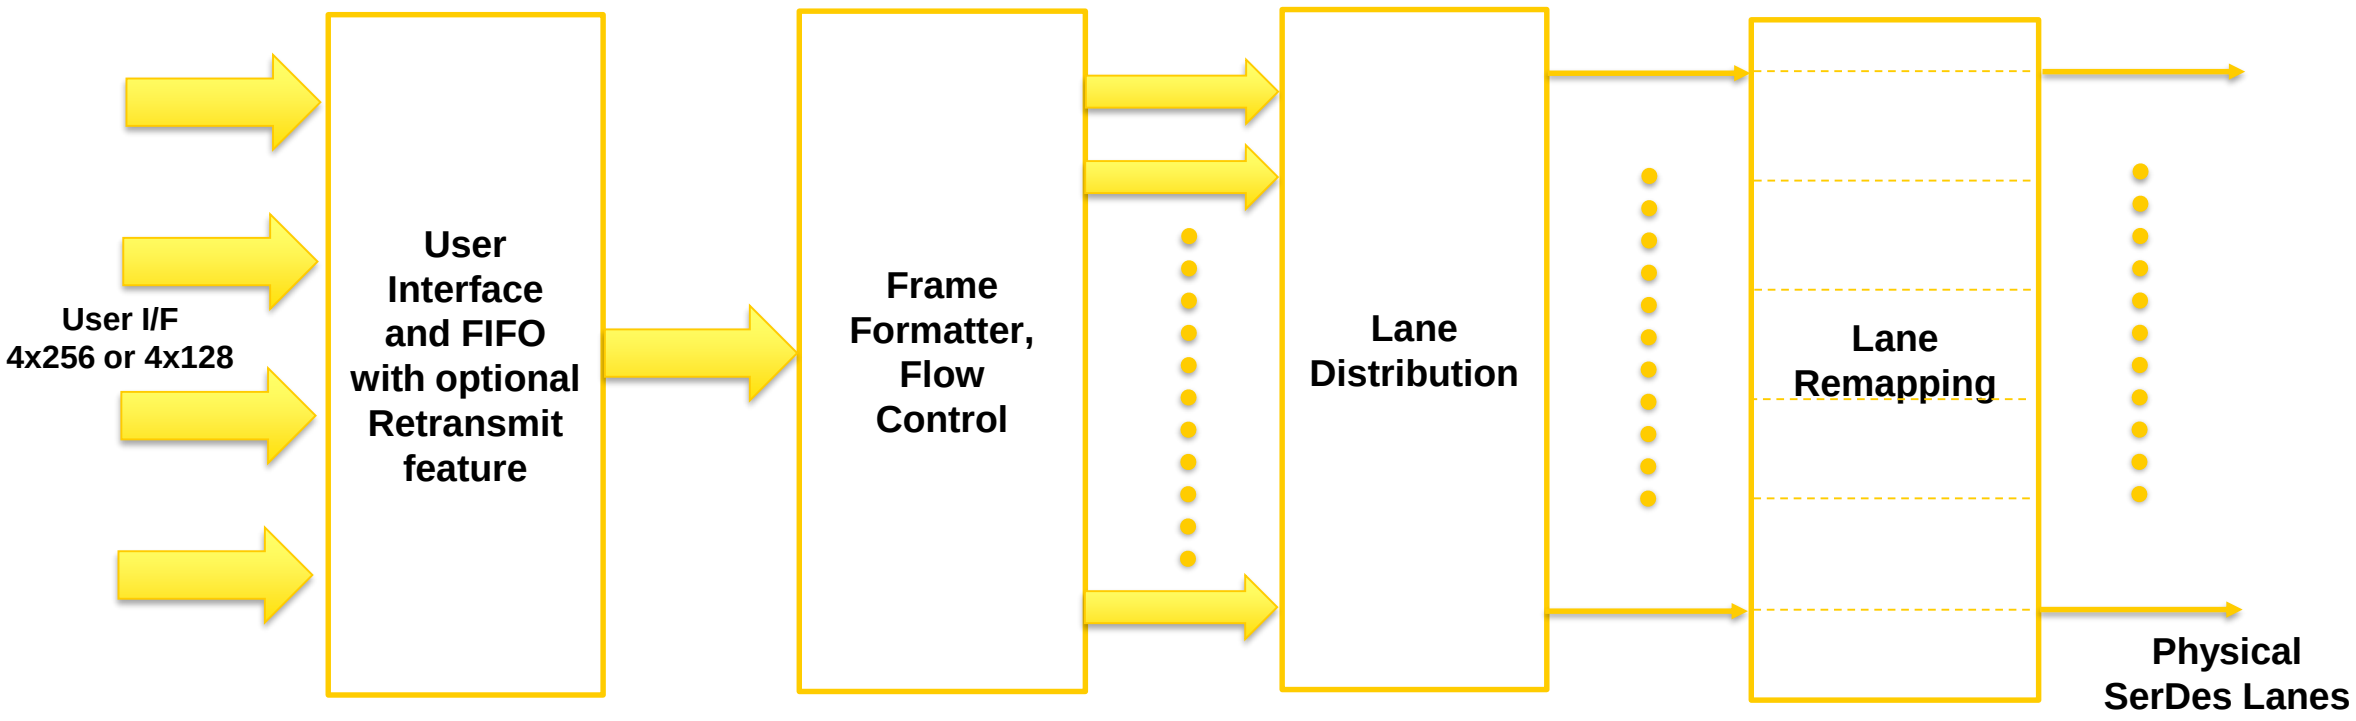

SiFive Interlaken IP Core – Configurability - Multi Core Configuration – Dual Core

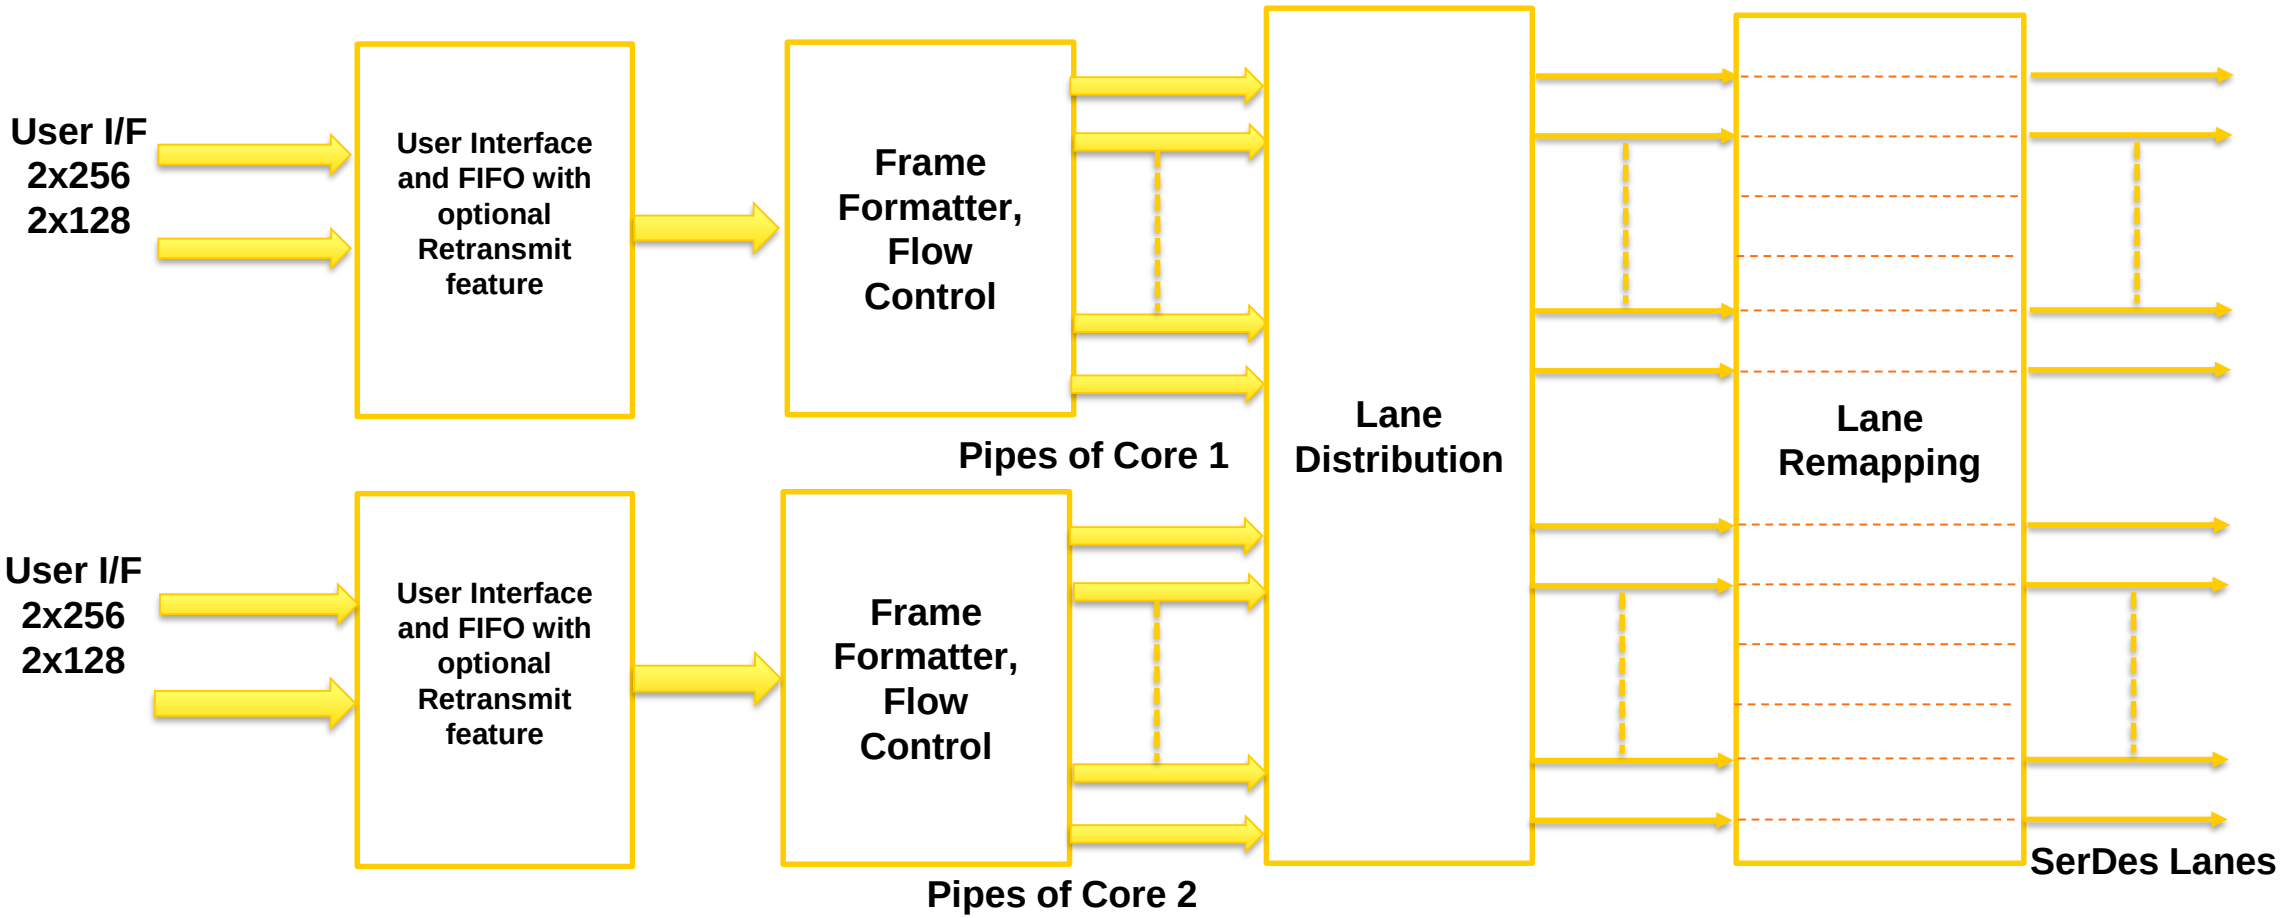

Lane Bifurcation

- A single Interlaken IP instance can be configured in-system to support different Interlaken interfaces
- Following example shows 1x1.2Tbps, 2x600Gbps or 4x300Gbps, leading to a more area efficient and flexible implementation

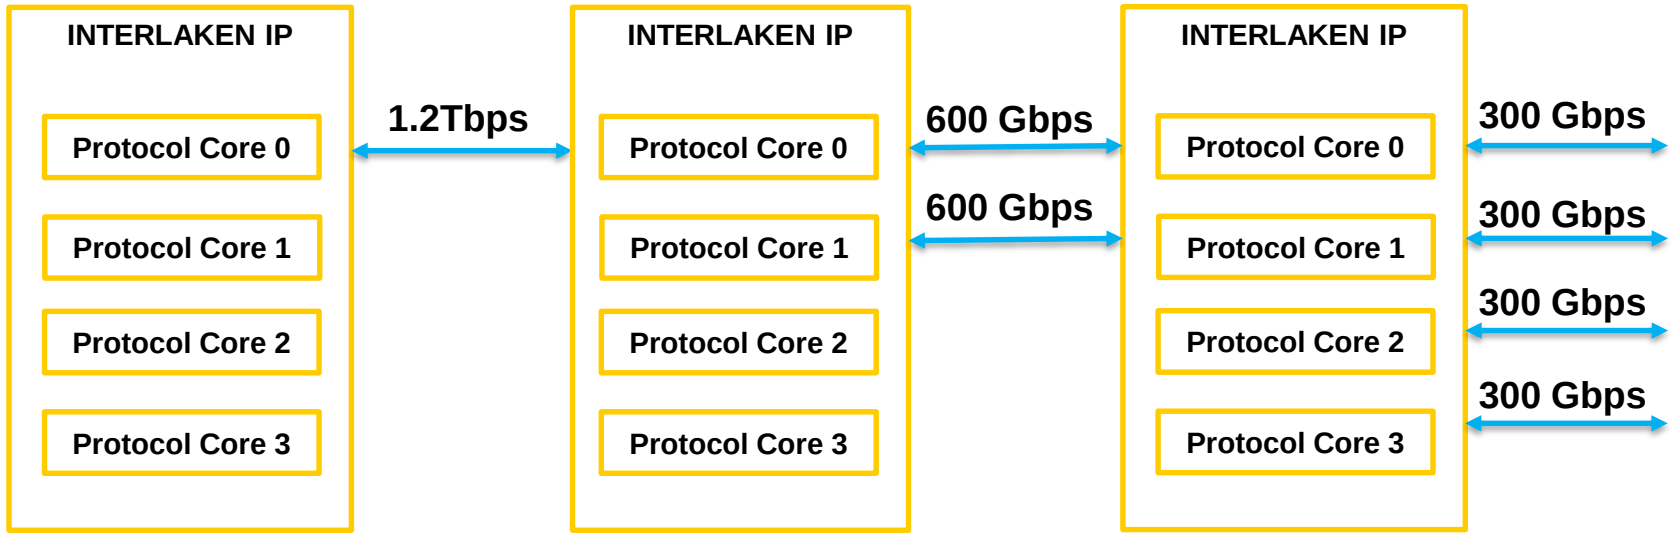

SiFive Interlaken Low Latency IP – Key Features

- Up to 30% latency reduction as compared to High-Bandwidth ILKN core.
- Up to 256 Gbps bandwidth support
- Up to 8 SerDes lanes with custom width support
  - Scalable up to 256Gbps with SerDes rates from 3.125Gbps to 32Gbps
- Flexible User Interface options
  - 64b: 1x64b, 2x64b, or 4x64b
- Support for 256 (extendible up to 64K) logical channels
- Lane resiliency support (independent Lane enable/disable)
- Lane bifurcation support (1:2 and 1:4)
- Interlaken Retransmit Extension support
- Interlaken-LA protocol support
- Simultaneous In-band and Out-of-Band flow control with programmable dual calendar
- Built-in error detection/injection for testability

Legacy Implementation of Retransmit

- The legacy implementation of RT works on the “Store and Forward” principle
- Incoming burst is fully stored in the RX RT buffer and CRC24 check is done
- This mechanism is followed for every incoming burst

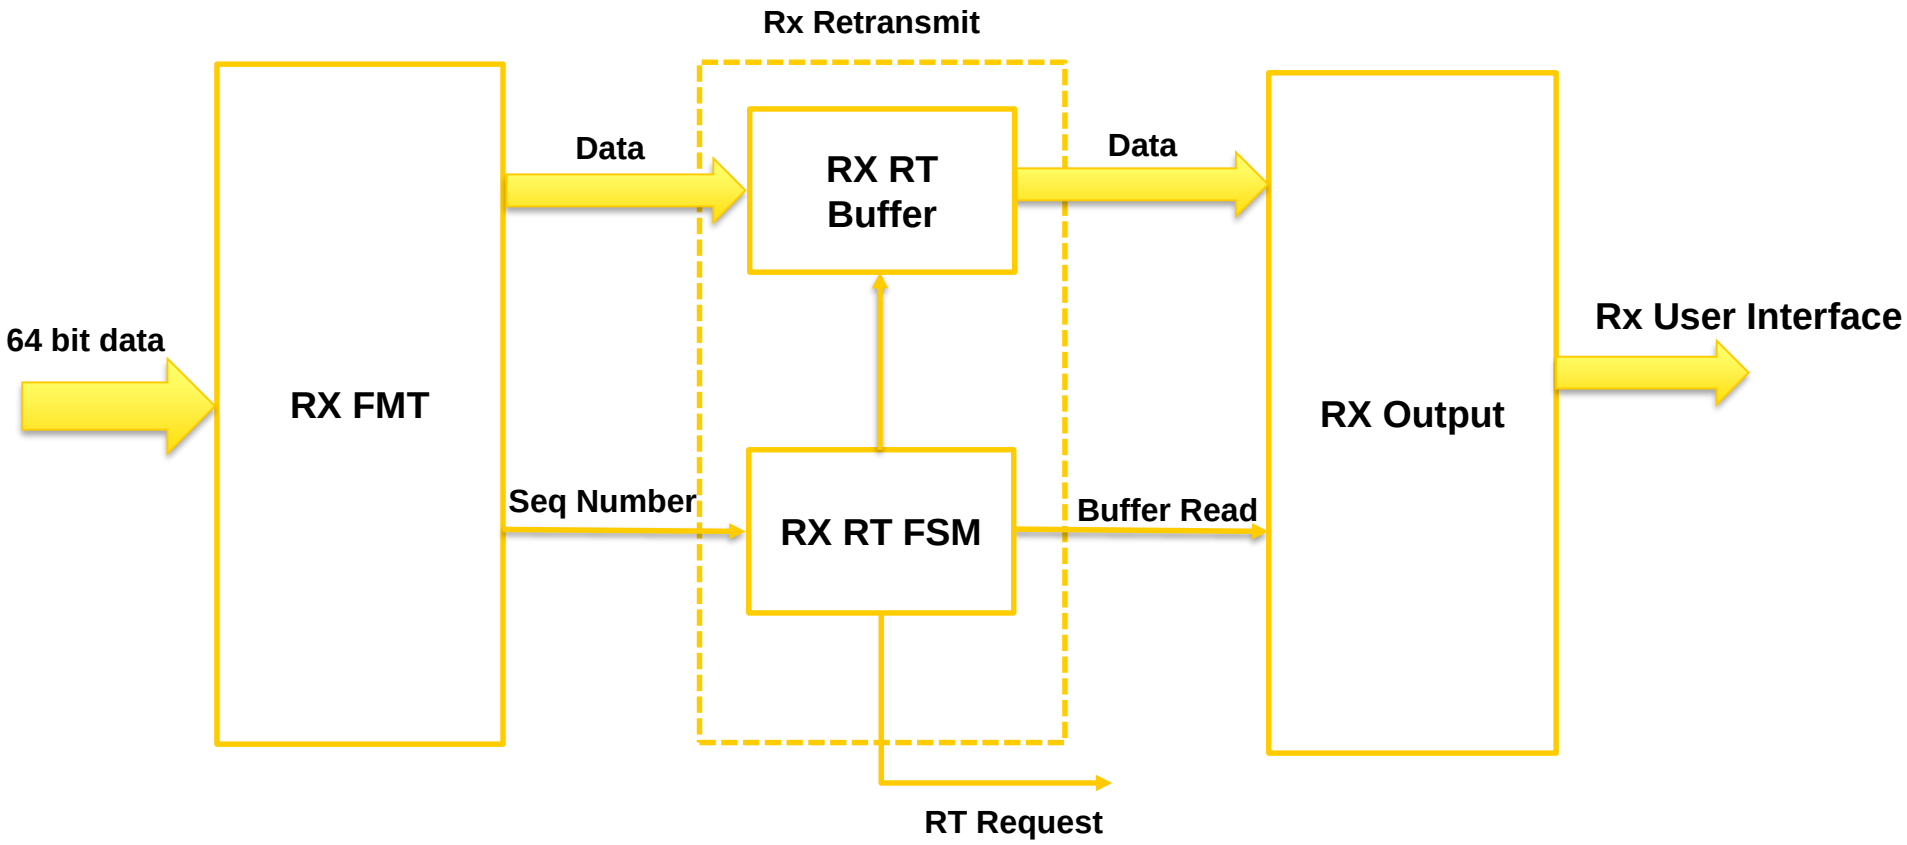

Cut-through Implementation of Retransmit

- The figure above shows the improvement over the “Store and forward” scheme.
- The incoming burst is directly forwarded to the RX UI. If the CRC24 error is seen the “burst_ignore” indication is sent over the RX UI till next good burst is received

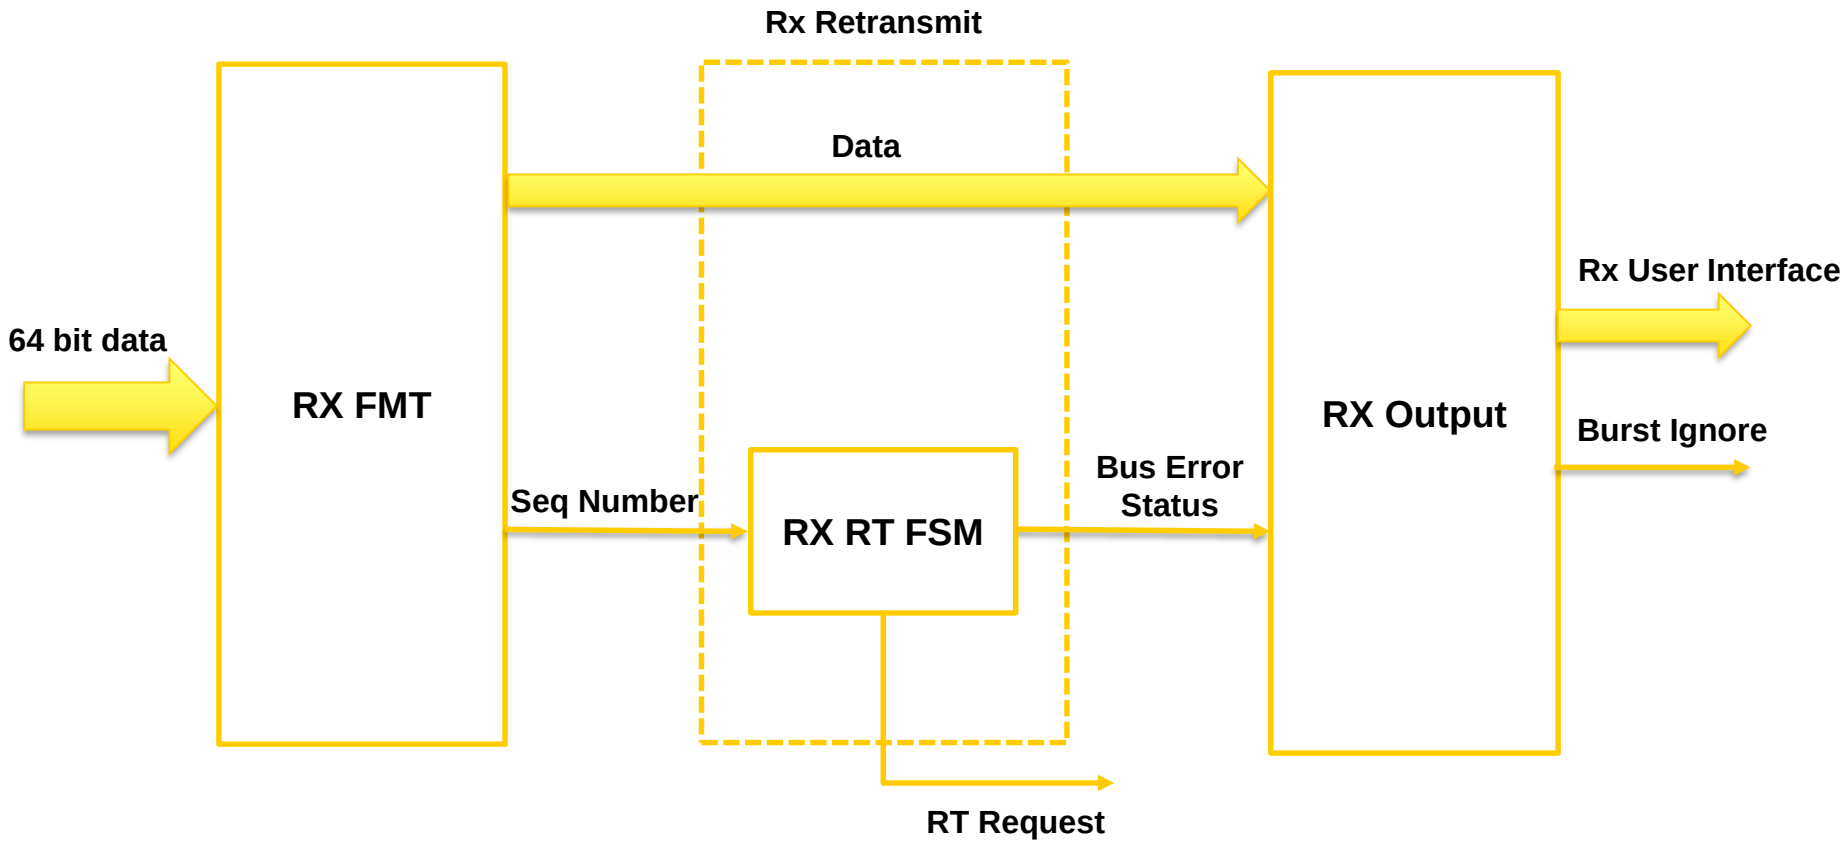

Deliverables

- Project Control Document configuration file (filled by customer)
  - The ILKN core has two types of configurations:
    - Hardware configuration is captured through a “Project Control Document”
    - Some items are software configurable through the memory mapped registers
- Documentation: -
  - Interlaken IP Product Brief
  - Interlaken IP Specification (NDA)
  - Memory-Mapped Register Manual
  - Design Verification Plan
- Synthesizable RTL
- Scripts for synthesis and static timing
- Assertions for the User-Interface and Config Registers
- Test simulation environment

Thank You