



AMD Together We Advance Financial Services

Alastair Richardson
Global Strategic Business Development AMD

The Alveo™ X3 Series

Accelerating Your Trading Strategies

Plug-and-Play Low Latency NIC

- Deterministic latency for reliable trade execution
- Plug-and-play with existing SW Stack

CPU Offload for Hybrid Trading

- Offload compute-intensive functions closer to wire (NIC)
- Improve hit-ratio, transaction cost-analysis (TCA)

Adaptable Accelerator for Custom Solutions

- Develop specialized Fintech solution in FPGA logic
- Customize in C/C++ or design in RTL for ultimate HW flexibility



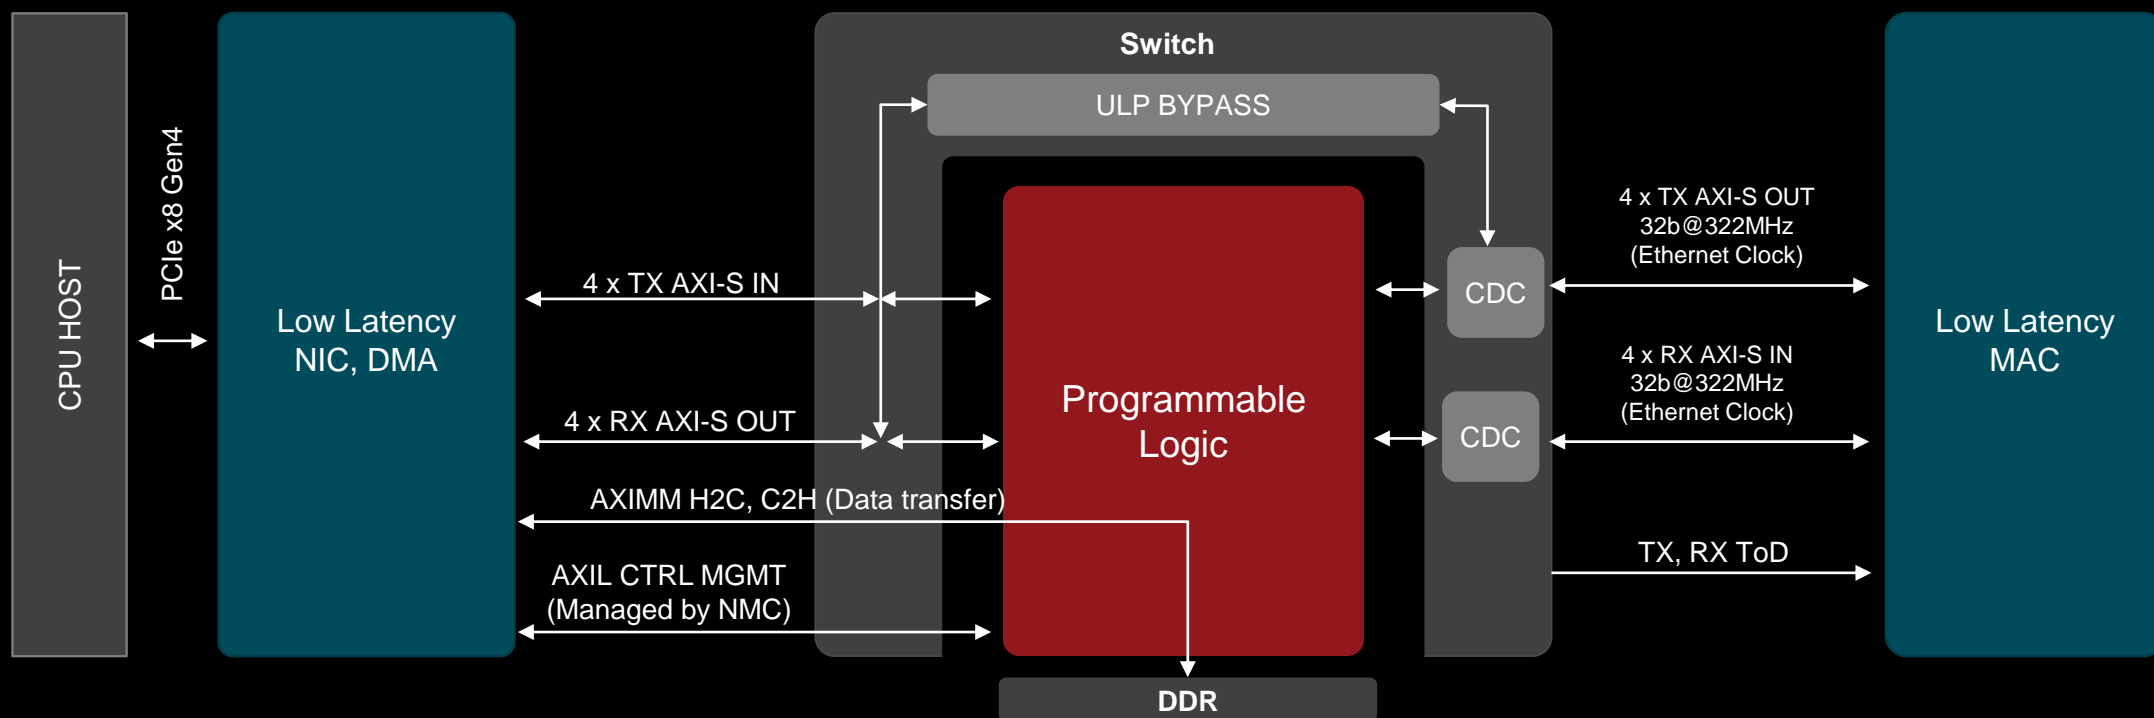
Interface	4x 10/25GbE
Logic Density	400K ¹ – 1M ² LUTs
Form Factor	HHHL
Max Power	75W

1: X3522 Low Latency NIC with hybrid upgrade package

2: X3522PV Accelerator Card

Alveo X3522 Hybrid Mode Low Latency Shell*


- Low latency shell remains static, providing functionality of the low latency NIC
- One Physical Function (PF) of NIC can be used by application to configure design in programmable logic
- Design with Vitis™ Unified Software Environment



* Subject to change

EASY AI INFERENCE ACCELERATION WITH UNIFIED INFERENCE FRONTEND (UIF)

Fast AI inference is important for your business. The acceleration offered by various devices (CPUs, GPUs, FPGAs, ...) is attractive, but you can't afford to learn all the tool chains/programming methods needed to port your models. You want to *train once, and run the inference on multiple hardware platforms*



AMD 

Disclaimer & Attribution

Timelines, roadmaps, and/or product release dates shown in these slides are plans only and subject to change.

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

©2022 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Athlon, CDNA, EPYC, Infinity Fabric Radeon, RDNA, ROCm, Ryzen, Ryzen Threadripper, Xilinx, the Xilinx logo, Alveo, Artix, Kintex, Spartan, Versal, Vitis, Virtex, and Zynq and combinations thereof are trademarks of Advanced Micro Devices, Inc. Microsoft is registered trademark of Microsoft Corporation in the US and other jurisdictions. SPEC®, SPECrate®, SPECint and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.