



# Crossfield Technology & Nuclear Physics SBIR/STTR

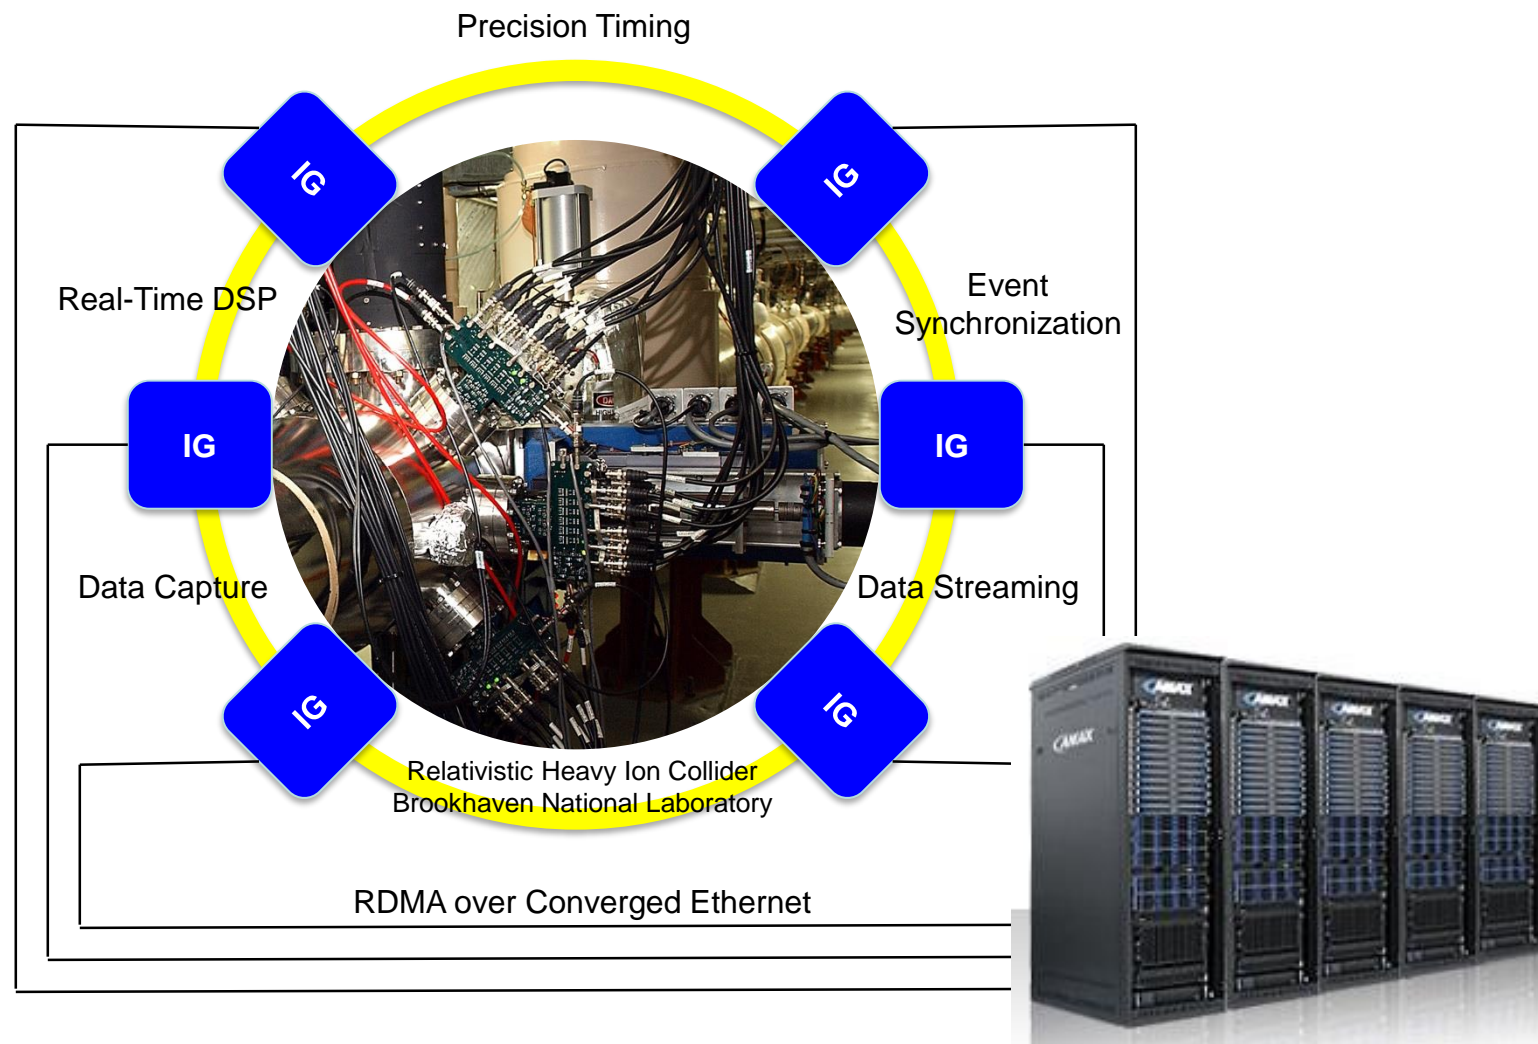
Software-Driven Network Architecture for Synchronous Data  
Acquisition -- DE-SC0015151

# Presentation Outline

- Program Overview
- A little about us and how we got here
- Challenges to Overcome
- Planned Solution

# **APPLICATION HIGH ENERGY PHYSICS SYNCHRONOUS DATA ACQUISITION**

# High Energy Physics Synchronous Data Acquisition



# Software-Driven Network for Synchronous Data Acquisition

- DoE SBIR Phase II Program
  - Software-driven instrumentation for event building in high energy physics
  - FPGA-based instruments capture sensor data from GPS ADCs
  - Sensor data processed in real-time, providing pulse-height and pulse-width analysis, pulse time-of-arrival measurement, and other instrumentation functions traditionally performed in fixed, dedicated instruments
  - Processed event data streamed to HPC for software-based Event Building
- Products Being Developed
  - RDMA over Converged Ethernet (RoCE) IP Core for Stratix 10 SoC FPGAs
  - Precision Timing IP Core (PTP/SynchE and support) Core for Stratix SoC FPGAs

# COMPANY OVERVIEW

# Company Overview

- Started in 2003 as a technology company bringing together a broad range of experience in high-speed networking, integrated circuit design and wireless system design
- Commercialize technologies developed through
  - BAA (SBIR & RIF) programs
    - DoD, DoE, NASA, NIST
  - And various commercial customers
- Crossfield has developed multiple generations of Instrumentation Gateways, non-volatile storage, and adapter modules
- Provide software services to capture, analyze and display sensor data streamed from instrumentation gateways

# Core Technologies




Wireless Gateways



Instrumentation Gateways




Networking Cards



FMC Modules

Real-Time Embedded Systems & Software  
High-Speed Networking  
High Performance Computing  
Synchronous Data Acquisition  
Wireless Instrumentation Systems  
Hardware in-the-Loop (HWIL) Simulation



VPX Modules



# **INSTRUMENTATION GATEWAYS (IG)**


# What are Instrumentation Gateways?

- IGs turn simple devices into smart networked edge devices that allows data to be gathered from or distributed to many synchronized points
- Key Features:
  - High-speed gateways provide advanced networking capabilities, such as Remote Direct Memory Access (RDMA) over InfiniBand or Ethernet
  - IEEE 1588v2 Precision Time Protocol and Synchronous Ethernet provide network time synchronization at the nanosecond level
  - Significant edge processing to “shape” high volumes of data
- The gateways are needed in embedded, desktop, rackmount and OpenVPX form factors

# Real Time Data Acquisition and Distribution

## 56G Instrumentation Gateway

Enables rapid implementation of demanding data acquisition, real-time simulation, process control and other applications that stream data across networks



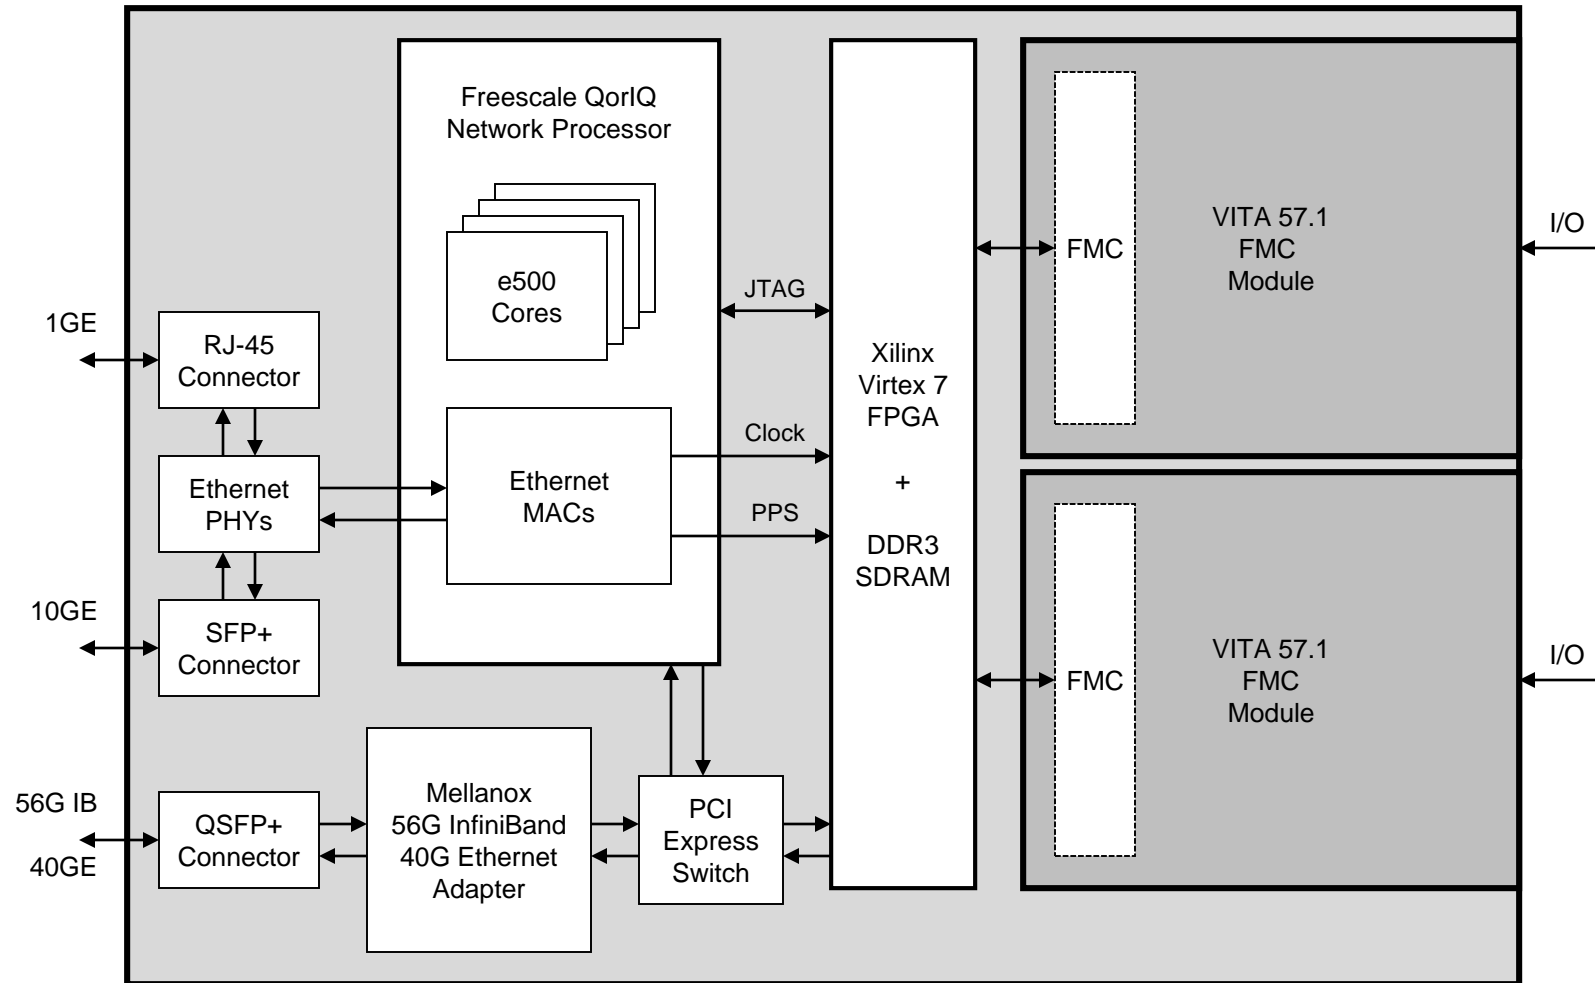
- Two FMC interfaces to high performance FPGA
- Hardware IEEE 1588v2 PTP and ITU-T SyncE
- Data plane – FPGADirect™ across 56G IB or 40G RoCE
- Control plane – 1GbE or 10GbE
- Integrated Web server

# Long History of Instrumentation Gateways Crossfield

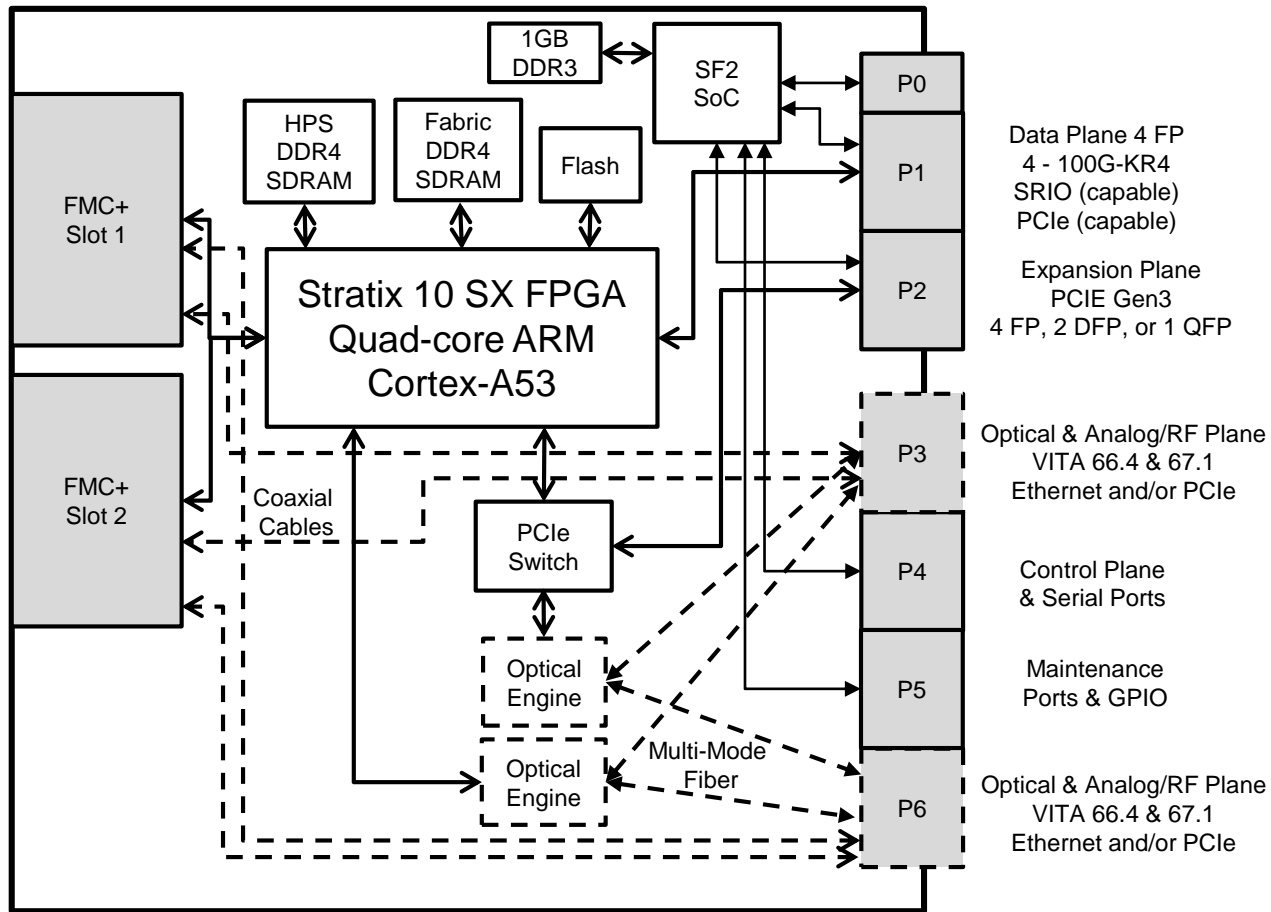
- Concept of Instrumentation Gateways started in 2005 with a MDA SBIR Phase I
  - Focused on solving a Synchronous Data Distribution problem
  - Initial architecture was around PCIe AS until the chip vendors went away
  - Transitioned to iWARP (Internet Wide-area RDMA Protocol)
    - Working with one of my startup companies
  - InfiniBand or RoCE turned out to be sufficient and easier
- Culminated with AFRL CRP Real-Time Hyperspectral Scene Generation
  - Started 2010 as is just now coming to an end
  - Crossfield essentially complete in 2015
- Transitioned to Synchronous Data Acquisition in 2012
  - Interest from DoE and Navy SIGINT
  - Same architecture delivers either or both
- DoD driving smaller platforms which require us to consolidate into single FPGA
  - Investigation started there fits perfectly with DoE topic

Nuclear Physics Software and Data Management - Software-Driven Network Architectures for Data Acquisition

# 40GE - 56G IB IG with Dual FMC Slots



# 6U HOST & SOSA Stratix 10 SX IG

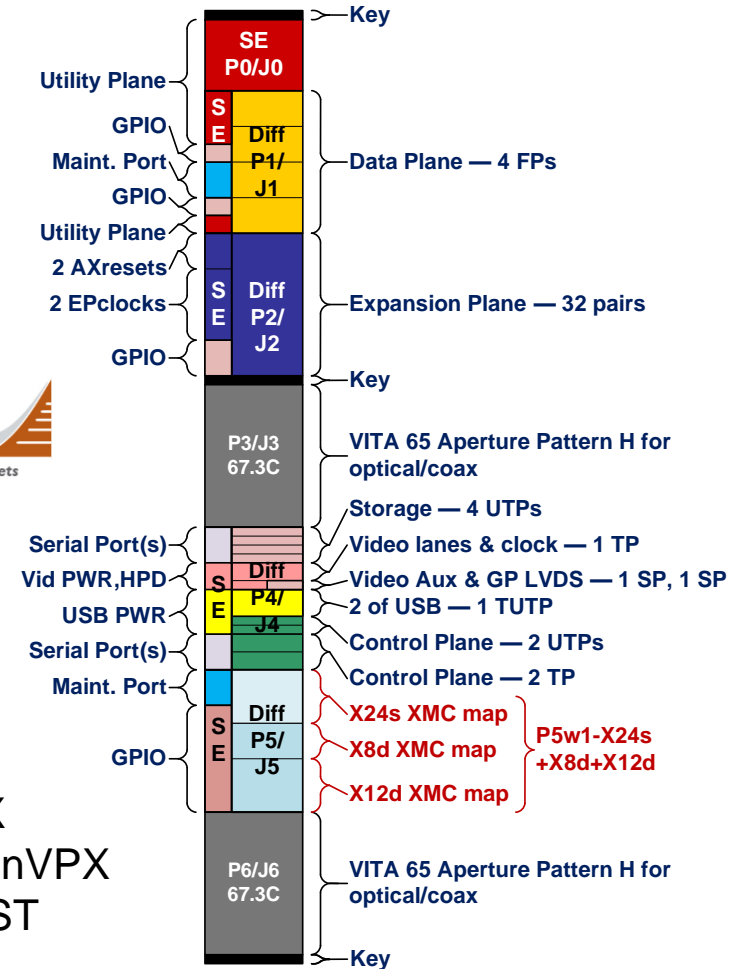


Dashes denote optional placements

- Data Plane 4 FP  
4 - 100G-KR4  
SRIO (capable)  
PCIe (capable)
- Expansion Plane  
PCIe Gen3  
4 FP, 2 DFP, or 1 QFP
- Optical & Analog/RF Plane  
VITA 66.4 & 67.1  
Ethernet and/or PCIe
- Control Plane  
& Serial Ports
- Maintenance  
Ports & GPIO
- Optical & Analog/RF Plane  
VITA 66.4 & 67.1  
Ethernet and/or PCIe





VITA 46 VPX  
VITA 65 OpenVPX  
VITA 84 HOST

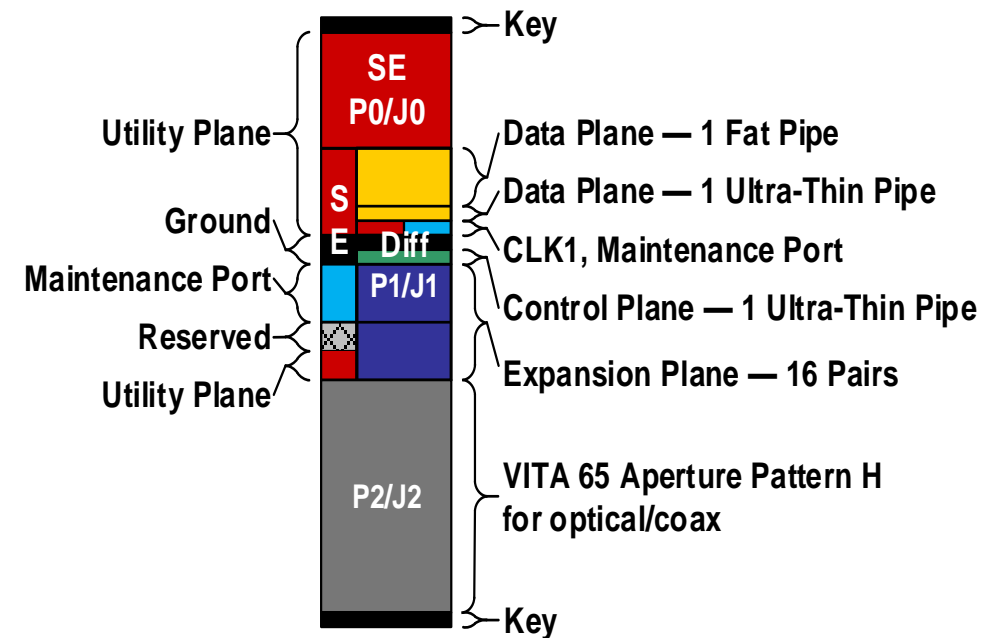
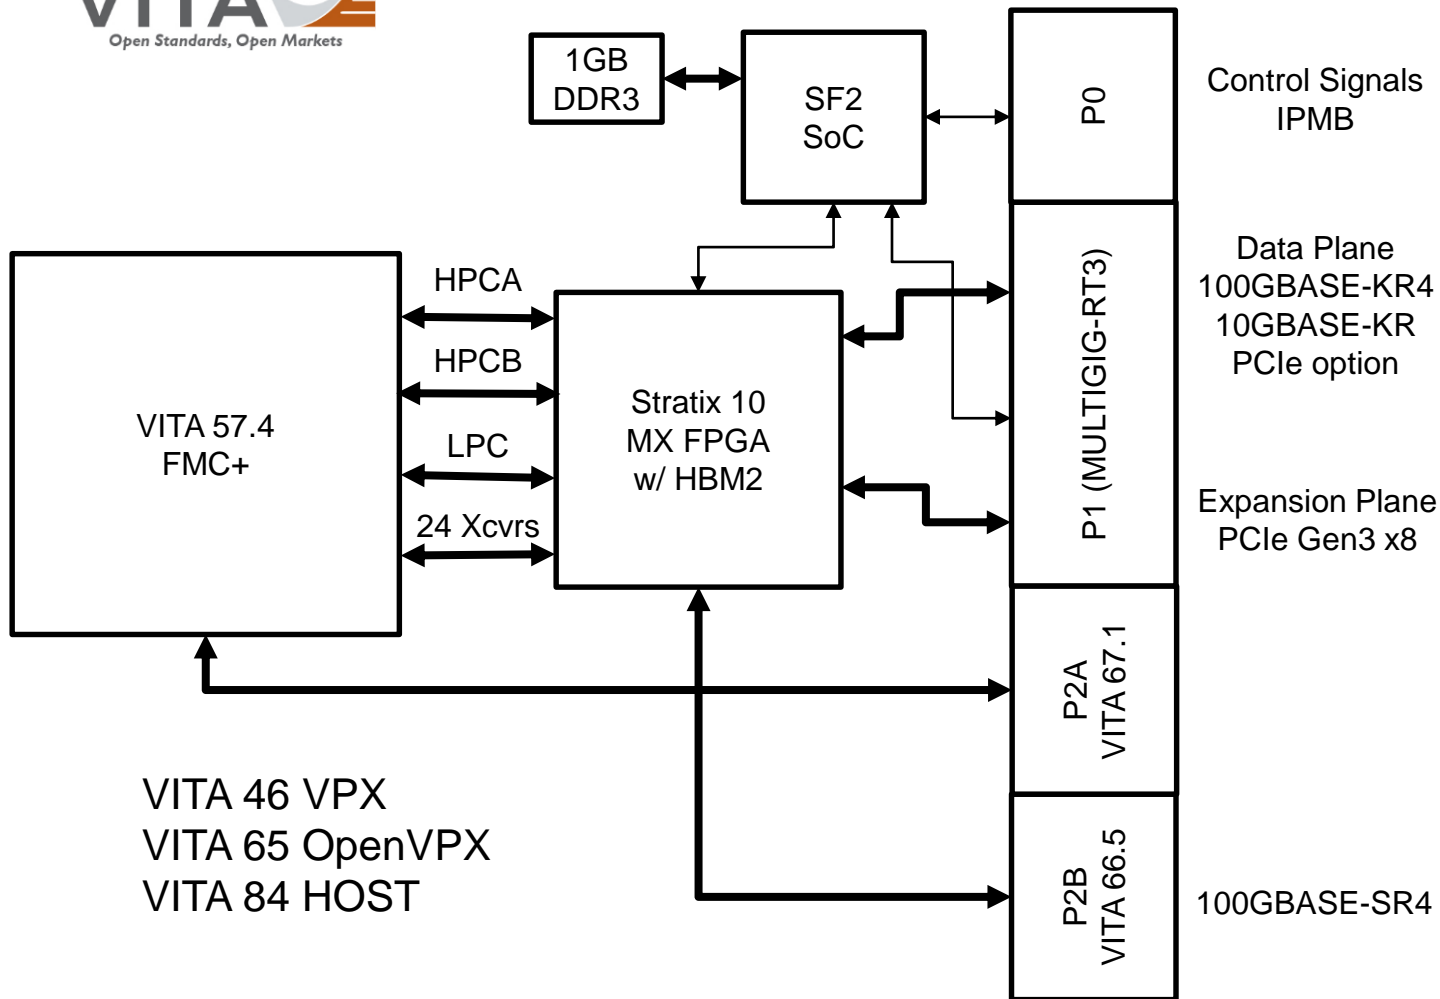


MOD6-PAY-4F1Q1H4U1T1S1S1TU2U2T1H-12.6.3-n

# 6U VPX Stratix 10 SX FPGA SoC 3D Rendering




# 3U HOST/SOSA/CMOSS Stratix 10 MX IG





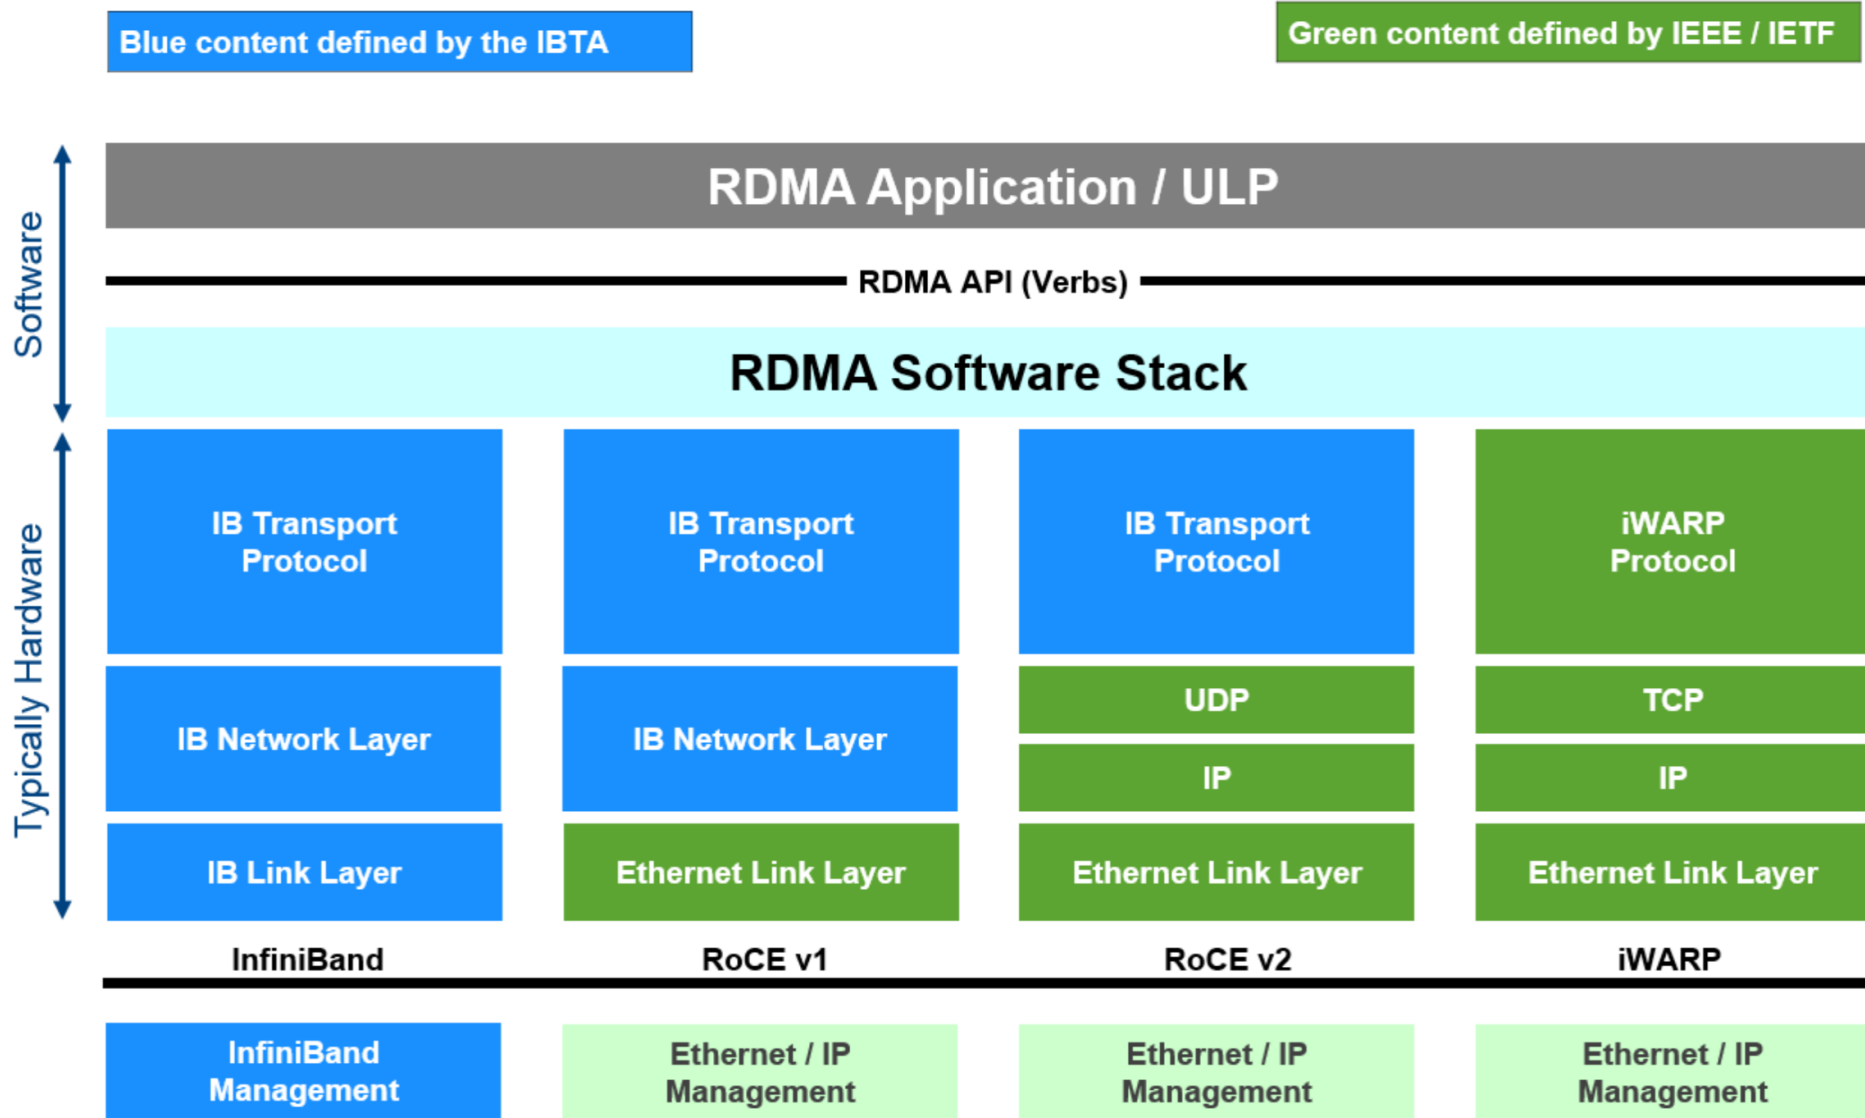
# Integrated Memory Enables 3U Form Factor



3U VPX Stratix 10 MX RF Interface

# FPGA IP CHALLENGES AND OPPORTUNITIES

# RDMA Networking Choices

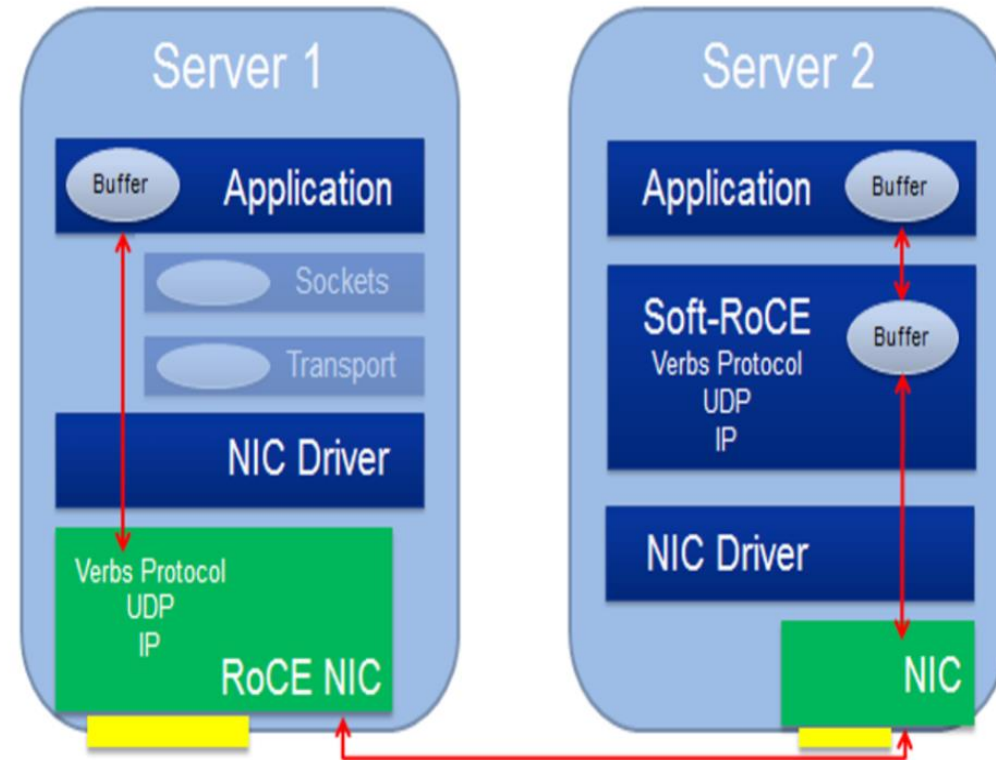


# FPGA RDMA IP Challenges

- InfiniBand
  - Subset of functionality is readily available
    - Unreliable Connection (UC) only
    - Expensive -- ~\$750K
    - Incomplete
      - Transport service and Link layer only (no general purpose DMA function)
- iWARP
  - IP not available on open market and likely much to complex to develop
- RoCE
  - Mostly complete IP in development at Xilinx
    - Price unknown
  - Nothing available for Stratix FPGAs
- However, there is a Soft RoCE
  - RDMA Transport in a Software Implementation

# Soft RoCE

- Full software implementation of RoCEv2
- Support in Linux readily available
  - On Github --  
<https://github.com/SoftRoCE>
- Developed & Verified for X86
  - Main contributors
    - IBM, Mellanox, System  
Fabrics Works



Soft-RoCE implements the packet processing otherwise managed by the RoCE NIC.

[http://www.roceinitiative.org/wp-content/uploads/2016/11/SoftRoCE\\_Paper\\_FINAL.pdf](http://www.roceinitiative.org/wp-content/uploads/2016/11/SoftRoCE_Paper_FINAL.pdf)

# Soft RoCE Performance Questions

- Published performance characteristics for Soft RoCE on x86 processors
- Performance is acceptable but at what cost of CPU utilization and power?
- What about embedded ARM cores?

## Results : Analysis

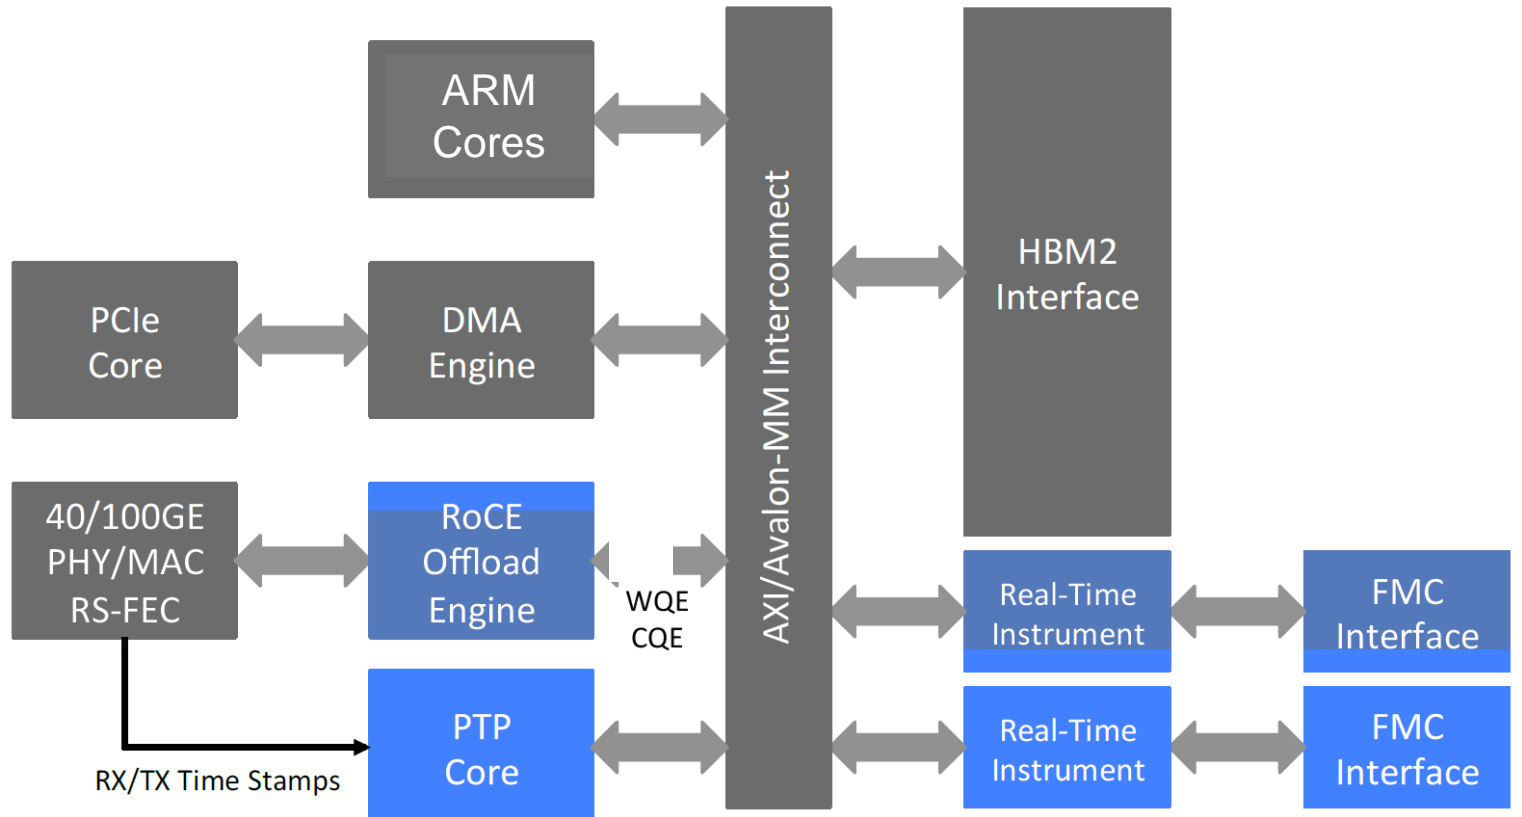
Peak Values	IB QDR	RoCE	Soft RoCE	No RDMA
Latency ( $\mu$ s)	1.96	3.7	11.6	21.09
One-way BW (MB/s)	3024.8	1142.7	1204.1	301.31
Two-way BW (MB/s)	5481.9	2284.7	-	1136.1

- RoCE performance gains over 10GbE:
  - ▣ Up to 5.7x speedup in latency
  - ▣ Up to 3.7x increase in bandwidth
- IB QDR vs. RoCE:
  - ▣ IB less than 1 $\mu$ s faster than RoCE at 128-byte message.
  - ▣ IB peak bandwidth is 2-2.5x greater than RoCE.

[https://www.lanl.gov/projects/national-security-education-center/information-science-technology/\\_assets/docs/2010-si-docs/Team\\_CYAN\\_Implementation\\_and\\_Comparison\\_of\\_RDMA\\_Over\\_Ethernet\\_Presentation.pdf](https://www.lanl.gov/projects/national-security-education-center/information-science-technology/_assets/docs/2010-si-docs/Team_CYAN_Implementation_and_Comparison_of_RDMA_Over_Ethernet_Presentation.pdf)

# Acceleration Provides Solution

- Soft RoCE port to ARM Cortex-A53 complete on Aria 10 FPGA
- Profiling code in development
- Expectation is that 100GbE can't be saturated by CPU complex without help
- First accelerator will focus on wire saturation while minimizing CPU overhead



# THANK YOU

## Contact Information:

Terry Hulett

[terry.Hulett@crossfieldtech.com](mailto:terry.Hulett@crossfieldtech.com)

512-413-5413