MERCURY SYSTEMS™

INNOVATION THAT MATTERS™

# Advanced Fabric Innovations to Maximize the Performance of HPEC Systems

**Marc Couture**
**Director, Product Management**

**Arnold Sodder**
**Consulting Systems Architect**

**Embedded Technology Trends**
**January 22, 2013**

# Processing Density shows no sign of stopping

- AVX Core count, Shader Processor count, going up as nm geometry goes down
- Intel
  - Mobile Class has four AVX/AVX2-enabled cores + GT2/3 GPUs
  - Server Class moving from 8 cores to 10 and even 12 per device
- GPUs
  - AMD
    - 960 processors in 40nm 6970M
    - 1280 processors in 28nm 7970M
- Two of the above devices, Intel or GPU, can fit onto a single 6U OpenVPX module for an immense amount of compute power


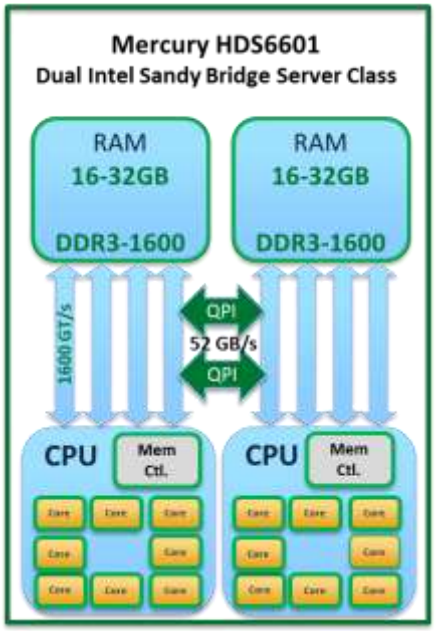



Hard to become compute bound, but I/O bound is easy

# Maximizing Capability per SWaP-C Platform

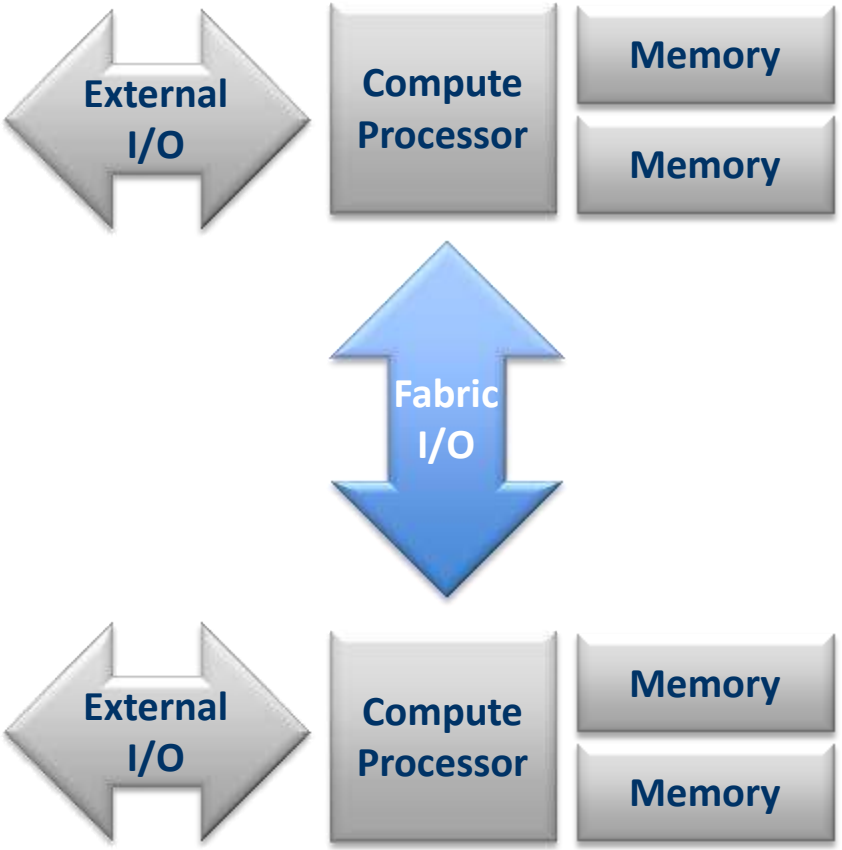


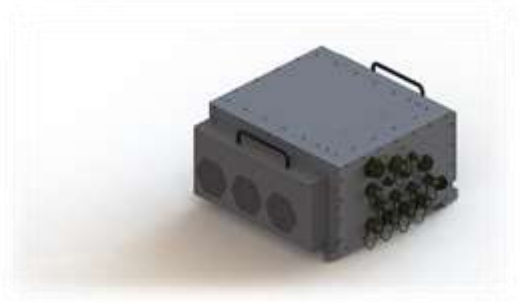

Balancing Compute, Memory, I/O (External & **Fabric**) Bandwidth

# Advances in Backplane Technology Needed for Next-Gen Interconnects

- Current OpenVPX backplane based on standard RT2 connector topping out at 6.25 Gbaud due to increasing:
  - Channel Loss
  - Channel Reflection
  - Channel Cross Talk
- Signal Integrity loss becomes exaggerated with increasing:
  - Slot count
  - Dataplane & Expansion Plane lanes
  - Traffic on mezzanine and RTMs
  - Temperature extremes
  - Vibration

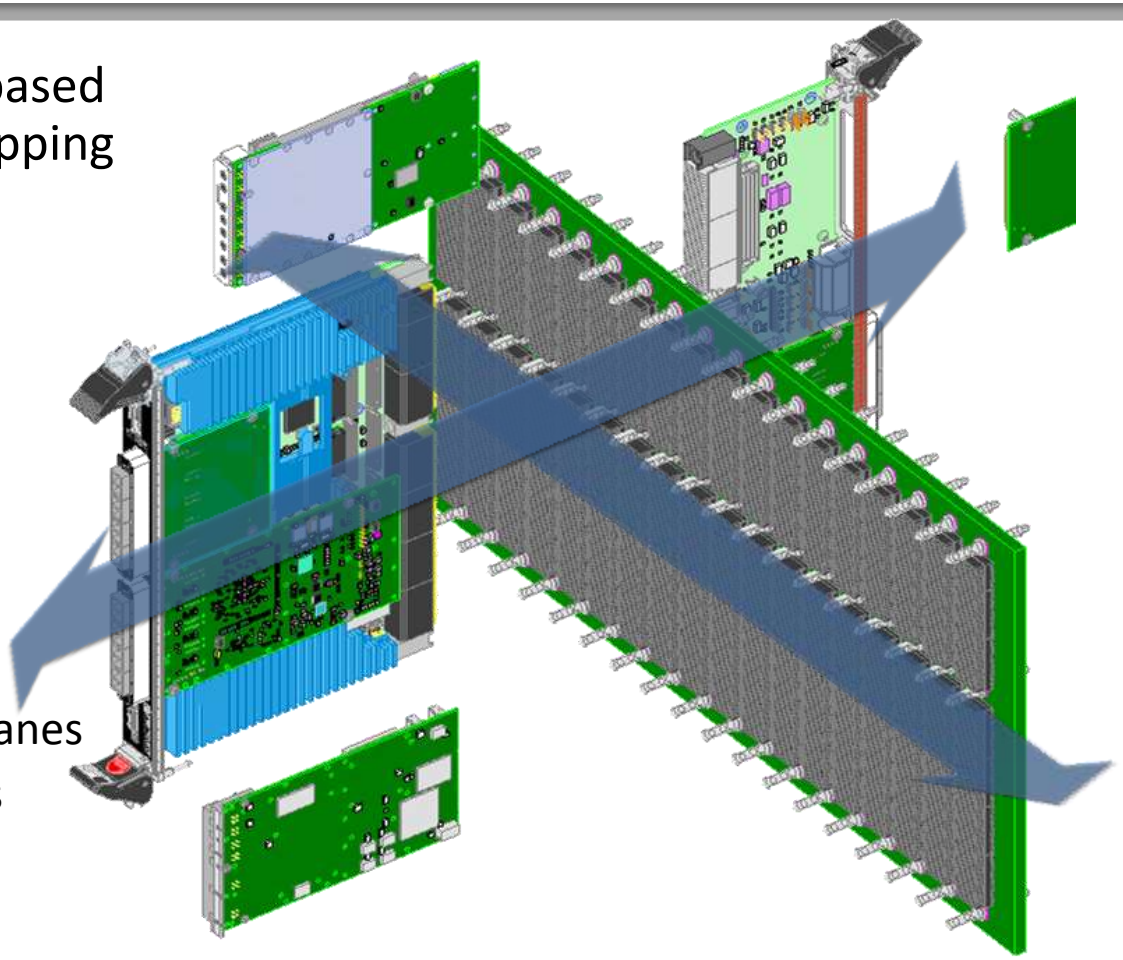

We need to speed up without degradation

# Next Generation Infrastructure

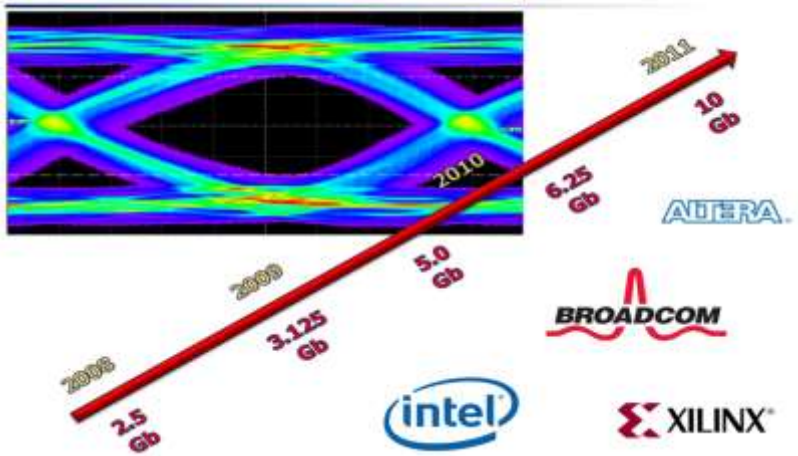


- Advances in connector and backplane technology will get us to 8.0 Gbaud and beyond, needed for Next-Gen Fabric Interconnects:
  - PCIe Gen 3
  - Infiniband QDR
  - 40 GigE

**Next Gen Connector Technology**

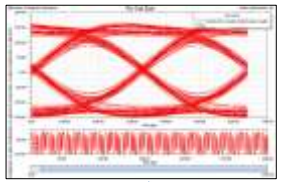

**Speed (beyond 6.25)**

**Rugged**
**(better vibration)**

The connector is the key to unlocking speed!

# The OpenVPX Multi-Plane Elevation

- Choices for the Data Plane mesh/switch fabric are varied
- The Expansion Plane is primarily PCIe

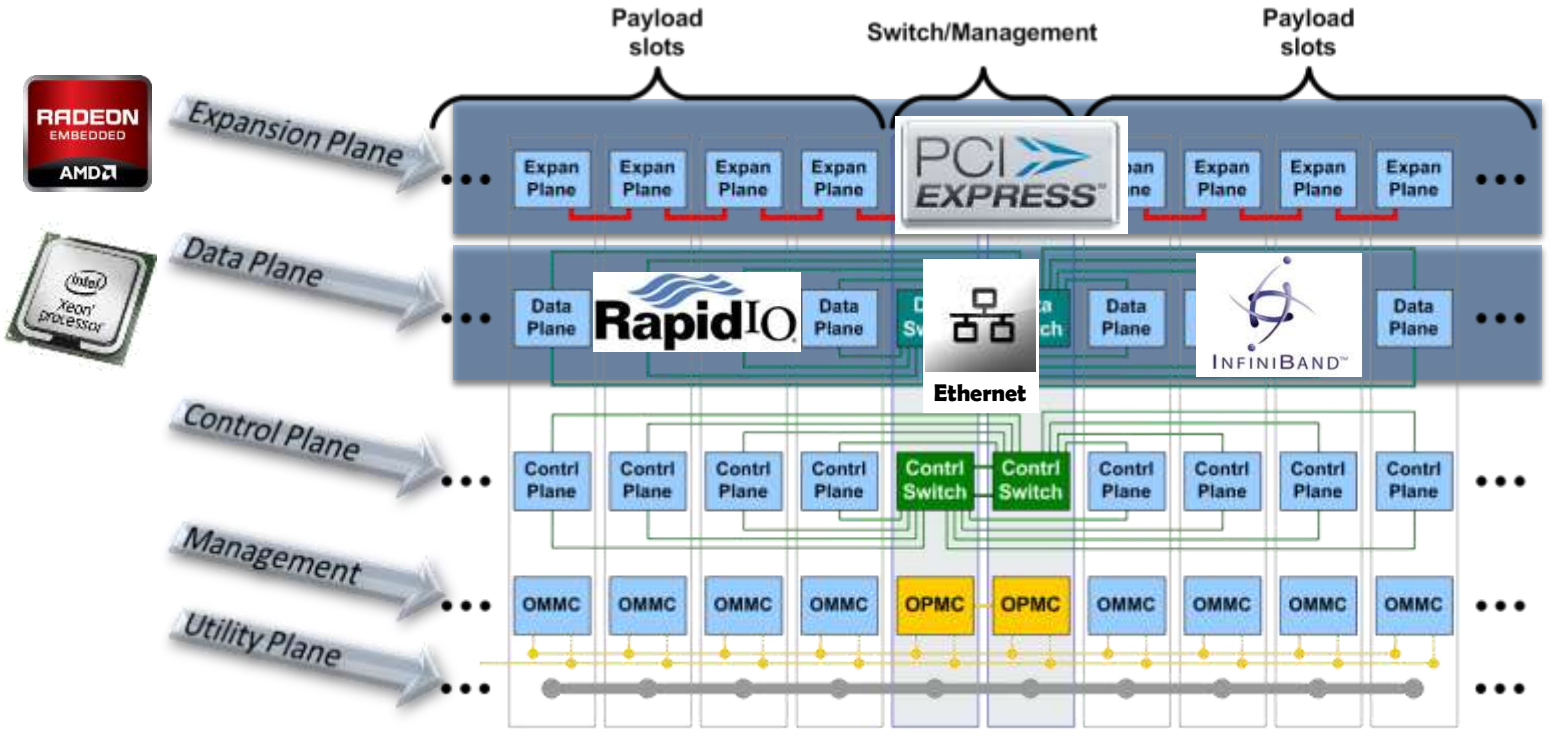


Enabling sophisticated interactions within the system

# Example Topologies, Switched Dataplane

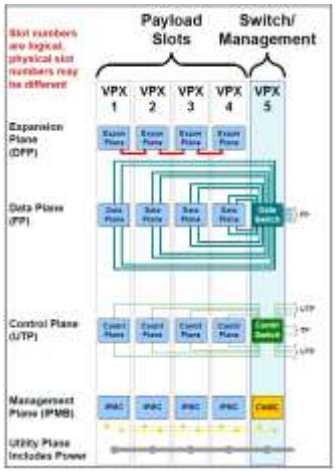


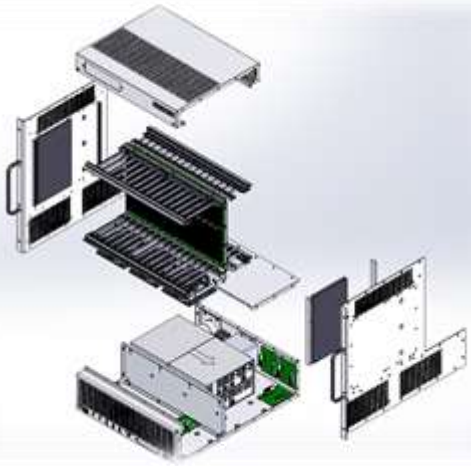

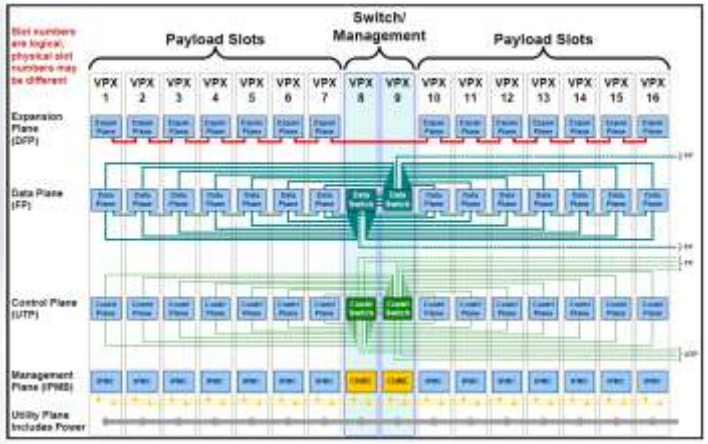


## Matching the problem with slots that scale effectively

# Payload Fabric Endpoints

- Currently utilizing Protocol Offload Engine Technology (POET™) IP on FPGAs for Serial RapidIO Gen 2 @ 3.125, 5.0, and 6.25 Gbaud/sec and 10 GigE
- Will be using Mellanox ConnectX-3 for InfiniBand at various rates and eventually 40 GigE
- Will intersect back with POET IP on FPGAs in the future

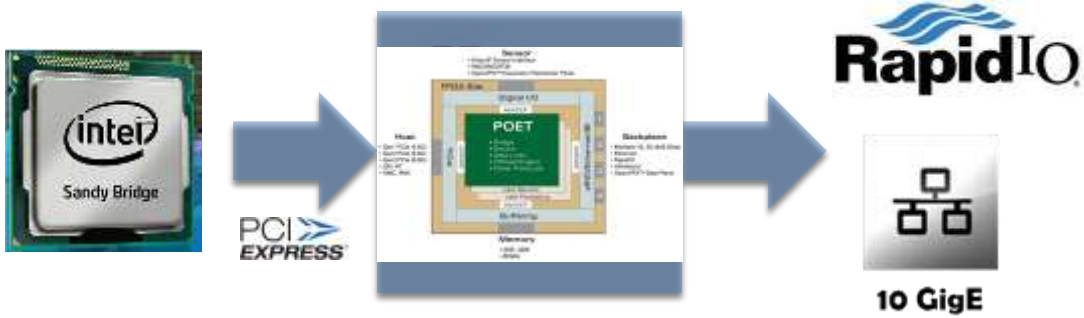


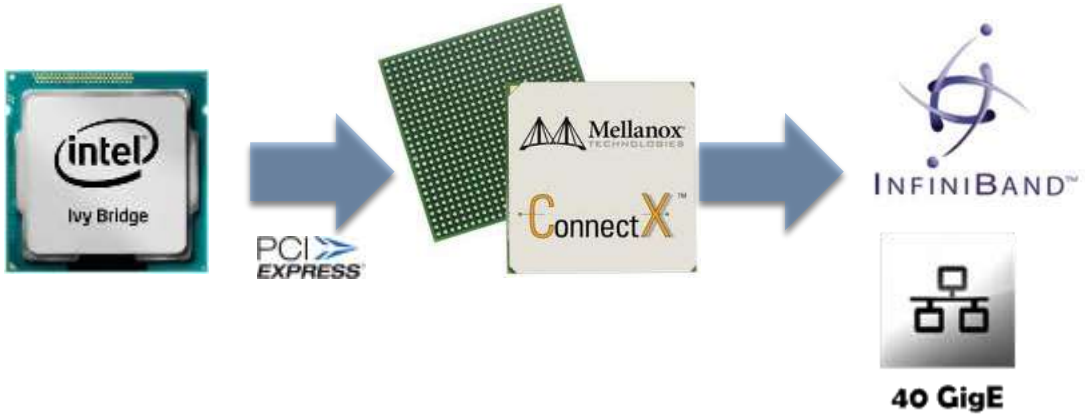


Next, we need to speed up the I/O to the processors

# Intel Server Class coupled with Mellanox

- Addresses the InfiniBand HPC market (e.g. Bladecenter) providing a path to embedded 6U OpenVPX
- Utilizes Mellanox ConnectX-3 for InfiniBand DDR, QDR
- Supports OFED, tightly coupled to InfiniBand
- Same analogy for 40GigE

Mellanox TECHNOLOGIES

ConnectX-3

## Mobilization/Ruggedization of the Server

# OpenVPX Switch Modules for HPEC

Dedicated Fabric Switches become increasingly important with increased slot-count, providing location-independent processing

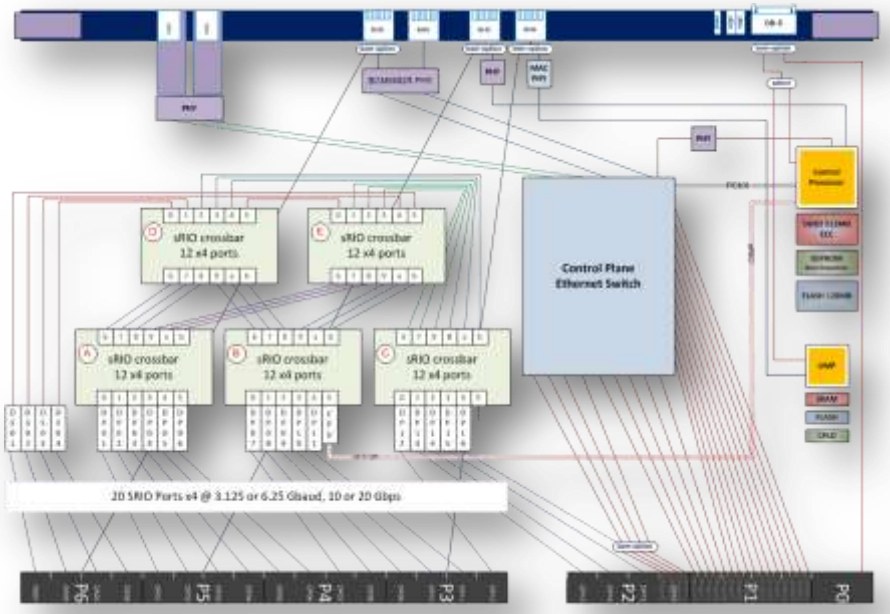


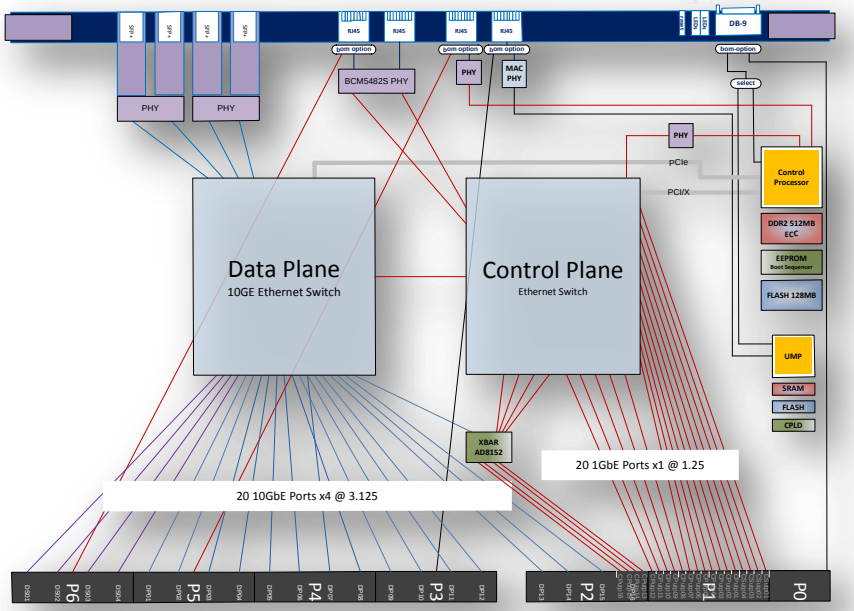


SRIO & 10GigE today, Moving to InfiniBand and 40GigE tomorrow

# Fabric Interconnect Enabling Multi-Dimensional Apps

- Multi-Channel Sensor Chains
- GHz Sample Clock alignment for coherency
- GBps Fabric Interconnect between chains
- High throughput, low latency, deterministic

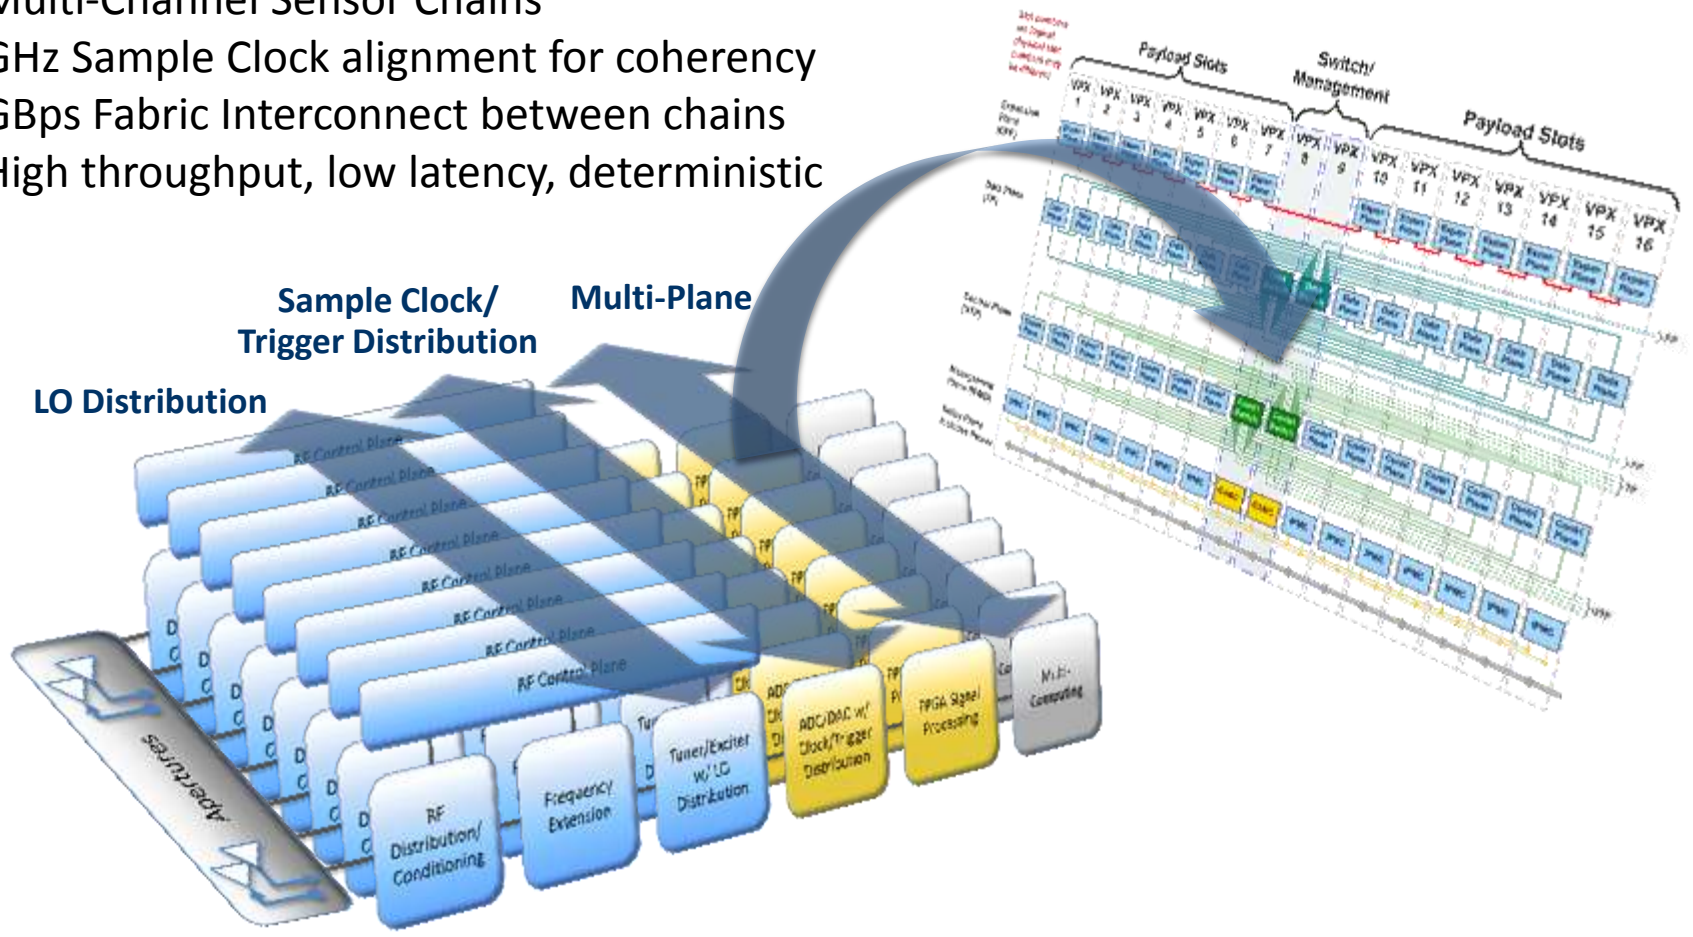


## Sensor I/O Integrated into the Switched Fabric

# HPEC Fabric Key Points

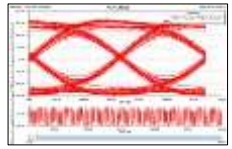

- The current VPX infrastructure is stalling out from an Signal Integrity aspect
- Next gen connector, backplane, and cooling technologies will kick OpenVPX into the next gear
- We'll be able to get maximum mileage out of these massive many-core HPEC compute engines
- And we'll be able to put them in tight spots (SWaP)

Breaking through to the next level with HPEC in OpenVPX

# Q & A

# Thank You