RapidIO for DAQ, Trigger and Data Analytics

Sima Baymani | CERN openlab Fellow

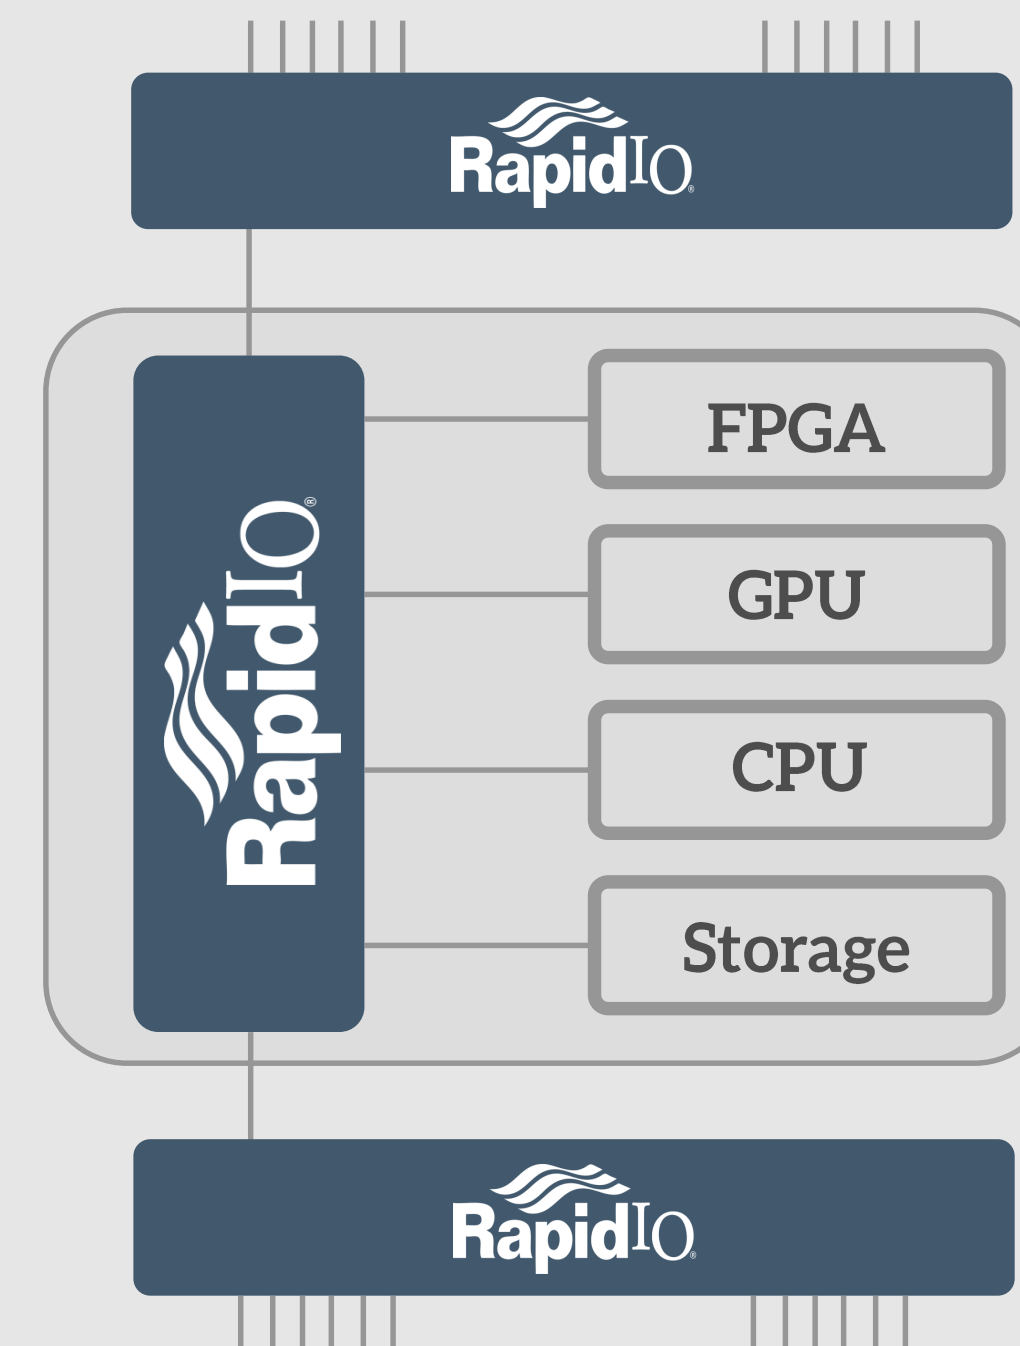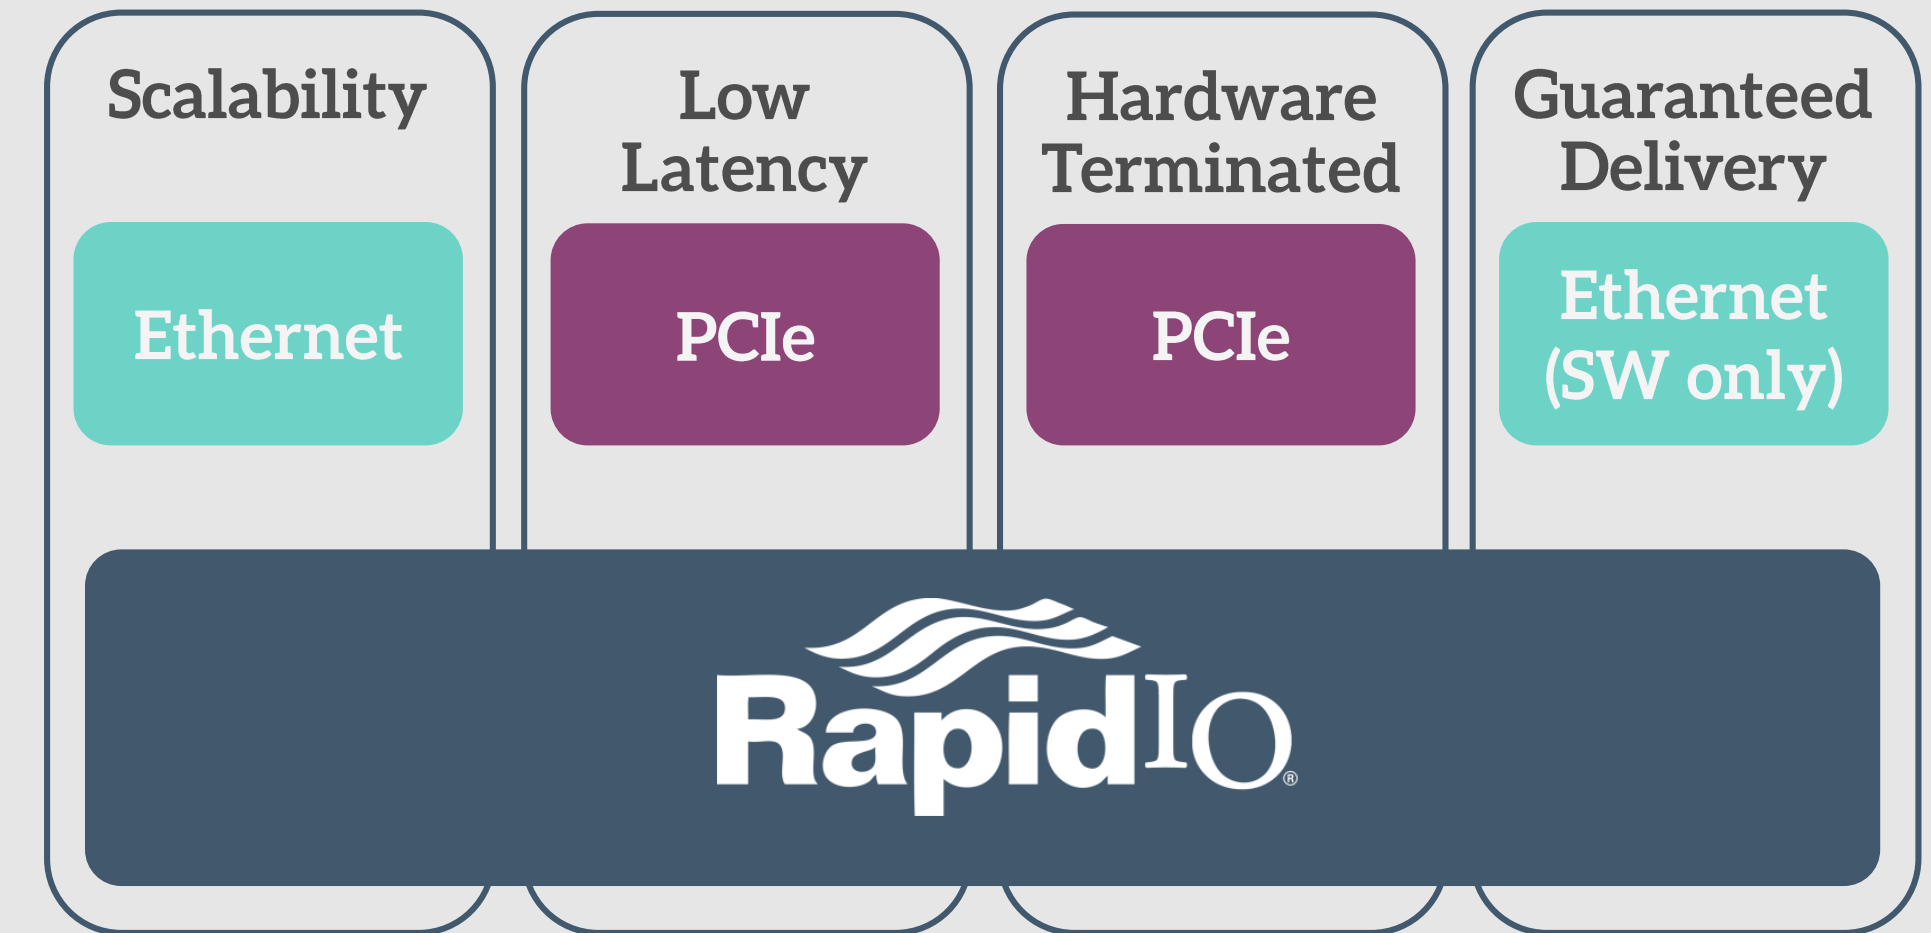Background image: Shutterstock

CERNopenlab

1

Openlab Partner: IDT

- Integrated Device Technology
- Partner since 2015
- HQ: San Jose, California
- ~1500 employees
- Hard real-time processing

Project Technology: RapidIO

- Created late 1990's
- Deterministic latency ~ 100 ns
- Combines strengths of PCIe and Ethernet
- Supports heterogeneous systems
- Open standard – rapidio.org

Scalability	Low Latency	Hardware Terminated	Guaranteed Delivery
Ethernet	PCIe	PCIe	Ethernet (SW only)

RapidIO

RapidIO

RapidIO
- FPGA
- GPU
- CPU
- Storage

RapidIO

Project Start

4 Day Introduction and Training

- Seminars
- Hands-on labs

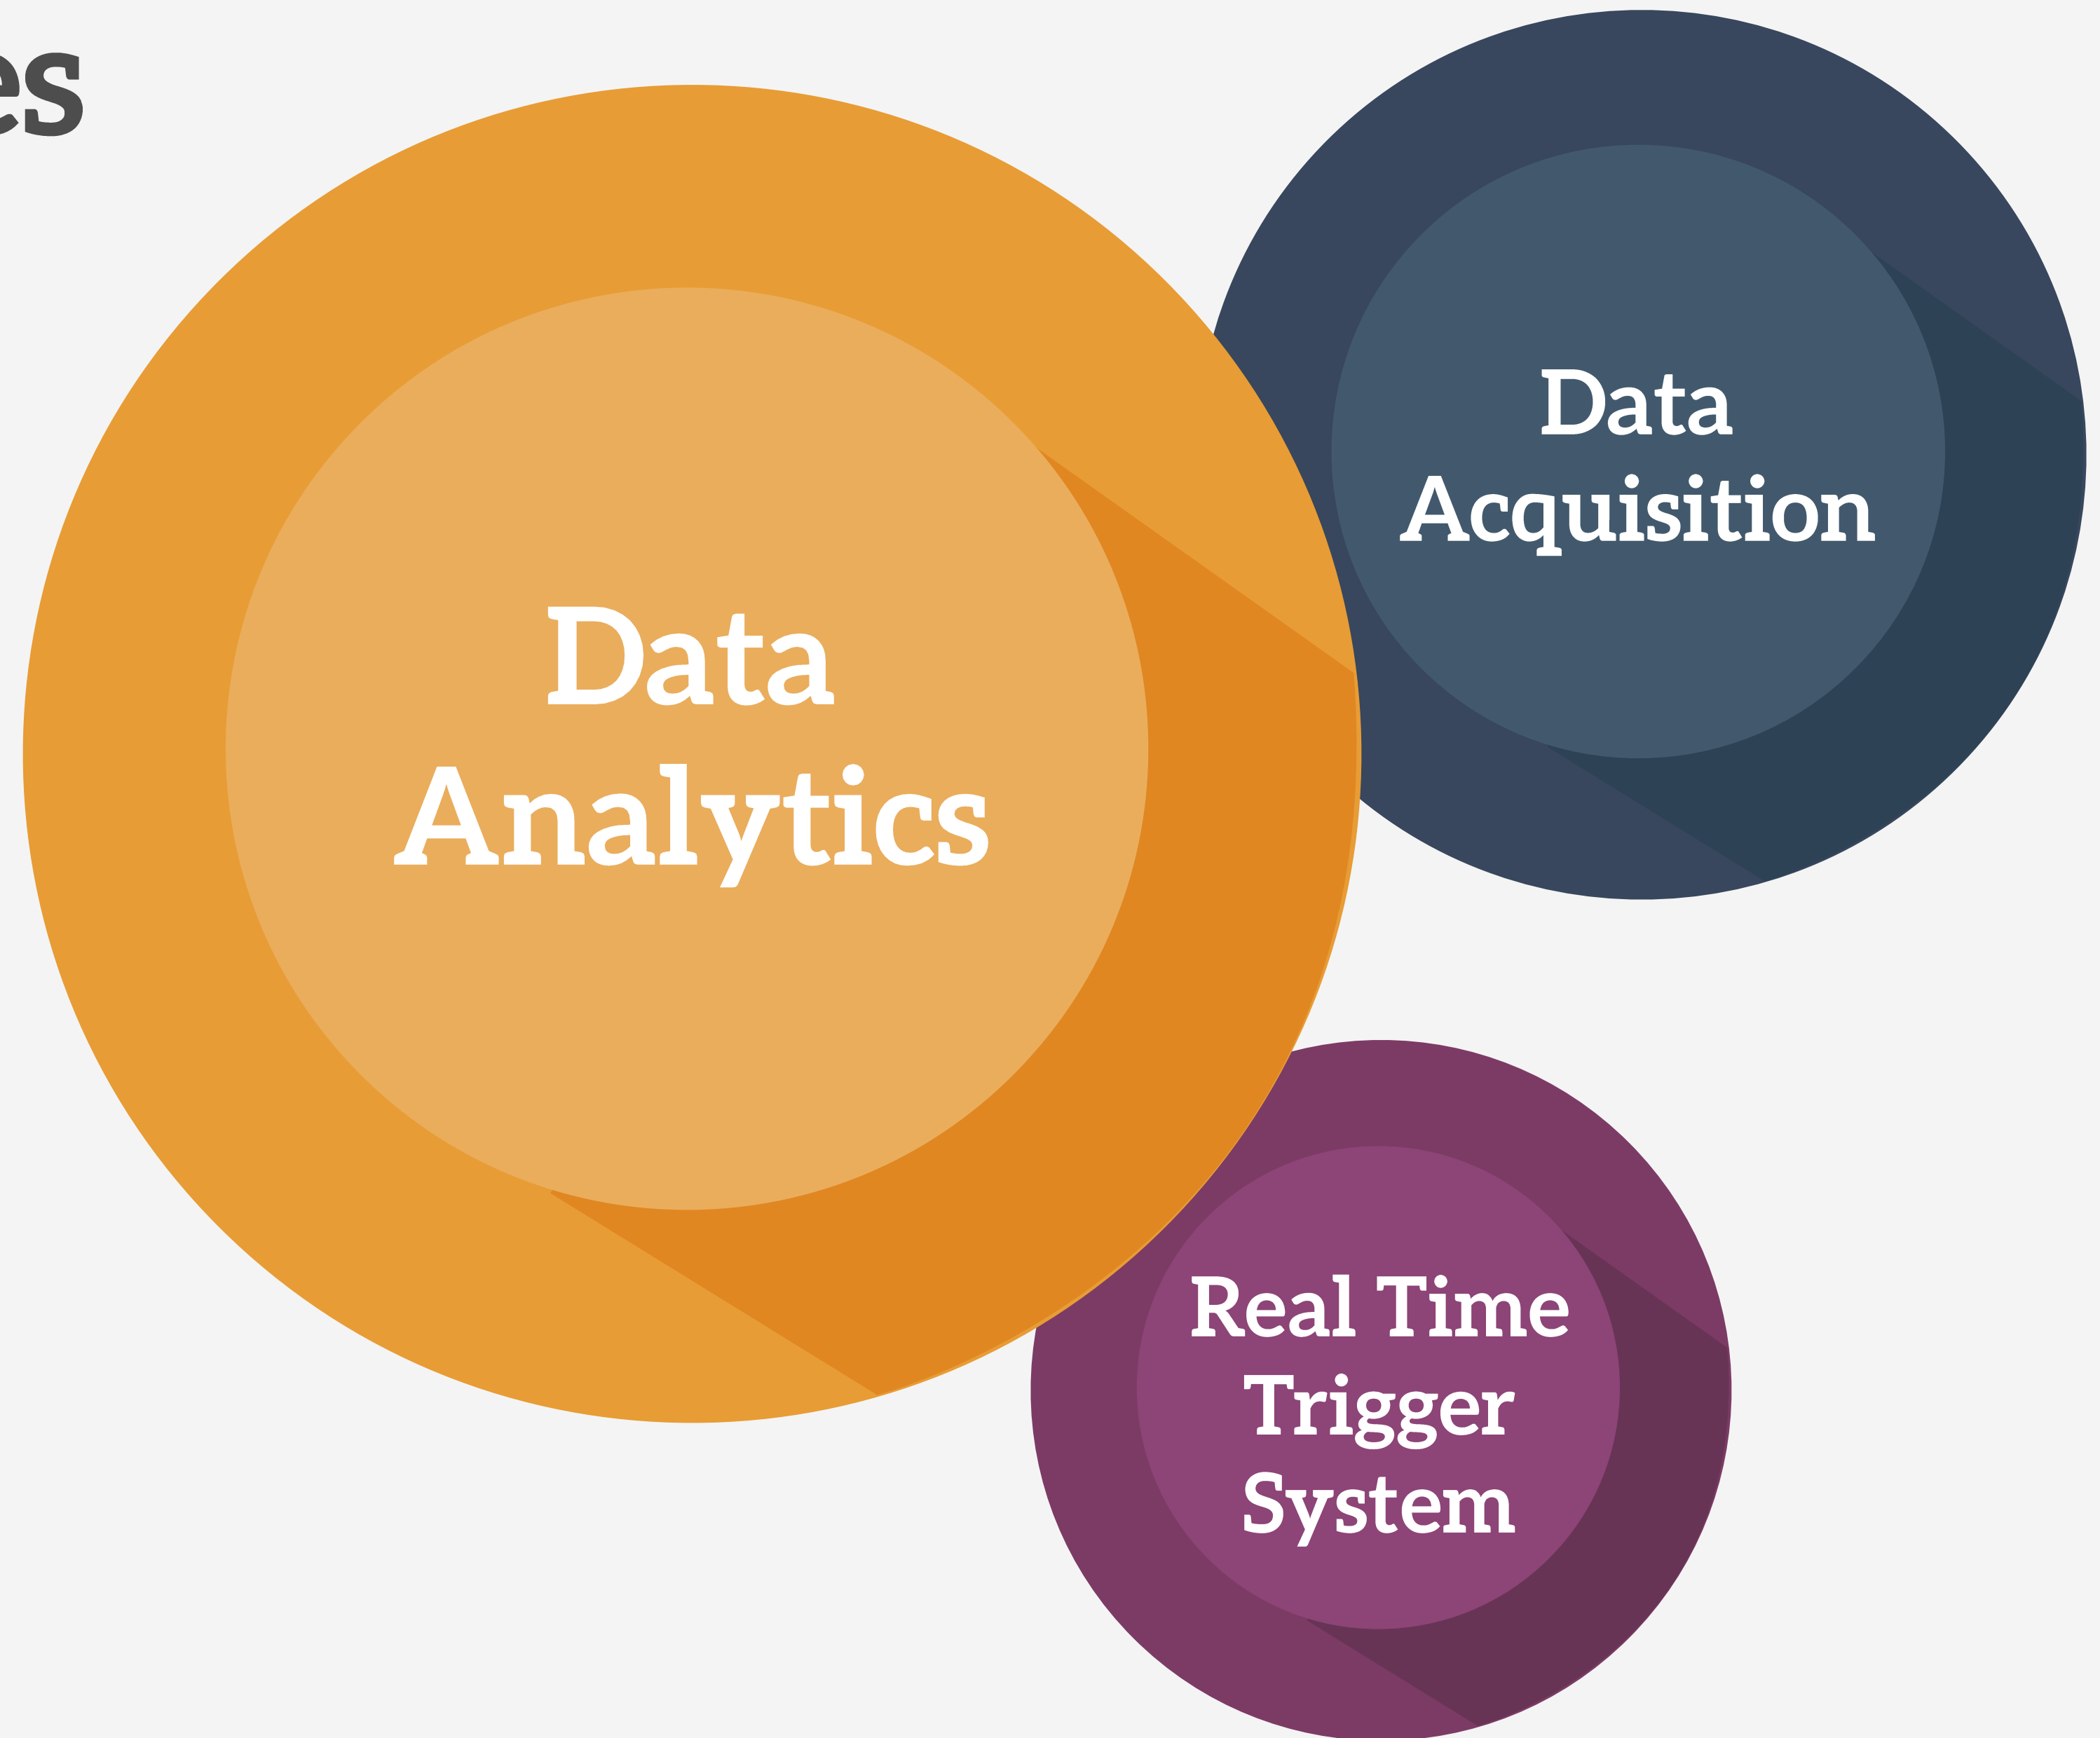

RapidExpress Bridge Card

Installation and Set-up of Cluster

- 4 nodes
- Each with a RapidIO NIC
- 1 RapidIO Switch

RapidExpress 8-port Switch

Project Use Cases

Exploring RapidIO

- Use cases
- Focus on different areas
- Evaluate suitability

Data Analytics

Data Acquisition

Real Time Trigger System

Data
Analytics

Use Case 1: Usage of RapidIO in a Data Analytics Environment

File Transfer Application

- Explore low level APIs
- Messaging API – protocol communication
- DMA API – data transfer
- Unicast - Multicast

Hadoop

- Hadoop deployed on 4-node lab cluster
- Using RIONET driver – emulating TCP/IP over RapidIO
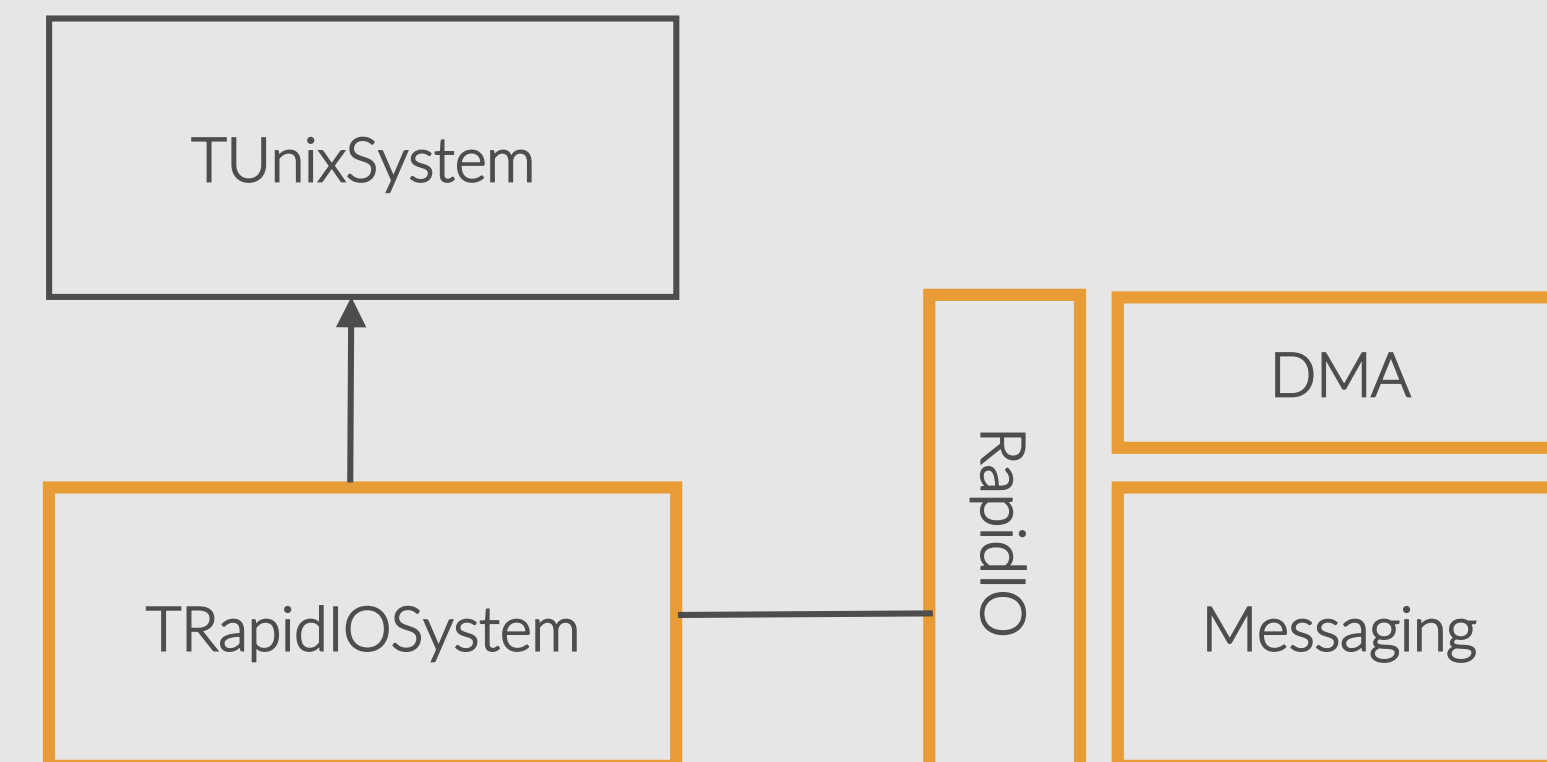- Nodes communicating solely using RapidIO

ROOT @ RapidIO

- ROOT over RIONET

- Native port of ROOT to RapidIO

  - Override core class in ROOT

  - Add RapidIO capabilities

  - Reuse base class functionality for

    filesystem and other services

- PROOF – distributed ROOT
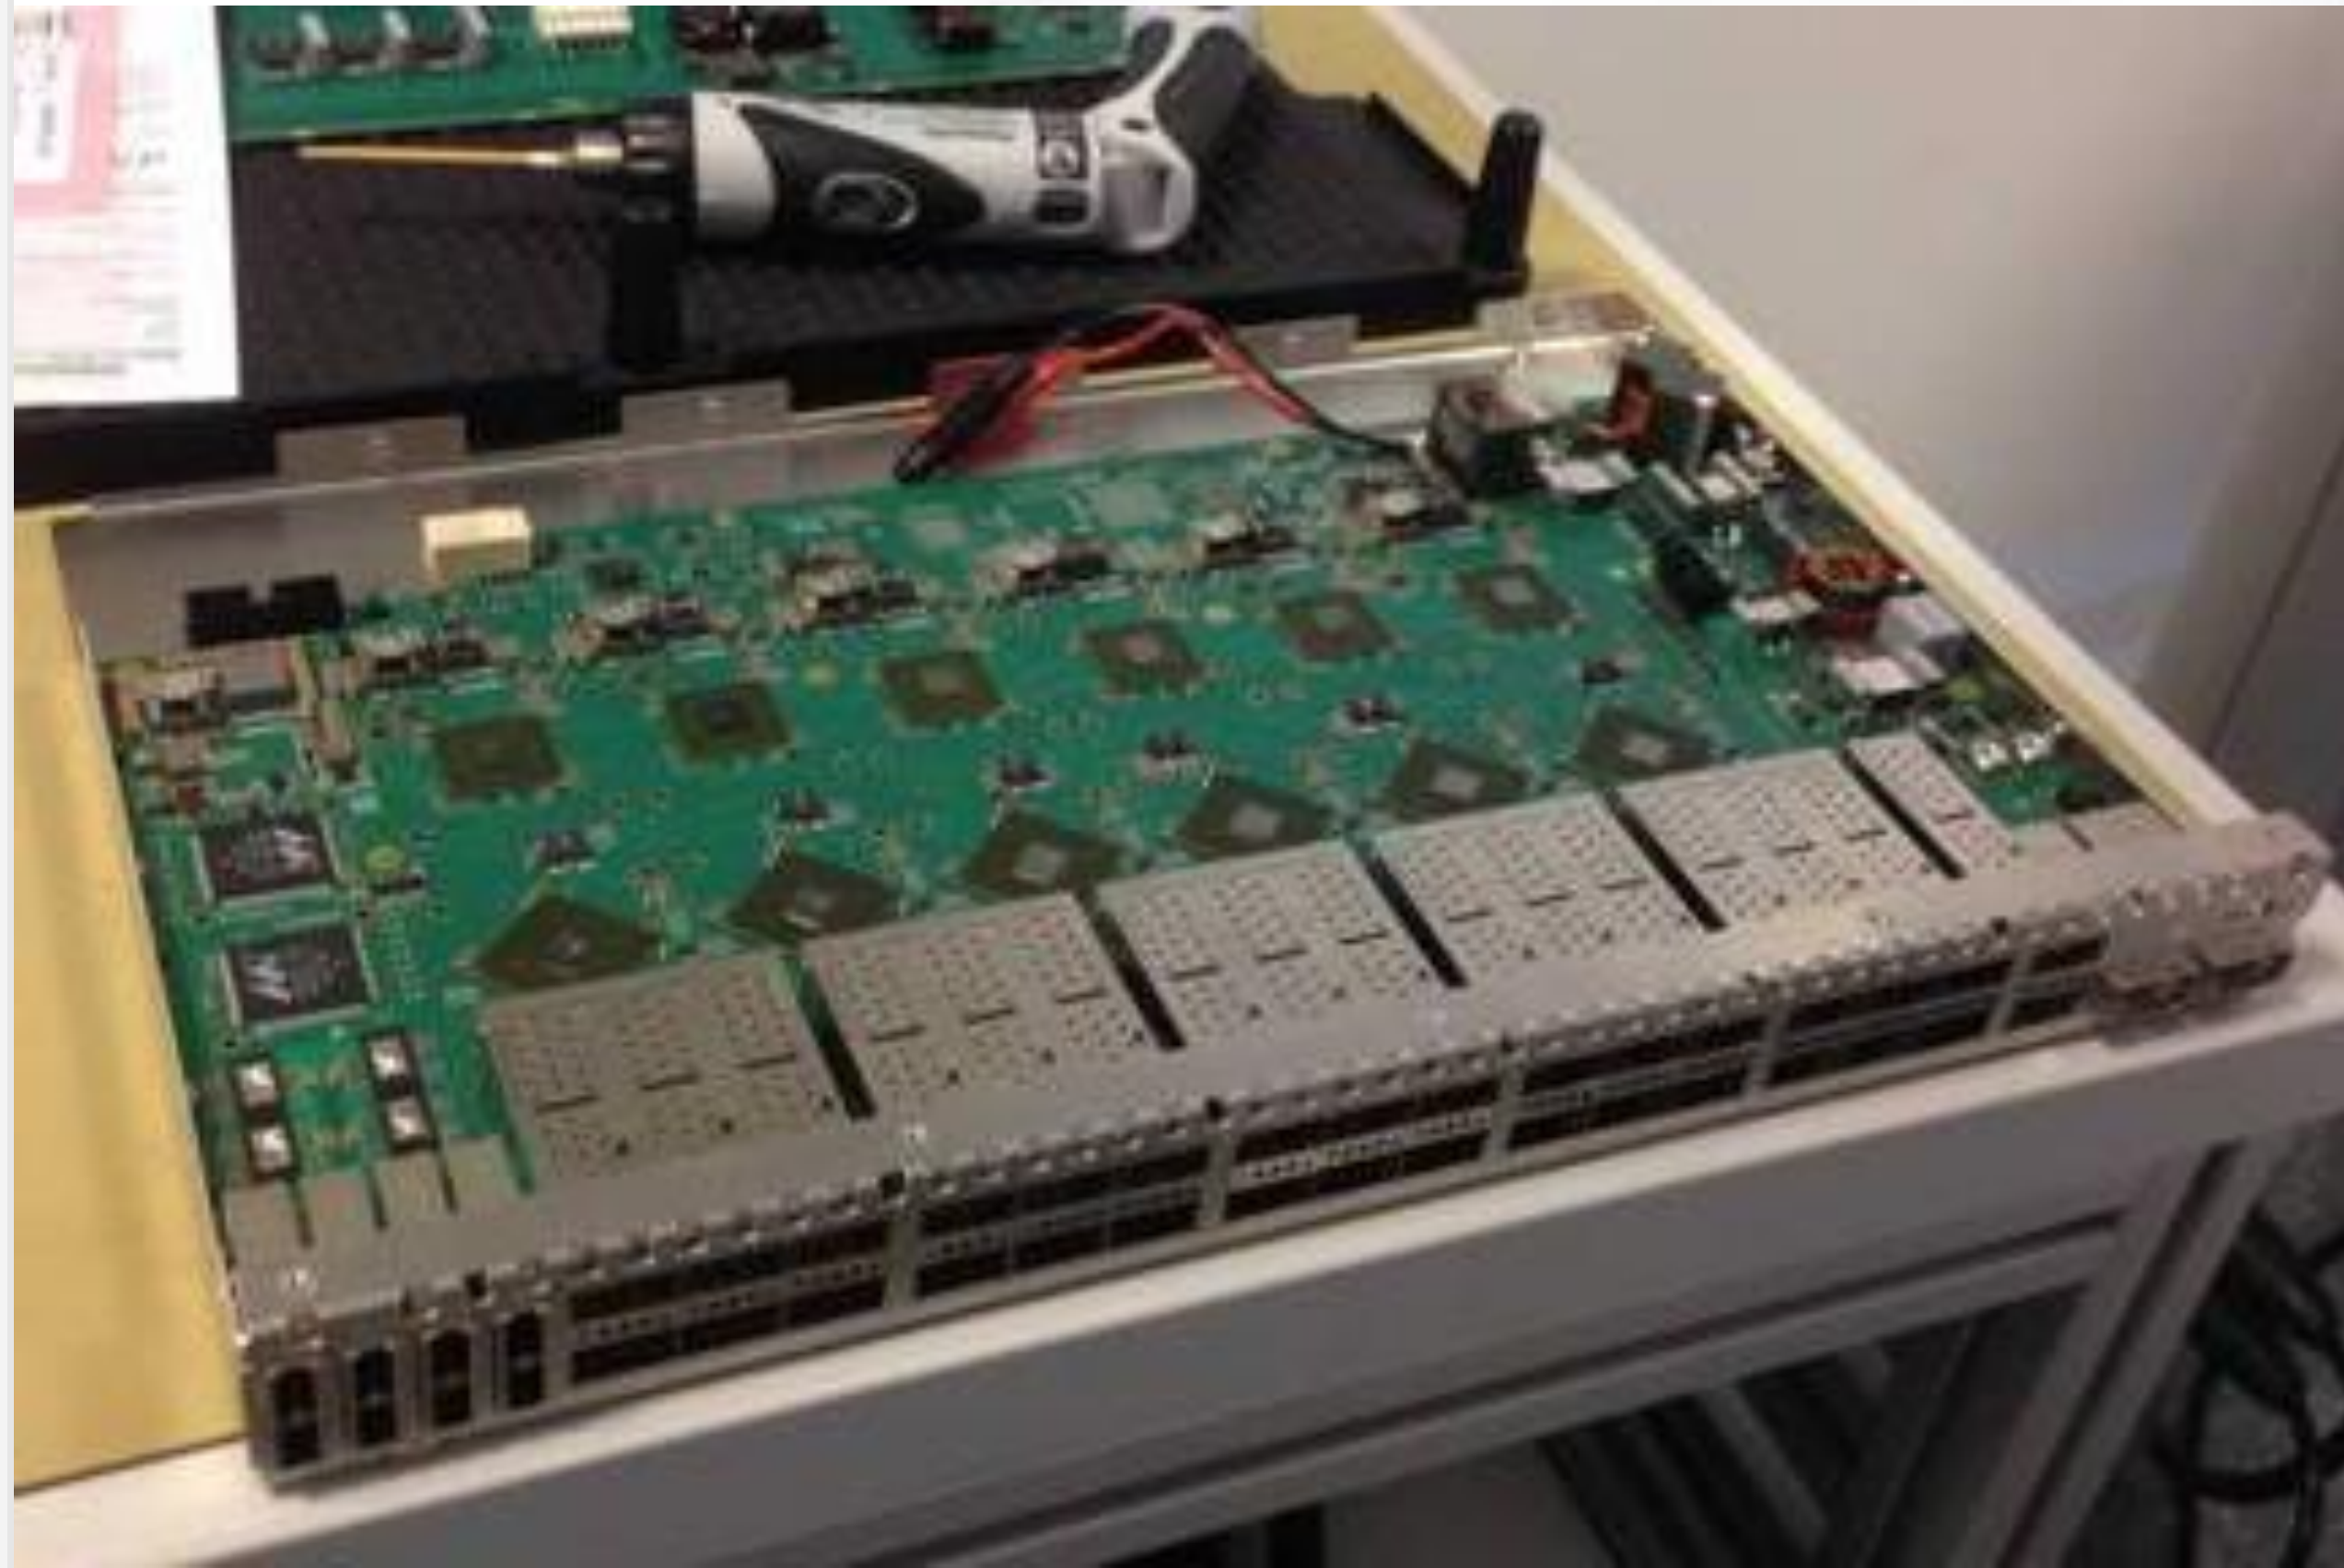
Architecture Overview
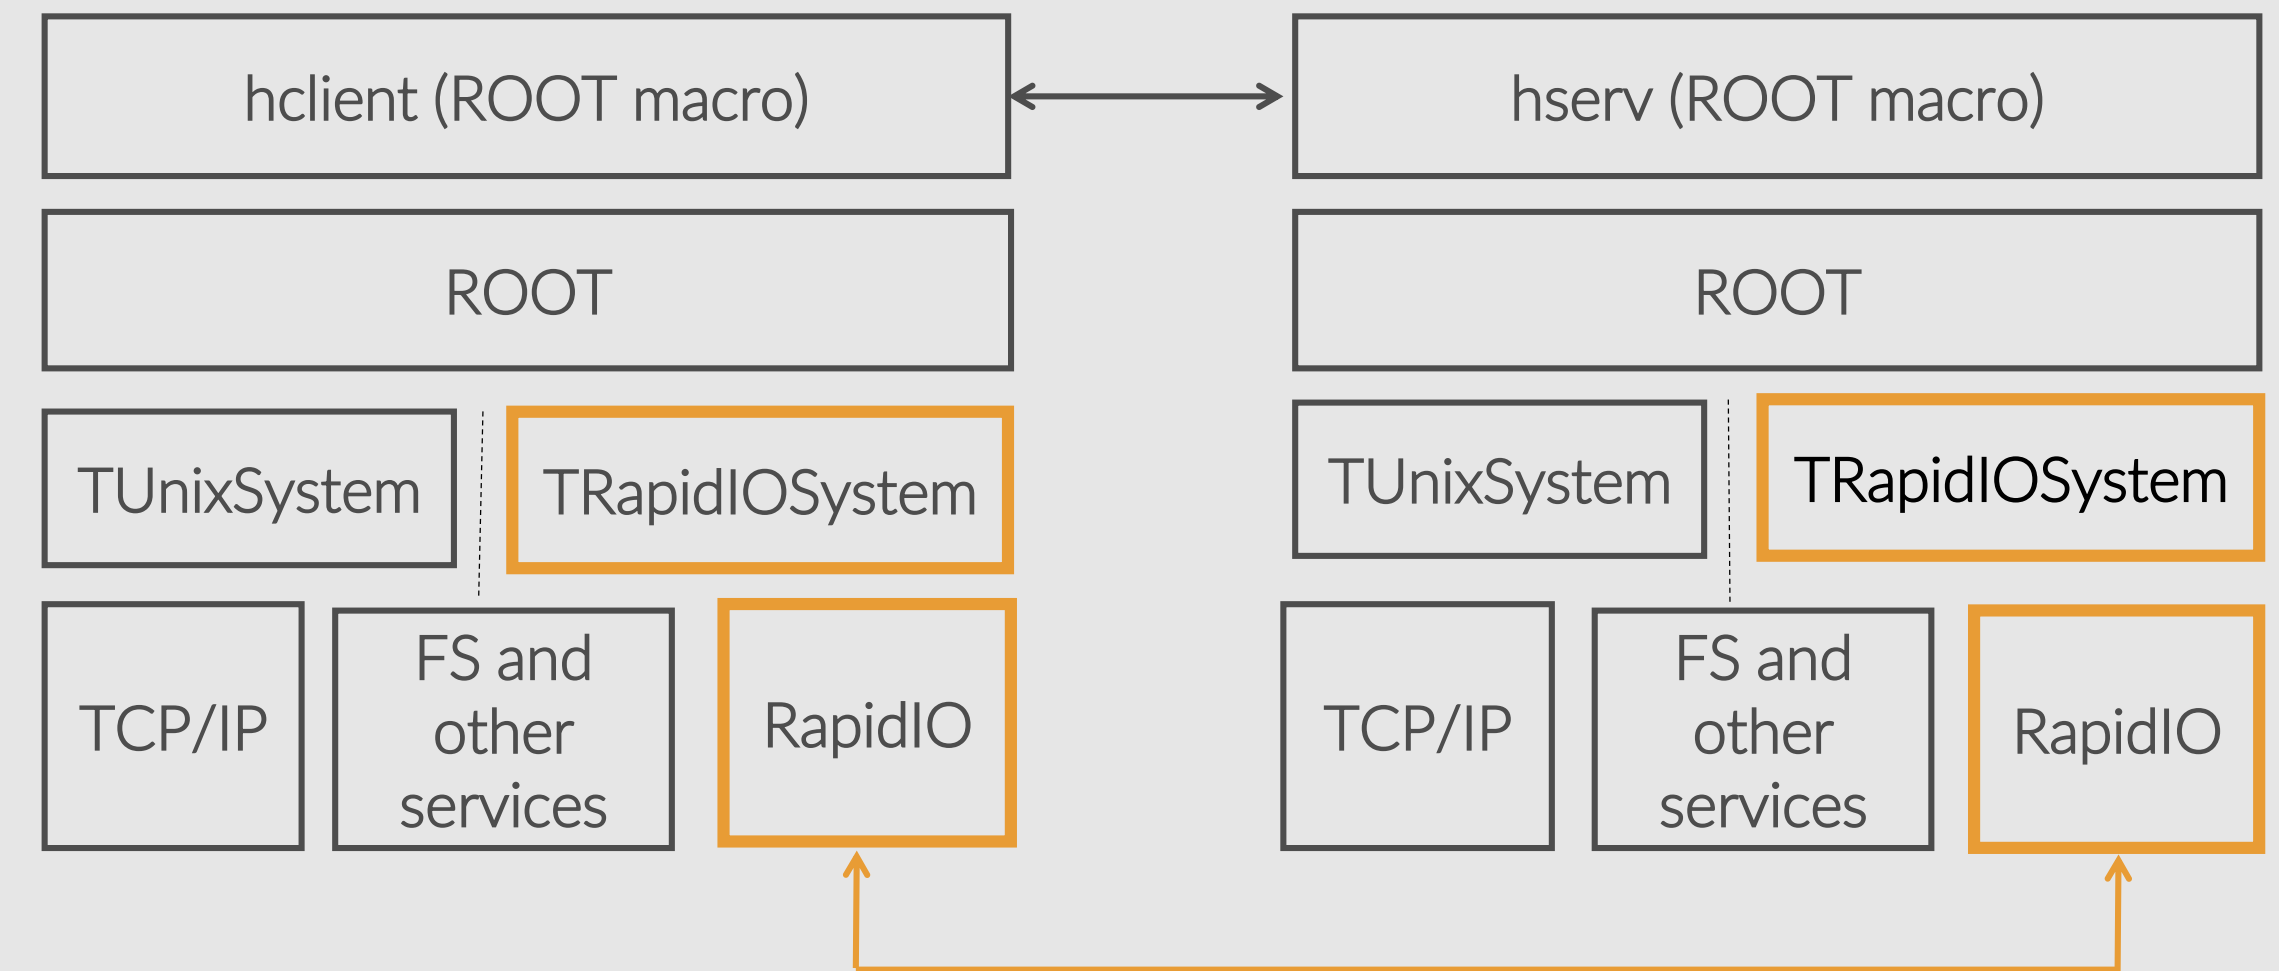
Class Overview

Project Activities

- Visit IDT Ottawa offices
- Project review
- In-depth training
- Target cluster
- Twiki
- Git repo
- JIRA - Kanban

Next Steps

Finalize ROOT proof of concept

PROOF port to RapidIO

Set up 19 inch rack servers and top of rack switch

RapidIO with enterprise analytics framework (Hadoop/Spark)

Benchmark applications

Q&A