

# Composable architectures in the box, in the rack, and rack-to-rack

Presenter: George Apostol, Founder and CEO, Elastics.cloud

# Elastics.cloud

## Profile

Founded: **2020**

Funding to date: **\$26m**  
Pre Series A

Patents filed: **5**

Total product development experience years: **1,500+**

Engineers worldwide: **65+**

Founders



**George Apostol, CEO & Founder**

- 35 years of experience designing system-on-chip (SoC), hardware, software, and systems
- Leadership and executive roles at Xerox/PARC, Sun, SGI, LSI Logic, Exar, PLX Technology, Samsung, TiVo, BRECIS
- Holds several patents for interconnect and interface design



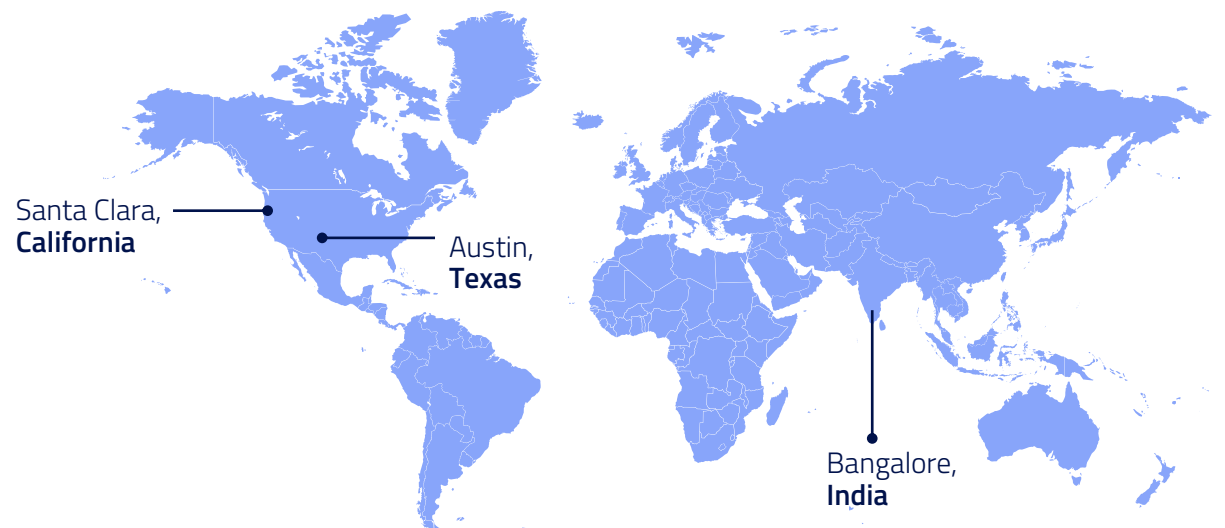
**Shreyas Shah, CTO, Chief Scientist & Founder**

- 25 years of experience in the design of semiconductor, system design, and architecture in fields of computing, networking, storage technologies, virtualization and Flash based storage
- Over 15 patents issued and numerous pending



Flash Memory Summit

Locations



# Industry/Market Pain Points


Addressed by Elastic.cloud



Flash Memory Summit


## Backward Compatible with PCIe

- Hybrid Switch with support of PCIe Gen1 – Gen5 and CXL 1.1/2.0
- CXL 3.0 features with COE




## PCIe/CXL Over PCIe/IB Cables: Bend Radius Limits Scalability

- COE (CXL Over Ethernet™) to scale inside the Rack and Rack2Rack across 32 racks in aisle
- Scalability and resources shared/ pooled across 4,000+ servers



## Slower Media & Disaggregated Memory

- Pre-fetching and caching to reduce the effect of slower media and dis-aggregated memory over the network
- Any CXL ports can be declared as cache port(s)
- Applications communicate with caches for faster responses



## Latency Affecting Application Performance

- Latency one NUMA hop: No change in applications with memory on CXL with switch




## Extending Resource Sharing Across Multiple Fabric Elements

- Memory Pooling: Inside the Rack and Rack2Rack across 32 racks in aisle
- Storage, Networking, GPU accelerators, AI accelerators, Accelerator2directmem™



## GPU Accessing Large Memory Pools Not Enabled with CXL 2.0

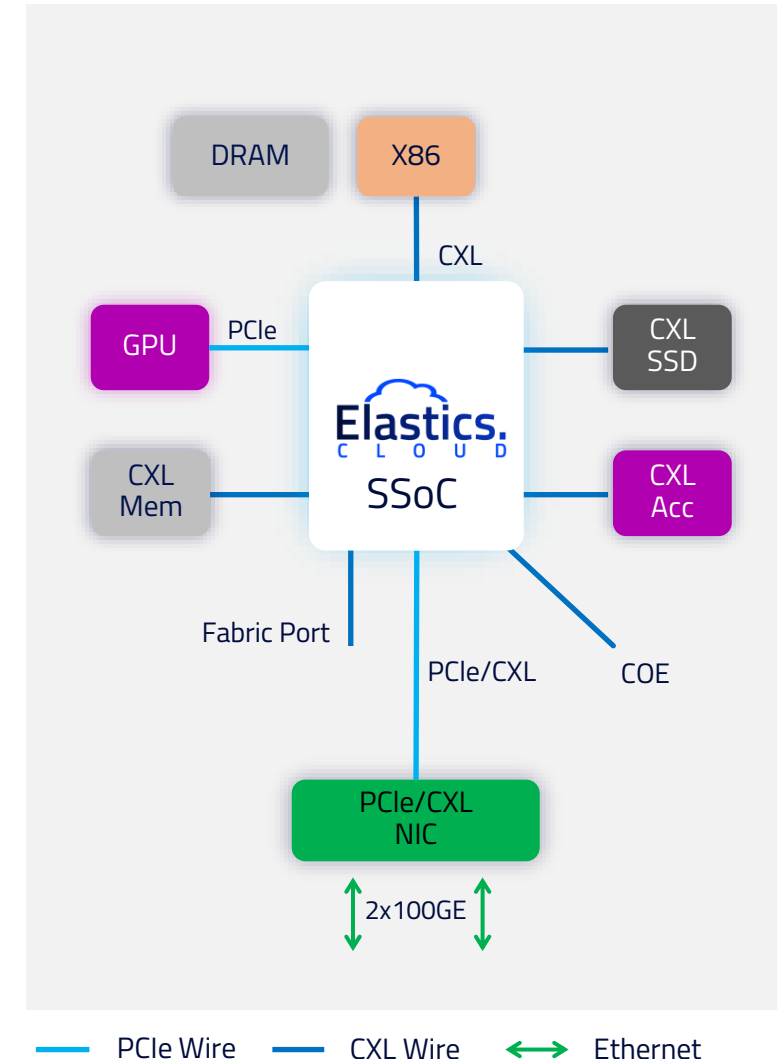
- GPU on PCIe can't access CXL-attached memory without going through CPU
  - Increases latency, limited bandwidth through CPU
- Elastic.cloud Accelerator2directmem™ enables GPU to directly access CXL-attached memory



# Elastics.cloud

## More than a CXL Fanout Switch

Feature	Benefit
256 lanes of PCIe 5.0/CXL 2.x hybrid switch 32 virtual hierarchies across multiple switches	Flexibility for diverse hardware configurations in a multi-tenant environment
Embedded compression and encryption accelerators	Intelligent management of data flows between compute nodes, storage and accelerators
Up to 800 Gb (two 400 Gb, four 200 Gb or eight 100 Gb) COE	Patent-pending technology to establish low latency, highly scalable interconnect using existing network infrastructure. Extends interconnect beyond CXL limit of 4,000 devices while maintaining low latency by bypassing TCP/IP stack
Enterprise class RAS	Enhance CXL Consortium-defined RAS features with advanced error protection, handling and logging, and providing "solder ball to solder ball" data path error detection and correction where feasible
Low Latency switch fabric, AnyPort2AnyPort	Architecture supports more low-latency apps

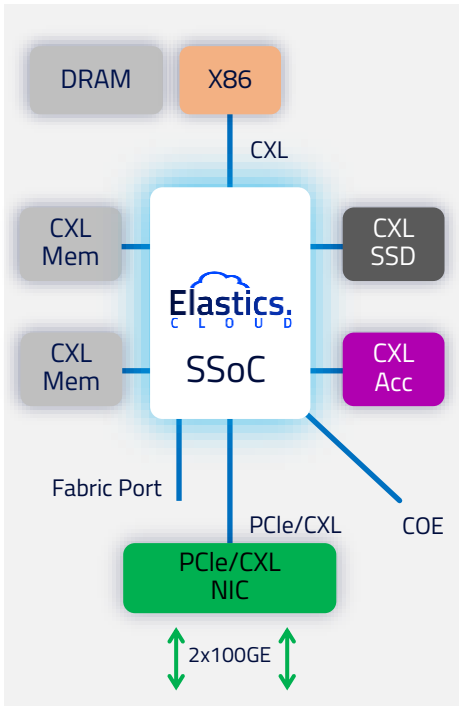


# SSoC Memory and Resource Expansion and Pooling Use Cases

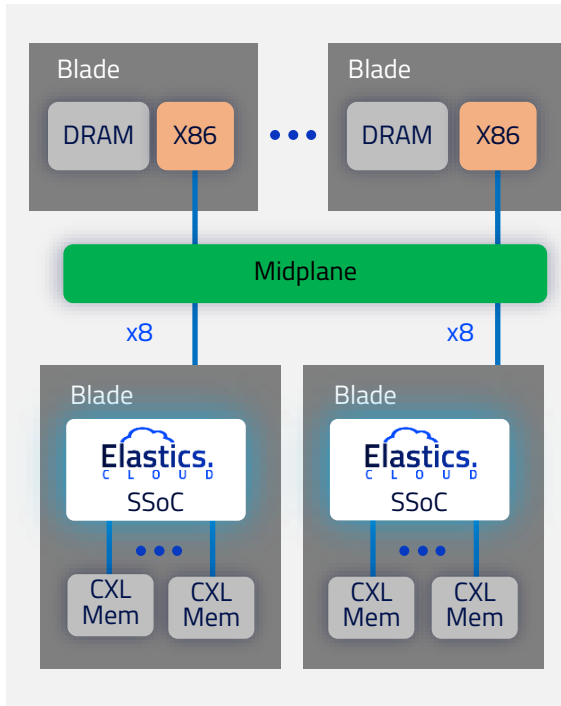


Flash Memory Summit

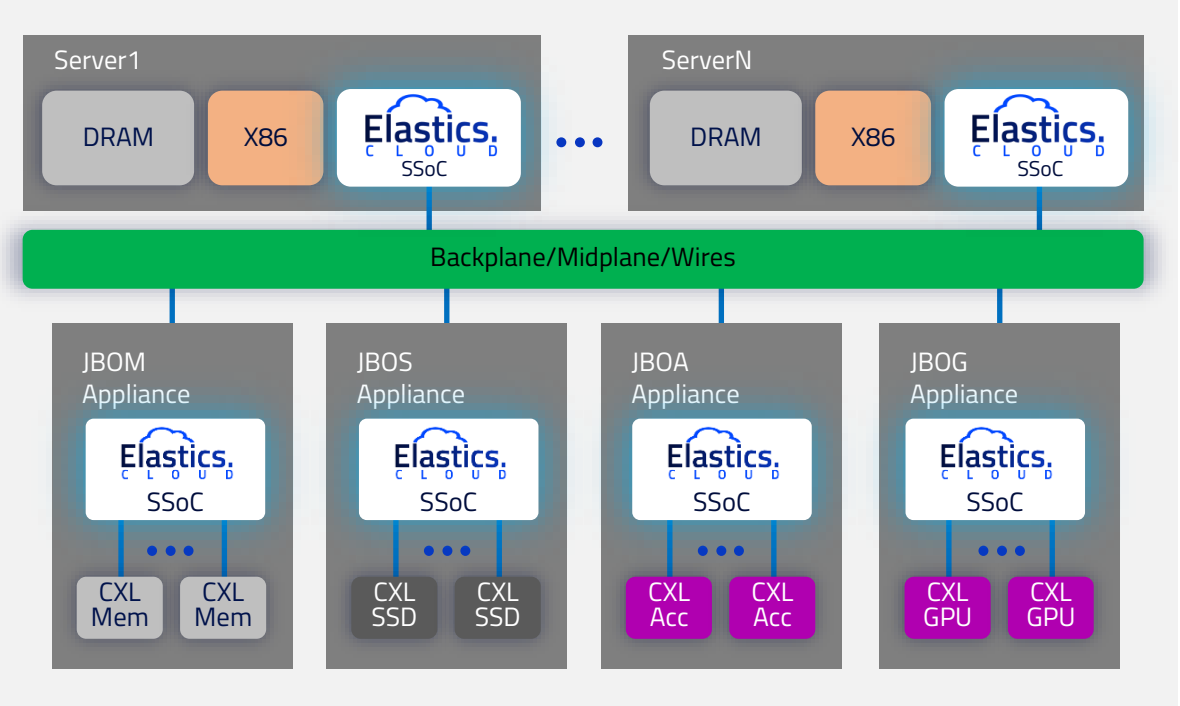
## 1. Memory Expansion in Single Socket Server



## 2. Memory Pooling in Single Chassis



## 3. Resource Pooling in Disaggregated Rack



— PCIe Wire — CXL Wire ↔ Ethernet

- SSoC enables pooling of any resource type
- Disaggregation and pooling allow dynamic composition of server resources for a given workload
- Multiple hosts can share resources from resource pools
- Managed by a single management host/BMC

# CXL Over Ethernet (COE) Implementation

- COE: CXL protocols over Ethernet
- Advantages
  - Seamless support for existing Ethernet infrastructure
  - Intra-rack and inter-rack reach
  - Supports existing Ethernet switches in the market and future low latency Ethernet switches
- Support for DCBx, QCN and replay buffers on each side of COE implementation for lossy Ethernet network


# COE Implementation in Rack Systems

Composability within the rack

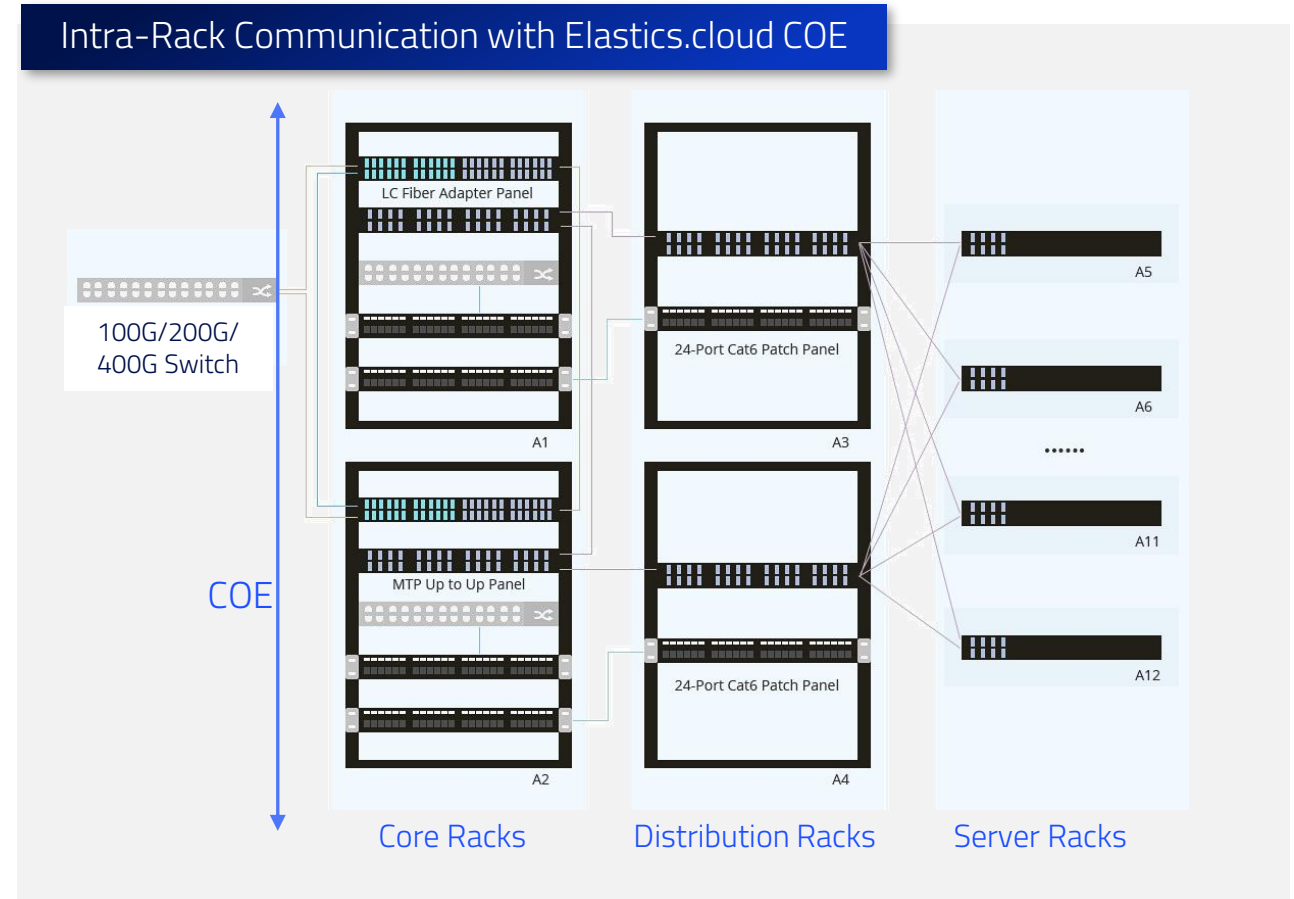


Flash Memory Summit

## Components Inside the Rack



## Intra-Rack Communication with Elastics.cloud COE



Ethernet wire eliminates cable bend radius limitations  
COE enables resource sharing/memory pooling within the rack



# COE Implementation in Rack Systems

Composability Rack2Rack



Flash Memory Summit

GPUs & Memories Used for  
ML/AI Training Clusters



Elastics.cloud's COE technology enables sharing of GPUs & Memories

Inter-Rack Communication with Elastics.cloud COE

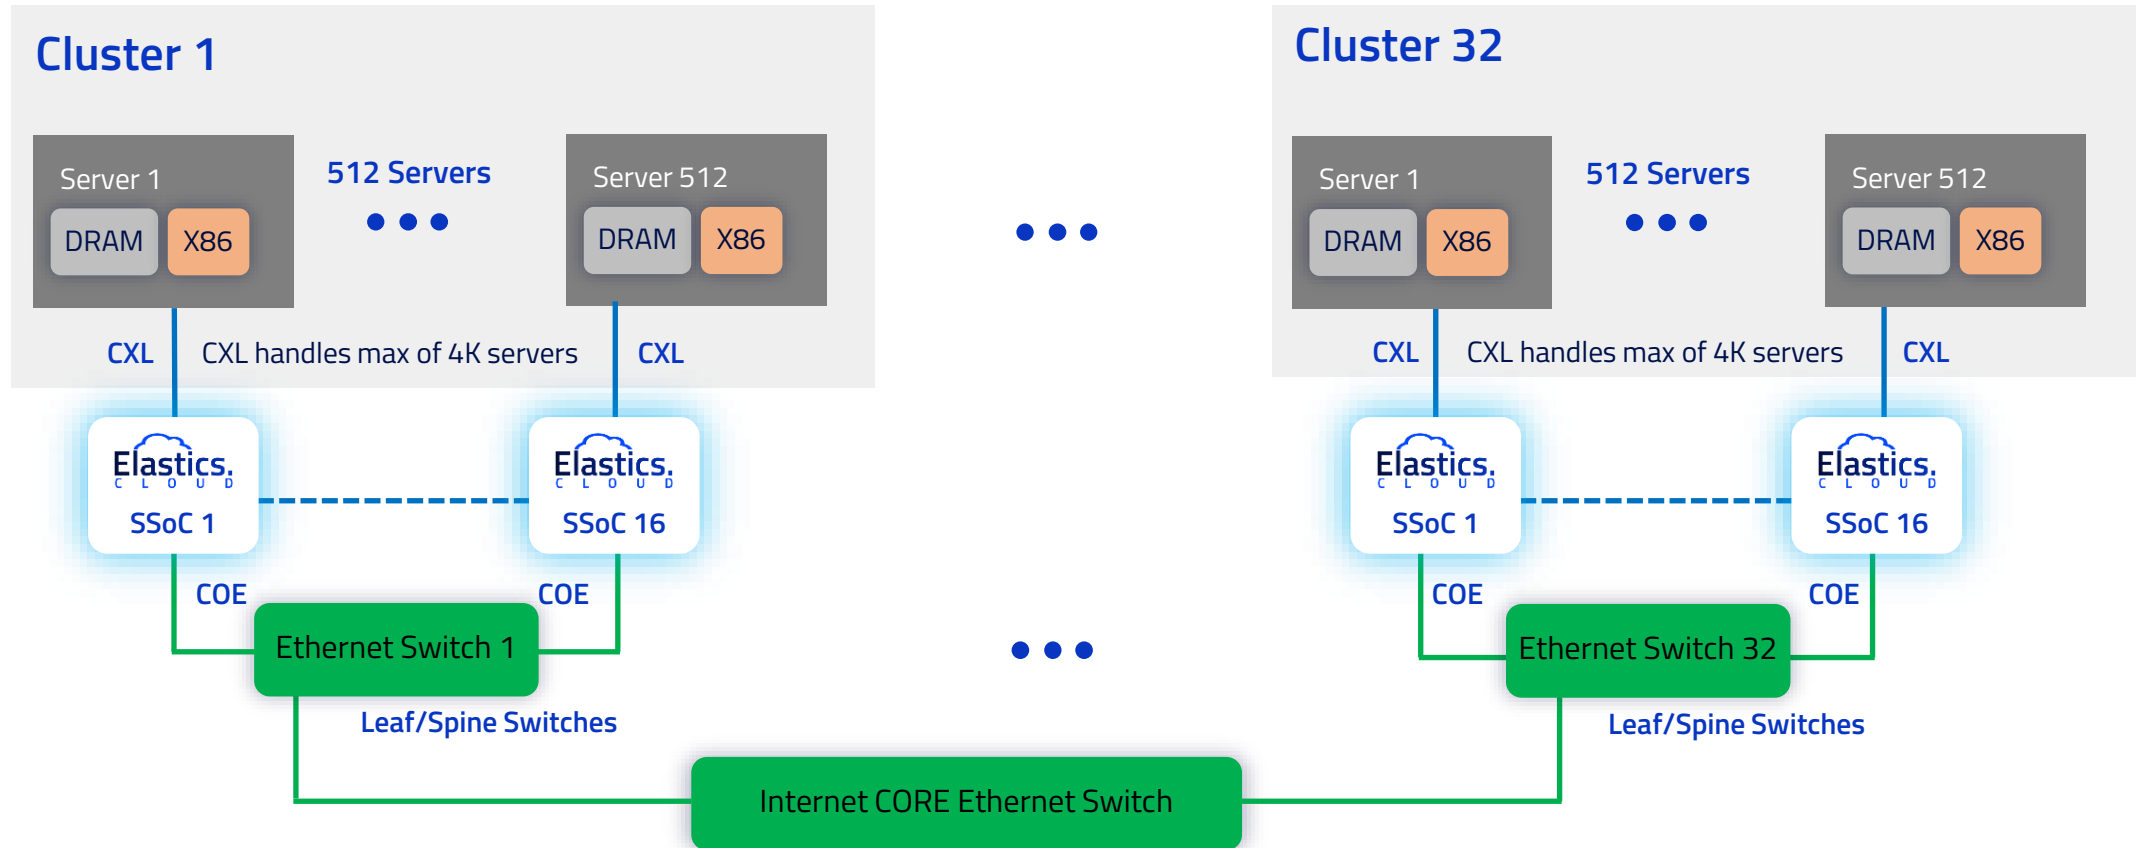


Elastics.cloud's technology reduces the effect of slower media and disaggregated memory over the network



# Cluster Connectivity Using CXL and COE

16K servers connected (32 clusters of 512 servers)



COE™ extends to hundreds of thousands of servers seamlessly

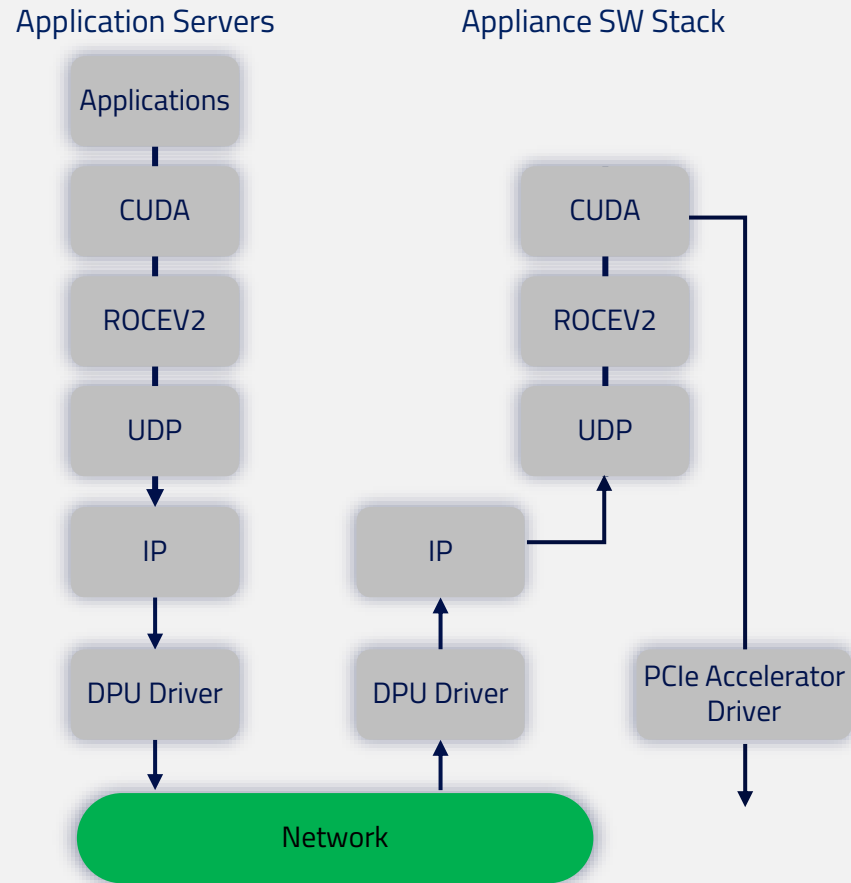
# ML/AI Training Software Stack with COE

Eliminates the need for ROCEV2/V3

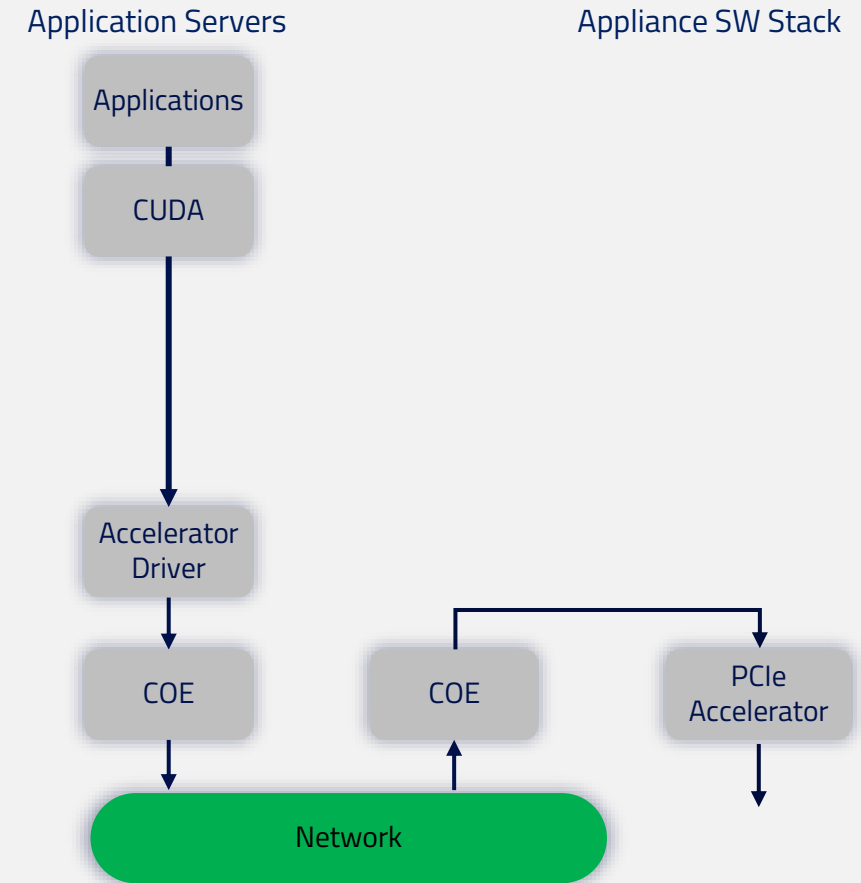


Flash Memory Summit

## Current Implementation



## With Elastics.cloud



Reduces complexity of software  
Significantly reduces latency through software stack

# FM API Manages Resources:

in the Box, in the Rack, and Rack2Rack

Switch management is done through the BMC/host using standard FM APIs

Multiple COE connected switches can be managed by single host/BMC

REST API with Elastics.cloud extensions supports multiple switches connected over COE

FM API with Elastics.cloud extensions can be easily integrated with orchestrators to create specific configurations for dynamic workloads e.g., Kubernetes

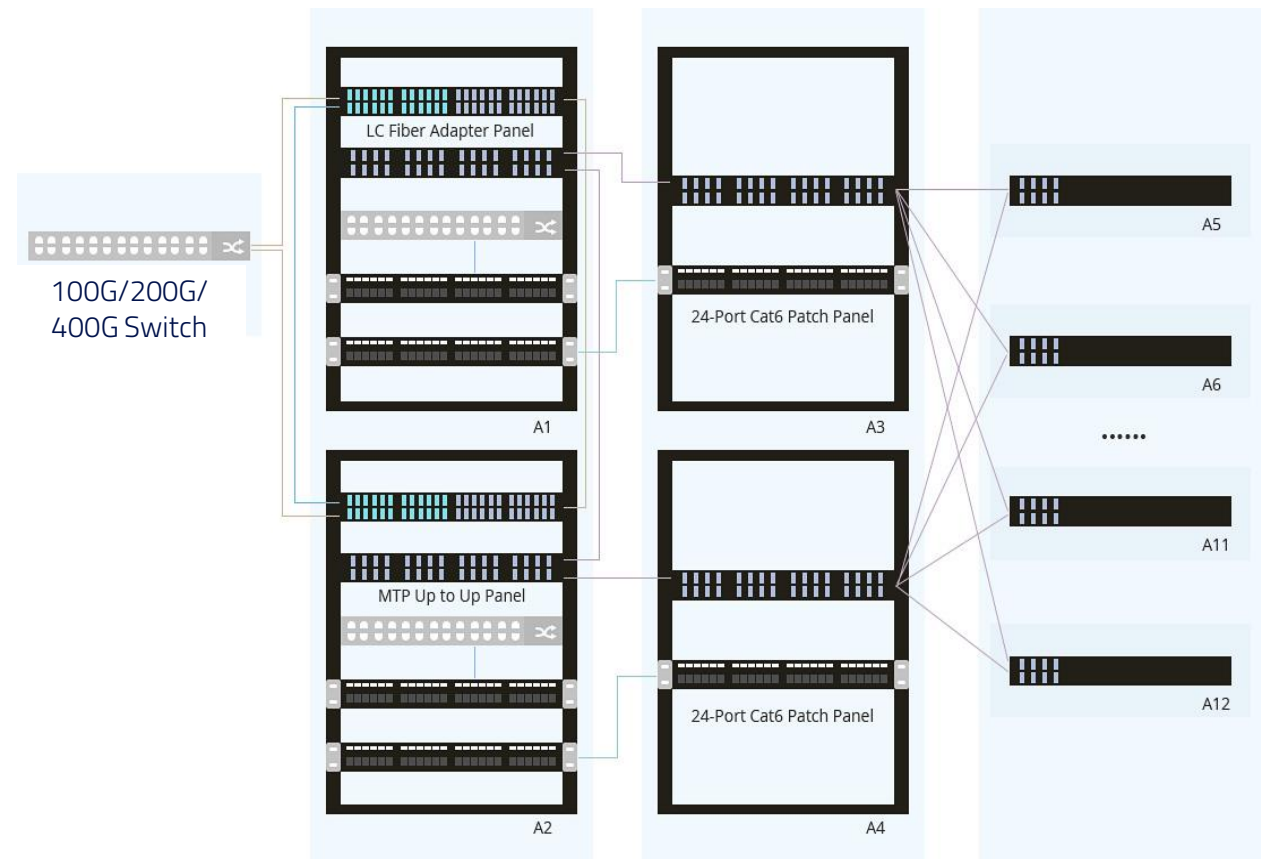
16K servers connected (32 clusters of 512 servers) via CXL and COE can be managed by Elastics.cloud's solutions

Elastics.  
CLOUD  
SSoC



Flash Memory Summit

FM API/Elastics.cloud extension orchestration tool



# Summary



Elastics.cloud is highly focused on bringing CXL-based composable infrastructure to market



Key features of Elastics.cloud's solutions:

- Backward compatibility with PCIe Gen5
- Lowest ball-to-ball latency in industry
- COE scales CXL inside the rack and rack-to-rack
- Solves industry pain points



Elastics.cloud's solutions enable system level optimizations

- System level build of JBOMs, JBODs, JBOx
- Dynamic assignment of pools of resources (CPU, memory, storage, accelerators) connected via CXL
- Simplified management infrastructure