SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

CXL 3.x and beyond: Converged Fabric for Heterogeneous Architectures

Presented by

Shreyas Shah – Cofounder and CTO, Elastics.cloud

Agenda

- Data growth and need for heterogeneous architectures
- Compute Express Link (CXL) introduction
- Heterogeneous architectures with CXL
- Elastics.cloud mission
  - Elastics.cloud's Next Gen Architectures
- CXL Switch use cases
- Applications on current FPGA platform
- Summary and next steps

Data Growth and Heterogeneous Architectures

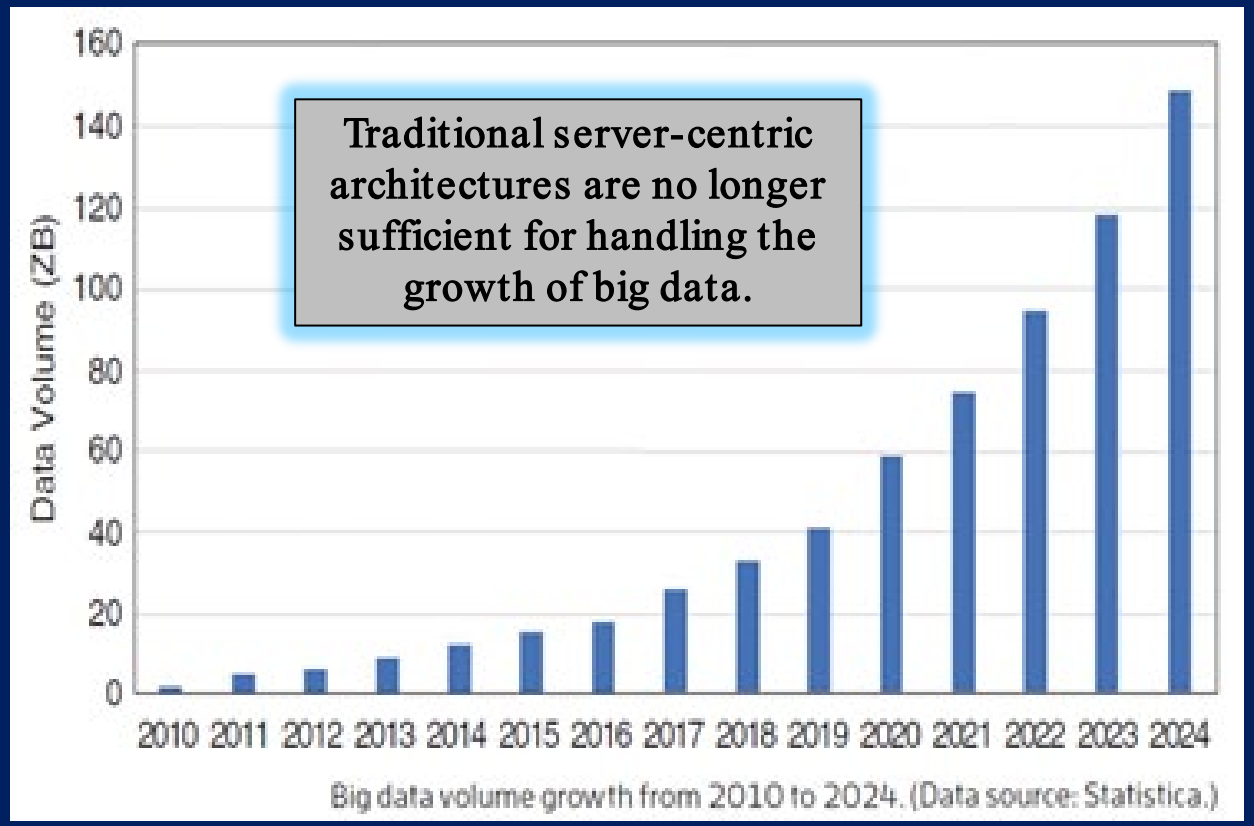

Growth in big data requires additional computing & memory resources

Processing

- CPUs
- GPUs
- AI ASICs
- Smart NICs

Storage

- Smart SSDs
- Smart Memory

Networking

- Ethernet

100 GE → 400 GE → 800 GE → 1.6T → 3.2T → 6.4T

Compute Express Link (CXL) Introduction

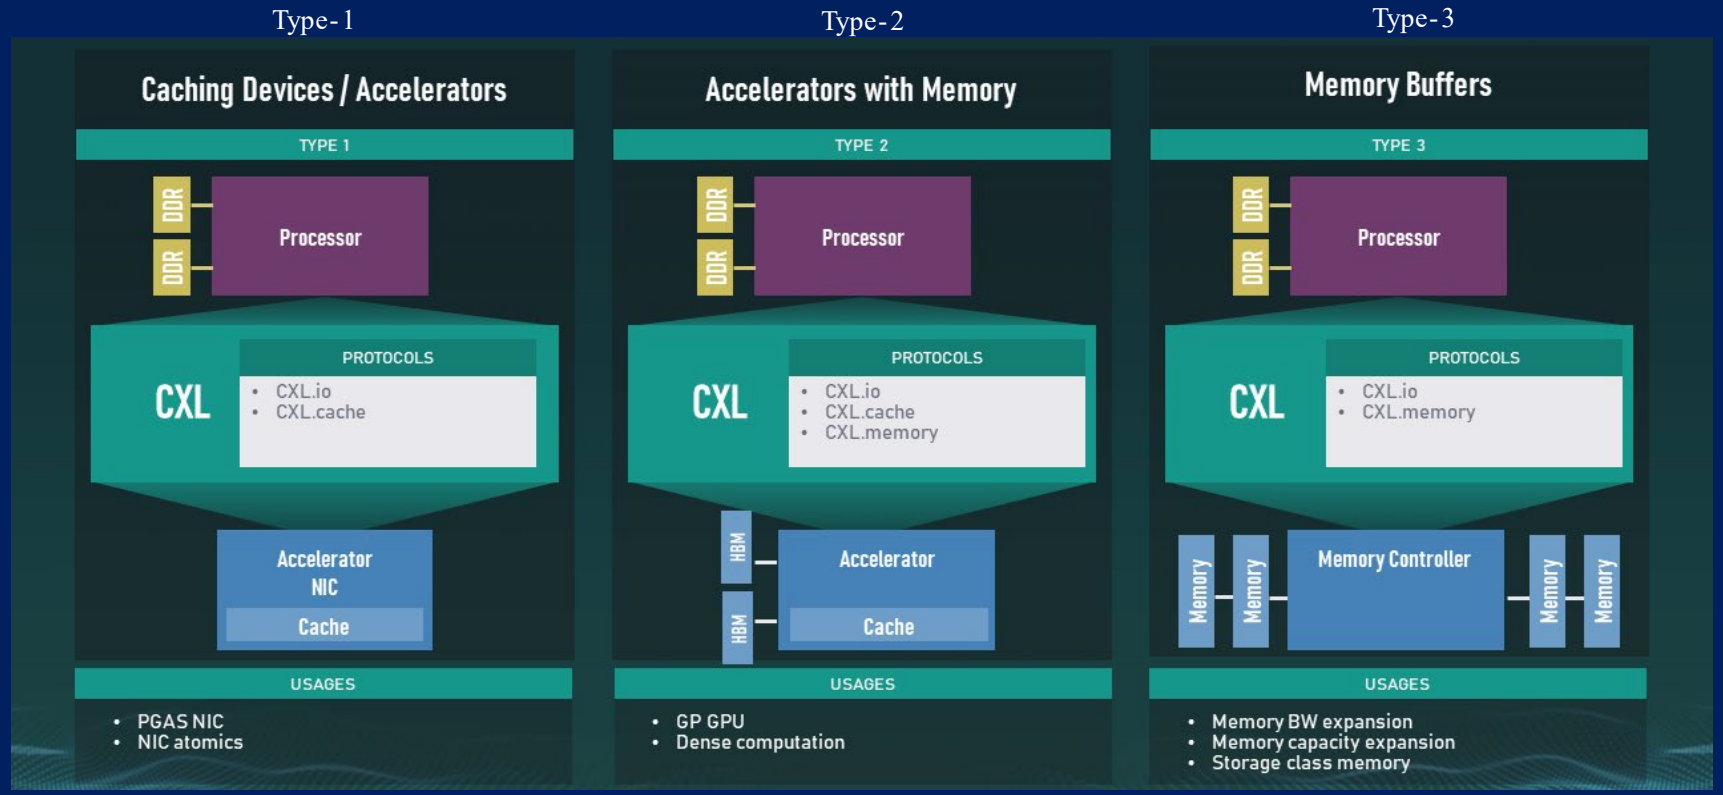

- CXL (Compute Express Link) is a protocol to connect accelerators, memory and devices to processor
- Reduce TCO, increased performance, layered on top of PCIe electrical
- Lower latency and higher bandwidth
- No ordering maintained like PCIe
- Common address map across all the devices – removes the CPU as a central processing in server architectures – helps with heterogeneous architectures

CXL Feature Introduction

Features	CXL 1.0/1.1	CXL 2.0	CXL 3.x
Release date	2019	2020	2023
Max link rate	32GTs	32GTs	64GTs
Flit 68 byte (up to 32 GTs)	✓	✓	✓
Flit 256 byte (up to 64 GTs)			✓
Type 1, Type 2 and Type 3 Devices	✓	✓	✓
Memory Pooling w/ MLDs		✓	✓
Global Persistent Flush		✓	✓
CXL Integrity and Data Encryption (IDE)		✓	✓
Switching (Single-level)		✓	✓
Switching (Multi-level)			✓
Direct memory access for peer-to-peer			✓
Enhanced coherency (256 byte flit)			✓
Memory sharing (256 byte flit)			✓
Multiple Type 1/Type 2 devices per root port			✓
Fabric capabilities (256 byte flit)			✓

- CXL 1.0/1.1 – introduction to CXL technology in market
- CXL 2.0 – supports switching
- CXL 3.x – composability and sharing of components for heterogeneous architectures

Heterogeneous Architectures with CXL

CPU-centric Architectures

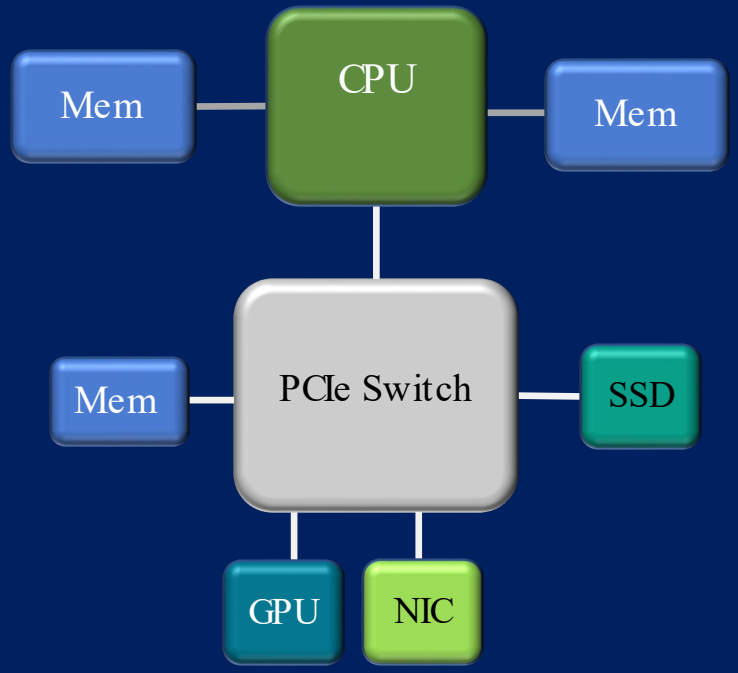

Datacentric Architectures

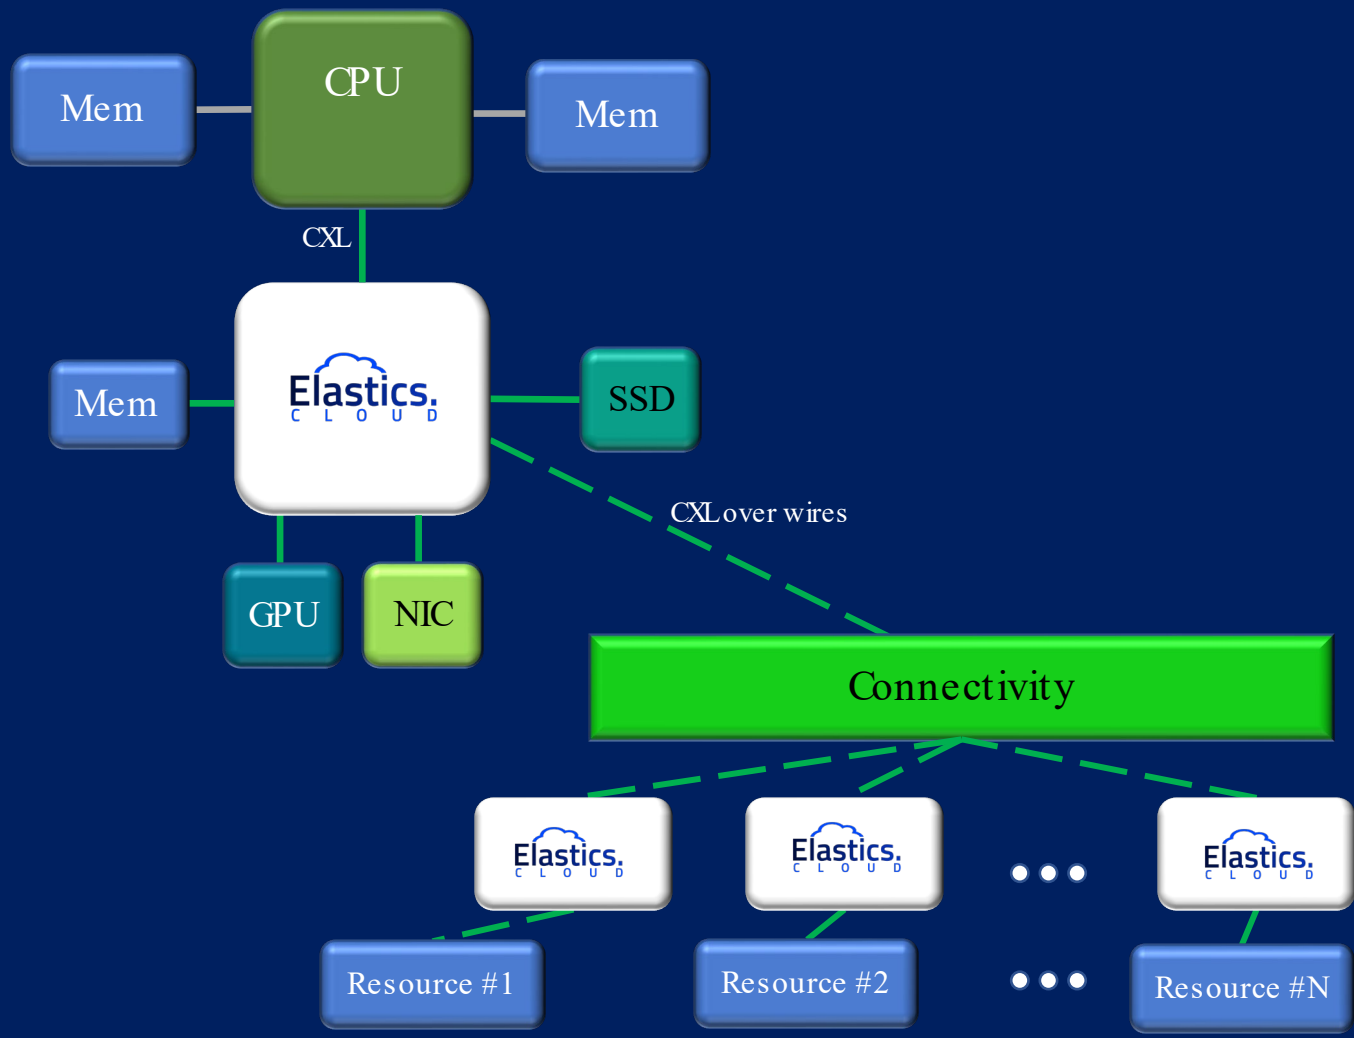

PCIe
DDR
CXL

Elastics.cloud Mission

Leveraging the Compute Express Link (CXL) interconnect to deliver new composable, heterogeneous architectures

Future Rack-Scale Architecture

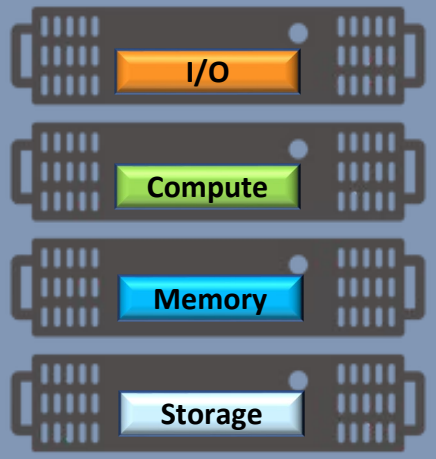

- Improved utilization of components
- Flexibility to right-size resources to match workloads
- Component expansion with high-speed, low-latency interconnect
- Multi-host, shared component enablement
- Intelligent buffering and traffic management
- Software-managed orchestration

Composable set of pooled, shared and disaggregated compute, network, memory, & storage resources

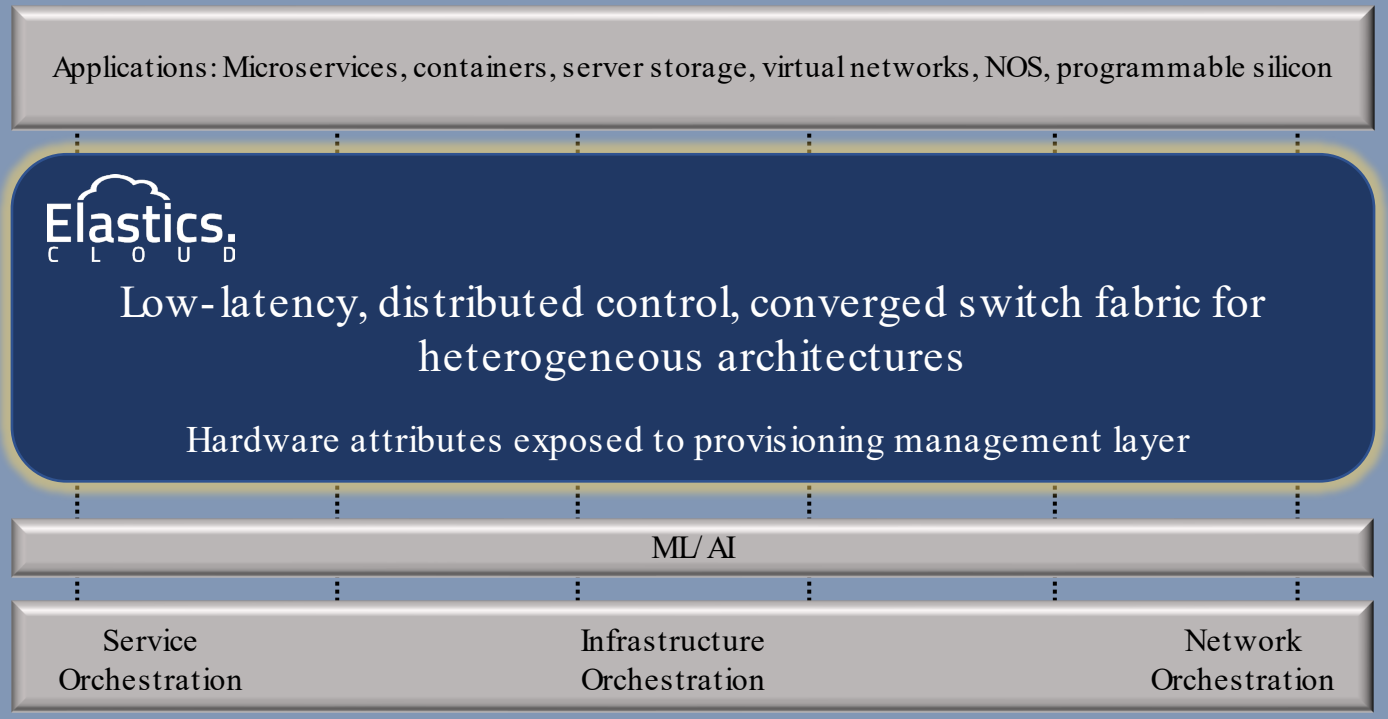

CXL based components intelligently composed with software orchestration to optimize performance and TCO

Elastics.cloud's Next Gen Architectures

Intelligent, highly performant switching at the intersection of compute, memory, storage, and networking

Industry Challenges

- Siloed servers in data centers due to unique memory requirement per application
- Underutilized CPU cores due to number of memory controllers on CPU socket
- Stranded memory, accelerators, processors, storage, and other server resources

Elastics.cloud's Solution

Total Cost of Ownership

- Resource pooling
- Elimination of stranded resources

CXL-Enabled Composability

- Resource disaggregation and pooling
- Type 1, 2, & 3 devices
- CXL 1.1, 2.0, PCIe

Elastics. CLOUD

Performance

- Memory expansion
- Low-latency
- Intelligent control

Scalability

- Server resource expansion
- Resource pooling within the rack
- Resource pooling rack-to-rack

Memory and Resource Expansion And Pooling Use Cases

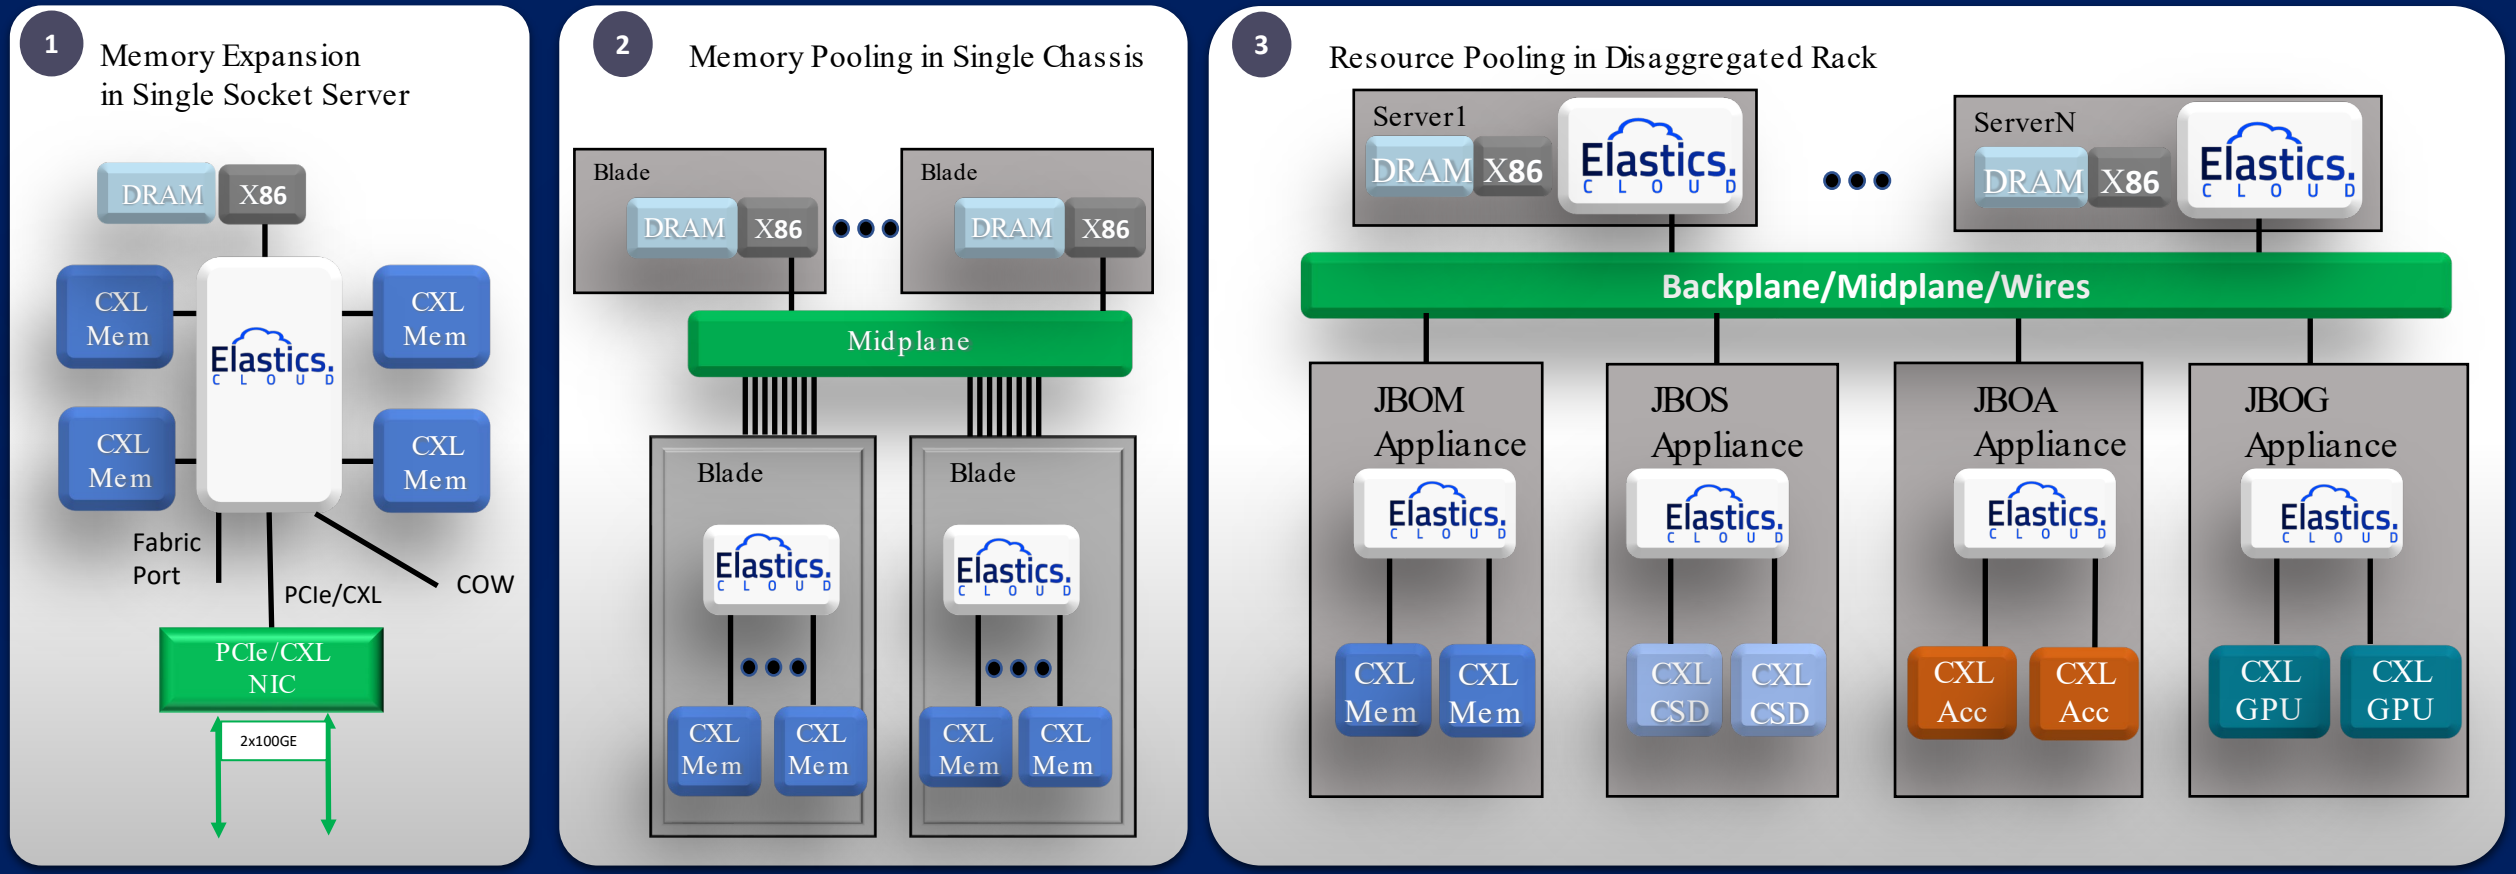

- SSoC enables pooling of any resource type
- Disaggregation and pooling allow dynamic composition of server resources for a given workload
- Multiple hosts can share resources from resource pools

Applications on Current FPGA Platform

Redis Labs

Latency up to 22x lower

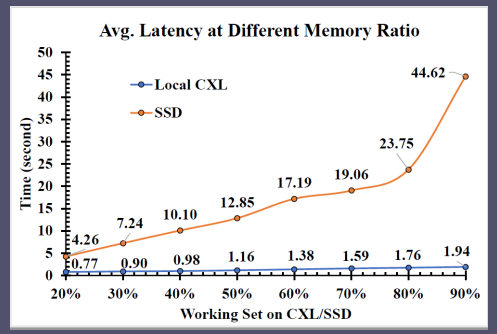

Ops/sec up to 23x higher

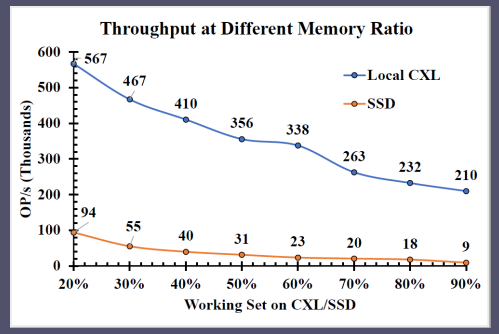

Bandwidth up to 24x higher

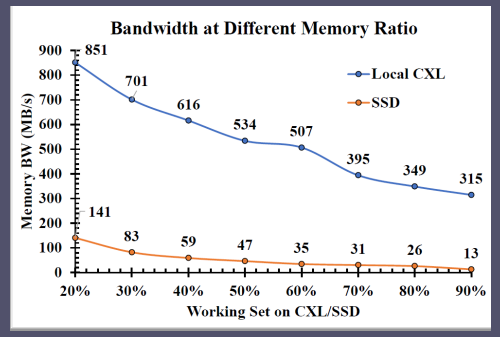

RocksDB

Ops/sec up to 12.6x higher

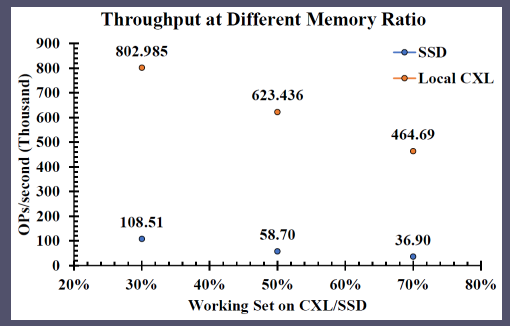

Completion time up to 1.74x lower

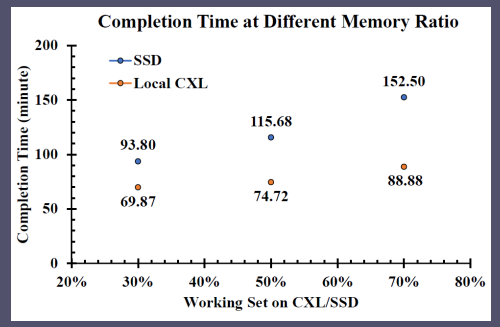

- In-memory databases benefit from CXL-based architecture
  - Redis Labs
  - RocksDB
- Datasets increase in size, swap from SSD to CXL - 20%to 90%
- Performance difference gets higher as more memory moves from CPU to CXL

Summary and Next Steps

- CXL is widely accepted as interconnect of choice for all server components
- CXL is a true converged fabric connecting CPU, GPUs, AI ASICs, storage, NICs
- CXL evolution from 1.1 to 3.x
  - Backward compatibility with PCIe electricals
  - CXL 3.x has necessary features for resource sharing across multiple servers, reducing TCO at scale
- Composable architectures require CXL
  - CXL enables low-latency interconnect across heterogeneous components
  - Ensures connectivity of server resources within rack and rack-to-rack
- For more information, contact us: sales@elastics.cloud

SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Please take a moment to rate this session.

Your feedback is important to us.

Post-Summit, visit www.snia.org/cms-summit for additional content.

