



PERSISTENT MEMORY

PIM SUMMIT

JANUARY 24, 2019 | SANTA CLARA, CA

## New Interconnects

Moderator: Andy Rudoff, SNIA NVM Programming Technical Work Group and Persistent Memory SW Architect, Intel



**SNIA** PERSISTENT MEMORY  
**PM** SUMMIT  
JANUARY 24, 2019 | SANTA CLARA, CA


CCIX: Seamless Data Movement for Accelerated Applications



Millind Mittal  
CCIX Consortium

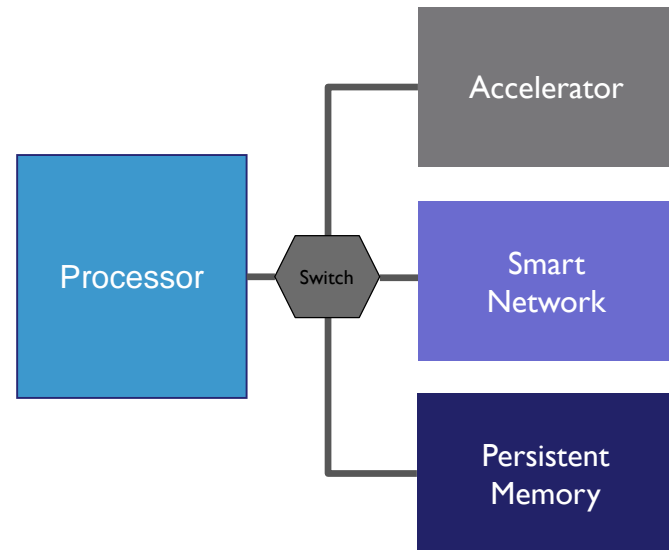
# CCIX: Cache Coherent Interconnect for Accelerators

- A new class of interconnect for accelerated applications
  - ◆ Accelerators are viewed as NUMA sockets(peers to processors) instead of IO interface connected devices
- Mission of the CCIX Consortium is to develop and promote adoption of an industry standard specification to enable coherent interconnect technologies between general-purpose processors and acceleration devices for efficient heterogeneous computing.



# Accelerators Trends

- More work being offloaded to accelerators
  - ◆ ML, video/audio, analytics
- Accelerators are shifting from devices with function offload to devices that augment host processor resources
  - ◆ Self hosted accelerator
  - ◆ System memory expansion (DRAM and Persistent Memory) with near-memory processing
  - ◆ Compute Augmentation
- Increased inter-node data transfers
  - ◆ Resource pooling
  - ◆ Resource dis-aggregation



# The CCIX Consortium

- Incorporated in 2017
- 50 Members & Growing
- Complete Ecosystem
- Join CCIX Now!  
([www.ccixconsortium.com](http://www.ccixconsortium.com))

Promotors			
			
			
Contributors			
			
			
			
			
			
			
			
Adopters			
Arteris, Inc.  Baikal Electronics  Bitman Technologies, Inc.  Chengdu Higon Integrated Circuit Design Co., Ltd.  Ericsson AB  Guizhou Huaxintong Semiconductor Technology Co. Ltd.  INVECAS, Inc.  Nokia Solutions and Networks Oy  Parade Technologies, Inc.  Phytium Technology Co., Ltd.  PLDA  Shanghai Zhaoxin Semiconductor Co., Ltd.  Silicon Laboratories Inc.  SmartDV Technologies India Private Ltd.  Socionext Inc.  Viavi Solutions, Inc.  Wave Computing			

# Availability of CCIX Base Specification 1.0



NEWS

## CCIX™ Consortium Enables Next Generation Compute Architectures with the Availability of Base Specification 1.0

*Specification Release Allows Companies to Deliver CCIX Production Devices*

**Beaverton, Ore – June 18, 2018** – The CCIX™ Consortium today announced the release of the CCIX Base Specification 1.0. The CCIX specification enables a new class of high performance, low latency cache coherent interconnect for the next-generation cloud, artificial intelligence, big data, database and other datacenter infrastructures.



*“The CCIX ecosystem is providing the industry with a flexible interconnect to deliver true peer processing in cache coherent topologies with improved performance over existing interconnect technologies. We are beginning to see the availability of the first products with support for the CCIX 1.0 specification and expect this strong adoption to continue growing with the availability of the production version of the specification.”*

–Gaurav Singh  
CCIX Consortium Chairman



# THANK YOU



PERSISTENT MEMORY

PIM SUMMIT

JANUARY 24, 2019 | SANTA CLARA, CA

## Gen-Z Fabric

Kurtis Bowman

Director of Architecture and Technology, Dell



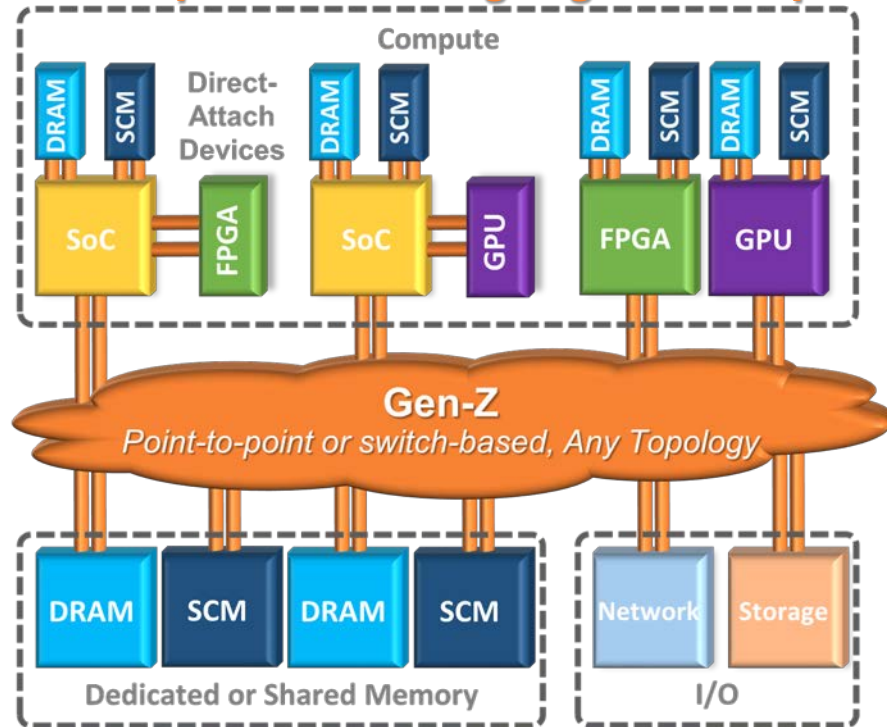
# About The Gen-Z Consortium

- The Gen-Z Consortium launched in October of 2016 to create an open, industry standard for a high speed, low latency, scalable, memory centric fabric
- There are currently 60+ member companies covering all of the disciplines required to create a vibrant ecosystem
- Gen-Z has shown demos of memory pooling with multiple servers over the last year
- Gen-Z members have released design IP and silicon vendors have started designs for Gen-Z devices
- All released and selected draft specifications are available on [www.GenZConsortium.org](http://www.GenZConsortium.org) for public review and comment


# Gen-Z Connects Disaggregated Components

- High Performance
  - High Bandwidth, Low Latency, Scalable
  - Eliminates protocol translation cost / complexity / latency
  - Eliminates software complexity / overhead / latency
- Reliable
  - No stranded resources or single-point-of-failures
  - Transparently bypass path and component failure
  - Enables highly-resilient data (e.g., RAID / erasure codes)
- Secure
  - Provides strong hardware-enforced isolation and security
- Flexible
  - Multiple topologies, component types, etc.
  - Supports multiple use cases using simple to robust designs
  - Thorough yet easily extensible architecture
- Compatible
  - Use existing physical layers, no OS modifications required
- Economic
  - Lowers CAPEX / OPEX, unlocks / accelerates innovation

## Gen-Z speaks the language of compute

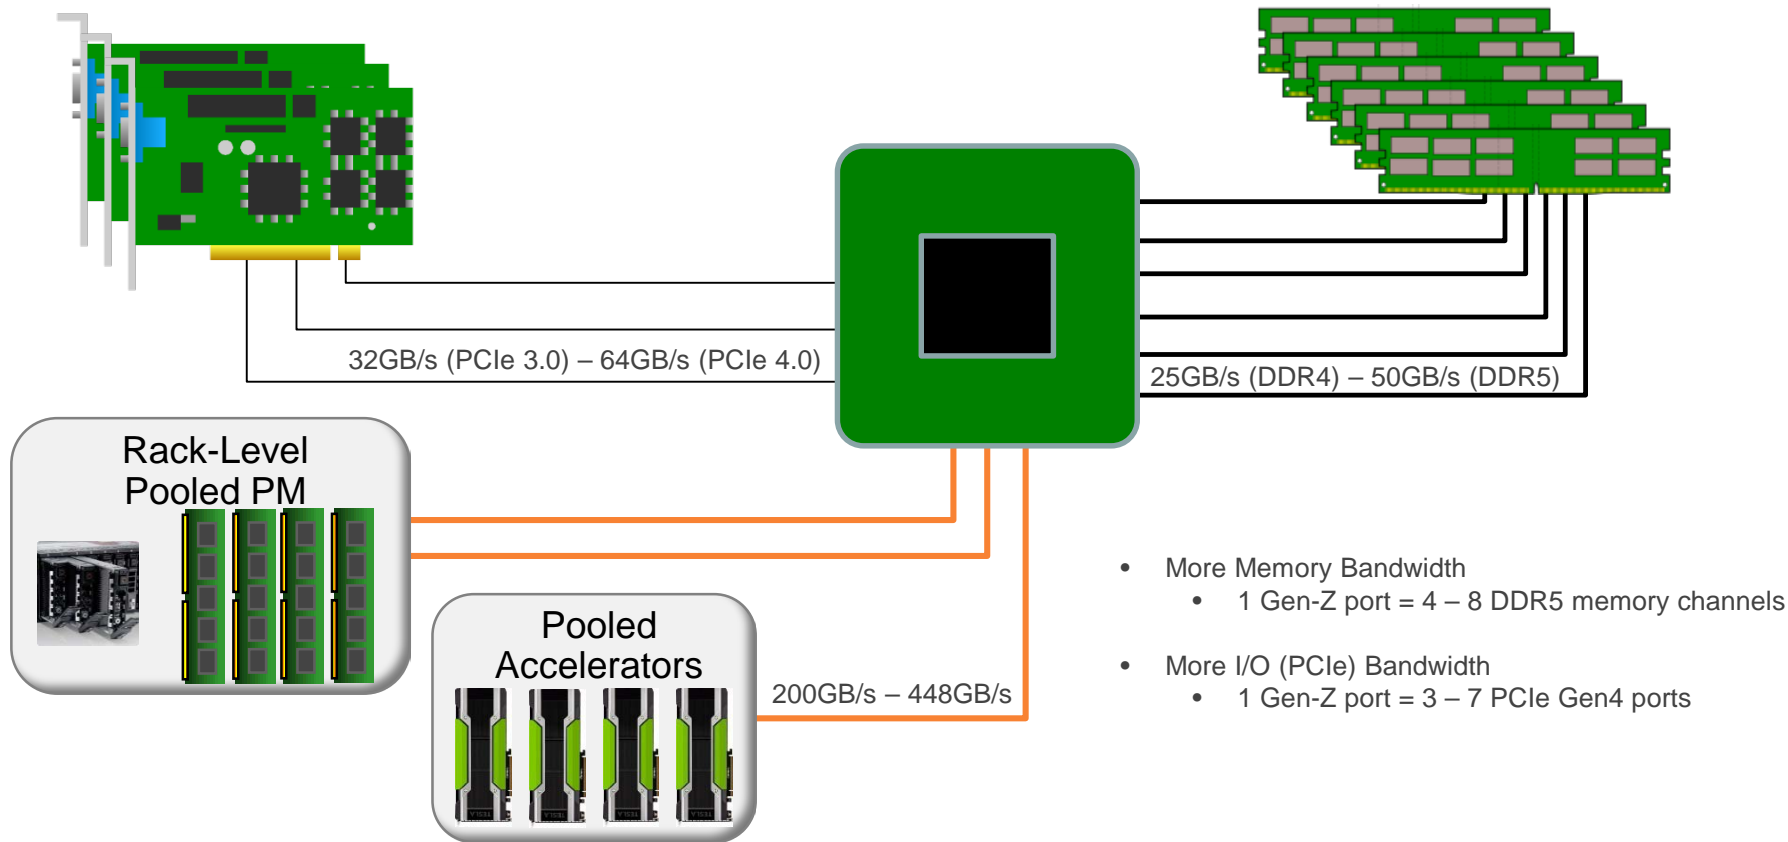


# Gen-Z Allows Memory Innovation



- 2-8 High-speed Serial Links
- 4-8 Memory Channels
- Low Latency, High performance
- 17-25 GB/s / Channel
- Split Memory Controller
- 288pins / DIMM
- Asynchronous Interface
- Synchronous Interface
- Processor is media agnostic

# Added Bandwidth Is Critical



# Thank You



PERSISTENT MEMORY

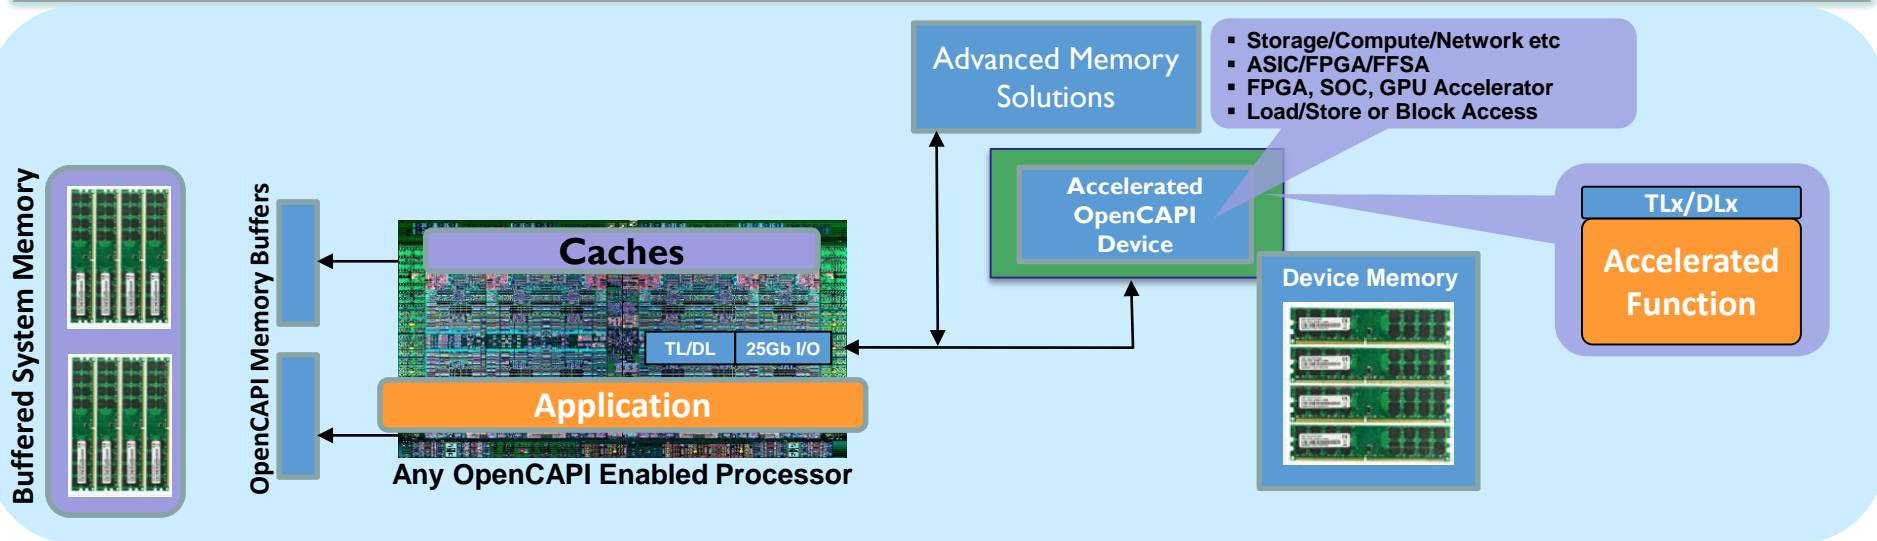
PIM SUMMIT

JANUARY 24, 2019 | SANTA CLARA, CA

OpenCAPI

Steve Fields, Fellow, IBM

# Open CAPI Key Attributes


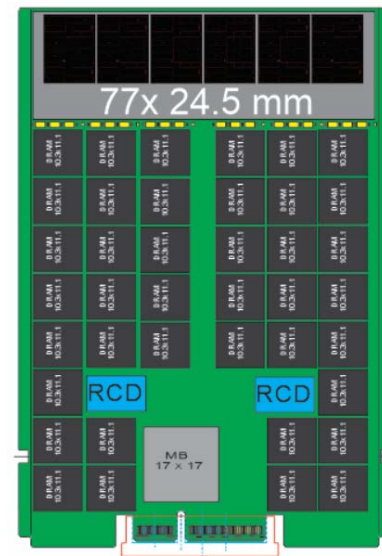




- Optimized for High Bandwidth and Low Latency in point to point Host-to-Device topology
- Architecture agnostic bus – Applicable with any system/microprocessor architecture
- Virtual addressing – device operates in application user space, eliminates communication overhead
- Coherent caching of system memory by OpenCAPI device
- CPU coherent device memory (Home Agent Memory) for both traditional DRAM and Persistent Memory
- Minimal OpenCAPI design overhead (FPGA less than 5%)



# OpenCAPI Advantages for Memory

- Enables idle latency close to that of direct-attach DDR RDIMM
- >2x bandwidth for <1/2 the CPU pins compared to direct-attach DDR interfaces
- Provides media-agnostic memory slot where controller is packaged with media
- Enables system-level flexibility for multi-tiered memory configurations
- Enables common electrical interface for memory, accelerator and coherent I/O device attach





Questions?