Welcome to

DesignCon® 2026

WHERE THE CHIP MEETS THE BOARD

Conference
February 24–26, 2026
Santa Clara Convention Center

Expo
February 25–26, 2026

LPDDR6 Memory: A New Choice for AI Data Center Applications

Frank Ferro, Group Director, Product Marketing, Memory IP
Cadence Design Systems, Inc.

SPEAKER

Frank Ferro

Group Director, Product Marketing, Memory IP, Cadence

fferro@cadence.com, www.cadence.com

Frank Ferro is the group director of product marketing at Cadence Design Systems responsible for memory and storage interface IP products. Frank joined Cadence after spending 10 years at Rambus, Inc. working on memory and SerDes IP products. In addition to Rambus, Frank spent more than 20 years at AT&T, Lucent, and Agere Systems. Mr. Ferro holds an executive MBA from the Fuqua School of Business at Duke University, an M.S. in computer science and a B.S. in electronic engineering technology from the New Jersey Institute of Technology.

Topics

- AI Market and Compute Demands
- LPDDR6 and LPDDR5X use in the Datacenter
- Cadence LPDDR6/5X Memory IP
- Summary

Data Traffic from AI Driving Datacenter Growth

- Global hyperscale datacenter growth continues due to accelerating data traffic
- Number of Hyperscale datacenters doubled over the last 5 years to >1,130 in 4Q24
- The capacity of these datacenters has more than doubled in less than 4 years
- Capacity is forecasted to double again in 4 years driven by generative AI compute demand

Hyperscale Datacenters

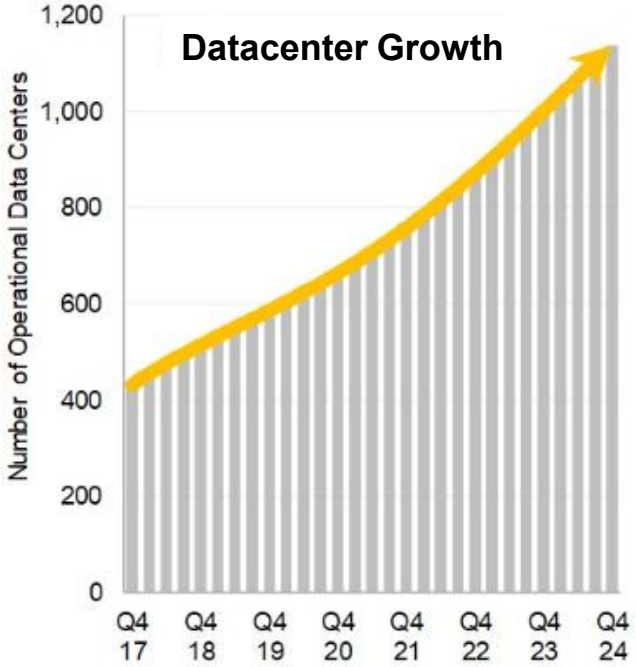

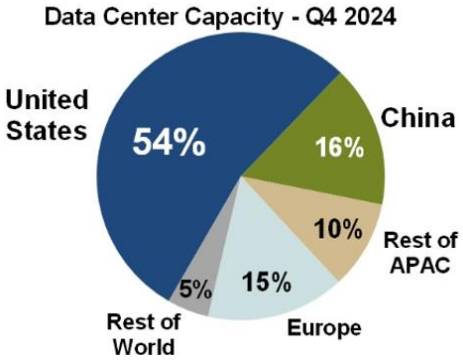

Synergy Research, March 2025

AI Model Sizes Continue to Increase

- Generative AI/ML training workloads are voracious in their demand for data and bandwidth
- AI/ML training is growing at 10X annually with GPT-3 @175 billion parameters and GPT-4 @1.8 trillion
- Model size grew at 410x over two years, the memory hardware grew at 2x over two years
- Development of custom accelerator to optimize performance and memory subsystem
- Need for greater memory and compute efficiency: reduced precision, quantization, and sparsity

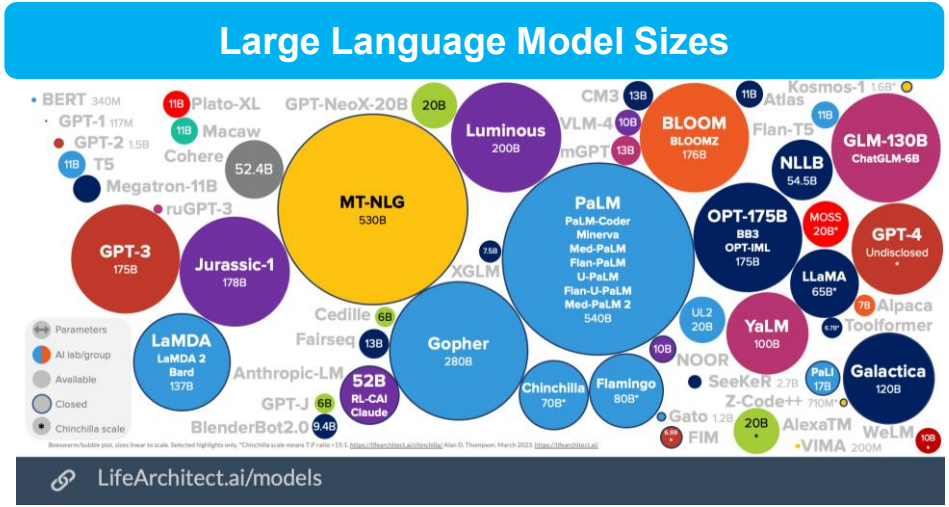

GPT-3 175B parameters
GPT-4 1.8T parameters

Minimum Memory Requirements for Popular LLM

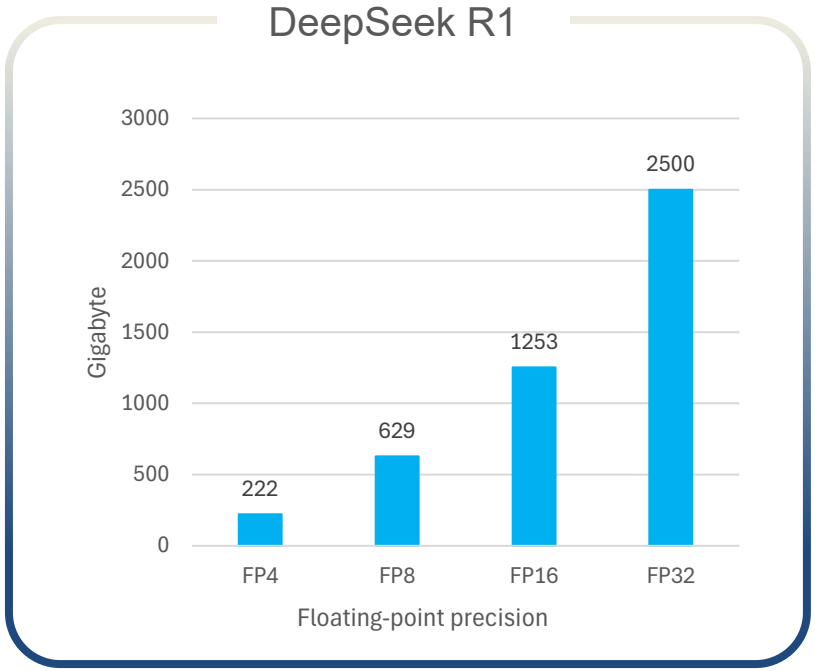

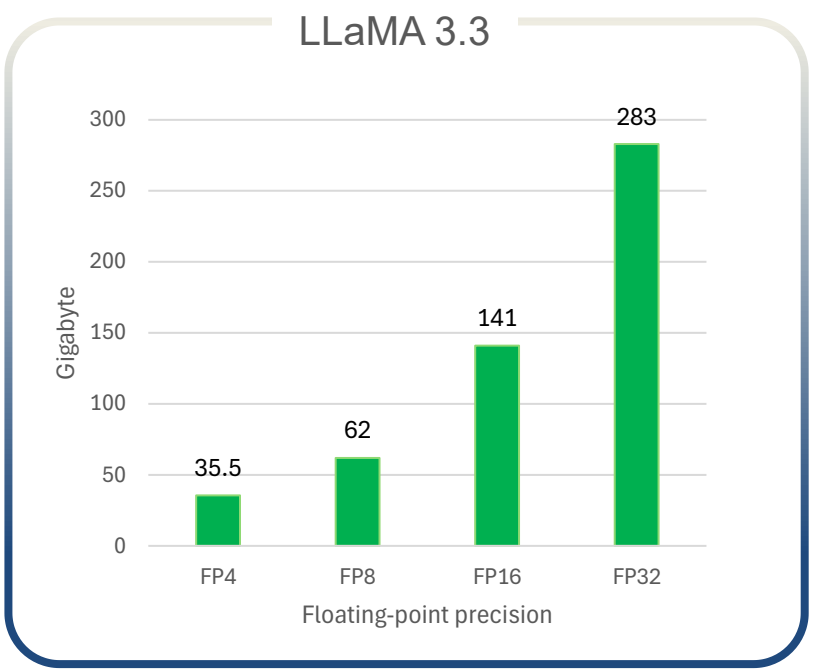

FP16 is a popular choice for LLMs

Data Center Rack Memory Trends

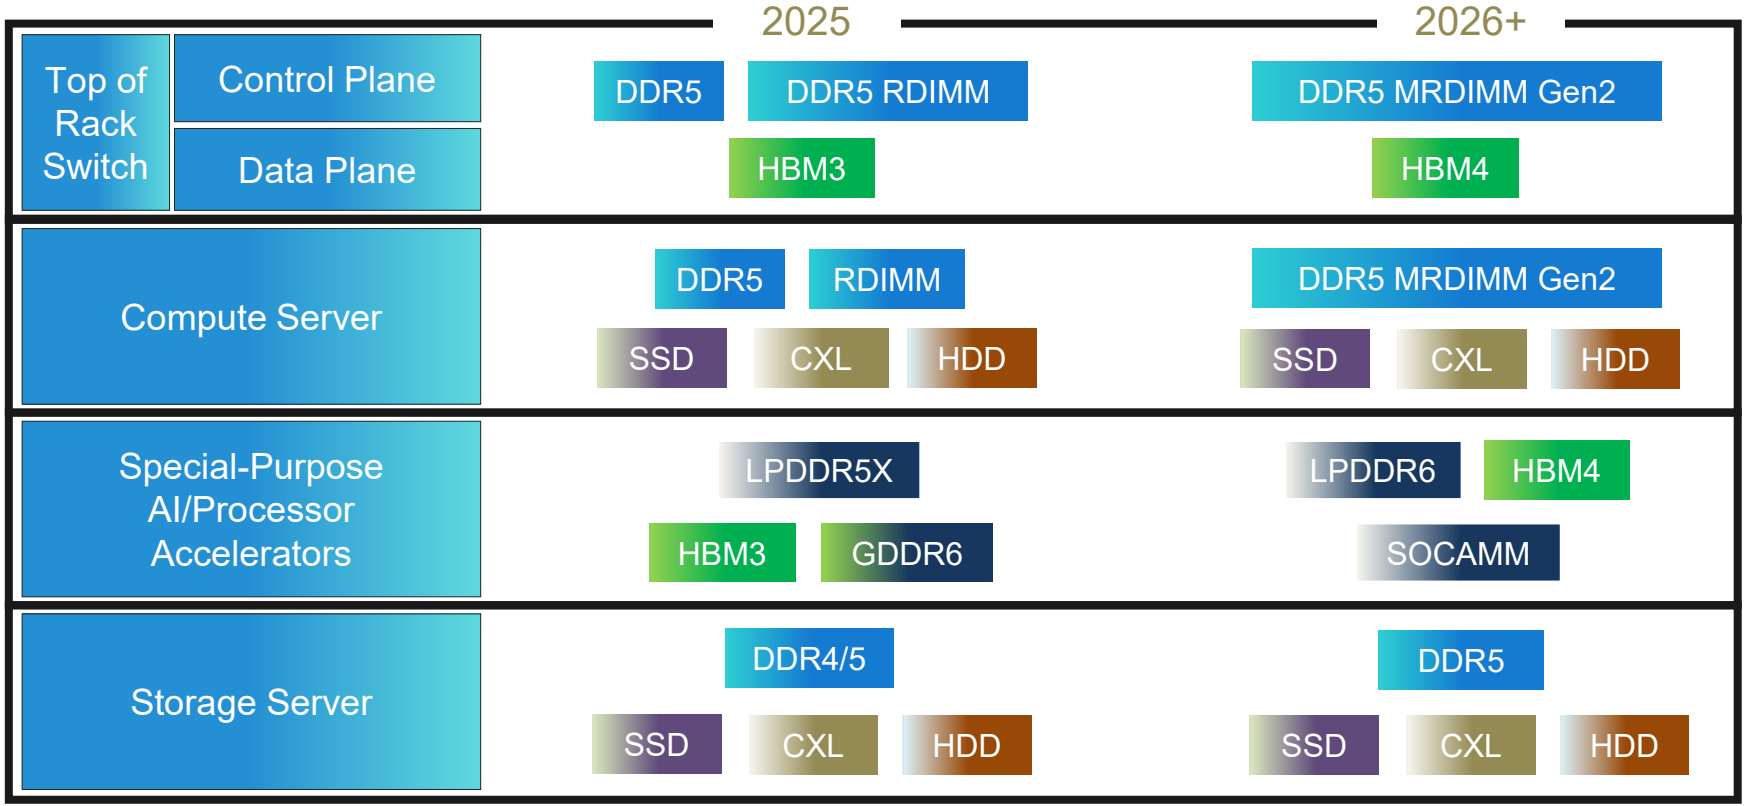

Dataflow in AI Factories

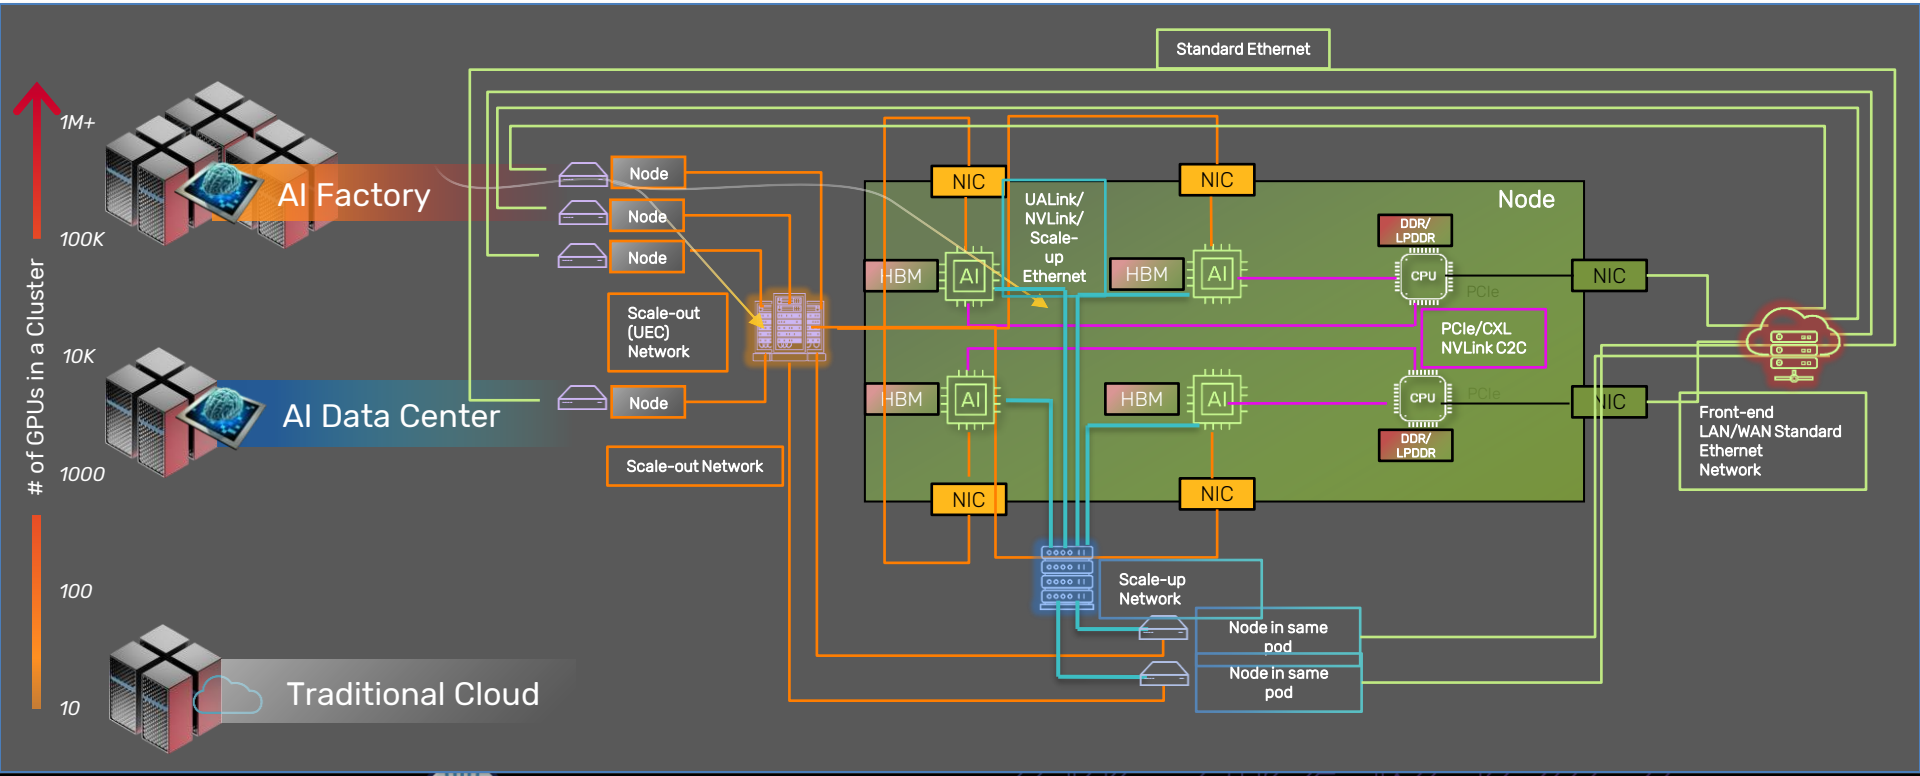

LPDDR6 and LPDDR5X use in the Datacenter

Optimized for low power, high performance with good memory capacity

Memory Choices for AI Training and Inference

- AI Training requires the highest bandwidth >1TB/s: HBM is primary choice
- LPDDR6/5X and GDDR6/7 offer high performance at a good price point for AI Inference
  - Cost effective using standard PCB and package
- LPDDR5X SOCAMM new form factor for CPU memory

Memory Protocol	Capacity	Cost/ Performance	Bandwidth	Power Consumption	Packaging Cost
DDR4/5					
LPDDR6/5X					
GDDR6/7					
HBM3E/4E					

Traditional DDR DIMM or PCB

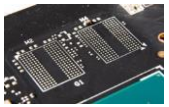

LPDDR/GDDR on PCB

HBM Stacked DRAM – 2.5/3D

Key: Best Good

LPDDR6 for AI in the Datacenter

Balancing performance, power, and capacity

- **High Performance for AI**
  - 691Gbps bandwidth per device
  - Dual sub-channel architecture for flexible operation while maintaining a small access granularity of 32 bytes
  - On-the-fly burst length control to support 32-bit and 64-bit data access
- **Better Capacity**
  - Static efficiency mode designed to support high-capacity memory configurations and maximize bank resource utilization
- **Low Power**
  - Lower operating voltage
  - Dynamic Voltage Frequency Scaling (DVFS-LP)
  - Support for both partial self and active refresh to reduce refresh power consumption
- **RAS (Reliability, Accessibility, Serviceability)**
  - Per Row Activation Counting (PRAC) to support DRAM data integrity
  - Carve-out meta mode, which is defined to enhance overall system reliability by allocating specific memory regions for critical tasks

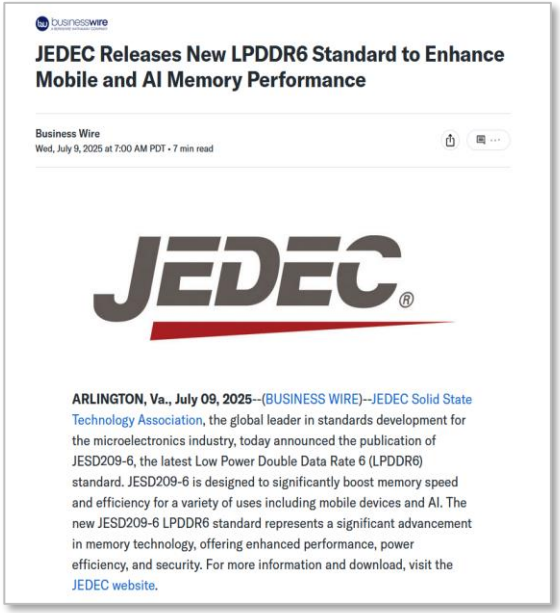

businesswire

JEDEC Releases New LPDDR6 Standard to Enhance Mobile and AI Memory Performance

Business Wire
Wed, July 9, 2025 at 7:00 AM PDT • 7 min read

JEDEC®

ARLINGTON, Va., July 09, 2025--(BUSINESS WIRE)--JEDEC Solid State Technology Association, the global leader in standards development for the microelectronics industry, today announced the publication of JESD209-6, the latest Low Power Double Data Rate 6 (LPDDR6) standard. JESD209-6 is designed to significantly boost memory speed and efficiency for a variety of uses including mobile devices and AI. The new JESD209-6 LPDDR6 standard represents a significant advancement in memory technology, offering enhanced performance, power efficiency, and security. For more information and download, visit the JEDEC website.

LPDDR5X in Datacenter Workloads

Micron Case Study*

- Lower memory power
  - >75% lower power (Multi-chase and POT3D)
- Higher memory bandwidth
  - >35% higher memory bandwidth (Llama3 8B)
- Better overall system power
  - 10% Better power efficiency (perf/watt) for Llama 3 (8B)
- Higher throughput
  - >5x inference through and 80% lower inference latency for Llama 3 (70B) (NVLink and Grace CPU also add to the performance)
  - 80% lower inference latency
- Moving to LP for next generation

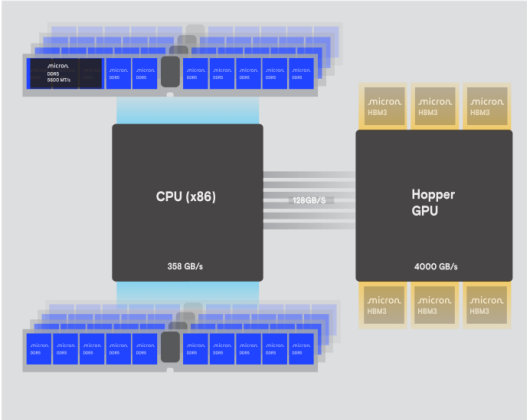

Figure 2: DDR5 system – DDR5 and x86 architecture

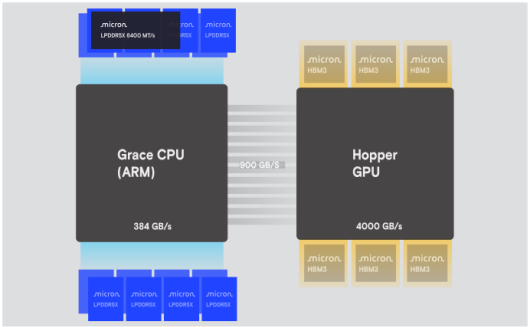

Figure 1: LPDDR5X system – LPDDR5X and ARM architecture

*The role of low-power (LP) memory in data center workloads, Micron technical brief, 2025
LPDD5X @ 6.4Gbps, total bandwidth: 384GB/s compared to DDR5

LPDDR6/5X SOCAMM

- High memory bandwidth
  - 2.5x the bandwidth of RDIMM*
- Lower Power
  - One-third the power of RDIMM*
  - Better thermal management
- Compact modular design
  - Scalability and serviceability
- High Capacity: 128GB per module (LP5X)
  - Four high density LPDDR memory stacks
  - Capacity will increase with LPDDR6

*According to Nvidia, Korean Herald, April 2025

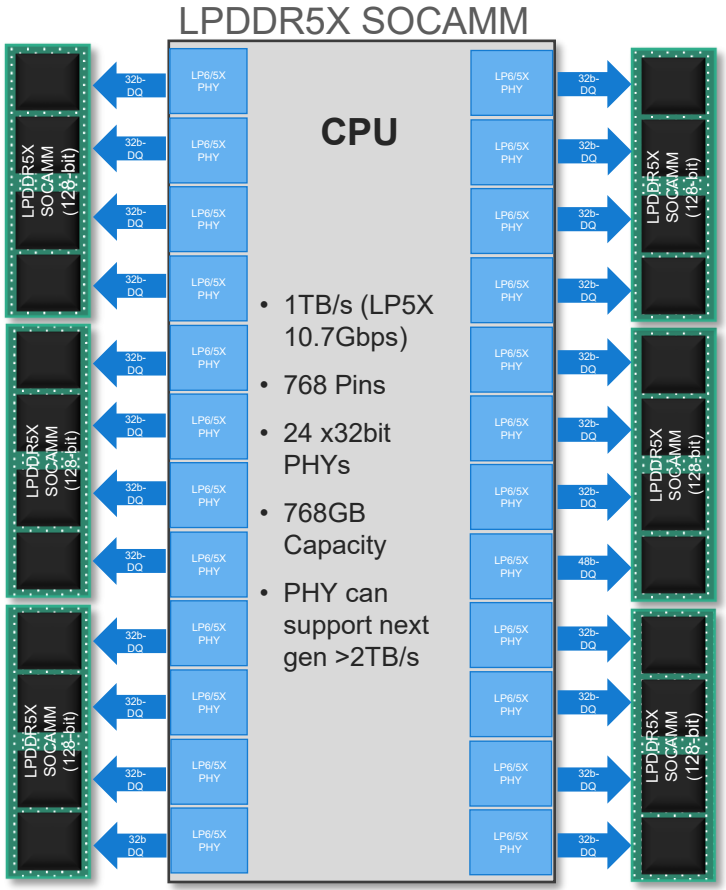

Cadence LPDDR6/5X Memory IP

High-performance solutions for advanced datacenter, networking and AI applications

Cadence Memory IP for AI

- **Focused on Performance**
  - Advanced mixed signal architecture: best signal integrity
  - System design expertise: provide customer reference design to ensure system performance goals are reached
- **PHY Architecture for longevity: Meet or exceed speeds of next standards**
  - Designed for DDR5 MRDIMM, LPDDR6/5X, HBM4E, GDDR7
- **First to market with the leading-edge memory IP**
  - LPDDR6/5X 14.4G IP announced in conjunction with the standard (July'25)
  - DDR5 12.8G MRDIMM announced in April'25
  - HBM4 12.8G IP announced in April'25 in conjunction with the standard
- **Large and experienced memory IP organization**
  - Systems, Silicon, Lab, Memory Controller and Support Teams

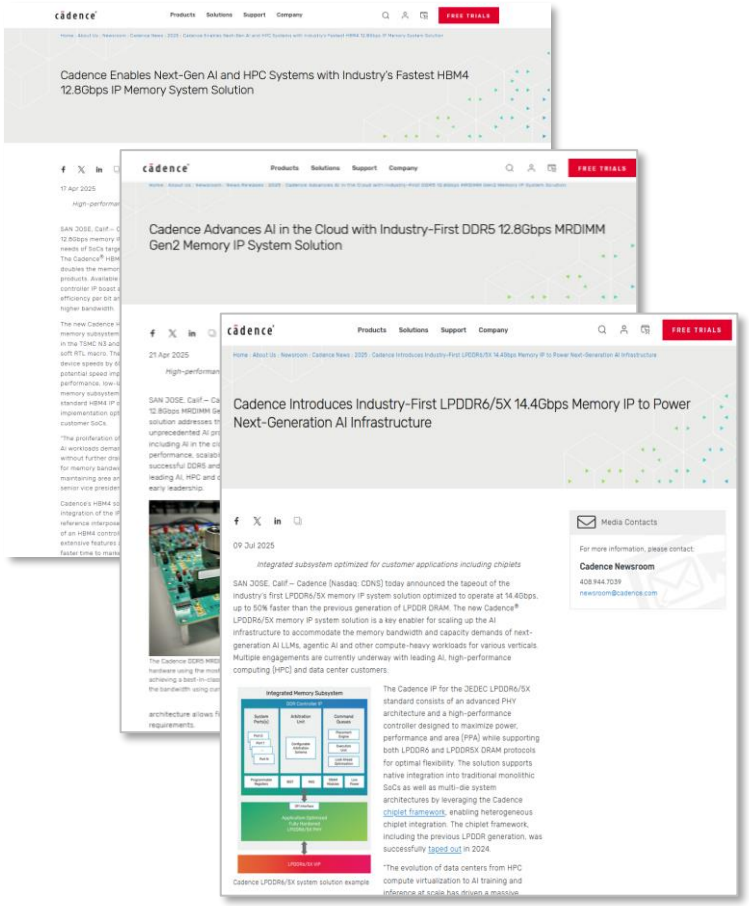

Complete Memory IP System Solution

Easy integration that is optimized for your application

- Fully hardened PHY
  - Optimized for each customer
  - Multiple floorplans, bump maps, packages, and PCB
  - Available in multiple foundries and nodes
- Controller delivered as soft IP
  - Optimized for customer data flow
- Multiple processor port options
  - AHB, AXI, CHI and Native
- Memory subsystem
  - Delivered as an integrated subsystem
  - Fully validated PHY and Controller
- Silicon validation
  - Rapid bring-up and remote debug software

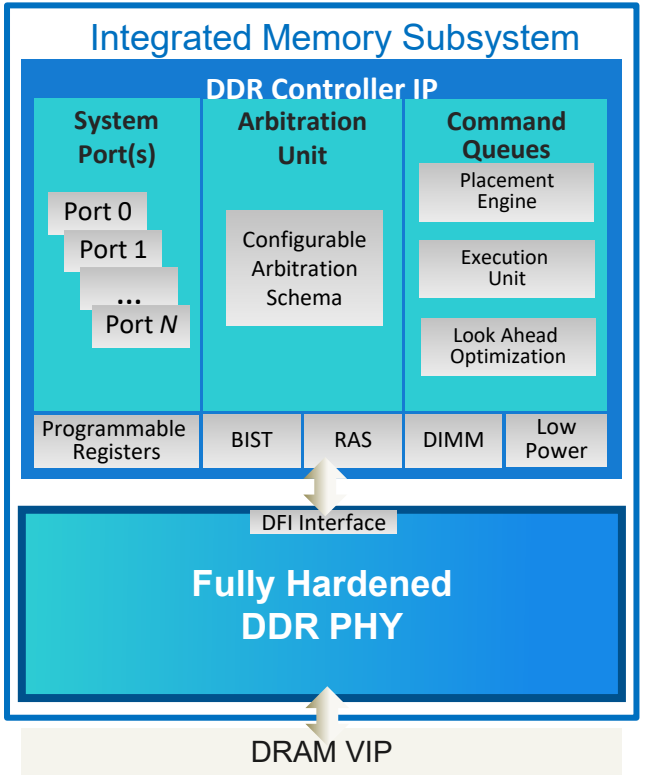

Cadence LPDDR6/5X 14.4Gbps PHY

- **Key Features**
  - 14.4Gbps data rate (no OD)
  - Two 12b Subchannels
  - Data and address pin configuration for 4R LP6
  - Per bit de-skew on the write read data and control path
  - Improved bump map to ease routing
  - Shoreline or depth optimized floorplans
  - NS and EW available

- **Silicon proven performance optimized architecture**
  - Custom high-speed IO/delay/logic within reusable hardened sub blocks
  - Low speed digital top level eliminates timing bottlenecks
  - Ultra short local high speed clock trees
  - Hardened high speed logic

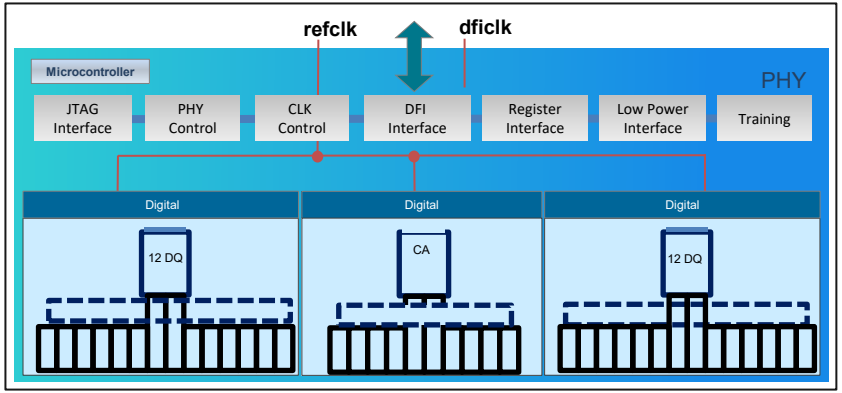

LP6 14.4G TX eye

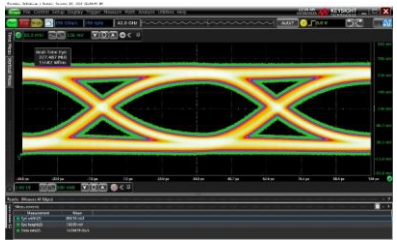

LP6 10.7Gbs RX eye

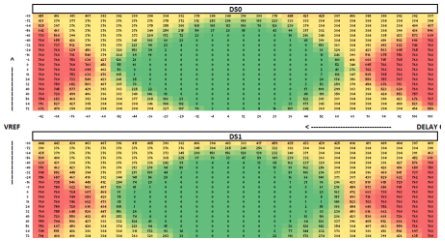

Cadence LPDDR6/5X Silicon Results

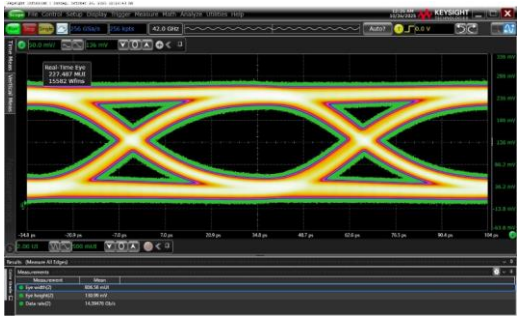

LPDDR6 14.4G TX Eye

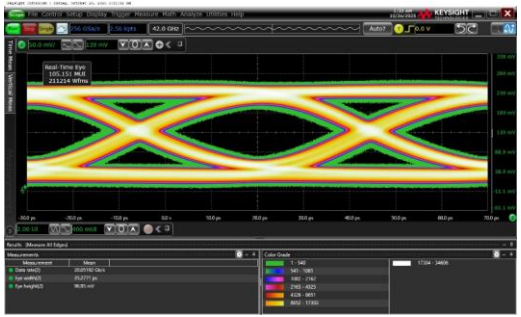

LPDDR6 20G TX Eye (over clocked)

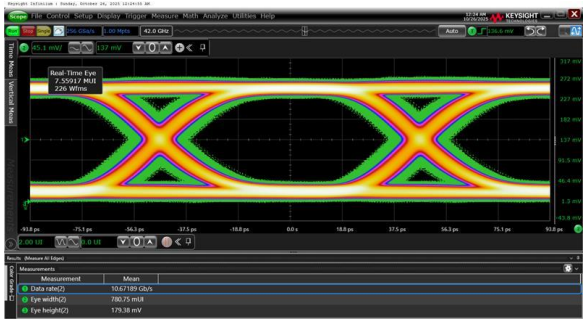

LPDDR5X 10.7G TX Eye

- LPDDR6/5X shows excellent results (LP6 POP DRAM)
- LPDDR5X testing with 10.7G DRAM

LPDDR6/5X Development Board

LPDDR6, LPDDR5X and LPDDR5X SOCAMM Testing

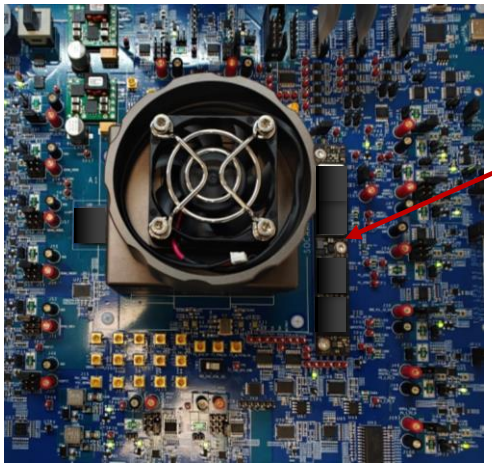

LPDDR6/5X Development Board
Supports LP6 POP, LP5X and LP5X SOCAMM

LP5X SOCAMM

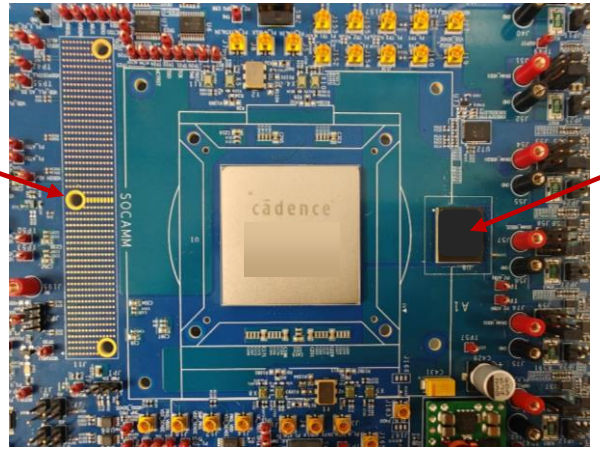

LP5X DRAM

LPDDR6/5X Development Board
(SOCAMM Removed)

LPDDR5X SOCAMM

Cadence Complete System Design Flow

Floor Plan/Bump Map

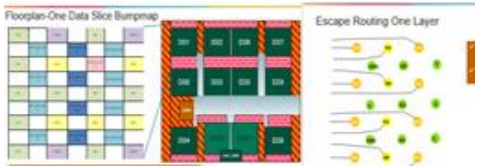

Achieve balance between area, pitch, escape routing, and decap

Package/Board Routing

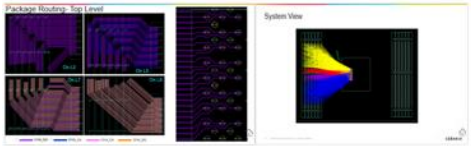

Optimal package board routing solution and stack-up

I/O Models

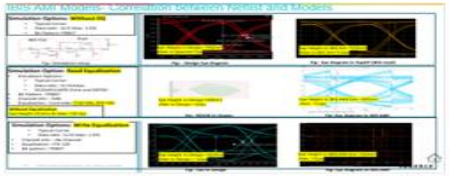

Best-in-class IBIS and PDN models

SI/PI Analysis

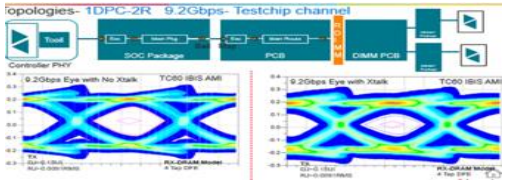

Signal and power integrity analysis in Cadence tools

Reference Platform

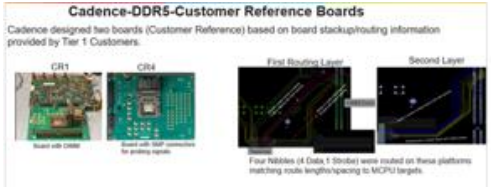

Enable customers with proven reference designs

Post-Silicon Validation

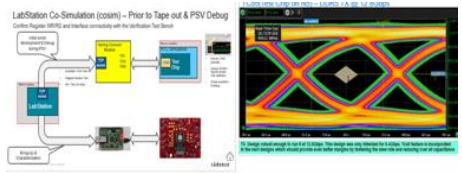

Fast silicon bring-up and validation

Cadence Silicon Bring-Up Software

- **Preconfigured for every customer design**
  - Preprogrammed PHY and CTRL register values
  - Brings up memory interface without SoC firmware
- **Independent operation**
  - Easy connect though SoC JTAG or debug interface
  - Requirement is to read/write registers in PHY and CTRL
- **Comprehensive debug and margining features**
  - 2D software eye diagram for every data pin
  - Step through the initialization and training steps
  - Pause at any bring-up step to verify or change results
- **Import settings to SoC firmware**
  - Bring-up code is open to read and modify
  - Easy export of register settings for SoC firmware
- **Remote support**
  - Support customer systems from Cadence labs

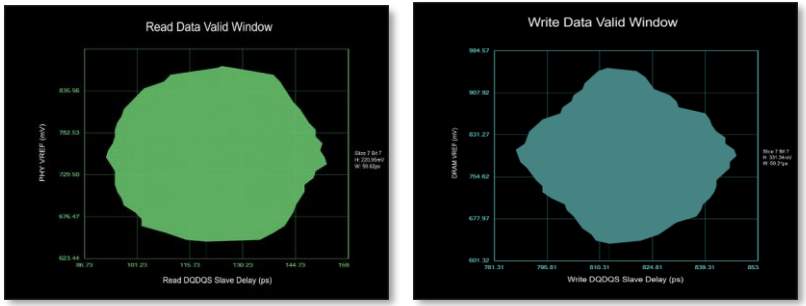

DDR5 12.8Gb/s MRDIMM Gen2
Read and write software eye diagrams

"We used the DDR IP bring-up software to try various IP settings and determine the *optimal DDR system initialization code* to be used in the firmware." – US-Based System OEM

"We used Cadence bring-up software enabled on a Windows laptop and were able to perform training and BIST to successfully write/read/compare DDR memory. Pretty good success for *just over an hour* of work." – US-Based Storage OEM

Summary

- Advanced AI applications are demanding far greater memory bandwidth
- LPDDR6 provides 691Gb/s per device and is the best trade-off of cost capacity and power for AI training and inference
- New LPDDR SOCAMM form factor for CPU memory saves power and boosts performance
- Fully-integrated memory subsystem solution PHY and Memory Controller
- Subsystem validation in hardware: test chip includes both PHY and Controller
- PCB and package reference design support
- Cadence is the only company to offer a complete memory IP portfolio for AI applications including DDR, LPDDR, HBM and GDDR

Thank you!

QUESTIONS?

Frank Ferro

Group Director, Cadence

fferro@cadence.com | cadence.com