

Course Description

This course outlines the architecture of AMD Versal™ AI Engine ML v2 (AIE-ML v2), a part of AMD Versal AI Edge Series Gen 2 devices, and explores the features and key architectural enhancements with this iteration of AI Engines.

This course provides an overview of both native and supported data types and highlights how to program the AI Engine and migrate older AI Engine designs. The enhancements included in AIE-ML v2, utilizing the DSP libraries, along with the compute capabilities and how to analyze performance, will also be demonstrated.

The emphasis of this course is on:

- Providing an overview of the new AI Engine (AIE-ML v2) architecture
- Describing the system design planning and application partitioning methodology
- Describing the AMD Vitis™ and AI Engine tool flow
- Providing an overview of the native and supported data types for functional implementation in AIE-ML v2
- Illustrating the programming model and the usage of memory tiles (shared buffers) for AIE-ML v2
- Utilizing the Vitis DSP library for AI Engines in implementing a matrix multiplication with multiple tiles
- Analyzing reports using the Vitis Analysis view and reviewing throughput and performance of different AI Engine architectures

What's New for 2025.2

- All labs have been updated to the latest software versions

Level – VER 1

Course Details

- 1 day instructor led training (online or in person)
- 8 lectures
- 4 labs

Price – \$800 or 8 AMD Training Credits

Course Part Number – AIE-MLV2-ARCH

Who Should Attend? – DSP users, software developers, system architects, and anyone who needs to accelerate their software applications using our devices

Prerequisites

- Familiarity with the C/C++ programming language
- Vitis tool for acceleration development flow
- Familiarity with basic signal processing concepts

Software Tools

- [Vitis Unified IDE 2025.2](#)
- Optional: MATLAB® tool (any latest version)

Hardware

- [Architecture: Versal adaptive SoCs](#)

After completing this comprehensive training, you will have the necessary skills to:

- Describe the AMD Versal AI Engine AIE-ML v2 architecture
- Follow the system partitioning and system mapping methodology
- Describe the AI Engine tool flow and utilize the supported data types for AIE-ML v2
- Program the AI Engines with the help of the programming model and describe the utilization of memory tiles

- Utilize the AI Engine DSP library and test a matrix multiplication design with multiple tiles for parallelization
- Analyze AI Engine designs using the Analysis view (Vitis Analyzer utility) and analyze the throughput and performance of different AI Engine architectures

Course Outline

- **AMD Versal Adaptive SoC: Architecture Overview**
Provides a high-level overview of the Versal architecture, illustrating the various compute resources available in the Versal architecture. {Lecture}
- **Introduction to the AIE-ML v2 Architecture**
Discusses the AI Engine AIE-ML v2 array architecture and its tiles. Also lists the key differences between the AIE, AIE-ML, and AIE-ML v2 architectures. {Lecture}
- **Versal AI Edge Series Gen 2: Application Partitioning**
Covers the system design planning and partitioning methodology. Also explains what application partitioning is and how an application can be accelerated by using various compute domains in the Versal device. {Lecture}
- **Versal AI Edge Series Gen 2: AIE-ML v2 Tool Flow**
Reviews the Vitis tool flow for the AI Engine and demonstrates the full application acceleration flow for the Vitis platform. {Lecture, Lab}
- **Supported Data Types for AIE-ML v2**
Provides an AI Engine functional overview and identifies the supported vector data types and high-width registers for allowing single-instruction multiple-data (SIMD) instructions. Covers the vector-based floating-point data types and block floating-point data types. {Lecture}
- **The Programming Model and AIE-ML v2 Memory Tiles**
Highlights the basics of AI Engine kernels and graph code as well as the data flow graph model. Also describes the memory tiles in the AIE-ML v2 architecture and illustrates AI Engine-ML v2 programming using both shared buffers and external buffers with the AIE-ML v2. {Lecture, Lab}
- **Overview of the AI Engine DSP Library with AIE-ML v2**
Provides an overview of the AI Engine DSP library, which enables faster development and comes with ready-to-use example designs that help with using the library and tools. Also describes how to use and configure the library functions. {Lecture, Lab}
- **Analyzing the AI Engine Design Reports and Performance**
Covers the different reports generated by the Analysis View in the Vitis Unified IDE and how to use these reports to optimize AI Engine designs. Also shows how to migrate a design to the new architecture as well as compare performance improvements after migration. {Lecture, Lab}

Register Today

Morgan Advanced Programmable Systems, Inc. (Morgan A.P.S.) delivers public and private courses in locations throughout the central US region; including Iowa, Illinois, Kansas, Minnesota, Missouri, Nebraska, North Dakota, South Dakota, and Wisconsin.

Visit morgan-aps.com/training, for full course schedule and training information.



■ You must have your tuition payment information available when you enroll. We accept credit cards (Visa, MasterCard, or American Express) as well as purchase orders and AMD training credits.

Student Cancellation Policy

- Student cancellations received more than 7 days before the first day of class are entitled to a 100% refund. Refunds will be processed within 14 days.
- Student cancellations received less than 7 days before the first day of class are entitled to a 100% credit toward a future class.
- Student cancellations must be sent [here](#).

Morgan A.P.S. Course Cancellation Policy

- We regret from time-to-time classes will need to be rescheduled or cancelled.
- In the event of cancellation, live on-line training may be offered as a substitute.
- Morgan A.P.S. may cancel a class up to 7 days before the scheduled start date of the class; all students will be entitled to a 100% refund.
- Under no circumstances is Morgan A.P.S. responsible or liable for travel, lodging or other incidental costs. Please be aware of this cancellation policy when making your arrangements.
- For additional information or to schedule a private class contact us [here](#).

Online or in person training with real hardware

- Morgan Advanced Programmable Systems, Inc. has set up a training VPN where engineer participants can take classes online using the same computers and devCards used during in-person training.
- Even better, and upon request, you can use these computers after hours on training days to experiment with labs. This is not possible for in-person training.
- Additionally, just like in-person training, the laptops and devCards, tools, OS, and licensing are set up in advance.
- In some ways, live online-training is better than in-person...for example, you can grant the instructor permission to look at your Vivado, PetaLinux terminal, or Vitis for extended periods of time if your lab is not going exactly as planned to a missed step.
- This is often more comfortable than two engineers crowding around a laptop screen.
- Taking remote training also allows you to learn some tips and tricks for working remote. Whether your devCard is in the lab

down the hall, or across the world via VPN, you can control your AMD based device quickly and efficiently.