

**PHISON**

**aiDAPTIV<sup>+</sup>**

# Affordable LLM Training & Inference




# > What Phison is Announcing for aiDAPTIV+ at GTC 2025



## Demo of World's 1<sup>st</sup> LLM Training Laptop

Maingear laptop PC built for AI powered by Phison aiDAPTIV+

- Full Model Fine-Tuning
- Enhanced Inferencing



## Support for NVIDIA Jetson IoT Devices

Edge computing and robotics powered by Phison aiDAPTIV+

- LoRA Model Fine-Tuning
- Enhanced Inferencing

## INFERENCE

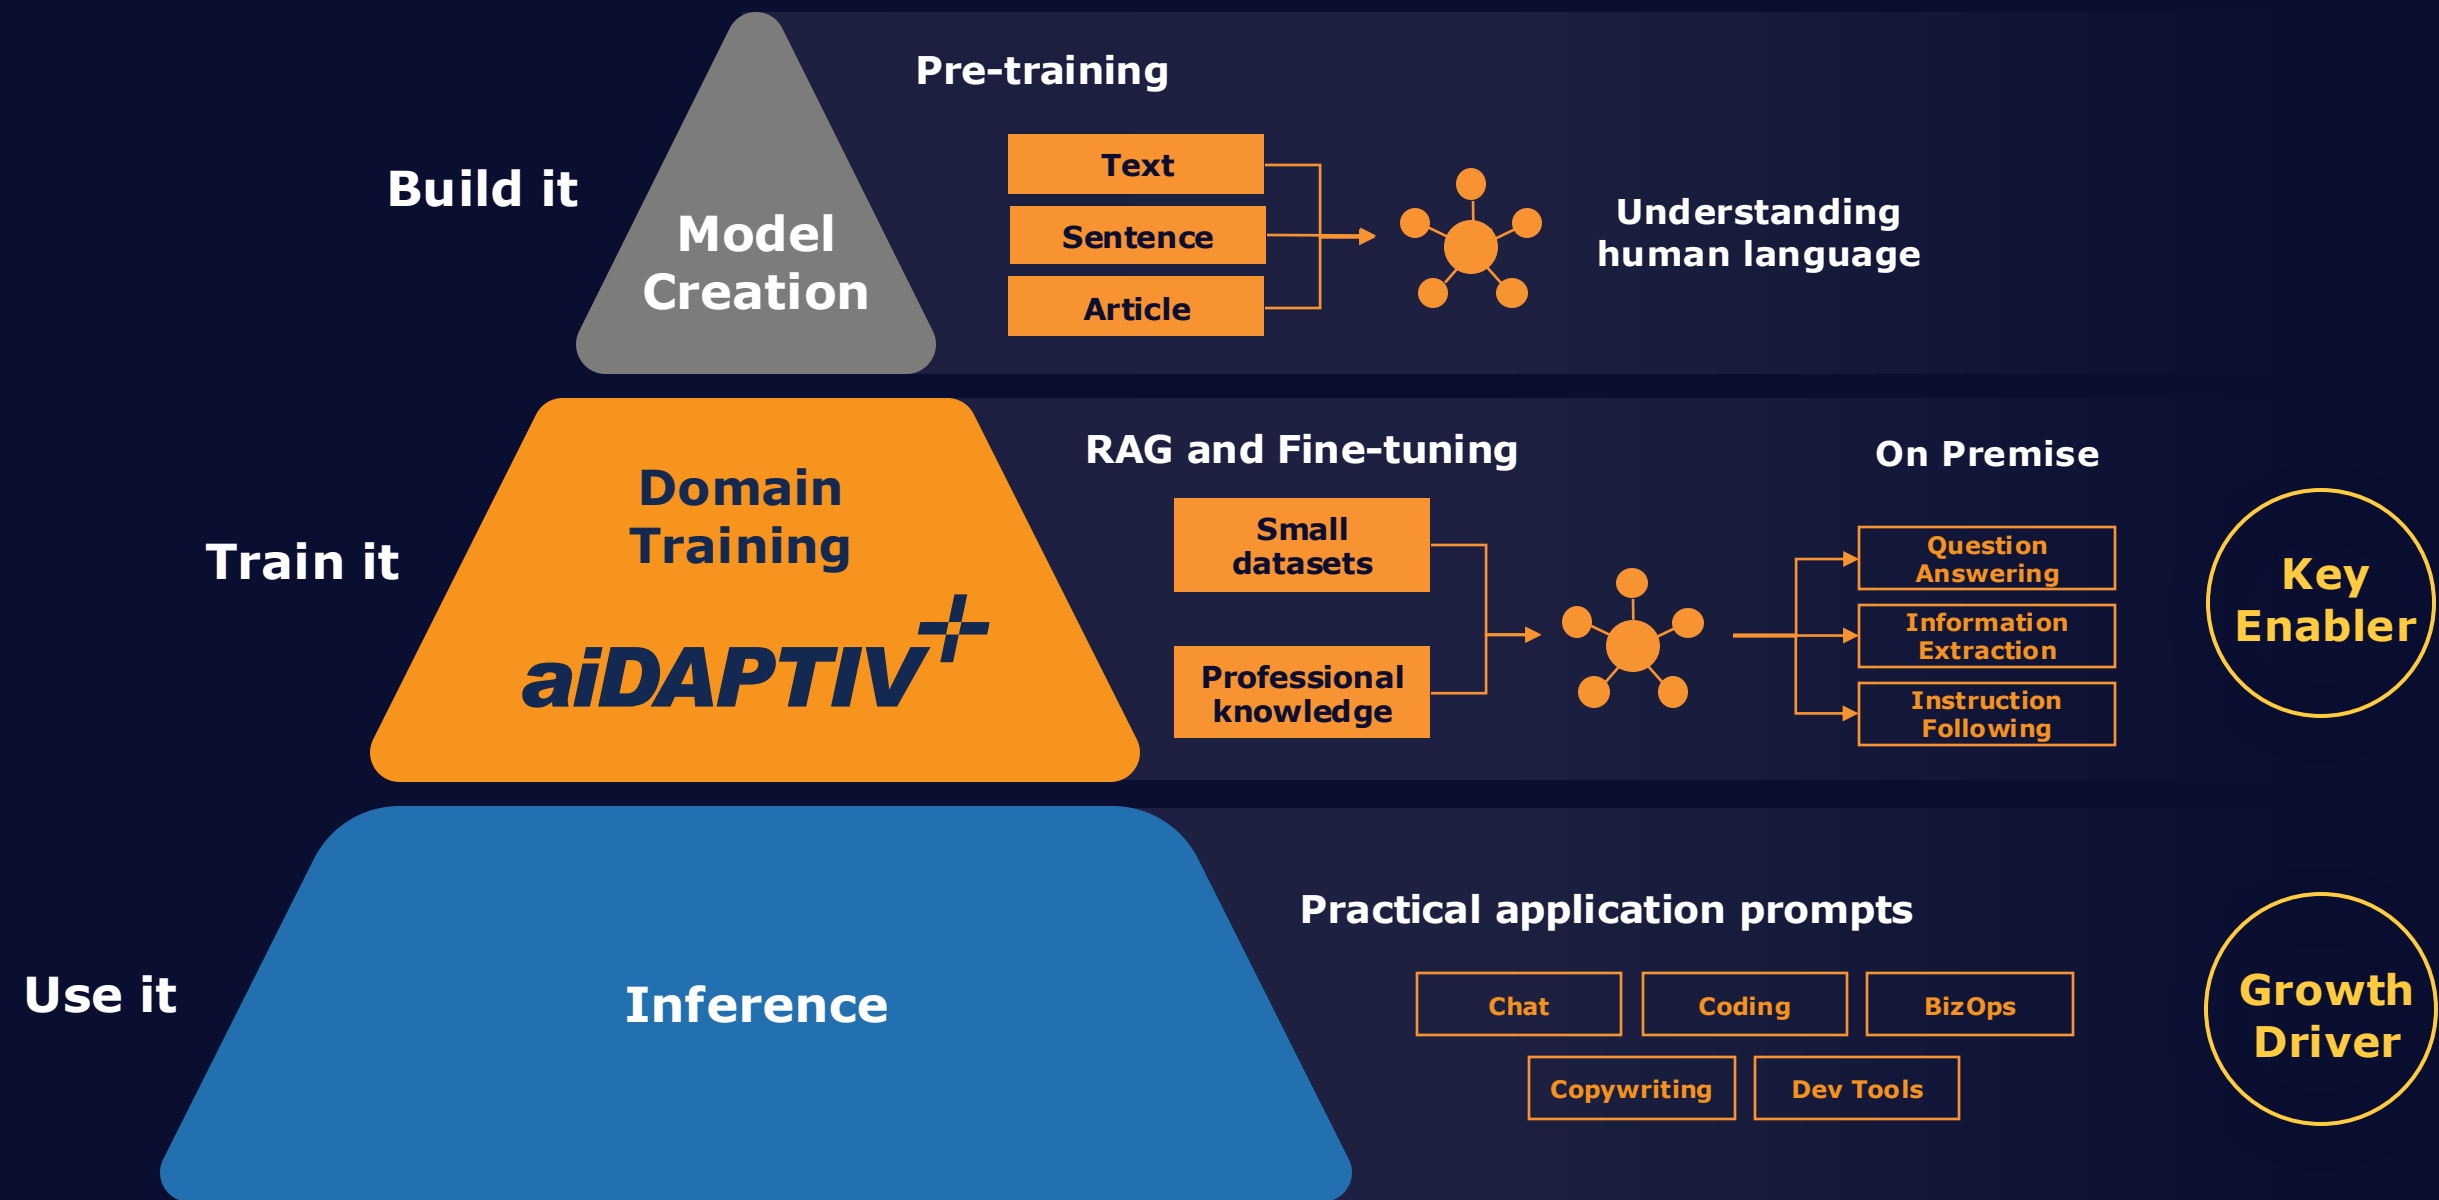
### Enhanced Inferencing

A better user experience

- Longer Token Lengths delivers more accurate results
- Faster Time to First Token recall for quicker research



# > Today's LLM Market Segment



GPU Cards Requirements

>1000

Massive GPUs for High Computing Power

10~100

Depends on Memory Size (GPU RAM ≥ 20x Model Size)

1

Minimal requirement

# > Phison aiDAPTIV+ Delivers Affordable Memory Scaling

## Current AI Algorithm Architecture

GPU + HBM/GDDR Scaling Limits

Lots of GPUs  
to Run Mainstream LLM




## Phison's aiDAPTIV+ AI Algorithm Architecture

Expands Memory Scaling with NAND Flash

# 8-10x Cost Savings


Extend GPU RAM



1: Based on NAND capacity

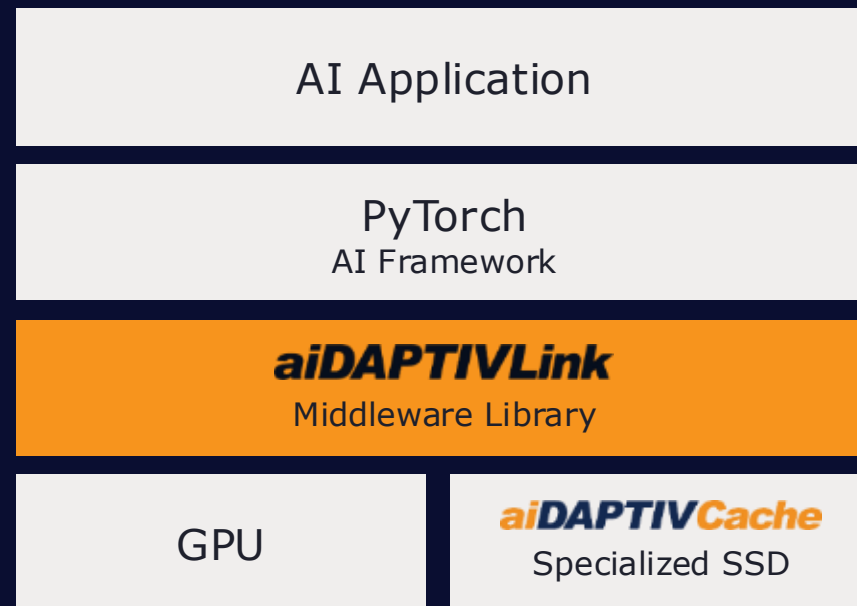
# > How aiDAPTIV+ Does It

Builds a Larger Memory Pool for LLM Training  
By Tiering **Expensive** VRAM and **Affordable** Flash Memory





# > Phison aiDAPTIV+ LLM Domain Training Solution



**aiDAPTIVLink**

**Middleware**

**Coordinates the swapping  
between HBM/DRAM and  
Flash Memory**



**aiDAPTIVCache**

**AI-Series SSD Family**

**Seamless Integration  
with VRAM/DRAM**



**aiDAPTIV+**


**Pro Suite UI**

**End-to-End User Experience  
Spanning Training to Inference**

# > **aiDAPTIV<sup>+</sup>** Workstation Scaling

Makes Possible Fine-Tuning Beyond VRAM Capacity

Single node 4x GPU configuration comparing GPU only and GPU with **aiDAPTIV<sup>+</sup>**



**Fixed  
VRAM  
Capacity**



Training Set Size  
HBM Pool (Usage%)  
Minimum GPU Count

140 GB  
192 GB (**73%**)  
4 / 4

260 GB  
192 GB (**120%**)  
4 / 6


1400 GB  
192 GB (**729%**)  
4 / 30

Notes:

1. Scaling is linear based on GPU count and model size
2. Training is based on 10M tokens

**System Configuration**

- RAM: 512 GB
- GPU: 4x RTX 6000 ADA
- GDDR: 192 GB





# > Train LLMs at Any Budget with **aiDAPTIV<sup>+</sup>**



**IoT Device**  
Up to 64B Parameter  
**LoRA** Model Training  
**\$500-1,000**



**Laptop PC**  
Up to 8B Parameter  
Full Model Training  
**\$2,000-3,000**



**Desktop PC**  
Up to 13B Parameter  
Full Model Training  
**\$3,000-4,000**




**Workstation PC**  
Up to 100B Parameter  
Full Model Training  
**\$5,000-50,000**



**Server**  
Up to 405B Parameter  
Full Model Training  
**\$50,000+**

# > **aiDAPTIV<sup>+</sup>** The Affordable Path to **1 Trillion Parameter Training**

The largest size model training at 1/30 (<4%) of the cost



**2U AI Training Appliance**

- NVIDIA GH200 Superchip x 1
- 480GB DRAM
- Phison aiDAPTIV+ MW/SW & aiDAPTIVCache 4TB SSD x 4




**\$3M**



**\$100K**




# Thank You!



***PHISON***






GTC 2025

# Enterprise SSD

CONFIDENTIAL

**PHISON**

# Phison Enterprise SSD Product Line-up



## High Performance

### X-Series

PCIe 4.0/5.0  
Up to 30.72TB  
1/3/60 DWPD



U.2 / U.3 / E3.S

## Data Center

### D-Series

PCIe 4.0/5.0  
Up to 122.88TB  
0.3/1/3 DWPD



E3.S / E3.L / U.2 / E1.S / M.2

## SATA

### S-Series

SATA III  
Up to 15.36TB  
0.4/1/3 DWPD



2.5"

## Boot Drive

### B-Series

PCIe 4.0 / SATA III  
Up to 960GB  
1 DWPD



2.5" / M.2

## Artificial Intelligence


### AI-Series

PCIe 4.0  
Up to 8TB  
Up to 100 DWPD



U.2 / M.2

# D200V/D205V Series Product Overview



## D200V/D205V QLC SSD

Phison’s D200V/D205V series utilize QLC technology to meet the high storage density demand driven by AI. With an impressive 122TB per SSD, D205V will help drive the trend of efficient data storage with improved space utilization and reduced power consumption.

Specification	Features	Capacity
<ul style="list-style-type: none"><li>Interface: PCIe 5.0 x 4</li><li>Protocol: NVMe 2.0</li><li>Form Factor: U.2 / E3.S / E3.L</li><li>DWPD: 0.3 DWPD</li><li>MTBF: 2.5 million hours</li><li>Warranty: 5 years</li></ul>	<ul style="list-style-type: none"><li>3D QLC NAND</li><li>Dual-port design</li><li>Power loss Protection</li><li>Namespaces: 128</li></ul>	<p><b>D200V:</b></p> <ul style="list-style-type: none"><li>30.72TB</li><li>61.44TB</li></ul> <p><b>D205V:</b></p> <ul style="list-style-type: none"><li>122.88TB</li></ul>



# X200 Series Product Overview



## X200 Series SSD

The best PCIe Gen5 performance, features, endurance, and economics for enterprise applications. The X200 shows Phison's dedication to developing advanced SSD technology to lead the industry in density, performance, and power efficiency for all mass-capacity storage providers.

### Specification

- Interface: PCIe 5.0 x 4
- Protocol: NVMe 2.0
- Capacity: Up to 30.72TB
- Form Factor: U.2 / E3.S
- DWPD: 1 and 3 DWPD
- MTBF: 2.5 million hours
- Warranty: 5 years


### Features

- Dual-port design
- Ultra-low latency
- Power loss Protection
- MF-QoS
- Namespaces: 128

### Performance

Seq. Read	<div></div>	14,800 MB/s
Seq. Write	<div></div>	8,700 MB/s
Ran. Read	<div></div>	3,200K IOPS
Ran. Write	<div></div>	930K IOPS

# X100 Series Product Overview



## X100 Series SSD

Phison’s X100 series SSD is highly customizable and marks the spot for enterprises demanding faster and smarter global infrastructures. Featuring best-in-class performance, the X100 enables enterprises to reduce the total cost of ownership through higher storage density, lower power consumption and higher performance.

Specification
<ul style="list-style-type: none"><li>Interface: PCIe 4.0 x4</li><li>Protocol: NVMe 1.4</li><li>Capacity: 2TB to 32TB</li><li>Form Factor: U.3/U.2, 2.5” x 15mm</li><li>DWPD: Up to 3 DWPD</li><li>MTBF: 2.5 million hours</li><li>Warranty: 5 years</li></ul>

Features
<ul style="list-style-type: none"><li>Dual-port design</li><li>Ultra-low latency</li><li>Power loss protection (PLP)</li><li>End-to-end data path protection (E2EDPP)</li><li>Phison 5th Gen LDPC ECC engine</li><li>Self-encrypting drive (SED) &amp; FIPS 140-3<sup>1</sup></li><li>Optimized for 24/7 enterprise workload</li></ul>

Performance		
Seq. Read	<div></div>	7,400 MB/s
Seq. Write	<div></div>	6,900 MB/s
Ran. Read	<div></div>	1,750K IOPS
Ran. Write	<div></div>	470K IOPS

<sup>1</sup> Based on customer’s requirement.

# Industry's Most Advanced PCIe 5.0 Enterprise SSD (7.68TB, 1 DWPD)



**PHISON**

X200P



**SAMSUNG**

PM1743



**KIOXIA**


CM7-R




**SK hynix**

PS1010\*

Sequential (MB/s)



4K Random (KIOPS)



\*PS1010 performance is not based on 7.68TB