


Product and Solution Introduction

Sandy Riza

ALTOS Product Manager



About Altos Computing



2017

Acer established ALTOS Computing Inc.

1997

Altos launched first x86 server

1990


Acer acquired Altos Computer

1977

Altos Computer Founded
in the United States

Products & Solutions Portfolio

To quickly build an AI and Cloud Computing environment, accelerate the digital transformation, and create new business models.



Altos aiWorks & aiGeni
A Flexible and efficient AI computing solution



AltosCloud VDI/IDV/VOI
Achieve centralized and simplified Desktop OS image management



GRAID
NVME Optimization
Optimize NVMe Performance

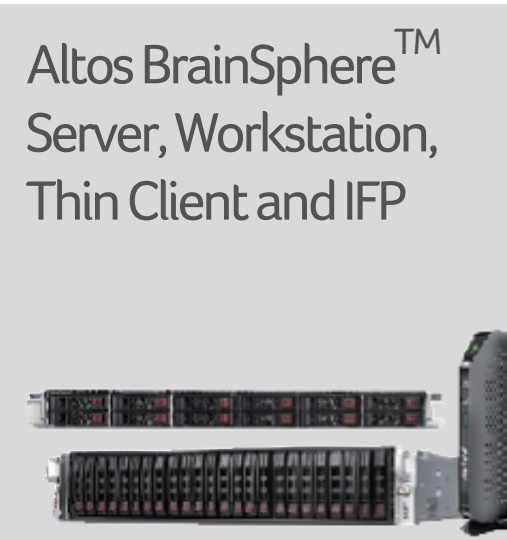


Data Analytics
Simplified Big Data Platform



Altos Security
ALTOS Next Generation Firewall

Altos BrainSphere™
Server, Workstation,
Thin Client and IFP



Storage
*NAS/SAN/DAS/FAS



QSAN


Altos BrainSphere™ Server/Workstation/Solution

Workstation P Series



T3-33S

T3-150



P10 F10 SE
P130 F10 SE



P5 F10
P5 F10 SE

Tower Server T Series



T3-5F5

T1-F6



T15 F6

Rack Server R Series



R680 F7



R785 F6



R380 F7



R3-8F6



R3



R385 F6



R385 F5

Storage Server S Series



R3-S4 F6



R3-S3 F6

IFPD



ALTOS IFP 65" 75" 86" 98"

Solution




ALTOS Next Generation Firewall



ALTOS Data Analytics

Altos - AI Server and Solution



Altos BrainSphere™ AI Server Portfolio

ALTOS
an Acer Group Company



Altos BrainSphere™ R785 F6

- 8 PCIe GPUs: up to H200 NVL / L40S/ RTX PRO 6000 Server Edition
- Dual AMD EPYC™ 9005
- 4U, MGX, air-cooling



Altos BrainSphere™ R880 F7

- 8 GPUs: B200
- Dual Intel Xeon 6
- 8U, HGX, air-cooling



NVIDIA GB300 NVL72

- 72x NVIDIA Blackwell Ultra B300 GPUs
- Direct Liquid Cooling



Altos BrainSphere™ R685 F5

- 8 PCIe GPUs: up to H100 / L40S / RTX 6000 Ada
- Dual AMD EPYC™ 7003/7002 series
- 4U, air-cooling



Altos BrainSphere™ R680 F7

- 8 PCIe GPUs: up to H200 NVL / L40S / RTX PRO 6000 Server Edition
- Dual Intel Xeon 6
- 4U, air-cooling



Altos BrainSphere™ R385 F6

- 2 PCIe GPUs: up to NVIDIA L40S, RTX 6000 Ada, RTX PRO 6000 Blackwell Max-Q
- Dual AMD EPYC™ 9005 / 9004
- 2U, air-cooling



Altos BrainSphere™ R3-8F7

- 2 PCIe GPUs: up to NVIDIA L40S, RTX 6000 Ada, RTX PRO 6000 Blackwell Max-Q
- Dual Intel Xeon 6
- 2U, air-cooling



Altos BrainSphere™ R3-8F6

- 4 PCIe GPUs: up to NVIDIA L40S, RTX 6000 Ada
- Dual 5th/4th Gen Intel Xeon Scalable processors
- 2U, air-cooling

Training

Inference

Altos BrainSphere™ AI/LLM Solution portfolio



GB300 NVL72

**Extreme
Scale HPC &
Gen AI Cloud**




8 x B200 NVL



R880 F7

**Large Scale Training/
LLM Fine-tuning**
Data sovereignty
And security compliance



8 X H200 NVL



R785 F6



R680 F7

**Enterprise Inference & GenAI
Microservices**
Improved employee productivity, optimized
internal processes



2 x NVIDIA RTX PRO 6000 Blackwell Max-Q



R385 F6



R380 F7

Entry AI / Edge Inference
Lower AI costs,
Accelerated Implementation

Altos BrainSphere™ R680 F7



ALTOS
an Acer Group Company



Ready Out-of-the-Box

Pre-installed with LLM, delivering a fully integrated software–hardware solution.

Rapid Deployment

Provides a model fine-tuning deployment guide to accelerate on-prem implementation.

Professional Services

A dedicated expert team supports industry-specific scenarios, offering on-site services.

Secure & Reliable

Local on-premise deployment ensures enterprise data never leaves the premises.

VRAM 1.1 TB	LLM Llama-3.1-405B	Memory 2TB
AI Compute 32 petaFLOPS of inference performance	Throughput 570 tokens/s	Concurrent Users 28

Product solution	Altos BrainSphere™ R680 F7 4U GPU
CPU	Dual 6th Gen Intel Xeon
GPU	8x NVIDIA H200 NVLwith NVLink™ interconnect
Network interface	On-board 2x10GbE + 1x1GbE management with NCSI; recommended to add 100–200GbE or InfiniBand NICs for external API traffic
AI compute	FP8 Tensor ≈ 32 PFLOPS, FP16/BF16 ≈ 16 PFLOPS
Throughput	Offline ≈740 tokens/s, Server ≈570 tokens/s (FP8).
Software bundle	Altos aiWorks inference / RAG suite + NVIDIA TensorRT-LLM / Triton Inference Server + vLLM.

Altos BrainSphere™ R880 F7



Enterprise



Education



Healthcare

ALTOS
an Acer Group Company



Ready Out-of-the-Box

Pre-installed with the LLM, delivering a fully integrated software-hardware solution.

Professional Services

A dedicated expert team supports user-specific scenarios

Optimal Cost

Deploying a hundred-billion-parameter model, enabling rapid and cost-effective deployment.

Secure & Reliable

Local on-premise deployment ensures enterprise data never leaves the premises

VRAM

1.4 TB

LLM

Llama-3.1-405B

Memory

4TB

AI Compute

72 petaFLOPS of training performance

Throughput

1,280 tokens/s

Concurrent users

64

Product solution	R880 F7 8U all-in-one, B200x8 training node.
CPU	Dual 6th Gen Intel Xeon
GPU	8x NVIDIA B200 SXM, NVLink + NVSwitch interconnect.
Network interface	On-board 2x10GbE + management; optional 200/400GbE or IB NICs.
AI compute	FP8 ≈ 72 PFLOPS, FP16 ≈ 36 PFLOPS.
Throughput	Llama-3.1-405B ≈ 1,280 tokens/s (Server Inference).
Software bundle	Altos aiWorks training suite + CUDA / PyTorch / vLLM.

Challenges Faced by AI Development Environments

Software Installation and Upgrade

Updating the software used for teaching and learning to keep pace with the latest trends and innovations.

Usage Optimize for maximum efficiency

The simultaneous operation of numerous software programs hinders optimal performance.

Security and Management

Stay vigilant against attacks or hacking attempts to ensure data integrity and security.



Altos aiWorks

AI Readiness Platform for Developer

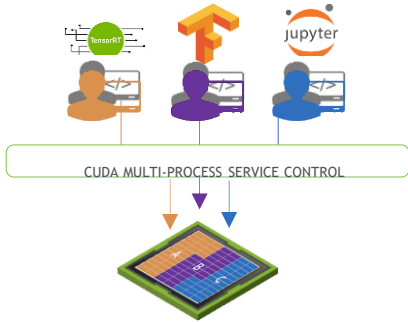


Flexible Allocation of Computing Resources

Multi-Process Service (MPS)




For AI Teaching & AI model development



- Software defined-Share GPU with multiple processes.
- Support NVIDIA GPU

Multi-Instance GPU (MIG)

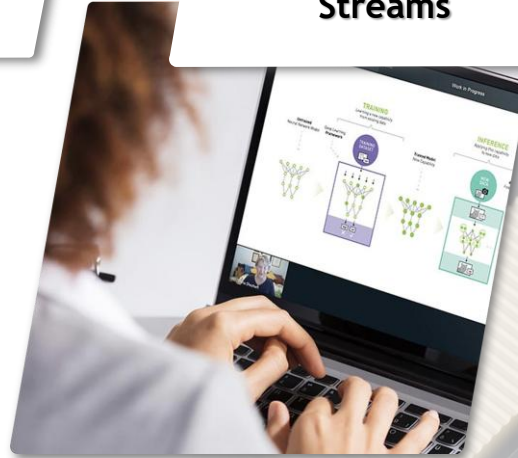


For AI model development & fine-tuning




- Physical partition-Separate GPU into multiple instances.
- Ensure isolation & QoS of each instance.
- Supports Nvidia H200, H100, A100, Pro 5000, Pro 6000, A30, B200

Streams



For AI model training & fine-tuning



- Allocate more than one GPU
- Support NVIDIA GPU

GPU SHARING POLICY

Altos Server & aiWorks – Integrated AI Orchestration Platform

AI Development Ecosystem Layer

<p>IDE Tools</p> 	<p>Training Framework</p> 	<p>Serving Tools</p> 
--	---	---

Pre-loaded AI development tools & AI frameworks

aiWorks GUI Control Plane

Multi GPU Computing	GPU-Fractioning	Batch Job	Single Sign-On
Multi Node	Multi-Tenancy	Scheduling	Custom Image

Simplified AI workload deployment

Orchestration Layer



AI Resource Analysis

AI Infrastructure

Compute	Storage	GPU	Network
---------	---------	-----	---------

GPU Resource Management

Infrastructure Engine Layer

					
R380 F7	R680 F7	R880 F7	R385 F6	R685 F5	R785 F6

Centralized AI Resource Management

```
# Setup GPU-Operator MIG add Nvidia Device Plugin Helm Chart Repo
helm repo add nvdn https://nvidia.github.io/k8s-device-plugin
helm repo update
kubectl create ns nvidia-device-plugin

# Prepare MPS Configmap Kubectl apply -f custom-mps-config.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  name: custom-mps-config
  namespace: nvidia-device-plugin
data:
  v100: |-
    version: v1
    sharing:
      mps:
        renameByDefault: true
        resources:
          - name: nvidia.com/gpu
            replicas: 8

# Helm Setup Nvidia Device Plugin
helm upgrade -i nvdn nvdn/nvidia-device-plugin \
  --version=0.16.1 \
  --namespace nvidia-device-plugin \
  --create-namespace \
  --set config.name=custom-mps-config

# Use MPS
kubectl label node <node> nvidia.com/device-plugin.config=v100 --overwrite
kubectl label node <node> nvidia.com/mps.capable=true --overwrite
```



GPU-Specific Settings

GPU 0: GPU Sharing Policy: Dedicated MIG

1/2	Memory: <input type="text" value="0.0"/> GB	<input type="text" value="0.0"/> GB	— +
1/3	Memory: <input type="text" value="25.0"/> GB	<input type="text" value="25.0"/> GB	
1/7	Memory: <input type="text" value="17.0"/> GB	<input type="text" value="17.0"/> GB	
	Memory: <input type="text" value="7.0"/> GB	<input type="text" value="7.0"/> GB	
<input checked="" type="checkbox"/> MPS	5% Memory: <input type="text" value="3.0"/> GB	<input type="text" value="4.0"/> GB	
	10% Memory: <input type="text" value="5.0"/> GB	<input type="text" value="6.0"/> GB	
	20% Memory: <input type="text" value="10.0"/> GB	<input type="text" value="11.0"/> GB	
	25% Memory: <input type="text" value="15.0"/> GB	<input type="text" value="16.0"/> GB	
	50% Memory: <input type="text" value="25.0"/> GB	<input type="text" value="26.0"/> GB	

GPU 1: GPU Sharing Policy: Dedicated MIG

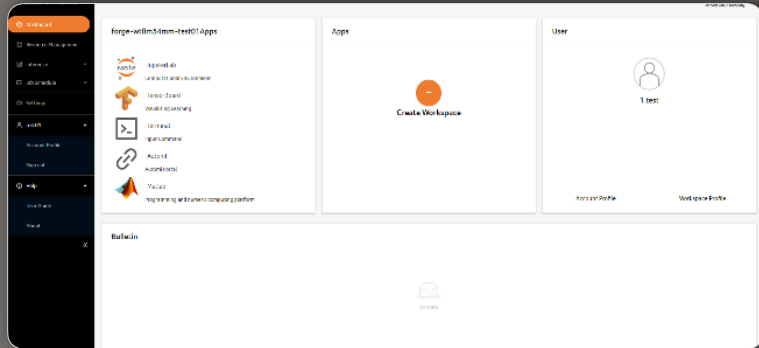
1/2	Memory: <input type="text" value="0.0"/> GB	<input type="text" value="0.0"/> GB	—
1/3	Memory: <input type="text" value="0.0"/> GB	<input type="text" value="0.0"/> GB	
1/7	Memory: <input type="text" value="0.0"/> GB	<input type="text" value="0.0"/> GB	
<input checked="" type="checkbox"/> MPS	5% Memory: <input type="text" value="2.0"/> GB	<input type="text" value="3.0"/> GB	
	10% Memory: <input type="text" value="3.0"/> GB	<input type="text" value="4.0"/> GB	
	20% Memory: <input type="text" value="6.0"/> GB	<input type="text" value="7.0"/> GB	
	25% Memory: <input type="text" value="8.0"/> GB	<input type="text" value="9.0"/> GB	
	50% Memory: <input type="text" value="13.0"/> GB	<input type="text" value="14.0"/> GB	

Easy GPU Setup with No-Code Tools

Configure MPS and GPU settings

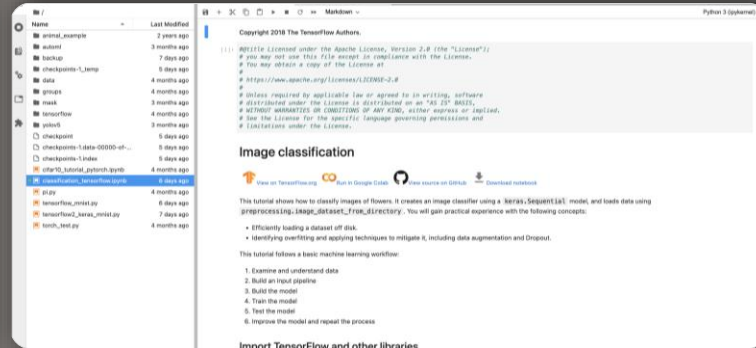
User Interface

Manage Tools and Resources in the Workspace



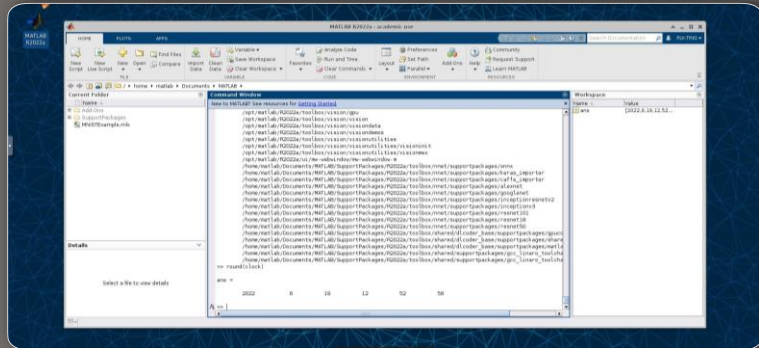
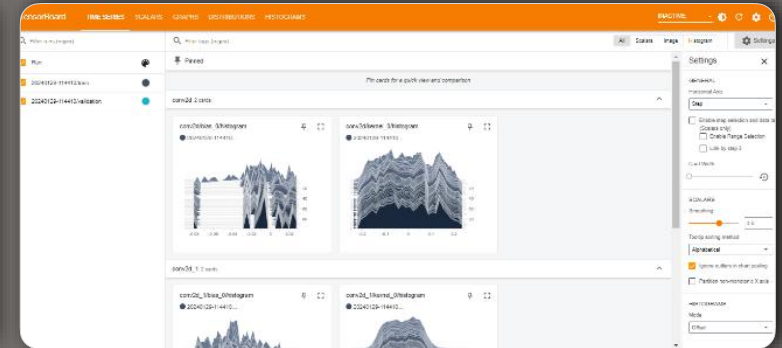
JupyterLab

Pre-configured Development Platform



TensorBoard

Track Metrics and Visualize Models



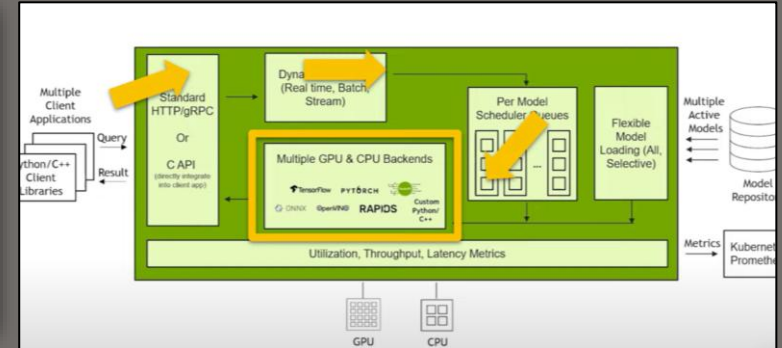
MATLAB

Containerized Service: Use with MATLAB License Key



Terminal


Command Line for Running Ubuntu Commands



NVIDIA triton

Inference Server for AI Model Deployment

Altos - AI Workstation



Building the Ideal Workstations for AI Development & Design

High-performance processors × Professional GPUs × Flexible memory configurations



Altos BrainSphere™ GB10 F1	Altos BrainSphere™ T3-150	Altos BrainSphere™ T3-33S
<ul style="list-style-type: none"> • Arm 20 core, 10 Cortex-X925 + 10 Cortex-A725 • 128GB LPDDR5X, 273 GB/s • 1000 TFLOPS(FP4) 	<ul style="list-style-type: none"> • Intel® Core™ Ultra (Series 2) • Up to DDR5 192GB • 2x RTX Pro 6000 Workstation Max-Q, up to 7000 AI TFLOPS 	<ul style="list-style-type: none"> • Intel® Xeon® W-3400/W-3500 • Up to DDR5 512GB • 2x RTX Pro 6000 Workstation Max-Q, up to 7000 AI TFLOPS
Up to ~7B parameters (Fine-tune AI model)	Up to 7B-30B parameters (Fine-tune AI model)	Up to 40B-70B parameters (Fine-tune AI model)
	Rapid CAD/CAE visualization Small-component simulation	Medium- to high-resolution CAD/CAE simulation

Altos aiGeni: GUI Provides ease of use and efficiency

1 No Setup Needed

Run one installer and be ready in **under 1 minute** - no manual setup or extra packages.

1

2 Data Stays Private

All your data stays on-premises - safe, private, and compliant.

2

3 Automatic Snapshots

Background snapshots every **15/30/60 minutes**; **one-click rollback** with no reinstall.

3



6

6 Refined UI/UX

Responsive layout with a clean, consistent design for faster onboarding and use.

5

5 Preloaded AI Tools

JupyterLab, VS Code, TensorBoard, Ollama, Automatic1111, WhisperUI - **single web entry, click-to-use.**

4

4 Real-Time Resource Panel

Built-in gauges for **CPU, GPU, memory, disk, and network** - no external agents required.

Altos aiGeni Portal Introduction

The screenshot displays the Altos aiGeni portal interface. The main content area shows the status of a container named 'llm-project-1'. The status is 'Running', indicated by a green dot. Below the status, there is a note: 'Auto Saving Environment in: 15 min' and 'Keep Last Backups: last 5'. The 'Port Mapping' section shows '8888 -- NodePort: 31423'. The 'Command' section shows the command: 'jupyter lab --ip=0.0.0.0 --allow-root --no-browser --NotebookApp.token=""'. The 'Volume' section shows the path: '/home/altos/jupyter -- /home/joyan/work'. The 'Events' section shows 'No events'. The 'Logs' section shows 'Waiting for logs...' and an 'OPEN TERMINAL' button.

The right sidebar displays hardware resource monitoring for 'Predator PTN16-51':

- GPU Utilization: 0%
- GPU VRAM: 341 / 7805 MIB
- GPU Power: 1.8 / 60.0 W
- GPU Temperature: 37 °C
- CPU Usage: 3.7%
- Memory Usage: 18868 / 31763 MIB
- Disk Usage: 136 GiB / 937 GiB

Check Container Status

Auto backup your environment in every 15 min

Global File system location

Show logs to help user know the detailed information of system

IDE tools & AI tools

Hardware Resource Monitoring



GRAID

ALTO S

an Acer Group Company

**SupremeRAID™
NVMe SSD Solution**


The Challenge

Traditional RAID is a bottleneck

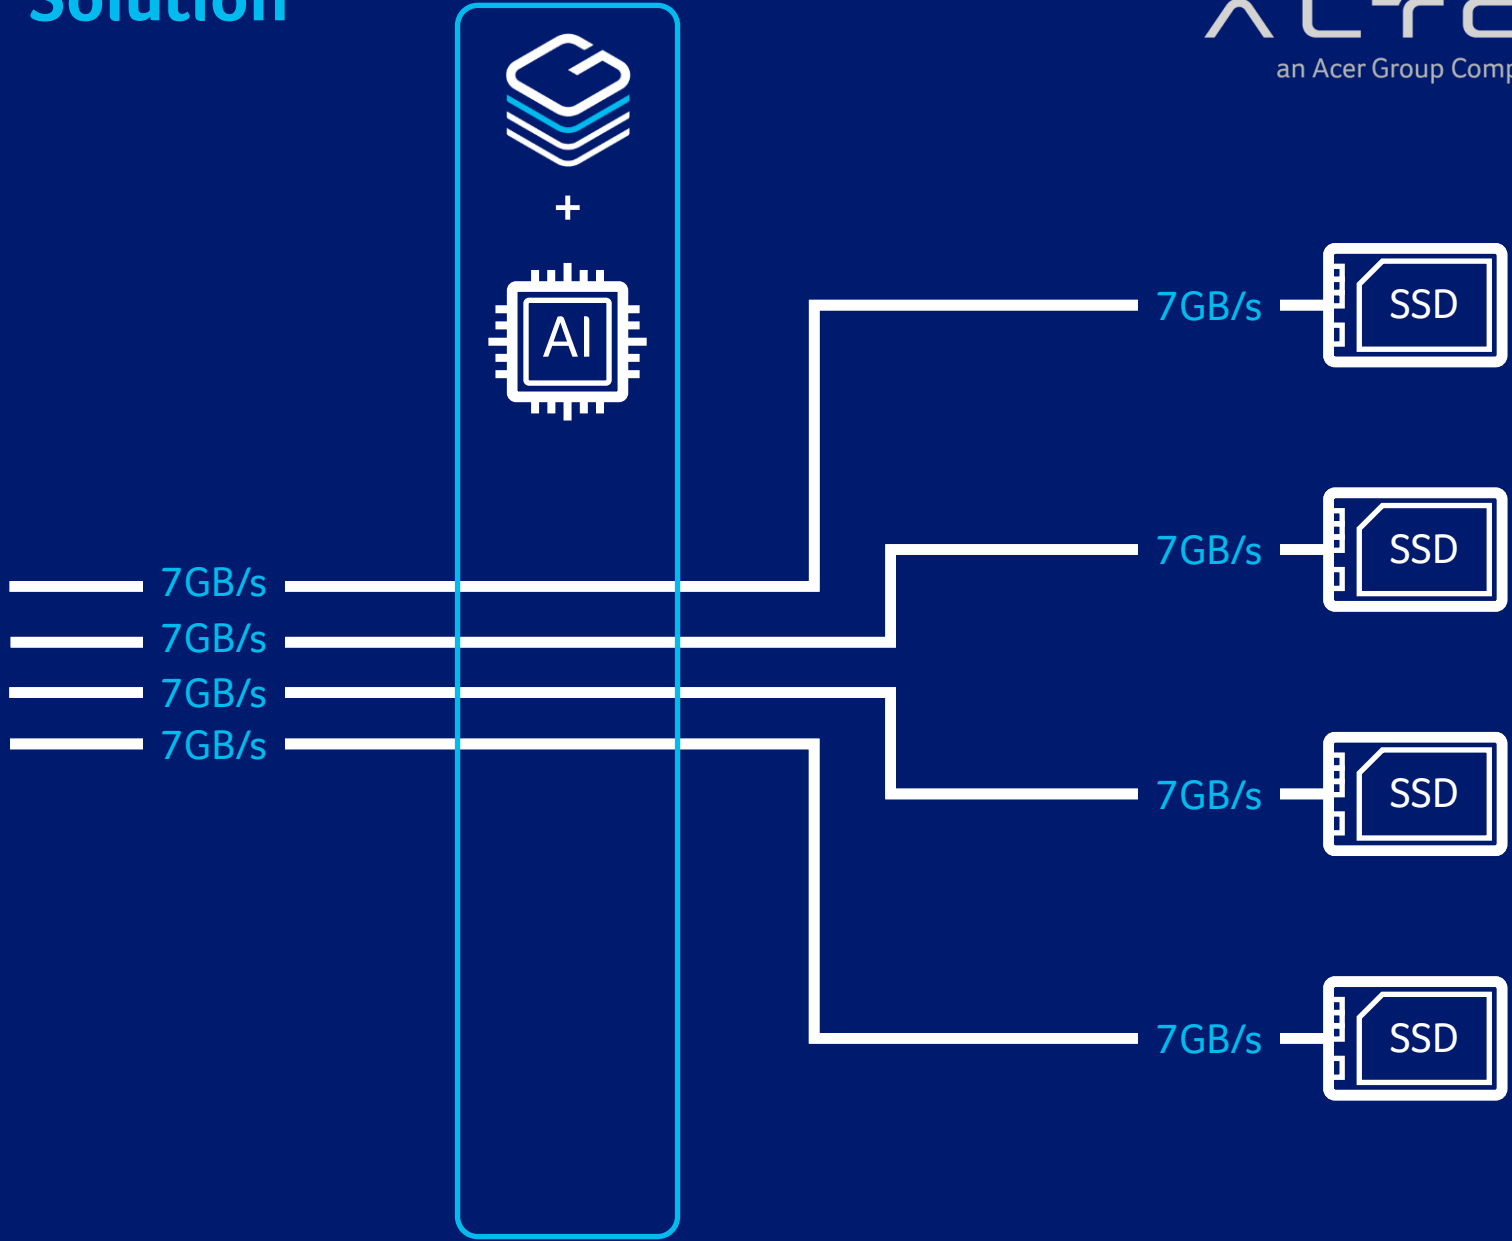

- **Flash storage** performance is evolving too fast to be fully utilized by existing storage architectures
- **Software-based RAID** competes with other applications for CPU power, slowing the whole system down
- Advances in NVMe technology are moving so fast that **ASIC-based RAID** can't deliver the full performance capability of the SSDs



Why traditional RAID is a bottleneck?






With SupremeRAID™ Solution



Analyze 25TB Network Packet Data

Best Performance!

	5unit	1unit	1unit
			
Configuration	CPU: Xeon 4310 *2 Mem: 256GB HDD: 8TB SAS *10	CPU: Xeon 6330 *2 Mem: 512GB HDD: 3.2TB NVMe *20	CPU: Xeon 6330 *2 Mem: 512GB HDD: 3.2TB NVMe *20
Solution	Ceph	OS Raid (Raid 5)	Altos High-Performance NVMe SSD Storage (Raid 5)
Performance (Processing time)	94hrs	54hrs 1.74 x 40% time saving	10hrs 9.4 x 89% time saving
Performance IOPS	11.7K	83.9K	105.6K

Key benefit

- Originally, analyzing 25TB network packet data took 94 hours with five systems using 320TB SAS .

- Shifting to a single system with 60TB NVMe SSD in RAID 5 reduced the time to 54 hours due to NVMe's faster speed.

- By implementing Altos High-Performance NVMe SSD Storage Solution and using a single system with a 60TB NVMe SSD in RAID 5, the analysis time was further **reduced to just 10 hours.**



Get Creative With Touch

Introducing
ALTOS IFP Series | Smart Interactive Panel
65" / 75" / 86" / 98"

ALTOS
an Acer Group Company

Interactive Flat Panel (IFP)

The built-in player has powerful functions and a user-friendly interface, providing you with a seamless writing experience.



Interactive and Functionality




	IFP Smart Panel	Interactive Projector	Smart TV
Screen Size	65" - 98" Inch	80 Inch Max	65" - 85"
Function	Educational Activities & Teleconference Presentations & Online Meetings Connected to Various Devices Digital information boards Compute with Dual OS Entertainment	Presentation	Entertainment
Resolution	4K	FHD	4K - 8K
Screen Protection	7 MOHS - Hardness 9	No	Temperd Glass
Multiple Touch	50 Touch Point	No	No
Stylus Support	Yes	No	No
Connectivity	Wifi & Bluetooth	No	Wifi & Bluetooth
Ports	20 Port	5 Port	5 Port
Life time	50.000 Hours Panel	1000 Hours Lamp	10.000 Hours Panel

Product Walk Around

Front View



1. Type-C
2. HDMI in 1
3. Touch USB-B
4. Public USB 3.0
5. Public USB 3.0
6. Power button
7. Microphone or Camera with Microphone



open pluggable specification (OPS)

Side View




1. AC Power Switch
2. Power Cable Input



1. Public USB 3.0
1. RS232
2. HDMI in 3
3. HDMI in 2
4. DP in
5. Type-C
6. Touch USB-B
7. Public USB 2.0
8. Android USB 2.0
9. Public USB 3.0
10. Android USB 3.0

Bottom View



1. MIC in
2. Audio out
3. SPDIF out
4. HDMI out



Standing Bracket

ALTOS
an Acer Group Company

Healthy Learning

Health and learning go hand in hand, and a healthy learning environment is the basis for focused learning. Acer antibacterial series provides teachers and students with a worry-free learning solution to create a healthy learning environment.





Full Featured Eye Care

Designed for prolonged periods of usage, this product has passed TÜV Rheinland certification for user safety. It uses flicker-free technology to eliminate annoying screen flickering problems. BlueLightShield Pro and anti-glare features can reduce the burden on the eyes and effectively reduce visual fatigue.



TÜV /Eyesafe



Flickerless



BlueLightShield Pro



Anti-Glare



Creative Teaching

ALTOS IFP series has built-in professional writing software and many projection functions. The intuitive usage provides teachers with more flexible teaching methods.



15W
USB Type-C




10/100/1000 Mbps
Ethernet Port



Cloud OTA
Update

Intuitive writing software makes learning easier.



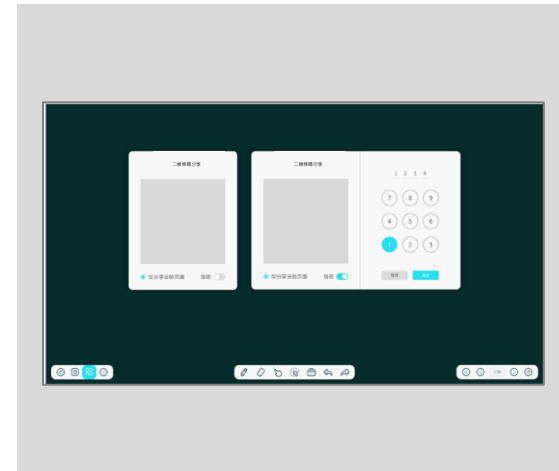
Multi-frame Writing

Divide the screen into 2, 3, and 4 writing areas without interfering with each other; convenient for different students to answer questions simultaneously.



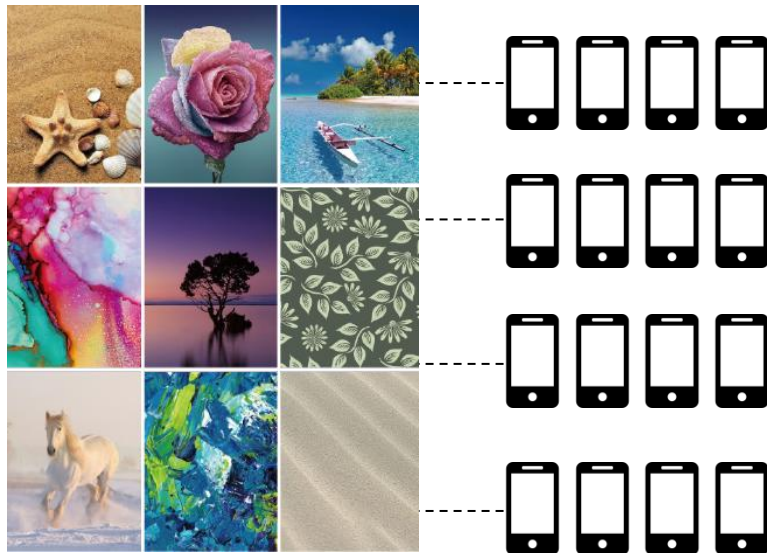
Free Hand Drawing

Provides a borderless writing and drawing space so that your creativity is not limited. Zoom in and out with one click, and expand the canvas infinitely, convenient for teaching and discussion.




Cloud Storage

Supports the storage and download of cloud data, as well as file sharing through QR code, so that teachers and students can retrieve the latest teaching materials anytime, anywhere.



Joint Collaboration

The collaboration function of Eshare Pro can work on up to 16 screens at the same time, which is convenient for discussions or for using different applications at the same time.



Limitless Projection

Eshare Pro allows you to enjoy content projection and the sharing of portable device screens. Easily exchange ideas and collaborate creatively with participants.

Interactive Features



VOTE

Through the App, users can use the voting option on their own phones by scanning a QR code. This is especially convenient with large groups of people, eliminating the need to count votes manually or using external software.



SCREEN RECORDING

Use the App to screen record. Easily record full classes and conferences for future playback and review.



SYNC WITH YOUR MOBIL PHONE

Educators do not have to stay at the front of the classroom all the time to use the IFP. Oftentimes, teachers have to walk around the classroom to check on student work



4K2K UHD LED Backlit Display

Acer Altos IFP is a series of UHD LED-backlit displays. It has 4K2K (3840x2160) ultra-high resolution and clear imagery. Its power consumption is lower than that of traditional CCFL monitors, and its service life is longer. It is suitable for multimedia applications that require fine picture quality and rich visual communication content.



LED Backlight

10
Bit

10 Bit

10Bit (107 million) Colors

Acer Altos IFP series' 10-bit (107 million colors) professional color presentation presents crystal clear imagery, providing viewers with a fantastic visual experience.



IR Touch

Infrared (IR) Technology

The Acer Altos IFP Series adopts infrared (IR) touch technology, which greatly reduces touch blind spots and improves touch accuracy. The panel brightness is 350 ~ 400nits, which is the ideal brightness for indoor environments.



Touchpoint

50 Touch Point

50 Point Touch Display

Acer Altos IFP series allows you to easily browse content with the touch of a finger. The IFP series supports 40-point* touch, and can recognize various gestures.



IFP

Built-in IFP (Interactive Whiteboard)

Through the IFP (Interactive Whiteboard) function, online sharing and collaboration can be provided, which is helpful for interactive teaching and for improving meeting efficiency.



18/7 Operation

18/7 Operation

The Acer Altos IFP series is specially designed for long-term use, with a service life of up to 50,000 hours and an automatic cooling system to ensure that the equipment can effectively maintain 18/7 uninterrupted long-term operation.

Altos Customers in Indonesia



Universitas Terbuka
Making Higher Education Open to All



POLITEKNIK NEGERI BANDUNG
Assuring Your Future



PERPUSTAKAAN NASIONAL
REPUBLIK INDONESIA

A night view of Earth from space, showing city lights and a glowing blue network overlay of nodes and lines. The text "THE BEST IS YET TO COME" is centered in white.

THE BEST IS YET TO COME