

XPerf ClusterReady™ Validates AI Racks 10× Faster: Delivers Industry-Leading AMD GPU Cluster Results with MiTAC Computing

Rack validation and optimization software that generates tokens per

second output under AI workloads — closing the gap between L11/L12 testing and the revenue performance operators depend on.

“

Partnering with MiTAC, we validated and optimized an MI325X rack in days and proved it outperforms AMD's published benchmarks by 20%. That's the standard ClusterReady™ sets.”

*Alex Carter, Found/CEO of
XPerf Inc*

ROUND ROCK, TX, UNITED STATES, June 18, 2026

/EINPresswire.com/ -- XPerf Inc. announced its deployment with [MiTAC Computing](#), ClusterReady™ demonstrated that the AMD Instinct™ MI325X AI air-cooled racks—featuring MiTAC Computing's G8825Z5 servers—achieved performance 20% higher than AMD's publicly reported FLUX.1-dev benchmark results under the tested configuration, and produced the first reported Llama 3.1 8B pre-training performance results on MI325X. Across the full 4-node rack, XPerf's network optimization delivered a 10× improvement in multi-node communication bandwidth, enabling 95.8% scaling efficiency on Llama 3.1 8B pre-training. Full MLPerf® Training V6.0 benchmark

results have been submitted.

Conventional L11/L12 rack tests verify integration and functional readiness but cannot replicate live sustained AI workloads — a rack that passes every standard test can still fail in production, eroding the ROI operators depend on. ClusterReady™ stress-tests racks using a combination of inference and training AI workloads alongside XPerf's proprietary test suite — exercising GPUs, memory, interconnects, network fabric, and thermal systems under conditions that mirror live production. ClusterReady™ measures tokens per second output under sustained load, giving a concrete, revenue-linked performance baseline for every rack. When issues are detected, XPerf's proprietary algorithms identify the root cause and, uniquely among validation tools, prescribe specific remediation or optimization steps, reducing manual engineering effort.

Key features:

- 10× Faster Validation — XPerf's proprietary algorithms and automated workflow compress full

rack validation from weeks to hours or days, so racks reach revenue-generating production in a fraction of the time.

- Stress Test Suite — combines inference and training AI workloads with XPerf's proprietary test suite to stress GPUs, memory, interconnects, network fabric, and thermal systems, exposing hardware and firmware issues that conventional methods miss.
- Root Cause Analysis with Prescribed Solutions — the only validation tool that identifies the root cause of detected issues and prescribes specific remediation and optimization steps, cutting engineering triage time and getting racks ready faster.
- Revenue-Linked Performance Baseline — every rack ships with a ClusterReady™ validation report documenting tokens per second output across inference and training workloads — providing a performance baseline for revenue generation.

"XPerf delivered excellent results on cluster-level validation and benchmarking within days — work that typically takes weeks. Our AMD MI325X AI air-cooled rack validated seamlessly and achieved 20% stronger performance on industry-relevant AI workloads than the standard AMD MI325X rack, highlighting MiTAC's engineering expertise in building high-performance AI infrastructure."

— Raymond Huang, GM, MiTAC Computing Technology Corp., USA

"The gap between assembling a GPU rack and knowing its performance has cost this industry too much for too long. Partnering with MiTAC, a leading global AI server manufacturer, we validated and optimized a MI325X rack in days and proved it outperforms AMD's published benchmarks by 20%. The same speed applied to the validation of a MI350X rack. That's the standard ClusterReady™ sets."

— Alex Carter, Founder & CEO, XPerf Inc.

Availability — ClusterReady™ is commercially available and designed for rapid deployment in supported environments. To learn more or schedule a ClusterReady™ validation demo, visit xperf.ai/clusterready

About XPerf Inc.

Founded by ex-Intel engineers with deep experience deploying GPU clusters at scale, XPerf Inc. is

XPerf ClusterReady™
Validate AMD Instinct™
MI325X AI Racks
10x Faster

20% Stronger AI workload performance
with validation in days, not weeks

XPerf ClusterReady™ Validate AMD Instinct™ MI325X
AI Racks 10x Faster

an AI infrastructure software company building the control plane for heterogeneous data centers. Its product — ClusterReady™, which validates AI racks and clusters 10× faster than L11/L12 workflows using real AI workloads, and OpsPilot™, which autonomously predict and resolves GPU cluster failures on-premise — cover the full cluster lifecycle from pre-production validation to autonomous operations. XPerf Inc. is a member of NVIDIA Inception, CNCF, and the gBETA accelerator program. Learn more at xperf.ai

About MiTAC Computing

MiTAC Computing Technology Corp., a subsidiary of MiTAC Holdings, specializes in AI, HPC, cloud, and edge computing. MiTAC Computing employs rigorous methodologies to ensure uncompromising quality—across barebones, systems, racks, and cluster levels—fully achieving performance and integration. With a worldwide presence and end-to-end capabilities—from R&D and manufacturing to global support—MiTAC Computing provides agile, customized platforms for hyperscale data centers, HPC, and AI applications. Learn more at www.mitaccomputing.com.

Trademarks: AMD, the AMD Arrow logo, AMD Instinct, and combinations thereof are trademarks of Advanced Micro Devices, Inc. FLUX.1 is a trademark of Black Forest Labs GmbH. NVIDIA is a registered trademark of NVIDIA Corporation. Intel and Intel Gaudi are trademarks of Intel Corporation. Llama is a trademark of Meta Platforms, Inc. Other names are for informational purposes only and may be trademarks of their respective owners.

Forward-Looking Statements: This press release contains forward-looking statements regarding XPerf's products and related system validation results. Neither XPerf Inc. nor MiTAC Computing makes any guarantee of future performance outcomes. These statements are based on current expectations and are subject to risks and uncertainties that could cause actual results to differ materially. Performance results are based on internal testing conducted by XPerf Inc. in collaboration with MiTAC Computing; actual results may vary based on system configuration, software versions, and workloads. Performance results referenced herein are specific to the MiTAC AMD Instinct™ MI325X rack used in this test environment and are not a guarantee of performance for any other MI325X deployment or configuration. XPerf undertakes no obligation to update these statements after the date of this release.

Alex Carter

XPerf Inc.

press@xperf.ai

Visit us on social media:

[LinkedIn](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/920283785>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire,

Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.