



Clockwork.io Raises \$31M as LinkedIn, Together AI and WhiteFiber Adopt Its Resilience Software to Stop Wasting GPU-Hours

Descrizione

COMUNICATO STAMPA â?? CONTENUTO PROMOZIONALE

LinkedIn prevents tens of thousands of GPU-hours of downtime monthly; new TorchPass innovations preserve AI workload progress without code changes and speed reinforcement learning.

PALO ALTO, Calif., Oct. 5, 2026 /PRNewswire/ â?? Clockwork.io, whose fault-tolerance software keeps AI training, reinforcement learning and inference workloads running through infrastructure failures, today announced \$31 million in new funding, production deployments at LinkedIn and Together AI, and expanded adoption by WhiteFiber.

The company also introduced two new capabilities for its TorchPass solution, each capturing the state of a running distributed AI job. Multi-node platform snapshots, an industry first for training, save an entire running job across every node without changes to the training code and preserve it for recovery. Fast, asynchronous application checkpoints, taken in the background while the job runs, accelerate reinforcement learning. They deliver updated model weights to the inference replicas that generate rollouts, the examples the model learns from, so those replicas spend less time waiting or working from a stale model.

Fault tolerance has become a requirement for AI at scale

Large distributed AI workloads can span thousands of GPUs that must stay in sync: one failed GPU, dropped link or frozen server can stall the entire job. Meta reported unexpected interruptions averaging roughly one every three hours during a 54-day period of Llama 3 training on 16,384 GPUs.

The typical response is to reload a checkpoint, a saved copy of the jobâ??s progress. Recovery can take up to 90 minutes, leaves healthy GPUs waiting, and requires the job to repeat work completed since that checkpoint. Customers pay for idle GPUs and repeated computation, and models take longer to complete. As jobs grow, each restart puts more GPU time at risk.

At this scale, keeping useful work running through failures is an infrastructure requirement. Clockwork.io meets it with a fault-tolerance suite that platform teams deploy as a layer between the hardware and the workload. LinkPass reroutes traffic around a failed link so the job never sees the fault. TorchPass moves work from a failing GPU to a healthy one so training continues instead of rolling back. Both are in production. TorchPass's new platform snapshots, announced today, capture the state of a running distributed job so the whole job can be restored when a failure is too large to migrate around.

“Failures are inevitable at AI scale. Losing hours of useful work to them should not be,” said Suresh Vasudevan, CEO of Clockwork.io. “Fault tolerance is a goodput multiplier: it keeps GPUs doing useful work instead of waiting for recovery or repeating work already done. We built our software alongside enterprises and cloud providers operating some of the largest GPU fleets, so it handles the failures they actually see. That protection belongs in the infrastructure enterprises and cloud providers rely on every day.”

Adoption expands across enterprises, hyperscalers and neoclouds

Enterprises running their own GPU fleets, hyperscalers and neoclouds are adopting Clockwork.io for the same reason: more of their GPU-hours go to useful work.

LinkedIn has deployed LinkPass network fault tolerance across its AI infrastructure fleet and prevents tens of thousands of GPU-hours of downtime each month.

“At AI infrastructure scale, a single network issue should never sideline healthy GPUs or interrupt running workloads. Before Clockwork.io, one InfiniBand NIC flap could remove an eight-GPU server from service, while a switch port flap could drain a second server, doubling the impact to 16 GPUs,” said Raghu Hiremagalur, SVP, CTO Infrastructure, LinkedIn. “Clockwork.io helped transform that operating model. Its network fault-tolerance technology automatically reroutes traffic onto healthy paths, allowing jobs to continue uninterrupted while link, optic, cable, or NIC faults are repaired. In aggregate, Clockwork.io prevents tens of thousands of GPU-hours of downtime per month across our fleet. By turning what were once disruptive operational incidents into manageable maintenance events, Clockwork.io has helped improve infrastructure utilization and operational efficiency.”

Together AI is bringing TorchPass to market as a service on its GPU Clusters. At the PyTorch Conference, the two companies will demonstrate a live multi-node training job continuing through injected network and GPU failures without restarting.

“Our customers grade us on goodput, the share of their GPU-hours that actually move the model forward,” said Pavneet Ahluwalia, Product Lead, Together AI. “Node repair already detects faults and provisions replacement capacity automatically. Clockwork.io's TorchPass and LinkPass build on that foundation and are designed to keep jobs moving through GPU faults and link failures, preserving progress. We are bringing them to market as the next layer of resilience in the platform.”

WhiteFiber (NASDAQ: WYFI), an existing customer, is expanding its use of Clockwork.io software across its growing global GPU-as-a-service footprint.

“Pressure-testing a cluster's reliability before it reaches production is critical, because a customer who inherits a hidden fabric fault pays for it later in failed jobs and lost GPU-hours,” said Tom

Sanfilippo, Chief Technology Officer, WhiteFiber. “Marginal optics, misconfigured NICs, and links that pass a basic test but degrade under load can slip through. Clockwork.io’s automated fleet audit validates every link and node at once, localizes faults in minutes, and lets us correct them before acceptance. We bring clusters up faster, and a customer’s first training run lands on a fabric validated end-to-end, not just powered on. With market demand growing as rapidly as it is, getting validated capacity to customers quickly is critical to our business, and it is why we are expanding Clockwork.io across our clusters.”

New TorchPass capabilities put workload protection in platform teams’ hands

Clockwork.io extended TorchPass beyond GPU migration with two capabilities platform teams have not had: a snapshot of a whole distributed job that they can take themselves, and application checkpoints fast enough to run in the background.

For training, platform snapshots save a running job’s execution state across all of its nodes so the job can be restored after an interruption. Platform teams and AI infrastructure engineers deploy it for supported workloads without waiting for application owners to modify their code or add checkpointing logic. Enterprise teams can protect training jobs across their fleet with one mechanism, and cloud providers can protect customer jobs whose code they do not control. Where teams checkpoint at the application level, TorchPass’s fast checkpoints can be taken more often, so less progress is lost and less computation repeated after a failure.

For inference, large models run across two or more servers, so one bad link can take down a whole replica and cut off a user session or agent task mid-stream. LinkPass keeps those multi-server replicas serving through link failures.

Reinforcement learning depends on both training and inference. Copies of the model generate rollouts, the trainer learns from them, and updated weights must reach the copies before they can generate with the latest version. TorchPass’s application checkpoints carry the updated weights to the rollout replicas sooner, while LinkPass keeps those replicas serving through link failures.

“Cluster fault tolerance used to be a training problem. It is now an inference problem too,” said Dylan Patel, Founder, CEO, and Chief Analyst at SemiAnalysis, whose ClusterMAX ratings benchmark GPU cloud providers. “In our ClusterMAX, TorchPass cuts training goodput loss from 14% to under 3% for a gold-rated neocloud. Reinforcement Learning (RL) ties the two together: inference replicas generate rollouts, the trainer learns from them, and the updated weights go back to the replicas. Clockwork.io keeps replicas serving through link flaps and network failures. Its extremely fast checkpoints accelerate weight transfer back into the rollout fleet, so neither direction stalls the run. One fault-tolerance layer under training, inference, and RL is where this has to be solved.”

Enterprise platform teams and cloud providers can contact Clockwork.io to evaluate the software for their workloads or explore partnership opportunities.

The round and what it funds

The round was co-led by Premji Invest, Wing Venture Capital, and Seligman Ventures, with participation from existing investors NEA and e& Capital. It brings Clockwork.io’s total funding to \$73 million. The company will use the capital to accelerate the rollout of its fault-tolerance suite across

training, inference and reinforcement learning, expand enterprise adoption, and scale delivery through cloud partners.

“The one thing that scales perfectly is unreliability: put enough GPUs in one machine and something is always failing,” said Greg Papadopoulos, Venture Partner, NEA. “The old playbook: stop the job, reload a checkpoint, makes no sense at today’s scale. Clockwork treats failure as the normal state: TorchPass migrates training off a failing GPU live, and now snapshots an entire running job with no code changes. It’s already saving tens of thousands of GPU-hours a month. We first backed Clockwork.io in 2021 and are thrilled to keep supporting them as they define the performance layer of the AI cluster.”

About Clockwork.io

Clockwork.io pioneers Software-Driven AI Fabrics®, a programmable layer between hardware and workload that makes GPU clusters observable, fault-tolerant, and fully utilized across any accelerator, network, or cloud. AI workloads need the whole cluster to act as one machine, yet failures and bottlenecks idle GPUs. Clockwork.io’s FleetLens platform recovers that lost capacity: nanosecond-accurate telemetry pinpoints the GPU, node, or link slowing or stalling a job, validates clusters before launch, and, with LinkPass and TorchPass, keeps workloads running through infrastructure failures and degradations, from training to inference. SemiAnalysis has independently benchmarked TorchPass as recovering from failures faster than checkpoint-restart and leading open-source frameworks. LinkedIn, Together AI, WhiteFiber, Wells Fargo, Nebius, NScale, and DCAI trust Clockwork.io to power their AI infrastructure. Learn more at www.clockwork.io.

View original content to download multimedia:<https://www.prnewswire.co.uk/news-releases/clockworkio-raises-31m-as-linkedin-together-ai-and-whitefiber-adopt-its-resilience-software-to-stop-wasting-gpu-hours-302897588.html>

Copyright 2026 PR Newswire. All Rights Reserved.

COMUNICATO STAMPA – CONTENUTO PROMOZIONALE: Immediapress – un servizio di diffusione di comunicati stampa in testo originale redatto direttamente dall’ente che lo emette. L’Adnkronos e Immediapress non sono responsabili per i contenuti dei comunicati trasmessi

–

immediapress/pr-newswire

Categoria

1. Comunicati

Tag

1. ImmediaPress

Data di creazione

Ottobre 5, 2026

Autore

redazione

default watermark