

Whitepaper

AI Anywhere ... Where is Anywhere?

Sponsored by:



Whitepaper

Sponsored by Hammerspace

AI Anywhere ... Where is Anywhere?

Introduction

Analysts and business leaders increasingly agree that AI is critical to market competitiveness. As a result, enterprises are accelerating efforts to run AI anywhere (and everywhere).

To create and deploy the most effective, efficient, and competitive AI solutions, enterprises are testing and developing solutions using the multitude of available AI technologies. For each of these approaches, AI engines demand access to as much data from as many data sources as possible.

Making data available and managing access to the data for AI is the biggest challenge facing IT leadership in every enterprise today.

Enterprise data is typically stored across many data silos – including highly distributed data sources at the edge, multiple on-premises datacenters and public cloud providers, application-specific datastores, and across different geographies, just to name a few.

One approach is to aggregate data into a small number, or ideally a single, data location. It requires new storage capacity. But adding capacity often creates new data silos, which in turn increases the complexity of managing data.

This approach also creates massive data gravity for large silos of data that may not be local to GPUs used for AI.

Not only does this approach carry significant CAPEX and OPEX, but it also presents environmental and sustainability ramifications as new capacity will require additional power and cooling.

For AI workloads, most enterprises do not need to invest in new capacity. They can use available

capacity on existing storage systems. However, existing storage systems often lack the features and performance needed.

Storage infrastructure is only one part to delivering AI anywhere. More importantly, it is necessary to focus on the data itself.

AI engines deliver the most relevant and accurate results when they can draw from as much enterprise data as possible at the highest performance possible. As noted earlier, enterprise data is highly distributed. Managing the data itself while maintaining the performance required for existing and AI applications is very complex.

If an enterprise chooses the strategy to aggregate its data to meet AI demands, it can undermine the performance needed to effectively feed the GPUs and AI agents used for AI workloads.

Latency can be introduced, which may impact customer response times, due to conflicts with other applications and metadata updates that occur along the data path.

Summary

- ❁ Enterprises seek to deploy AI anywhere, which encompasses both customer-facing and internal applications across all geographies.
- ❁ To maximize the return from AI, AI-engines and GPUs must be fed with massive amounts of data at the highest performance possible.
- ❁ However, capital and operational costs, as well as sustainability must equally be considered.
- ❁ Neuralytix believes the approach Hammerspace has designed to meet requirements, using its data platform technology, meets the demands needed to deliver AI anywhere.

Innovation is not organic

For AI, it's not just about raw performance in the traditional sense but also pipeline performance and operational efficiency. Delivering AI anywhere requires accessing data distributed at the edge, on-premises datacenters, and the cloud, feeding AI engines at the highest performance with minimal latency at every stage of the AI lifecycle. These challenges bring us full circle to what the fundamental requirements are.

When embarking on an AI anywhere strategy, how does an enterprise:

- ❁ Make as much data available from as many data sources as possible,
- ❁ Avoid performance impacts,
- ❁ Bolster business continuity,
- ❁ Minimize additional financial investments,
- ❁ Meet environmental and sustainability goals, and
- ❁ Maintain scalability?

Critically, how does an enterprise meet these needs without significant growth in OPEX due to added operational complexity and staff?

AI Anywhere Means Delivering Timely AI-Everything Everywhere

Today, upper management and boardrooms expect AI to be an accelerator for time-to-success. Success includes competitive advantage, operational efficiencies, and cost effectiveness.

Meeting these expectations demands the ability to deploy AI anywhere – ideally, to everything, everywhere. Accelerating time-to-success means delivering the benefits of AI online quickly.

But how does an enterprise deliver AI anywhere and everywhere rapidly while overcoming the challenges highlighted above?

Neuralytix favors a solution that allows data to stay *in situ* no matter where the data is located.

This approach overcomes many problems associated with delivery AI anywhere, including:

- ❁ Needless procurement of additional infrastructure, removing both CAPEX and OPEX,
- ❁ Environmental impact associated with duplicating infrastructure,
- ❁ Need to migrate data manually – removing the risk of data corruption or loss of context,
- ❁ Meet governance and regulatory compliance – keeping data where regulations dictate,
- ❁ Using real-time data – resulting in slow development, testing, and deployment of AI, and
- ❁ Impact on business application performance when data paths become over-saturated.

While an enterprise no longer needs additional infrastructure for replicated data, an incremental investment is still needed for use-case-specific resources to operate and store the output of AI solutions.

It is our position that solving the AI anywhere problem goes beyond a solution that can take advantage of data *in situ*. It must also have the following features:

- ❁ Ability to access data from edge, on-premises datacenters, and public cloud infrastructures,
- ❁ Non-disruptive integration into existing business applications,
- ❁ No agents or other client software which require extensive compatibility testing, and thus completely defeat the objective of accelerating time-to-success,
- ❁ Use existing and industry-standard storage protocols,
- ❁ Use automation to speed up data pipelines and ensure data is retained on the right tier to minimize cost and power consumption, and
- ❁ Solutions that are integrated with the other leading vendors in the AI ecosystem to accelerate deployment and reduce risk.

Hitting the Nail on the Head

Neuralytix research identifies Hammerspace, a data platform that unifies data across an enterprise, as a leading solution to address the requirements to meet the AI anywhere objective.

A key advantage of Hammerspace is its ability to locate, access, and manage data across all AI-relevant data and data sources spanning multi-vendor and multi-site on-premises storage to the public cloud – while keeping the data itself *in situ*.

Another key advantage of Hammerspace's platform is its use of industry standard data access and management protocols. It eliminates the need to install custom agents or client software on application servers. This allows Hammerspace to be integrated into an enterprise's existing infrastructure non-disruptively, while still providing parallel file system performance to accelerate AI applications.

Enterprises can continue to take advantage of the benefits and innovations of their existing storage infrastructure investments, simultaneously making the data AI-ready. This represents a significant cost advantage over solutions and approaches that require investments to duplicate or replicate existing storage infrastructure.

More importantly, the flexibility to unify access to data *in situ* across existing storage silos and orchestrate the data to compute resources anywhere enables rapid run test scenarios, proving the ROI of inferencing or agentic AI projects using cloud-based or GPU-as-a-Service resources before companies commit to capital investments in their own datacenters.

This automated data orchestration complements Hammerspace's ability to increase utilization and accelerate the performance of on-premises and cloud-based GPU clusters. It does so by aggregating local NVMe storage within servers to create an ultra-high-performance Tier 0 storage layer. Any underutilized local NVMe within the

storage pool further extends the cost advantage of Hammerspace.

To achieve this, Hammerspace creates a global namespace across the universe of data of an enterprise. At the same time, customers can also create independent namespaces for specific uses such as AI modeling.

In the context of AI, this can accelerate AI model design, testing, and deployments for AI teams to evaluate different AI engines and technologies simultaneously.

The combination of these benefits allows enterprises to quickly establish both applicability of AI engines, AI models, and return on investment (ROI).

Conclusion

Enterprises are deploying AI anywhere and everywhere. To serve the massive data and performance needs, AI requires dynamic cost-effective access to all its data without the limits of cost, latency, or operational complexity.

Neuralytix believes that the innovative solution and platform Hammerspace has developed makes it a leader in enabling the deployment of AI anywhere.

Neuralytix analysis shows that Hammerspace's industry standards-based platform extends beyond AI. Enterprises across all industries can use the same Hammerspace platform for AI and deploy it for other use cases.

About Neuralytix

Neuralytix® is a global industry analyst and consulting firm that provides industry-leading market analysis, innovative go-to-market strategies, and custom content solutions.

www.neuralytix.com

© Copyright 2025, Neuralytix Inc.

Neuralytix is a registered trademark. All rights reserved