


BROCHURE

AI Optimized Infrastructure Services

Building Robust
Foundation for
Enterprise AI at Scale







Introduction

AI is no longer a collection of use cases, it's a systemic shift in how enterprises operate, make decisions, and grow. Yet organizations find themselves at a critical inflection point: while AI capabilities are advancing rapidly, their ability to adopt and scale these innovations is constrained by underlying infrastructure, growing sovereignty regulations, Trust, Value of AI and Operational Readiness. This is consistent with results from the recent survey commissioned by LTM with Forrester Consulting.

Architecting Foundation Right for AI is not Optional ! Its a Strategic Imperative.

Business Problems

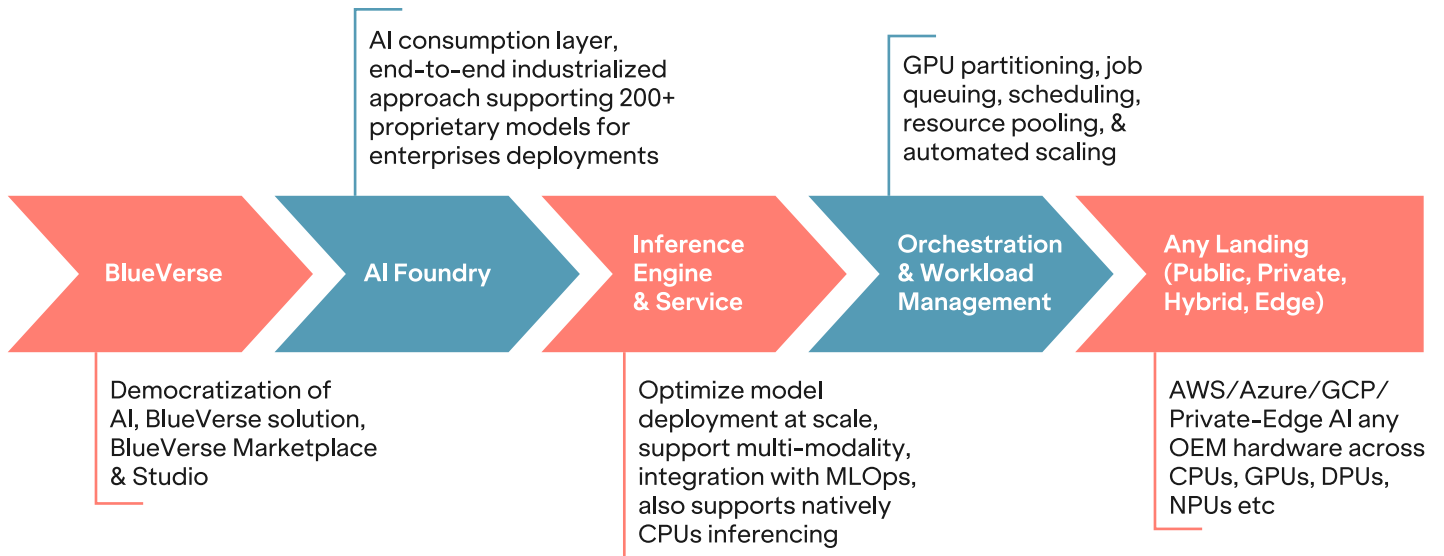
Enterprises face multiple barriers while scaling AI:

			
<p>Infrastructure volatility and cost inflation</p> <p>Hardware obsolescence, vendor lock-in, and expensive private AI options limit flexibility should change to Rapid Obsolescence, Vendor Lock-In and High Barrier for Private AI.</p>	<p>Divergent ecosystems of data and models</p> <p>Enterprise data systems built for transactions struggle to support distributed, high-volume, multimodal AI workloads should change to Operational fragmentation, Misaligned architectures and Poor Integration with legacy systems.</p>	<p>Compound AI infrastructure complexities</p> <p>GenAI and Agentic AI require multiple models, tools, and frameworks, creating integration and management challenges.</p>	<p>Return-on-AI and risk gaps</p> <p>Poor ROI models, compliance risks, and sub-optimal deployments drain millions in hidden costs. change to Unclear ROI Models, hidden risks and sub-optimal deployments drain millions in costs.</p>

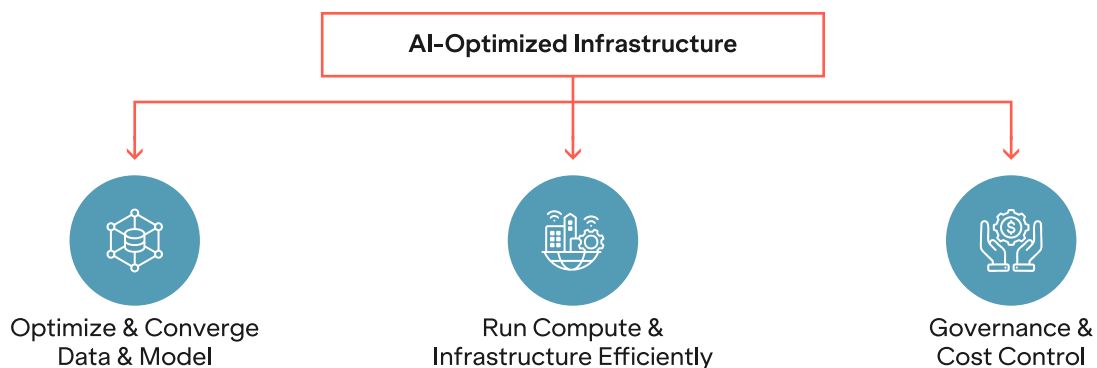
Our Solution

LTM's AI-Optimized Infrastructure Solution solves for business problems by going beyond Specialized Hardware, Network/I&O Efficiencies & storage systems, by providing a full stack solution which is Model, Modality, Hardware and/or Cloud Platform agnostic that enables to converge the two distinct worlds in today's enterprise architecture capability; traditional and AI workloads.

The solution is modular, composable and vendor-agnostic, delivering a resilient, reliable, cost- optimized, and unified platform for developers and business users to deploy, manage, and scale AI effortlessly.



Built on the belief that scaling AI is not just a model challenge, but also an infrastructure and economics challenge, the solution enables enterprises to unlock sustainable AI adoption by focusing on three levers:



Optimize & Converge Models & Data

- Smaller, efficient model architectures (quantization, pruning, fine-tuning)
- Smarter data strategies (active learning, preprocessing, caching)
- Bridging enterprise data with AI-native architectures (vector stores, multi-modality, distributed fabrics)

Run Compute & Infrastructure Efficiently

- Workload profiling and right-fit hardware mapping
- Orchestration, autoscaling & hybrid cloud strategies
- Advanced inference engines, dynamic GPU allocation, network I/O efficiency

Governance & Cost Control

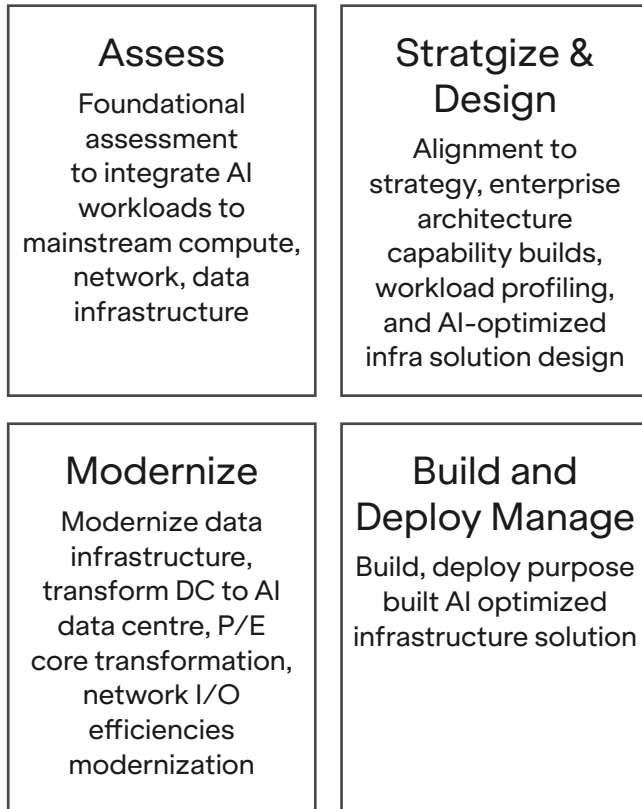
- Continuous observability & cost attribution
- Abstraction layers & hybrid deployment to reduce vendor lock-in
- Unified platform for accessible AI services (Model-as-a-Service, Token-as-a-Service) to mitigate ROI risks

Together, these levers and design principles of LTM's AI-optimized infrastructure solution empowers enterprises to scale AI economically while ensuring agility, resiliency, and reliability.



Our Approach to AI-optimized Infrastructure

LTM follows a structured, phased methodology that ensures AI infrastructure is functional as well as robust, scalable, and future-ready.



This approach ensures enterprises move from concept to consumption of AI with confidence, speed, and control.

Why LTM?

- | LTM AI-Optimized Infrastructure, any model in any hardware with any cloud (Public, Private, Edde or Hybrid) deployment archetypes, modular, composable and vendor-agnostic.
- | LTM AI Foundry supports 600+ hardware SKUs across Nvidia, Intel, AMD & Huawei and has pre-set deployment configurations for about 200+ model providers like OpenAI, Anthropic, Bedrock, Azure Foundry etc.
- | LTM BlueVerse an AI-native ecosystem that brings together intelligent agents, modular architecture, and built-in trust - providing everything enterprises need to move from AI concept-to-value, across today's systems and at scale tomorrow.
- | LTM curated best-fit ecosystem of partner, combining the agility and innovation of the startup landscape with the scale, stability, and proven performance of established products.
- | Commercial Dflexibility, LTM offers the AI-Optimized Intrastructure stack either in an as-a-service model or through OEM-specific deployment models, providing enterprises with choice and adaptability to align with their procurement and operational strategies.

Start your AI journey on infrastructure built for trust, scale, and speed.

Connect with us at cis.pulse@LTM.com.

LTM is a global technology services and consulting company and the business creativity partner to the world's largest and most disruptive companies. We bring human insights and intelligent systems together to help enterprises across industries rewire their business models, accelerate innovation, and drive AI-centric growth. With our integrated operations, transformation, and business AI services, we design and deliver solutions that create new productivity paradigms and new roads to value. Together with 87,000 employees across 40 countries and our global network of hyperscaler partners, LTM — A Larsen & Toubro company — owns business outcomes for over 700 clients, helping them to not simply outperform the market, but to Outcreate it.