


Delivering AI at Scale for Operators: Where Data Centers, Terrestrial Networks and Submarine Cables Converge

Harsha Ram

AI – THE 4th STAGE OF THE INDUSTRIAL EVOLUTION




STEAM

18th Century



Electricity

19th Century



Computing

20th Century



AI

21st Century

SOVEREIGN AI - WHY IT MATTERS

To protect India's data, accelerate AI innovation, ensure compliance, and build models that truly understand the Indian Context

**Protects
National Data &
IP**

**Ensures
Compliance &
Trust**

**Enhances Security
& Strategic
Autonomy**

**Optimized for
India's Language &
Context**

**Boosts Innovation
& Economic
Growth**

FROM CONTENT DELIVERY TO INTELLIGENCE ENABLEMENT

AI changes what users expect from the network



The last decade optimized the internet for video.
The next decade must optimize it for AI interaction.

Networks must evolve from delivering content to enabling intelligence.

AI SURGE IS REDEFINING INFRASTRUCTURE DEMAND

Accelerating Digital Transformation with High-Performance Compute and AI Workloads



AI Growth Trends

Global AI market

Projected Growth:
CAGR of **37.3%**, reaching **\$1.8 trillion**
by 2030.

India AI market

Projected Growth:
CAGR of **33.7%**, reaching approximately **\$6 billion**
by 2027.



Generative AI Workload Explosion

Large Language Models (LLMs) like OpenAI's GPT and Google's Gemini require **high-performance compute environments**.

Demand for hosting AI training and inference workloads is expanding across **globally**, driven by increased digital transformation initiatives.

THE NEW FOUNDATION OF INTELLIGENCE

AI doesn't start with models, it starts with the right infrastructure powering them

AI DATA CENTERS

GPU ready data centers
powering AI training & inference

High density racks with
liquid cooling

AI NETWORK

High Capacity, low latency networks
moving massive datasets in real-time

Long haul fiber & subsea
for global AI model distribution

TRADITIONAL VS AI ARCHITECTURE

	TRADITIONAL	AI-READY
Compute	• CPU centric, optimized for transactional workloads	• GPU/TPU centric, optimized for parallel processing
Storage	• Regional databases, structured data	• High-throughput, distributed storage for massive datasets
Networking	• Standard bandwidth, latency not critical	• Ultra low latency, high bandwidth (metro, long haul fiber)
Workloads	• Business applications, ERP, CRM	• AI model training/inferencing, large scale data analytics
Scaling	• Vertical scaling (bigger servers)	• Horizontal scaling (clusters of GPUs)
Cooling	• Air cooling sufficient	• Liquid cooling for high-density racks
Security	• Perimeter-based compliance for business data	• Data sovereignty, model governance, compliance for sensitive AI data

FOUNDATION OF FUTURE-READY AI NETWORKS



High Performance Core

- Multi terabit backbone
- High-capacity interconnects between AI compute clusters
- Carrier grade reliability with redundant routing paths
- Optimized for massive data throughput and GPU to GPU transfers



Low Latency Fabric

- Microsecond class transport for real-time AI workloads
- <1ms latency for AI model training & inferencing
- Metro & long-haul optimization for faster data movement
- Critical for autonomous, predictive, and real-time analytics



Secure by Design

- End to end encryption across all data in transit
- Zero trust segmentation to isolate workloads
- AI powered threat detection & mitigation
- Compliance ready for data sovereignty & governance



Unified DC-Cloud-Edge Ecosystem

- Seamless interconnect between on-prem, DC, cloud & edge nodes
- Direct cloud on-ramps to hyperscalers
- Edge AI enablement for localized, real-time processing
- Integrated with storage, compute, and GPU clusters



Dynamic & Orchestrated

- SDN enabled architecture for dynamic workload routing
- On-demand bandwidth for bursty AI training cycles
- Network Function Virtualization for service agility
- API driven orchestration for DevOps & AIOps integration

ARCHITECTURE IMPLICATIONS FOR PROVIDERS

AI changes traffic volume, traffic shape and the operating model

01 TRAFFIC GROWTH

AI becomes the next major driver after video.

02 TWO-WAY FLOWS

Prompts, voice, images and context increase upstream demand.

03 LATENCY MATTERS

Real-time AI needs consistent response, not bandwidth alone.

04 INFERENCE AT EDGE

Metro and edge compute move intelligence closer to users.

05 FIBER + SUBSEA

AI ecosystems need dense metro, long-haul and global capacity.

06 AI-NATIVE OPS

Prediction, optimization and remediation become automated.

PLAN FOR BANDWIDTH • ENGINEER FOR LATENCY • OPERATE WITH AI

INDIA AI INFRASTRUCTURE STORY

Fueling Regional Growth with Carrier-Rich, Low-Latency Network Backbone

Large AI campuses
in metros


Dense fiber networks across the
metro sites

Express long-haul routes
specifically designed for high
performance

10's of fiber pairs in
submarine cable systems
across multiple sites

Edge sites across high density
Tier 2 cities for inferencing
use cases

Secured Network Managed
Services with observability



**Tried.
Tested.
Trusted.**

Your digital infrastructure
partner for over 25 years



Network



Data Centers



Cloud



Cyber Security



Managed Services



CIO Choice 2025
for AI-Ready Data
Centers



Best Future-Proof
Data Center
Project



Award for
Sustainability



Innovation in
Data Center



India's 1st Nvidia
DGX Ready Liquid &
Air Cooled DCs



Best in
Services Industry



DC Services India
& Managed Cloud
Services APAC



Managed Network
Services - MQ

THANK YOU