cyfuture
Thinking Future. Moving Ahead.

Data Centre & GPU Capabilities

Enterprise infrastructure
for cloud and AI workloads



Cyfuture India Private Limited

October 2026

cyfuture.ai

Cyfuture at a Glance

25+ years of digital infrastructure, cloud, data centre & AI execution

cyfuture.ai

2001
Founded
roots to 1998

5,000+
Workforce

25+
Delivery centres
India + US/UK

4
Data centres

5 MW+
Existing Tier III capacity

1,100+
IndiaAI GPUs awarded

12+ Fortune 500 clients

Company Overview

Cyfuture is a next-generation digital infrastructure and IT solutions provider, blending cloud, AI and data services to empower enterprises, governments and startups. Headquartered in Noida SEZ, the company serves clients across India, the US and the UK.



Proven Operating History

- **2001** Cyfuture founded
- **2010s** Tier III DC expansion / MeitY empanelment
- **2024** 25+ delivery centres incl. US & UK
- **2025** Chennai data centre + IndiaAI Mission positioning

Infrastructure & Delivery

- 4 operating hyperscale, Tier III-certified facilities across India
- Noida, Jaipur, Bangalore and Chennai
- 1,000+ rack capacity across facilities
- 99.99% uptime SLA
- Enterprise-grade cooling, security and managed services

- **ISO & PCI DSS**
- **CMMI Level 5**
- **MeitY CSP**
- **IndiaAI Mission vendor**
- **L1 bidder across AMD, Intel and NVIDIA GPU categories**

Leadership

Anuj Bairathi Promoter / Chief Executive Officer	**Ravish Sharma** Director / Chief Operating Officer	**Munish Mahajan** Director / Chief Human Resource Officer
Ajai Rai Chief Technology Officer	**Rana Roy** Chief Business Officer	**Himanshu Gupta** Chief Financial Officer

Selected Marquee Clients

TATA · GAIL · NTPC · pwc · SBI · unicef · punjab national bank · IDBI BANK · Bharat Petroleum · DHFL · IndianOil · HSBC

Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

Credentials & Recognition

Certifications, awards, media coverage and key milestones

cyfuture.ai

Certifications



IndiaAI Mission Empanelment



CMMI Appraised



PCI DSS compliance V3.2



NSIC-CRISIL SEC 2B rating



VMware Cloud Verified



Member of NASSCOM

Awards



Rajasthan - State Award for Export Excellence



Entrepreneur of the Year 2022



Times 40 Under 40 2023



Best Organization for Women 2022



Great Place to Work



Entrepreneur India - Business Transformation Award 2023

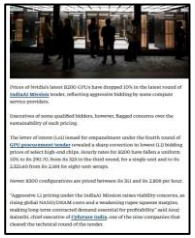
Recognitions















Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

Lines of Business

Connected services across cloud, AI, data centres and digital solutions.

cyfuture.ai

01 Enterprise Cloud

Secure, scalable cloud platforms for business workloads.

02 Enterprise AI Platform

AI-powered solutions that automate operations, enhance productivity, and accelerate business growth.

03 DC and DR Service

Data centre and disaster recovery services to support business continuity.

04 GPUaaS

On-demand GPU compute for AI training, inference and high-performance workloads.

05 Colocation Service

Secure, connected data centre space for critical IT systems.

06 AI Application Development

Custom AI applications that automate workflows, enhance productivity, and deliver intelligent business solutions.

07 Managed Support

Proactive IT support and monitoring that ensure seamless operations and business continuity.

One connected Cyfuture ecosystem: from secure data centres to cloud, GPU compute and AI-led digital solutions.

Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

Data Centres – Noida & Jaipur

Secure, scalable & reliable facilities supporting enterprise cloud, managed hosting, disaster recovery & business continuity

Noida Data Centre

Primary DC



NSEZ Phase II
Noida Special Economic Zone

35,000 sq. ft.
raised floor space

5 MW
power

ISO 22301
& PCI DSS certified

Services: colocation, managed hosting, cloud solutions, disaster recovery, high-speed connectivity

Infrastructure: dual power feeds, high-efficiency UPS, VESDA smoke and fire detection with suppression

Security: gated, monitored facility with stringent access controls

Sectors: e-commerce, healthcare, finance, education and government

Jaipur Data Centre

DR site



Sitapura IT Park
Jaipur

Tier III
certified facility

4.0 MW
power

ISO 22301
& PCI DSS certified

Resilience: dual power feeds, high-efficiency UPS, VESDA fire detection and suppression

Security: 24/7 surveillance and controlled access

Services: colocation, public, private and hybrid cloud, disaster recovery, managed IT

Sectors: mission-critical workloads across e-commerce, healthcare, finance, education and government

Upcoming AI Data Centre

Upcoming



Noida SEZ
upcoming AI campus

20 MW
liquid-cooled capacity

240 kW+
per-rack density

400G / 800G
networking ready

AI-first design: purpose-built for GPU thermal loads, not retrofitted from legacy colocation

Cooling: direct-to-chip liquid cooling with hybrid air / liquid zones

Next-gen ready: built for NVIDIA Blackwell and Vera Rubin class racks

Leasing: reserved blocks, anchor halls or whole-campus for AI labs, enterprises and sovereign programs

Cyfuture's AI and cloud expansion is built on top of operational, secure and service-rich data centre infrastructure in Noida and Jaipur, and extends to our upcoming liquid-cooled AI campus.

Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

Power & Physical Security – Noida

Redundant power and layered physical security for AI-ready colocation

cyfuture.ai

Power Infrastructure

Reliable | Redundant | High availability | AI-ready

1. **Dual Grid Feeds** — Feed A + Feed B from independent grid sources
2. **Substation & Transformers** — 11/33 kV incoming, stepped to 415 V
3. **MV/LV Switchgear** — Main supply, A & B paths
4. **UPS System** — N+N redundancy, A & B
5. **PDUs** — A-feed and B-feed distribution
6. **AI Compute Racks** — Dual-corded, fed from A + B

A Feed (Primary) + **B Feed (Secondary)**, independent end to end



Diesel Generators

N+1 redundancy; auto-start via ATS during grid outage.



DCIM / BMS Control

Real-time monitoring, load management and capacity planning.



Building Infrastructure

HVAC, fire detection & suppression, environmental monitoring.

Key Features

- **A+B redundant supply** — Dual grid feeds with N+1 infrastructure
- **N+1 resilient design** — High availability and fault tolerance
- **High power density** — Built for AI and HPC workloads
- **Real-time DCIM monitoring** — Power, energy and infra visibility
- **Scalable for expansion** — Flexible design to support growth
- **Energy efficient** — Optimised usage, low carbon footprint

Physical Security

People | Process | Technology

1. **Perimeter Security** — Fencing, guards & patrols, 24×7 surveillance, ANPR / barcode vehicle control



2. **Building Entry** — Visitor management, photo-ID check, access passes, turnstile / mantrap



3. **Access Control** — Multi-factor (card + biometric), role-based access, logs & audit trail



4. **Data Hall Security** — Locked racks / cages, in-hall CCTV, per-zone access, environmental monitoring



5. **Monitoring & Response** — 24×7 security operations, incident management, real-time alerts



Key Benefits

- Enhanced asset protection
- Controlled access
- Audit-ready environment
- Compliance support (ISO 27001, SOC)
- Business continuity
- Deterrence against threats

Redundant power and layered 24×7 security keep mission-critical and AI workloads protected at Noida.

Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

Compute Capabilities

Enterprise, cloud and AI compute, including NVIDIA B300 and RTX PRO platforms

cyfuture.ai

CPU Compute

Scalable x86 compute infrastructure supporting general-purpose, enterprise and mission-critical workloads.



- **Intel & AMD processors**
  Xeon Scalable | AMD EPYC
- **High core count**
  Up to 128 cores per CPU
- **Large memory capacity**
  DDR5, up to 4 TB per server
- **Flexible storage**
  Local NVMe / SSD / HDD

GPU Compute

GPU infrastructure for AI/ML training and inference, analytics, simulation and other accelerated workloads.



GPU Platforms

NVIDIA B300 | RTX PRO

- **High-speed interconnect**
  NVLink, PCIe Gen5, 400G/800G ready
- **Scalable configurations**
  Single node to multi-node clusters
- **Optimized for AI/ML, HPC**
  Training, inferencing, analytics

Virtualized Compute

Enterprise hypervisor platforms enabling secure, flexible and scalable virtual machines with resource isolation and multi-tenancy.



- **Multiple hypervisors**
  VMware vSphere | VCD, KVM, Hyper-V
- **Multi-tenant & isolated**
  Secure, isolated environments
- **VMs at scale**
  Compute, storage and network
- **High availability**
  Live migration support

Cloud-Native Compute

Containerized and Kubernetes-based environments for modern applications, microservices and DevOps workloads.



- **Kubernetes (TKG)**
  Enterprise container platform
- **Container as a Service**
  Deploy and manage containers
- **Microservices & DevOps**
  Supports CI/CD pipelines
- **Auto-scaling & HA**
  Scale with demand

HPC & Distributed

High-performance, clustered compute designed for parallel processing, scientific workloads and large-scale distributed applications.



- **High-speed interconnects**
  RDMA | RoCE | InfiniBand 400G/800G
- **Parallel workloads**
  Optimized for distributed compute
- **HPC frameworks**
  MPI, Slurm, etc.
- **Scalable clusters**
  From few nodes to 100s of nodes

Bare Metal Compute

Dedicated physical servers for workloads requiring direct hardware access, predictable performance and resource isolation.



- **Rack servers**
  Standard enterprise servers
- **Blade servers**
  High-density compute
- **Edge compute**
  Distributed and remote locations
- **Custom configurations**
  As per workload requirements

Flexible Compute Models: Dedicated, virtualized, clustered and on-demand compute configurations aligned to workload requirements.

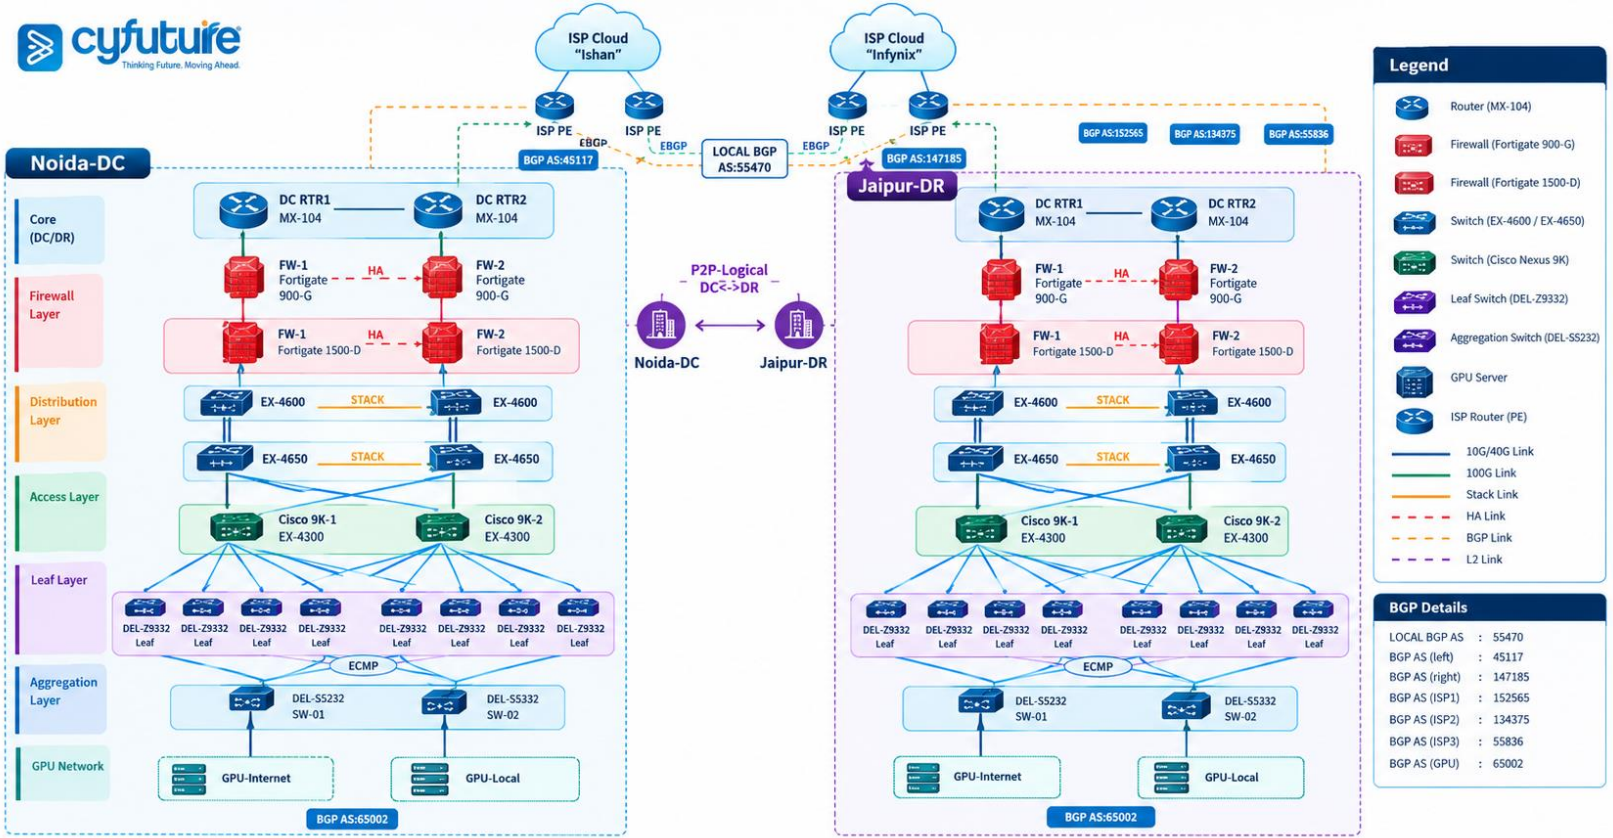


High Availability & DR: Compute across Noida Primary DC and Jaipur DR, supporting business continuity and resilient workload deployment.

Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

Network Architecture – Noida & Jaipur

Multi-vendor, redundant network design spanning Noida DC and Jaipur DR



Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

Enterprise & AI Storage Capabilities

Scalable, high-performance and secure storage for enterprise applications, cloud platforms and AI/ML workloads

cyfuture.ai

Storage Partners

Dell Technologies | HPE | NetApp | IBM | HITACHI | PURESTORAGE | VAST | WEKA | BeeGFS | MINIO | ceph | DataCore | COMMVAULT



Block Storage (SAN)

- High-performance block storage
- Low latency for mission-critical applications
- Supports databases, virtualization and enterprise workloads



File Storage (NAS)

- Shared file services for enterprise
- Ideal for collaborative and content-based workloads
- Scalable and highly available file storage



Object Storage

- S3-compatible object storage
- Massive scale for unstructured data
- Suitable for backup, archive and AI/ML datasets
- High durability and cost-efficient storage



Software-Defined Storage (SDS)

- Flexible and distributed storage architecture
- Supports virtualized and cloud environments
- Scalable and resilient design
- Unified storage across data centers



High-Performance Storage

- NVMe/SSD-based storage
- High throughput and low latency
- Designed for AI/ML, HPC and GPU workloads
- Supports large datasets, checkpoints and model training



Distributed Storage

- Horizontally scalable architecture
- High availability and resilience
- Supports large-scale and performance-intensive workloads
- Suitable for AI/ML and data-intensive applications



Backup & Archive Storage

- Data protection and retention
- Long-term archival storage
- Disaster recovery support
- Policy-based data lifecycle management



AI/ML Data Storage

- High-throughput storage for GPU clusters
- Optimized for large datasets and model repositories
- Supports distributed training and inferencing
- Scales with AI infrastructure growth

Scalable Architecture
Scale capacity and performance as per requirement

High Availability
Redundant and resilient design

Multi-Protocol Support
Block, File, Object

Cloud & DC Ready
Supports hybrid and multi-cloud environments

Flexible Deployment
On-premise, cloud or hybrid

Secure & Compliant
Enterprise-grade security and data protection

Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

SOC, NOC & Managed Services

24×7 security, infrastructure monitoring and managed operations across Noida and Jaipur

cyfuture.ai

Security Operations Centre (SOC)

24x7	96%	<2%
dual-site SOC	detection precision	false positives

- **Dual-site SOC**
  Noida and Jaipur, continuous security operations
- **Real-time telemetry**
  Elastic SIEM, 3 TB+ of logs daily
- **MITRE ATT&CK playbooks**
  Aligned detection playbooks for a stronger posture
- **Threat intelligence**
  MISP, Anomali, CERT-In IOCs, FS-ISAC feeds
- **Proactive hunting**
  Weekly hunt sprints and Purple Team drills
- **ML-driven detection**
  Custom models flag anomalies in real time

Network & Infrastructure Operations (NOC)

24x7	1 hr	99.982%
monitoring	outage response	DC uptime assurance

- **Infrastructure monitoring**
  Hardware, network, traffic, cloud and OS
- **Application monitoring**
  Application performance and URL monitoring
- **Layered SOC visibility**
  Physical-to-OS layer and application SOC
- **24x7 oversight**
  Continuous cloud infrastructure monitoring
- **Performance monitoring**
  Uptime, troubleshooting and optimization
- **Rapid response**
  1-hour response for service or system outages

Managed Services & Support

250+	100+	MeitY
cloud engineers	certified engineers	empanelled CSP

- **Managed SOC service**
  With managed migration, DLP endpoint and data encryption
- **24x7 support**
  Email, chat and phone for guidance and incident reporting
- **Instance & patch management**
  Provisioning, changes, updates and remote OS patching
- **Identity & access**
  IAM with role-based access control, plus MFA
- **Security oversight**
  Proactive monitoring and vulnerability assessment
- **Reporting**
  Regular MIS reports in agreed formats

Incident Response & BCDR

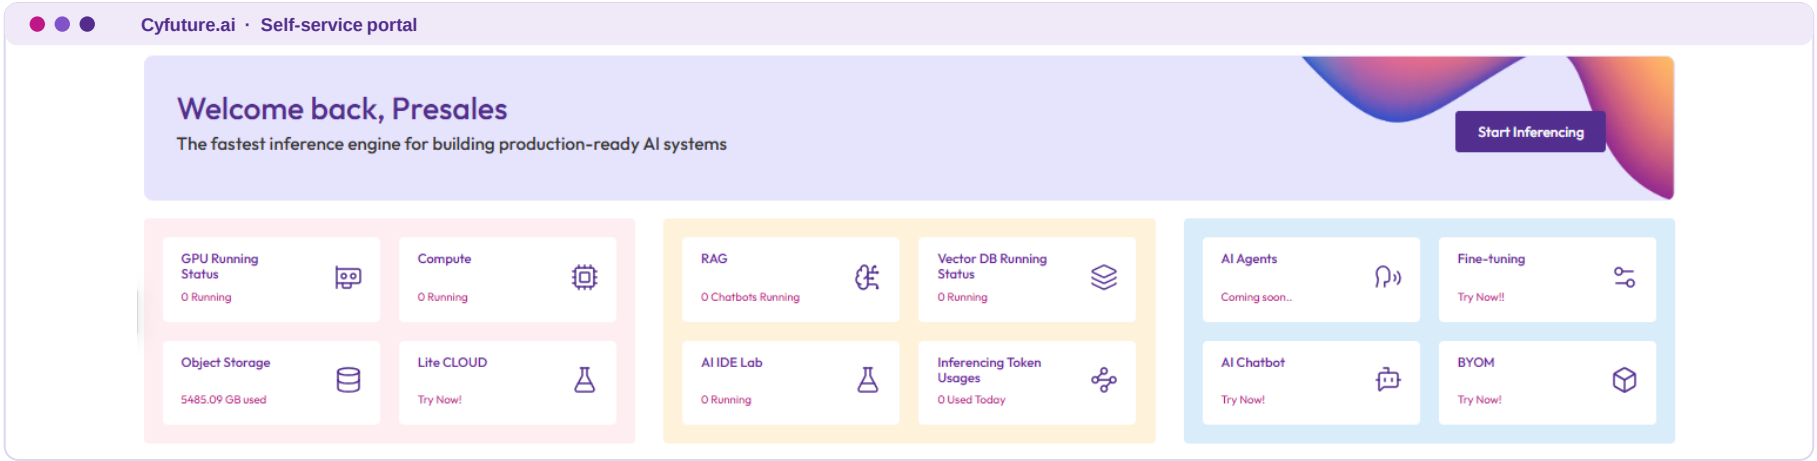
- **CERT-In compliant IR plan**
  1-hour incident notification window
- **3-2-1 geo-redundant backups**
  Verified daily for integrity
- **Isolated recovery**
  Clean-room restore after incidents
- **Quarterly BCP drills**
  RTO < 60 min, RPO < 15 min

Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

Self-Provisioning Portal

Create, manage and scale compute, storage, network and cloud resources from a single portal

cyfuture.ai



Self-Service Provisioning

Create and deploy compute, storage, network and cloud resources without manual intervention.

On-Demand Resource Allocation

Provision resources based on workload requirements with configurable capacity and specifications.

Automated Workflows

Streamlined provisioning, configuration and lifecycle management through automated workflows.

Centralized Management

Single portal to manage infrastructure resources, services, configurations and usage.

Real-Time Visibility

Monitor resource utilization, availability, performance and service status.

Role-Based Access

Controlled access based on user roles, projects and organizational requirements.

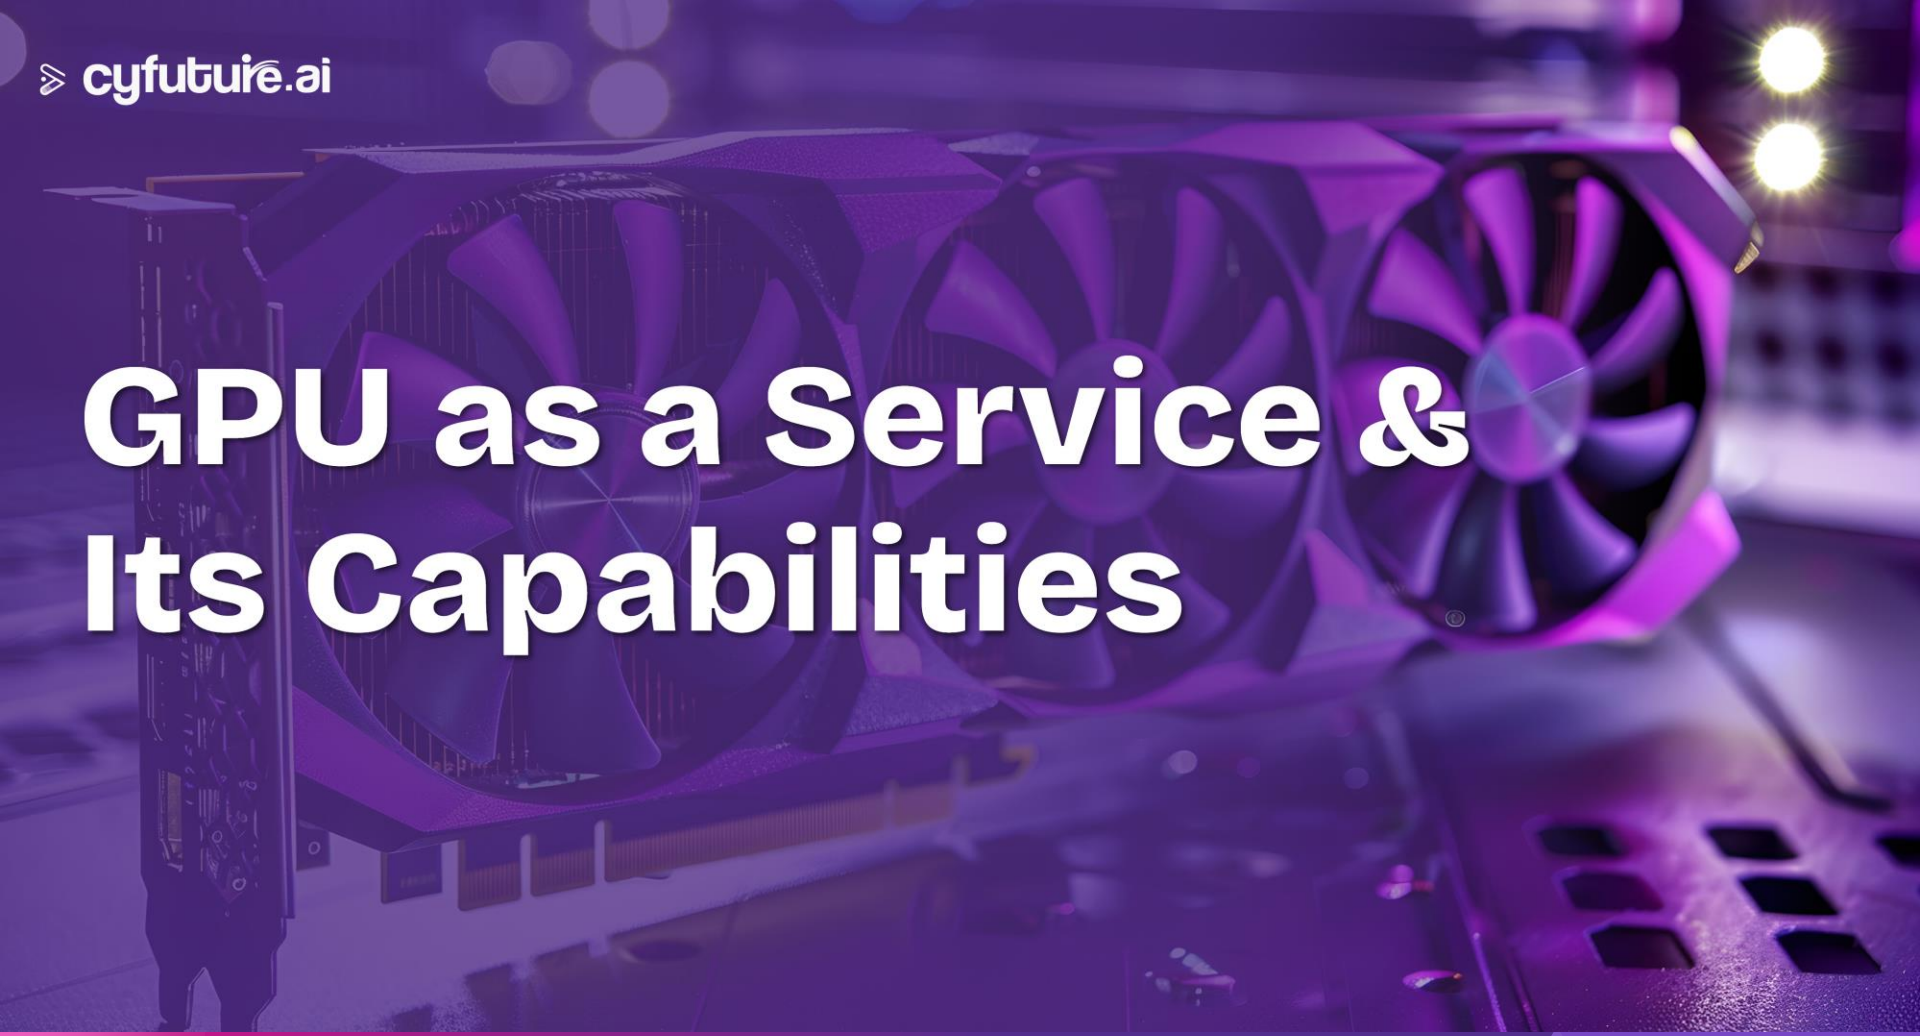
Usage & Cost Visibility

Track resource consumption and associated usage for better capacity and cost management.

Scalable & Flexible

Supports rapid expansion, modification and decommissioning of resources as business requirements change.

Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.



Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

GPUaaS – Bare Metal & Virtualized

Two ways to consume NVIDIA B300 capacity: dedicated bare-metal nodes or virtualized GPU access

cyfuture.ai

BaaS Bare Metal as a Service



Dedicated physical HGX B300 nodes with direct access to the GPUs, NVLink and network fabric, provisioned for the agreed contract period.

- **Full node:** 8 × B300, 2.3 TB GPU memory, 2 TB system memory, 30.72 TB raw NVMe
- **Direct fabric access:** ConnectX-8 NICs, RoCEv2 with GPUDirect RDMA
- **Dedicated capacity:** isolated resources and predictable performance
- **Starts small, scales out:** from a single B300 to 128 nodes (1,024 GPUs)
- **Validated handover:** GPU health, topology, network and RDMA checks

VM Virtualized GPU





Shared, multi-tenant access to GPU capacity through hypervisors and Kubernetes, with resource isolation between tenants.

- **Single-GPU entry:** start with one B300 and expand as workloads grow
- **Multi-tenant isolation:** enterprise hypervisors with resource isolation
- **Cloud-native scheduling:** Kubernetes or Slurm on dedicated control-plane nodes
- **Platform services:** GPU-as-a-Service, inferencing, fine-tuning, RAG
- **Packaging:** per-GPU memory and connectivity set in the final OEM design

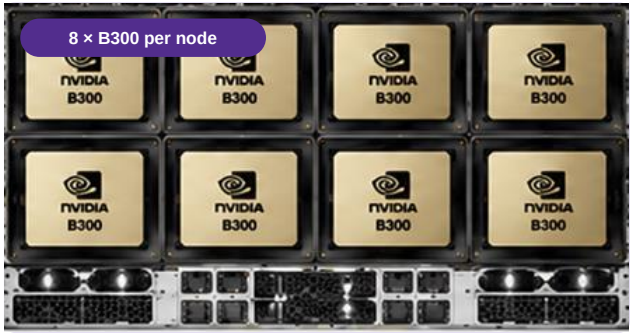
Cyfuture scope for either model

1 Consulting	2 Solution design	3 Deployment	4 Validation	5 Operations	6 Expansion
Workload discovery and architecture recommendation	B300 platform, connectivity and deployment architecture	Installation and configuration	GPU, network, RDMA and infrastructure checks	Monitoring, incident response, lifecycle and change management	Repeatable design for multi-GPU and cluster growth

Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

NVIDIA B300 – Platform Specifications

Blackwell Ultra accelerators on the HGX B300 platform: reference server specification



NVIDIA B300
Blackwell Ultra GPU architecture, 288 GB HBM3e per GPU

What B300 brings

- Blackwell Ultra GPU architecture for demanding AI workloads
- Large HBM3e memory for memory-intensive models
- High-bandwidth GPU-to-GPU communication through HGX B300
- High-speed networking options designed for distributed AI
- Accelerated data movement and collective communication

8 × B300	2.3 TB	2 TB	30.72 TB	16 × 400G
Blackwell Ultra SXM GPUs	total GPU memory	DDR5 system memory	raw local NVMe	RoCEv2 per node

Specification	HGX B300 reference node
Server platform	Supermicro HGX B300 / Pegatron AS801-2T1-8H2 (qualified OEM host)
GPUs	8 × NVIDIA Blackwell Ultra B300 SXM on the HGX B300 baseboard
GPU interconnect	5th-generation NVLink / NVLink Switch for GPU-to-GPU bandwidth
Processors	2 × Intel Xeon 6760P, 64 cores each (128 cores per server)
System memory	2,048 GB ECC DDR5 RDIMM (32 × 64 GB)
Local NVMe	4 × 7.68 TB U.2 (30.72 TB raw) plus 2 × 960 GB boot
GPU fabric	8 × ConnectX-8 SuperNIC, 2 × 400GbE RoCEv2 each, GPUDirect RDMA
DPU and services	NVIDIA BlueField-3 DPU, 2 × 200GbE service connectivity
Power and form factor	12 × 3000 W hot-swap PSUs (6+6); 8U rackmount, air-cooled

Where customers benefit

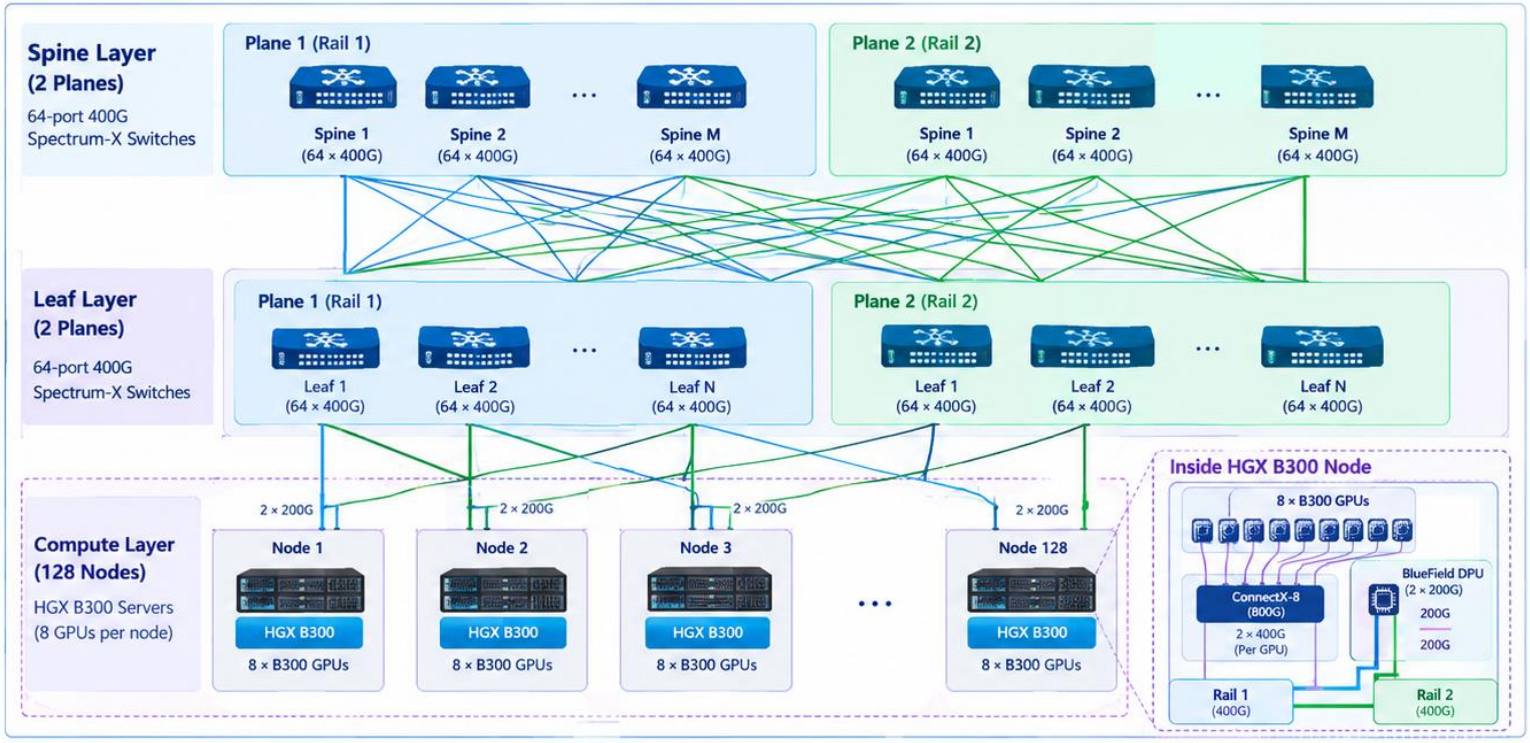
- **Generative AI** — LLM serving and real-time workloads
- **Training** — Fine-tuning and experimentation
- **Reasoning** — AI agents and automation
- **Simulation** — Compute-intensive workloads
- **Enterprise AI** — Analytics and data-driven work

Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

GPUaaS – Networking: Dual-Plane Leaf & Spine

Dual-plane, rail-optimized, 1:1 non-blocking 400G Ethernet fabric for 128 HGX B300 nodes (1,024 GPUs)

cyfuture.ai
FUTURE. BUILT TOGETHER



Reference Design

- **128 HGX B300 Nodes**
  1,024 B300 GPUs
- **32 Scalable Units**
  (4 Nodes per SU)
- **Dual-plane, Rail-optimized**
  Non-blocking Fabric
- **400G Ethernet (Spectrum-X)**
  64-port Switches
- **8 × B300 GPUs per Node**
- **Dual Port 200G BlueField DPU**
  (Per Node)
- **ConnectX-8 (800G) -> 2 × 400G**
  (Per GPU)

Network Segments

- Plane 1 (Rail 1)
- Plane 2 (Rail 2)
- Node / Server Link
- Management / OOB

Key Network Functions & Services

High Performance GPU East/West Traffic
2 × 400G per GPU
(800G total)

CPU / North-South Traffic
Converged network for management, internet and customer access

Storage Network
Dedicated storage connectivity with high throughput

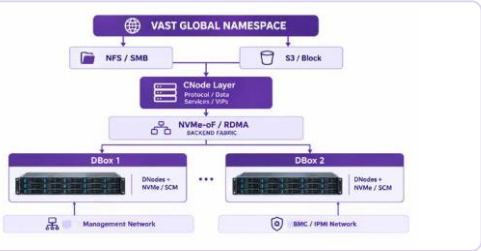
OOB Management
Physical/isolated network for management and monitoring

AI Storage – VAST, WEKA & BeeGFS

Parallel, scale-out storage platforms that keep GPU clusters supplied with data

cyfuture.ai

VAST Data

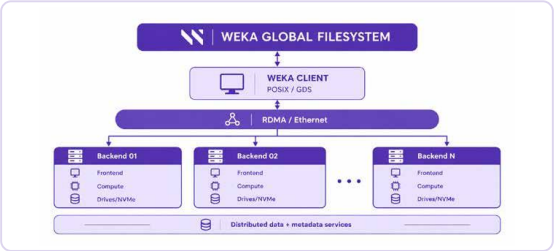


Architecture

Disaggregated Shared Everything (DASE): stateless CNodes run storage services while DBoxes hold the flash media, so compute and capacity scale independently.

- **Access:** global namespace over NFS, SMB and S3, with block access
- **Data path:** VIP pools spread client traffic and fail over transparently
- **Network:** separate frontend and NVMe-oF/RDMA backend fabrics

WEKA

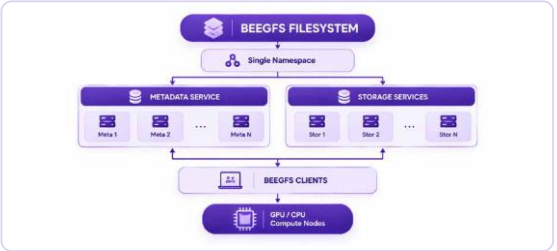


Architecture

Software-defined distributed file system on standard x86 servers with NVMe, with no centralized controller. Storage servers contribute drive, compute and frontend functions.

- **Access:** POSIX client, with RDMA and GPUDirect Storage in supported setups
- **Network:** Ethernet or InfiniBand, with multiple NICs for throughput and HA
- **Tuning:** NUMA-aware NIC assignment on GPU servers

BeeGFS



Architecture

Parallel cluster file system with POSIX semantics. Metadata servers are separate from storage servers, giving a single namespace across the cluster.

- **Striping:** files spread across storage targets so many clients share the bandwidth
- **Network:** InfiniBand or RoCE with multi-rail RDMA, or TCP/IP
- **GPU path:** GPUDirect Storage supported for GPU-oriented deployments

Best fit

- **VAST Data:** Large unified multi-protocol AI data platform
- **WEKA:** High-performance distributed AI/HPC file services
- **BeeGFS:** Flexible parallel file system for AI/HPC and research

Common design: independent 100G–800G storage fabric, redundant leaf paths and dual or multi-NIC clients, separate management / OOB plane, and sizing against aggregate GPU demand and checkpoint write bandwidth.

Copyright © 2026 Cyfuture India Pvt ltd. All Rights Reserved.

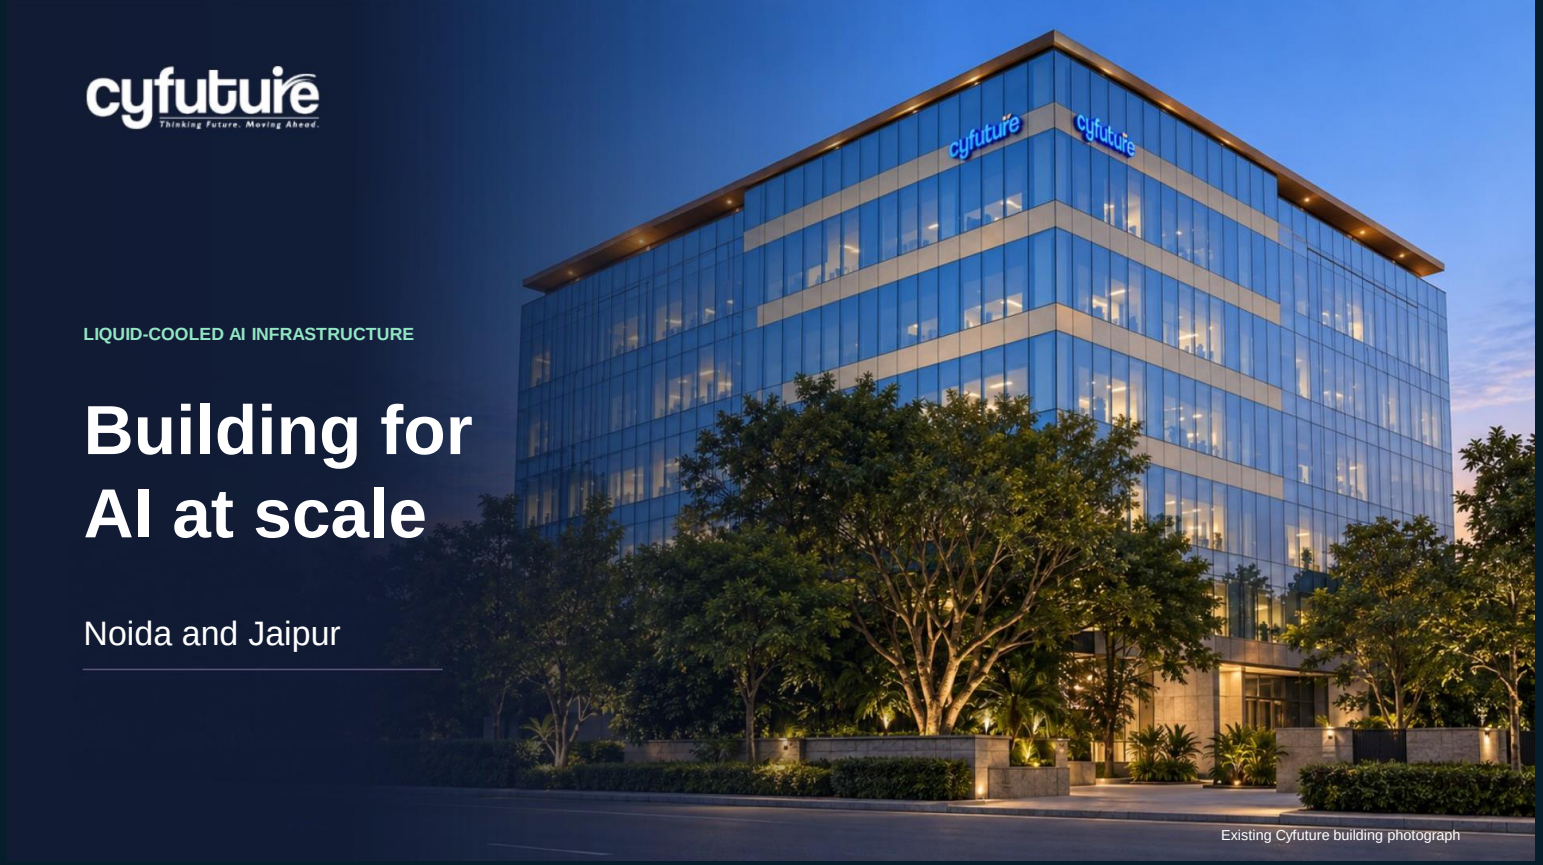
cyfuture

Thinking Future. Moving Ahead.

LIQUID-COOLED AI INFRASTRUCTURE

Building for AI at scale

Noida and Jaipur



Existing Cyfuture building photograph

20 MW

NOIDA

80 MW

JAIPUR

PLANNED DEVELOPMENT / OCTOBER 2026

01 / DEVELOPMENT OVERVIEW

Two locations. 100 MW planned.

A focused expansion for customers deploying dense AI infrastructure.

20 MW

NOIDA

80 MW

JAIPUR

Liquid cooling at the core

Direct-to-chip cooling with isolated facility and technology loops.

Flexible capacity for customers

Dedicated rack zones, private suites or larger capacity blocks, subject to final design.

20 MW Noida + 80 MW Jaipur is planned IT capacity. Cooling and electrical overheads require additional facility power.

02 / NOIDA FACILITY

20 MW liquid-cooled development

Noida, Delhi NCR

20

MW planned IT load

Dedicated capacity for AI training, inference and enterprise compute.

Power infrastructure

Multiple 33 kV dual transmission feeds in the proposed Noida design basis.

Liquid-cooled deployment areas

Platform-specific rack power, coolant interfaces and residual air cooling.

Customer deployment options

Dedicated zones and private suites, with connectivity and fit-out defined to scope.

Utility and facility specifications follow the brochure design basis and require confirmation for the selected 20 MW block.

03 / JAIPUR FACILITY

80 MW liquid-cooled development

Jaipur, Rajasthan

80

MW planned IT load

A second regional location for AI capacity and application resilience.

Transmission-level utility basis

132 kV connection in the proposed Jaipur design basis, with approved downstream distribution.

Cooling for local conditions

Heat rejection, water strategy and equipment derating sized for the actual site climate.

Regional deployment choice

Jaipur and Noida can support a multi-region strategy with application-specific recovery design.

Utility and facility specifications follow the brochure design basis and require confirmation for the selected 80 MW development.

04 / THE AI HALL

Designed around the AI rack



High-density power

Controlled liquid interfaces

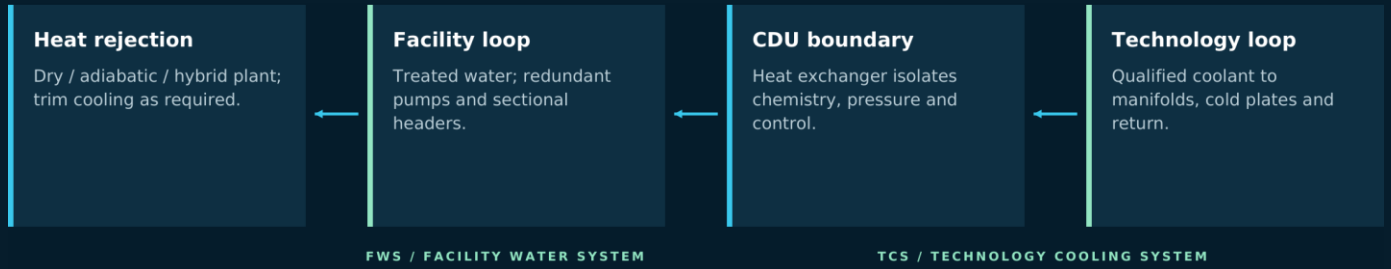
Service access

Illustrative concept visual from the supplied brochure. Equipment, layout and service clearances follow the final design.

05 / LIQUID-COOLING ARCHITECTURE

A controlled thermal path

Direct-to-chip cooling with separated facility and technology loops.



N+1 capacity target

Pump and CDU capacity by load block, with isolation points for maintenance and recovery.

Continuous protection

Leak sensing, flow and pressure monitoring, coolant sampling and contamination control.

Schematic from the brochure. Arrows show heat transfer towards ambient, not coolant flow. Final topology requires detailed design.

06 / RACK AND THERMAL PLANNING

Capacity tailored to the platform

Selected rack positions support different power and cooling envelopes.

DEPLOYMENT FAMILY	PROPOSED RACK ENVELOPE	INTEGRATION BASIS
B300 systems	60–70 kW	Qualified server configuration and heat split
AMD MI355X systems	80–150 kW	Server and rack configuration controls load
GB300 rack-scale systems	Up to 200 kW	Selected positions with approved liquid interface
Future rack-scale systems	Up to 300 kW	Enabling envelope at selected positions

25–45°C

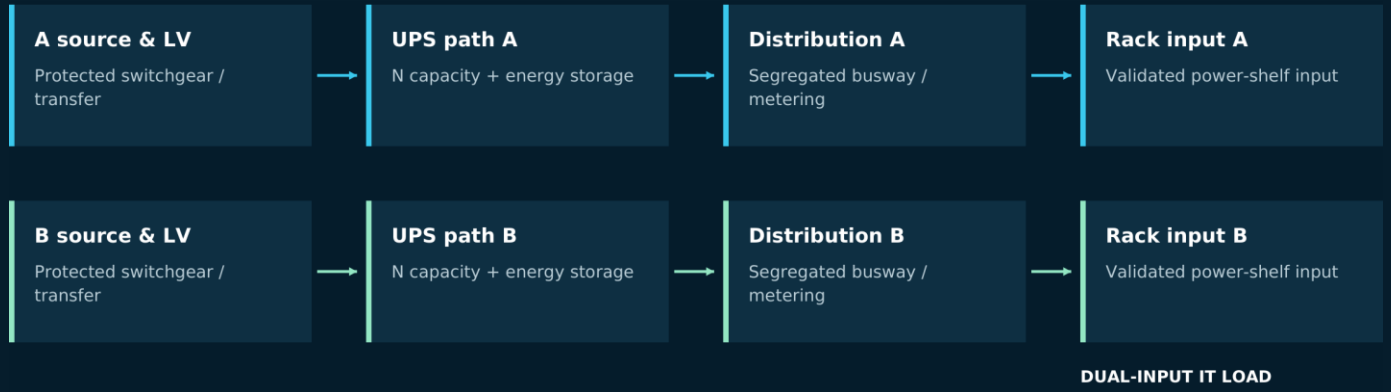
Indicative technology-loop inlet range only where the selected platform permits.

Proposed planning envelopes, not guaranteed platform consumption or a uniform entitlement at every rack.

07 / ELECTRICAL RESILIENCE

N+N UPS with segregated A/B paths

Each proposed UPS path supports the committed critical load.



N+1

Standby generation within each defined block

≥5 min

UPS autonomy target at end-of-life design load

24 h

On-site fuel autonomy target

Conceptual design basis. Component redundancy does not establish an end-to-end availability guarantee or Tier certification.

08 / BUILDING AND DEPLOYMENT INFRASTRUCTURE

Engineered for heavy AI equipment

Prefabricated structure and modular MEP infrastructure, with site-specific structural design.

6 m

Floor-to-floor planning basis

Overhead space for segregated power, network and liquid services. Final clear height follows coordinated design.

3,000 kg/m²

AI hall floor UDL target

4,000 kg

Wet rack planning envelope

6,000 kg

Heavy delivery route / lift target

All values are proposed engineering targets. Structural approval includes wet rack mass, local reactions and movement loads.

09 / CONNECTIVITY AND DATA

Infrastructure that keeps AI supplied

Compute, data services and management follow a defined integration boundary.

01	**Compute fabric**	400/800 Gb/s-class links, with topology and rail mapping matched to the selected platform.
02	**Storage and checkpoints**	Size capacity, throughput and metadata performance around real datasets and checkpoint behaviour.
03	**Campus connectivity**	Carrier-neutral approach, proposed diverse entries and meet-me rooms, with scoped cloud and WAN links.
04	**Protected management**	Separate administrative access and out-of-band control from tenant and operational technology networks.

Network, storage and GPU systems are customer-selected or separately scoped. Carrier availability and routes require validation.

10 / SECURITY AND OPERATIONS

Protection across the facility lifecycle

Proposed 24×7 facility monitoring and service coordination.

Controlled physical access

Perimeter controls, monitored entries, role-based access, visitor escorting and secure logistics.

Tenant and OT separation

Dedicated areas, auditable access and work orders, with protected building and cooling controls.

Facility command

DCIM, BMS and power monitoring, with alarm handling, capacity visibility and escalation procedures.

Disciplined maintenance

Documented SOPs, MOPs and EOPs, trained liquid operations and controlled customer-authorised work.

Site certification scope, response windows and service levels are evidenced and agreed for the specific facility.

11 / RESOURCE PERFORMANCE

Efficiency with measurable boundaries

Energy and water performance follows the actual site, load and thermal design.

PUE

Energy efficiency

Meter IT load, electrical losses and cooling separately.

WUE

Water accountability

Track site water use and compare heat-rejection options.

CO_2

Carbon visibility

Report electricity use and relevant emissions data.

Warm-water design

Higher permitted coolant temperatures can reduce compressor dependence. Site climate models establish the operating benefit.

No numerical PUE, WUE, renewable share or carbon-neutral commitment is represented in this presentation.

12 / CUSTOMER ENGAGEMENT

Capacity that fits the deployment

Lease a dedicated capacity block or discuss the full Noida or Jaipur facility scope.

ENGAGEMENT	TYPICAL SCOPE	AGREE AT DESIGN FREEZE
Dedicated capacity block	Rack zones, pods or halls with committed IT kW/MW	Power and cooling envelope, growth rights
Private AI suite	Segregated tenant area with tailored access and distribution	Security boundary, fit-out and service ownership
Full-facility engagement	20 MW Noida / 80 MW Jaipur, subject to final availability and agreement	Phasing, utilities, plant and handover boundary
Optional managed services	Remote hands, monitoring and scoped network / storage support	Hardware, licensing, support and change control

Capacity, energy, cooling, connectivity, fit-out and optional services are defined separately in the customer proposal.

13 / CUSTOMER DELIVERY

A defined path to operational handover

A commissioning and acceptance programme for each customer deployment.

01	**Requirements and interfaces**	Platform, MW, rack density, coolant and electrical boundary
02	**Design and approvals**	Coordinated drawings, capacity schedules and site acceptance
03	**Build and system tests**	Factory checks, installation, flushing and protection testing
04	**Integrated commissioning**	Utility loss, power transfer, cooling failures and thermal load
05	**Customer acceptance**	Workload validation, training, runbooks and signed handover

Site-specific delivery dates follow the approved project schedule, utility energisation, approvals and commissioning.

14 / SPECIFICATION SUMMARY

Noida and Jaipur at a glance

A common engineering approach, with site-specific utility and integration design.

PARAMETER	NOIDA	JAIPUR
Planned IT capacity	20 MW	80 MW
Utility design basis	Multiple 33 kV dual feeds	132 kV connection
Electrical resilience	N+N UPS; N+1 generation	N+N UPS; N+1 generation
Liquid cooling	Isolated FWS/TCS via CDUs	Isolated FWS/TCS via CDUs
Cooling redundancy target	N+1 by defined load block	N+1 by defined load block
Building height basis	6 m floor-to-floor	6 m floor-to-floor
Service readiness	Per approved project schedule	Per approved project schedule

Planning basis drawn from the supplied brochure. Targets require design approval and customer-specific acceptance.

CYFUTURE

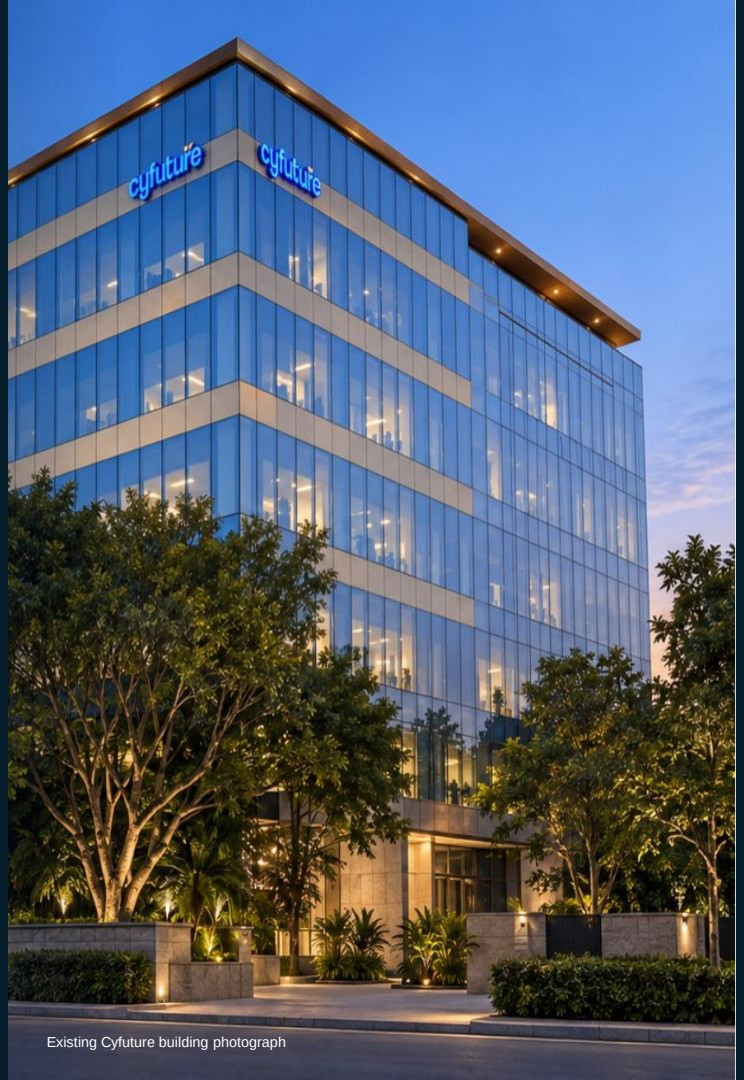
Build your next AI deployment with Cyfuture

20 MW NOIDA
80 MW JAIPUR

100 MW combined planned IT capacity

Technical and commercial engagement

cyfuture.com

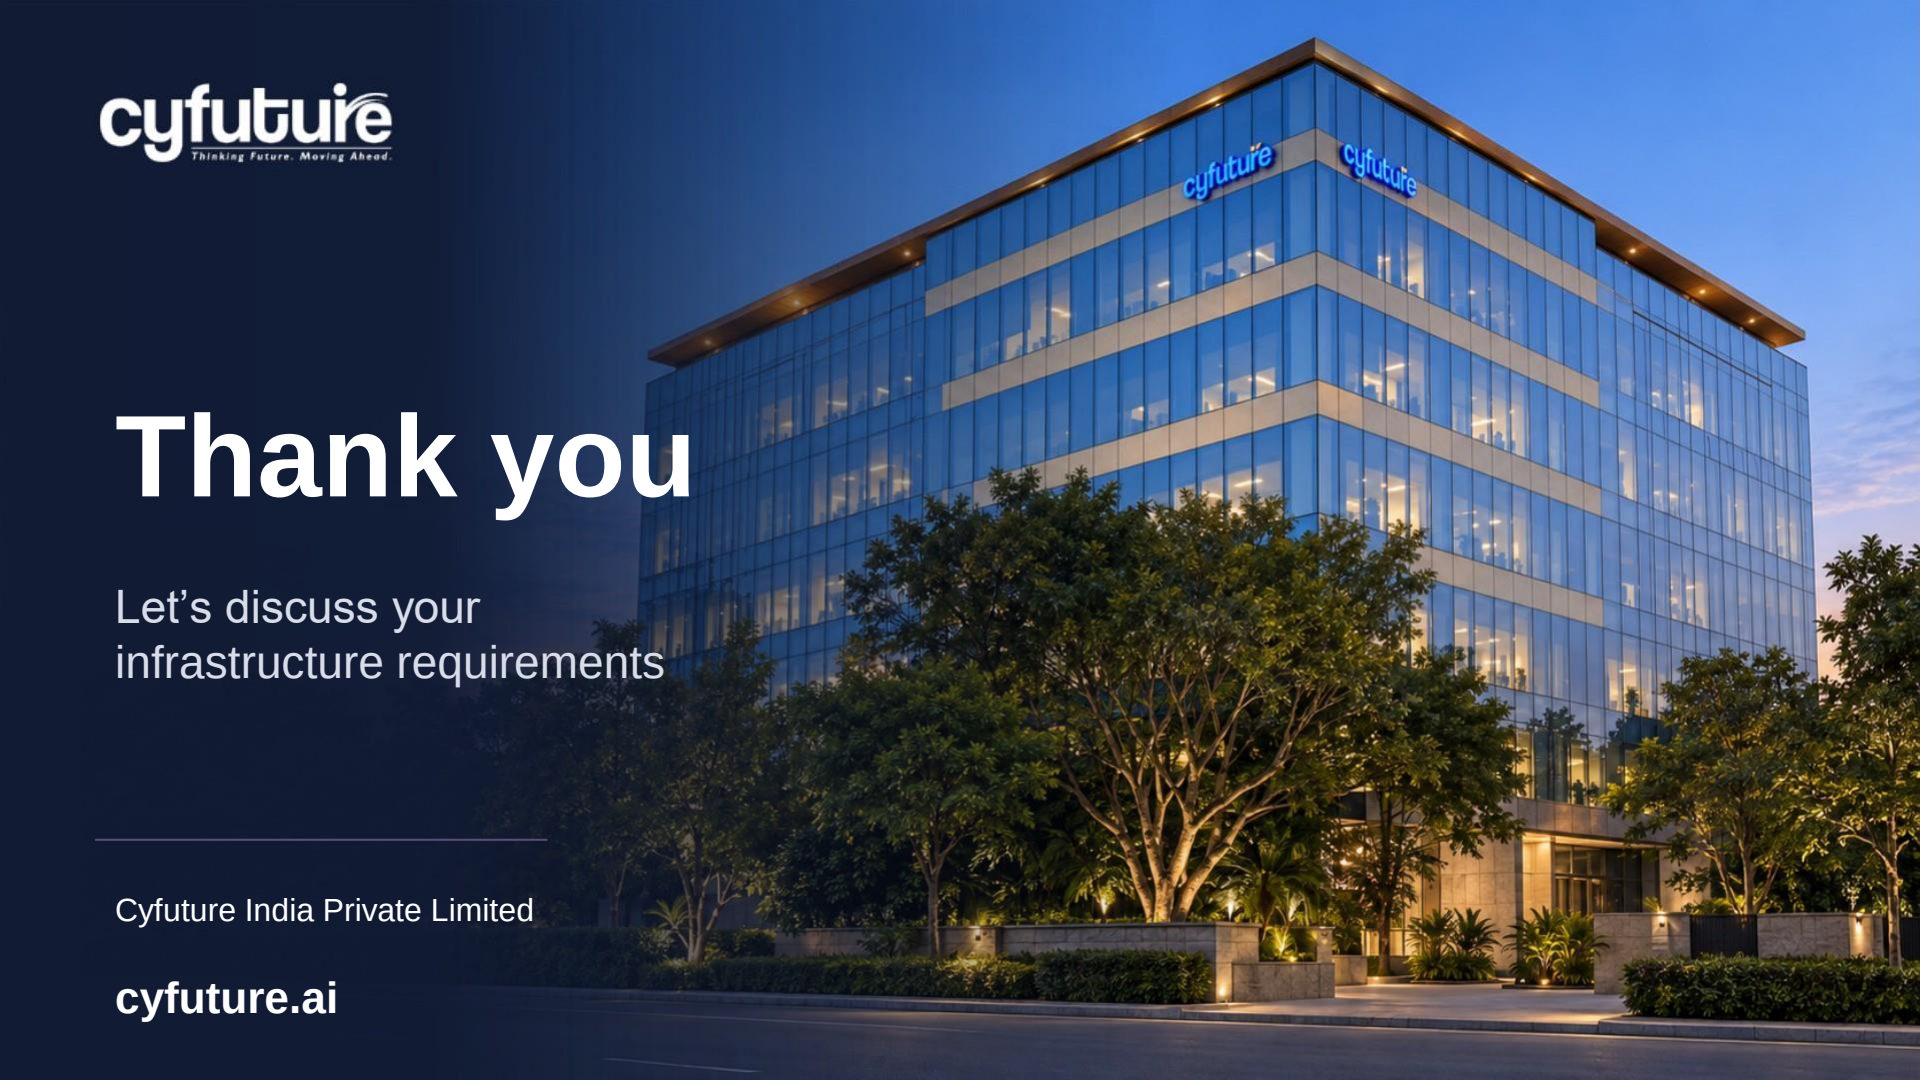


Existing Cyfuture building photograph

cyfuture
Thinking Future. Moving Ahead.

Thank you

Let’s discuss your
infrastructure requirements



Cyfuture India Private Limited

cyfuture.ai