cyfuture
Thinking Future. Moving Ahead.



NEXT-GENERATION AI INFRASTRUCTURE

100 MW

Built for the AI frontier.

LIQUID-COOLED DATA CENTRES

40 MW
JAIPUR-2 / MAR 2027

20 MW
NSEZ-2 / FEB 2027

40 MW
NSEZ-3 / MAR 2027

TECHNICAL & SALES BROCHURE | PLANNED PORTFOLIO | OCTOBER 2026

01 / THE OPPORTUNITY

The infrastructure behind your next breakthrough.

Three planned facilities. Two strategic regions. One integrated approach to power, cooling, connectivity and operations.

N+N

Power without compromise

N+N UPS paths and N+1 standby generation, with block-level capacity engineered around IT and critical cooling.

AI

Liquid-first design

Dedicated AI platform floors, direct-to-chip liquid cooling and complementary air-cooled halls.

6 m

Room to evolve

Prefabricated construction with a heavy-duty structural design and a 6 m floor-to-floor planning basis.

A buyer-ready technical conversation

Portfolio & locations 04-06 • Buildings & platforms 07-10
Reference design, power & cooling 11-15
Networks & storage 16-18 • Security & operations 19-21
Delivery & commercial engagement 22-24
Specification schedules & source notes 25-28

How to read this brochure

The site capacities, utility voltages, redundancy and dates follow Cyfuture's project brief. Detailed engineering values are proposed design targets. Final capacity, performance and service commitments are fixed in the approved design and customer contract.

RFS = ready for service. 100 MW is planned IT load, not total utility demand. No certification, platform endorsement or completed construction is implied.

02 / CYFUTURE

An Indian foundation. A global customer perspective.

Cloud, data centres, AI and managed services brought together by a business established in 2001.



ACTUAL CYFUTURE BUILDING PHOTOGRAPH

From infrastructure to outcomes

Cyfuture India Private Limited serves enterprises through cloud and infrastructure services, data centre solutions, AI and analytics, business applications and managed operations.

An established regional base

Existing operations in Noida and Jaipur provide the business context for this proposed expansion. The 100 MW programme is presented separately from the current operating estate.

A single commercial conversation

Engage on space, power, cooling, connectivity and optional managed services through one coordinated solution-design process.

Company background: Cyfuture About Us [1]. Photo supplied in Cyfuture B300 commercial proposal; it does not depict the three planned facilities.

03 / PORTFOLIO

100 MW of planned IT capacity.

Capacity stated at the IT load boundary. Regional and site allocations follow the Cyfuture project brief.

40 MW	20 MW	40 MW
JAIPUR-2	NSEZ-2	NSEZ-3

RAJASTHAN REGION

JAIPUR-2

40 MW IT

Planned RFS: March 2027

132 kV transmission

DELHI NCR REGION

NSEZ-2

20 MW IT

Planned RFS: February 2027

Multiple 33 kV dual feeds

DELHI NCR REGION

NSEZ-3

40 MW IT

Planned RFS: March 2027

Multiple 33 kV dual feeds

Scale by committed MW, not by a fixed rack count. Rack density, auxiliary IT, network and storage allocations determine the final deployable configuration within each site’s capacity.

Programme dates are planned RFS milestones, subject to utility energisation, statutory approvals, delivery and integrated commissioning.

04 / RAJASTHAN REGION

Jaipur-2 | Built around a 40 MW IT ambition.

IT Park, Sitapura, Jaipur • Planned RFS: March 2027

132 kV

Transmission-connected

Planned connection at 132 kV, with transformation, protection and downstream A/B distribution developed through the approved electrical design.

AI platform floors

Dedicated floor planning for liquid-cooled AI clusters, supporting scalable power and thermal blocks. Separate air-cooled areas accommodate suitable enterprise, network and storage equipment.

A Rajasthan engineering response

Heat-rejection plant sized against local hourly climate data and extreme design conditions. Water strategy, dust filtration, drainage and equipment derating are evaluated for the actual site.

Buyer diligence pack

Request the utility-sanction schedule, route drawings, land and building approvals, geotechnical report, structural capacity matrix, expansion phasing and site-specific service boundary.

40 MW IT requires additional facility capacity for cooling and electrical losses. Transformer and utility MVA are established by load-flow and power-factor studies.

Location, IT capacity, RFS and voltage: Cyfuture project brief. Grid capacity, protection settings and upstream independence require site-specific verification.

05 / DELHI NCR REGION

Two NSEZ facilities. 60 MW of planned IT load.

Noida Special Economic Zone • Capacity for separate deployments or a coordinated campus strategy.

20 MW

NSEZ-2

20 MW IT • Planned RFS February 2027
Multiple 33 kV dual transmission feeds, as defined in the project brief.

40 MW

NSEZ-3

40 MW IT • Planned RFS March 2027
Multiple 33 kV dual transmission feeds, as defined in the project brief.

Physical separation is engineered

Dedicated A/B paths, sectionalised distribution, separate meet-me rooms and diverse entry routes form the proposed basis. Upstream substations, cable corridors and flood exposures must be checked for shared risks.

Plan resilience at the right boundary

Two facilities in one region should not automatically be treated as independent disaster-recovery regions. For regional recovery, assess Jaipur plus Noida and validate application RPO, RTO and network behaviour.

The number and independence of upstream sources, sanctioned demand, SEZ customer eligibility and customs treatment are confirmed during diligence.

Engineered around the rack.



Illustrative AI-generated concept visualisation. Rack arrangement, pipework, equipment and clearances are subject to the final coordinated design.

06 / PREFABRICATED CONSTRUCTION

Factory precision. Site-specific structural strength.

A prefabricated delivery strategy for repeatable blocks, without compromising the engineering of high-mass AI floors.

Structure and envelope

Engineered steel or precast/composite structural modules with robust concrete/composite slabs. Fire resistance, corrosion protection, connections, vibration, deflection and seismic actions are designed for each site.

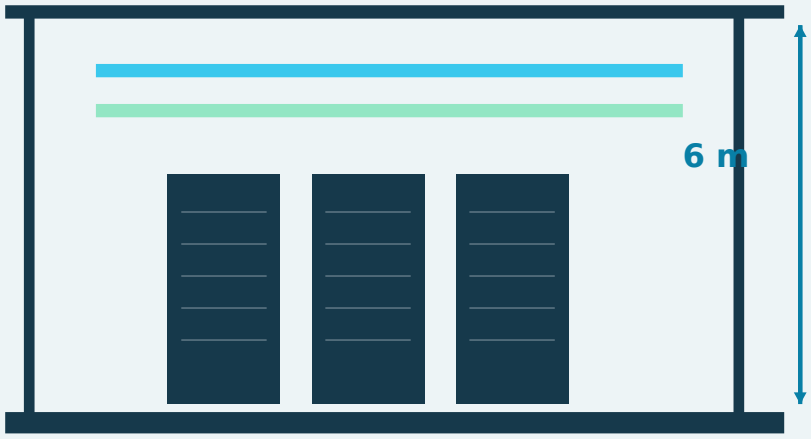
Modular MEP infrastructure

Factory-assembled electrical and cooling skids; pretested power rooms; repeatable pipe headers and busways; accessible isolation points. Maintain separate maintainable A/B service corridors.

Build, expand, maintain

Future blocks are isolated from live operations. Reserved plant bays, lifting zones, equipment replacement routes and maintenance clearances are part of the master plan.

ILLUSTRATIVE FLOOR SECTION



Structural floor-to-floor

Overhead segregation: A/B power, network and liquid services. Clear working height is established after coordination of slabs, pipework, busways and fire systems.

6 m is interpreted as floor-to-floor height. If 6 m clear internal height is required, the building section and total height must be increased accordingly.

07 / FLOOR & RACK ENGINEERING

Heavy AI requires four separate load checks.

Proposed structural design brief. These values are engineering targets, not a universal Vera Rubin standard.

DESIGN PARAMETER	PROPOSED BASIS	ACCEPTANCE REQUIREMENT
AI hall floor UDL	3,000 kg/m² (about 29.4 kN/m²)	Structural design of slab, beams, columns and foundations; all relevant load combinations.
GB300 / Rubin / Helios rack position	4,000 kg wet rack planning envelope	Local support, footprint, centre of gravity and simultaneous floor loading checked independently.
Local point and rolling loads	Individual checks required	Caster, levelling-foot, anchor, plinth and movement loads; UDL alone is insufficient.
Heavy equipment delivery	6,000 kg route / lift payload target	Verify packed assembly plus handling equipment; dimensions, turning radius and dynamic effects.
Air-cooled / support floors	1,500 kg/m² preliminary UDL target	Rack, battery, storage and other concentrated loads checked individually.

Rubin’s published approximate rack mass is 1,814 kg [6]; final wet/shipping mass is integration-specific. A 4,000 kg rack target is not a kg/m² floor rating.

08 / PLATFORM COMPATIBILITY

Dedicated floors for each compute generation.

Allocate floors by platform, power and cooling envelope. Floor counts and MW splits are set with customer demand.

FLOOR FAMILY	REFERENCE POINT	PROPOSED FACILITY ENVELOPE
Blackwell B300	DGX B300 reference: up to 15 kW/system; air-cooled configuration [3].	Dedicated B300 floor; 60-70 kW rack planning. Air containment or qualified liquid-cooled server configurations.
GB300 NVL72	Rack-scale liquid-cooled system; published maximum rack power up to 142 kW [4].	Up to 200 kW per selected rack position; exact liquid and residual-air split from integration schedule.
AMD MI355X	Up to 1,400 W/module with liquid cooling; reference maximum liquid inlet 43°C [7].	Dedicated MI355X floor; 80-150 kW rack planning band. Total server and rack design controls actual load.
Vera Rubin NVL72	Public reference supports a 45°C coolant inlet; final rack integration data required [6].	Dedicated next-generation floor; 300 kW enabling envelope at selected positions, subject to acceptance.
AMD MI455X / Helios	Rack-scale 72-accelerator blueprint with liquid-cooled integration [8].	Dedicated MI455X / Helios floor; 300 kW enabling envelope at selected positions, subject to acceptance.

Proposed rack envelopes are capacity reservations, not published platform consumption or a promise at every rack. Final server SKU, firmware and integrator schedule govern compatibility.

09 / NVIDIA REFERENCE ARCHITECTURE

From the AI factory blueprint to the facility.

Design principles informed by NVIDIA DGX SuperPOD, NVL72 AI Factory and public DSX facilities guidance.

01 / Scale-up domain

Rack-local accelerator interconnect. Treat the integrated rack as a tested compute unit.

02 / Scale-out fabric

High-bandwidth compute network across racks, with topology and rail mapping matched to the platform.

03 / Data services

Independent storage, tenant services, in-band management and protected out-of-band control.

Electrical integration

Resolve rack bus voltage, power shelves, AC/DC boundary, transient behaviour and power-cap settings before selecting distribution.

Thermal integration

Resolve inlet temperature, fluid quality, flow, differential pressure, heat capture and residual air load for every platform.

Digital validation

Use a coordinated 3D model plus electrical, thermal and network studies to validate the deployment before energisation.

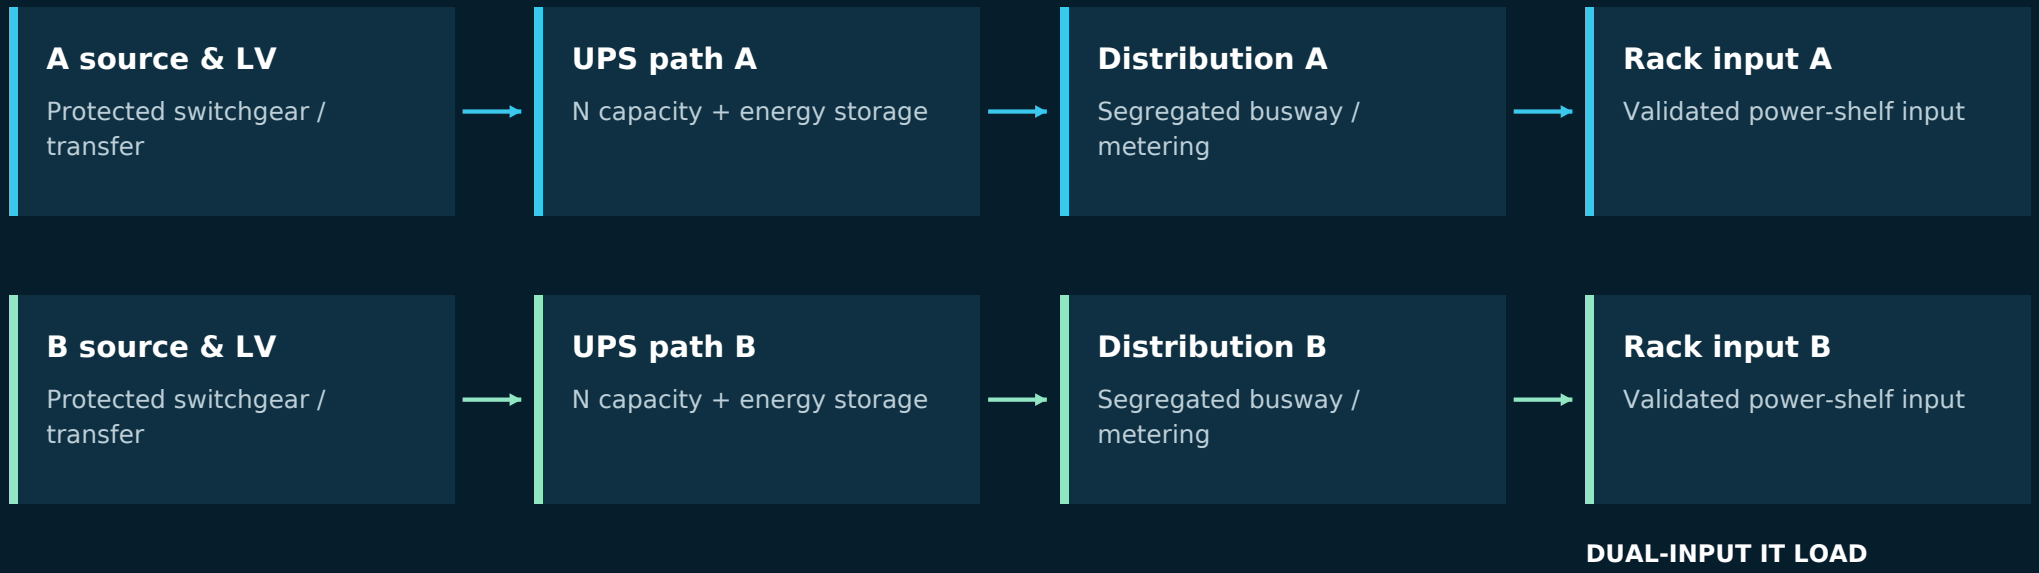
Reference-informed design is a starting point. Platform acceptance requires a documented compliance matrix against the applicable release and integrator requirements.

Sources [3]-[6], [9]. This brochure does not claim NVIDIA certification, endorsement, SuperPOD validation or access to restricted reference designs.

10 / ELECTRICAL RESILIENCE

N+N UPS. Two full-capacity paths.

Proposed A/B architecture: either UPS path can support 100% of the committed critical load.



Independence by design

Physical separation, selective protection, separate maintenance bypasses and fault containment. Confirm the customer’s rack power-shelf topology preserves the intended A/B resilience.

Ride through the transition

Proposed UPS autonomy: at least 5 minutes at end-of-life design load, validated against generator start, transfer and load steps. Include critical cooling pumps and controls in the continuity design.

Conceptual topology. Redundancy at a component level does not, by itself, establish end-to-end availability or any Tier certification.

11 / GRID & STANDBY POWER

Designed for the load step, as well as the steady state.

Utility interface, generator reserves and AI transients are engineered as one power system.

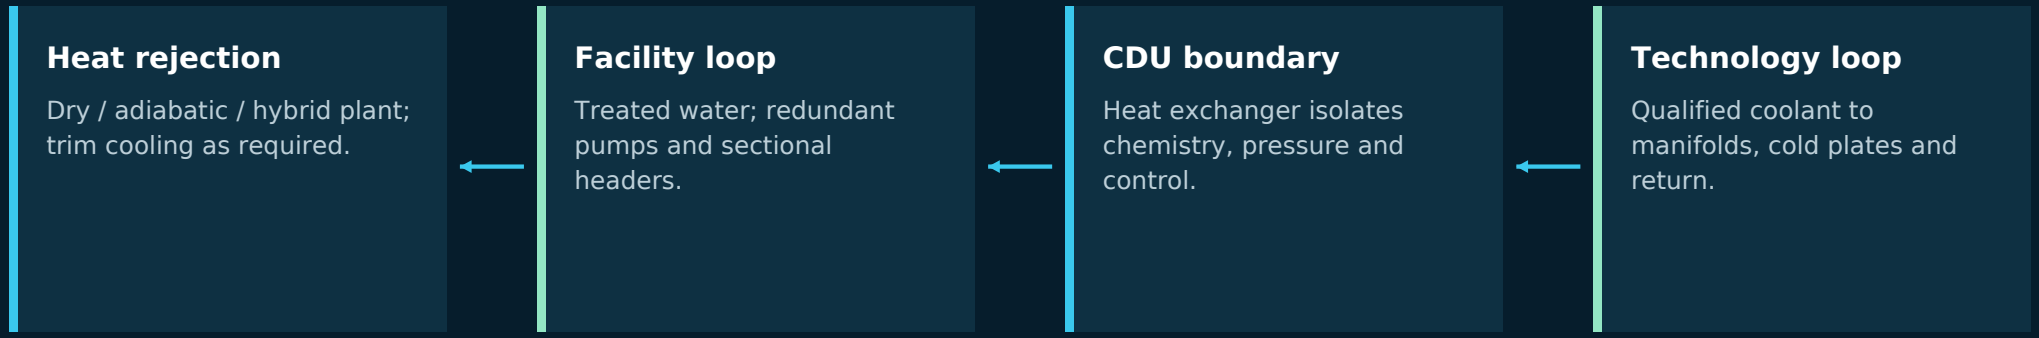
SYSTEM	PROPOSED SPECIFICATION
Jaipur utility	132 kV transmission connection; transformer configuration, protection and sanctioned capacity established by utility and design studies.
Noida utility	Multiple 33 kV dual transmission feeds; confirm separate routes and upstream common points through utility single-line drawings.
Standby generation	N+1 at the defined capacity block, supporting IT, critical cooling and essential auxiliaries after site derating.
Fuel resilience	24-hour on-site usable-fuel autonomy target at design emergency load; compliant storage and replenishment arrangements.
Power quality	Load-flow, short-circuit, discrimination, harmonic and arc-flash studies; power-quality monitoring and surge protection.
AI transients	Validate slew rates, synchronised load steps, UPS battery response, generator response and agreed platform power caps.

Illustration only: 100 MW IT at a PUE of 1.30 would draw 130 MW at the facility boundary. This is not a PUE commitment or a utility sanction.

12 / LIQUID COOLING

A controlled thermal path, from silicon to ambient.

Direct-to-chip cooling with separated facility and technology loops; maintainability designed into every block.



Capacity after a failure

N+1 pump and CDU capacity target per defined block. Separate distribution sections, isolation valves and duty sharing allow a unit to be removed without exceeding the committed thermal load.

Protect equipment and continuity

Continuous leak sensing, pressure and flow monitoring, strainers, coolant sampling and corrosion control. Size thermal buffers and UPS-supported controls/pumps for the tested failover scenario.

Arrows show heat moving from IT towards ambient, not coolant flow. Supply and return circuits are commissioned to the selected platform’s chemistry and pressure limits.

13 / THERMAL DESIGN

Warm-water AI. Air cooling where the workload needs it.

A mixed thermal strategy supports next-generation racks and the wider enterprise estate.

DESIGN TOPIC	PROPOSED BASIS
Temperature bands	Provide configurable technology-loop inlet setpoints, typically 25-45°C where the selected platform permits. 45°C is not a blanket requirement for every platform.
Design delta-T	10-15 K preliminary liquid-loop sizing band; final flow and pressure loss come from equipment curves and the approved thermal schedule.
Illustrative liquid flow	At 200 kW liquid heat and 10 K delta-T, water flow is about 287 L/min. At 300 kW, about 431 L/min. Verify coolant properties and heat-capture fraction.
Air-cooled halls	10-30 kW/rack planning band; higher densities only after airflow and containment validation. Proposed server-inlet operating target 22-27°C within equipment limits.
Residual air cooling	Allow for networking, storage, power-shelf and non-liquid components. Size the air plant to measured platform heat splits; do not assume 100% liquid capture.
Ambient and dew point	Model hot-weather operation for Jaipur and Noida. Keep coolant above local dew point or provide an approved condensation-control design.

Flow illustration uses $Q = m \times cp \times \text{delta-T}$, water cp about 4.18 kJ/kg·K and density about 1 kg/L. Sources [6], [9], [10]; detailed operating bands are proposed.

14 / NETWORK ARCHITECTURE

Four network functions. Clear boundaries.

Separate compute and protected management; storage and services may use validated converged fabrics with logical isolation.

01

Scale-out compute

400 / 800 Gb/s-class endpoint connectivity; non-blocking design target within the agreed cluster; rail and topology design matched to the platform.

02

Storage / data

High-throughput storage fabric, with RDMA where supported; capacity sized for ingest, checkpoints, recovery and concurrent training jobs.

03

Front-end / in-band

Tenant ingress, cluster services, orchestration and management networks; policy separation and redundant routed connectivity.

04

Out-of-band

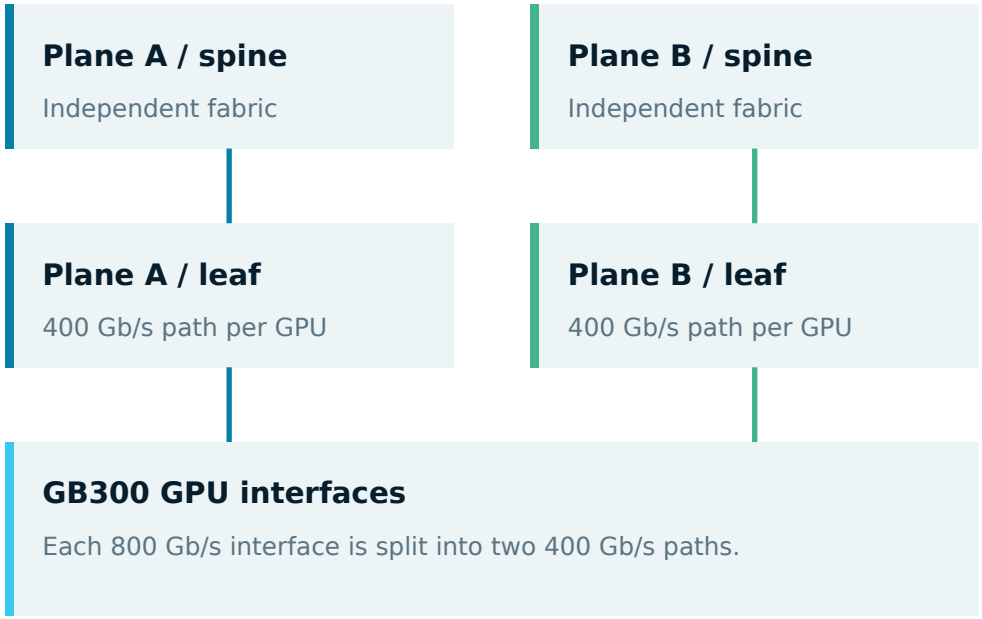
Separate device-management network, console access and protected administrative entry; independent recovery access.

References [3]-[5], [8]. Exact per-GPU, per-node and per-rack rates differ. 800 Gb/s is a network port class, not a guaranteed application throughput.

15 / FABRIC & CONNECTIVITY

An AI fabric with room to scale.

GB300 reference example: dual compute planes, with separate services and out-of-band management [5a].



A plane failure reduces available bandwidth. Port counts, rails, bisection capacity and switch ratios are fixed in the cluster design.

Carrier-neutral campus design

Proposed dual meet-me rooms and diverse entries, risers and cross-connect paths. Confirm carrier availability and upstream route diversity for the selected site.

Cloud, WAN and recovery

Private cloud links, internet, DCI, DDoS protection and encrypted WAN as scoped options. Contract bandwidth, latency, jitter and recovery objectives with each provider.

Fabric engineering

Use RoCEv2 with engineered ECN/PFC where required, or platform-appropriate InfiniBand. Test collectives, congestion, optics and failover. Segregate tenant traffic and administrative access.

Jaipur and Noida support a multi-region application strategy. Inter-city transport is specified separately from the local training fabric.

Conceptual reference interpretation. Reserve pathways for evolving optics, including potential 1.6 Tb/s deployments; supported reach and performance require final design.

16 / STORAGE & DATA SERVICES

Keep the accelerators supplied with data.

Storage is sized from workload behaviour, not just from the number of GPUs.

TRAIN

Hot training tier

Parallel file or high-performance object access with metadata capacity, small-file performance and direct data paths validated against the AI stack.

RECOVER

Checkpoint tier

Write bursts, checkpoint intervals and restart time determine sustained throughput, capacity and the required failure margin.

PROTECT

Durable data tier

Versioning, encryption, retention, immutable backup and geographically separate recovery copies as agreed in scope.

Acceptance based on real workloads

Agree dataset sizes, read/write mix, concurrency, metadata rates, checkpoint completion time and degraded-mode throughput. Run customer-representative tests alongside the GPU collective benchmark.

Management and visibility

Provide capacity, latency, error, fabric and thermal telemetry. Integrate tenant access controls, keys, audit logs and recovery procedures with the customer’s operating model.

GPU, fabric, storage and software systems are customer-selected or separately contracted managed services; they are not automatically included with colocation capacity.

17 / TRUST & SECURITY

Protect the campus, the control plane and the tenant.

Layered protection with clear customer and operator responsibilities.

Physical protection

Controlled perimeter, monitored entries, vehicle screening, mantraps, role-based access and visitor escorting. Secure loading and staging areas separate logistics from live operations.

Tenant separation

Dedicated suites or cages, documented access permissions, cabinet controls and auditable work orders. Camera coverage and retention are agreed within applicable requirements.

Operational technology security

Separate building, power and cooling networks from tenant networks. Use least privilege, MFA, monitored remote access, logging and controlled change windows.

Fire and life safety

Early smoke detection, zoned alarm systems, suitable suppression and code-compliant sprinkler/pre-action design. Separate electrical, battery and fuel risks with appropriate fire compartments.

Site-specific statutory compliance and certification plans are supplied during diligence. Existing corporate credentials do not automatically certify new facilities.

18 / RESOURCE PERFORMANCE

Measure the whole facility. Optimise the useful work.

A transparent energy and water strategy for dense AI operations.

PUE

Energy efficiency

Meter IT, electrical losses and cooling separately. Model annual PUE across actual load, weather and equipment configurations.

WUE

Water accountability

Track site water use and WUE. Compare dry, adiabatic and hybrid heat rejection against climate, water quality and availability.

CO_2

Carbon visibility

Report electricity use and relevant emissions data. Renewable sourcing options require contracted supply and evidence.

Warm water expands the options

Higher permitted coolant temperatures can reduce compressor dependence. Free-cooling hours are established by the site climate model; year-round chiller-free operation is not assumed.

Design for long operating life

Use variable-speed equipment, staged capacity, heat-reuse provisions where a viable local heat sink exists, and accessible modular components that can be maintained and replaced.

No numerical PUE, WUE, renewable share or carbon-neutral claim is guaranteed in this brochure. Performance commitments require a stated measurement boundary and operating conditions.

19 / OPERATIONS

Operational discipline, around the clock.

Proposed 24×7 facility monitoring and service coordination, with a documented escalation and change-control model.

CAPABILITY	CUSTOMER-VISIBLE DELIVERABLE
Facility command	DCIM, BMS and electrical monitoring integrated into alarm handling; dashboards for capacity, energy, temperature and critical conditions.
Remote hands	Authorised rack work, visual inspection and incident assistance under agreed response windows, permitted actions and customer approval rules.
Maintenance	Documented SOPs, MOPs and EOPs; preventive and predictive maintenance; switching control, spare parts and maintenance windows.
Incident response	Named severity definitions, escalation contacts, communications cadence, incident timeline, root-cause report and corrective actions.
Liquid operations	Trained handling, controlled coolant changes, leak response, isolation, flushing and contamination management with maintenance records.
Customer governance	Capacity forecasts, monthly service reviews, agreed audit evidence, access reviews and change notifications.

SOP / MOP / EOP = standard, maintenance and emergency operating procedures. Response and restoration commitments are defined in the service agreement.

20 / COMMISSIONING

Prove the failure cases before production goes live.

A staged acceptance programme for the building, the infrastructure and the integrated customer deployment.

01	**Design validation**	Approved drawings, load schedules, fluid specification, structural acceptance and coordinated model.
02	**Factory acceptance**	Electrical skids, control logic, cooling modules and protection settings tested against the design.
03	**Site acceptance**	Installation checks, flushing, pressure tests, protection injection and calibrated instrumentation.
04	**Integrated systems**	Utility loss, A/B isolation, generator load steps, pump/CDU failure, thermal ride-through and alarms.
05	**Customer workload**	Burn-in, sustained thermal load, network collectives, storage checkpoints and recovery validation.

Witnessed tests must reproduce agreed operating loads and failure modes. Test scripts, thresholds, exceptions and signed acceptance records become part of the handover pack.

21 / COMMERCIAL MODELS

Buy the capacity and support your programme needs.

A flexible offer structure for hyperscalers, model builders, sovereign AI and enterprise customers.

ENGAGEMENT	TYPICAL SCOPE	AGREE BEFORE CONTRACT
Powered AI colocation	Dedicated rack zones, pods, halls or floors; committed IT kW/MW; contracted cooling boundary.	Power and thermal envelopes; rack count; delivery phase; metering; growth rights.
Private AI suite	Segregated tenant area with custom access, network and liquid distribution requirements.	Security boundary; customer fit-out; dedicated/shared plant; service responsibilities.
Build-to-suit deployment	Customer-specific integration for rack-scale platforms and phased cluster capacity.	Design freeze; milestones; change control; equipment interface and acceptance.
Managed service options	Remote hands, managed network/storage, platform operations and monitoring, as scoped.	Hardware ownership; software licensing; support coverage; patching and incident boundary.

SLA schedule: define availability at the service boundary, temperature/flow/pressure bands, power quality, planned maintenance, response times, exclusions, measurement methods and service credits. No uptime percentage is represented as agreed until contracted.

Commercials separate capacity charges, energy, cooling, connectivity, cross-connects, fit-out and managed services. Prices and availability are provided in a customer-specific proposal.

22 / CUSTOMER ONBOARDING

From requirement to an accepted operating environment.

A structured discovery process reduces interface risk and makes the commercial proposal precise.

Tell us about the deployment

Platform and integrator; rack dimensions and packed/wet weights; maximum and steady-state kW; ramp behaviour; required MW and phasing; preferred site and target service date.

Define the thermal and electrical interfaces

Coolant, supply temperature, flow, pressure and liquid heat fraction; AC/DC boundary; connector schedules; UPS needs; rack A/B behaviour; transient and power-cap requirements.

Set the application requirements

Compute topology, storage, connectivity, migration and recovery; data residency, encryption/key ownership, assurance requirements and the split of customer/operator responsibilities.

YOUR DILIGENCE PACK

01 Site and utility evidence

02 Technical compliance matrix

03 Power and cooling capacity schedule

04 Structural and logistics acceptance

05 Carrier diversity and network design

06 Commissioning and handover plan

07 Service agreement and pricing schedule

Migration planning covers staging, lifting, rack integration, coolant filling, burn-in, change windows, data movement, cutover and a documented rollback path.

23 / SPECIFICATION SCHEDULE

The portfolio at a glance.

Project-brief inputs and proposed engineering targets are distinguished below.

ITEM	BASIS	STATUS
Total planned IT load	100 MW: Jaipur-2 40 MW; NSEZ-2 20 MW; NSEZ-3 40 MW.	Project brief
Planned RFS	NSEZ-2 February 2027; Jaipur-2 and NSEZ-3 March 2027.	Project brief
Utility interface	Jaipur 132 kV; Noida multiple 33 kV dual transmission feeds.	Project brief
Electrical redundancy	UPS N+N; generator plant N+1.	Project brief
Building strategy	Prefabricated structure / MEP modules; 6 m floor-to-floor height.	Brief + height interpretation
AI floor structure	3,000 kg/m² UDL; 4,000 kg populated-rack planning envelope.	Proposed; engineer to approve
Heavy delivery path	6,000 kg route/lift payload target, subject to packaged load and mover.	Proposed; route-specific
Ride-through and fuel	UPS ≥5 min end-of-life; fuel 24 h at design emergency load.	Proposed; test and approve
Cooling topology	Liquid first; isolated FWS/TCS through CDUs; N+1 block capacity target.	Proposed engineering basis

N denotes the capacity required for the defined load block under the stated design conditions. Ratings must include relevant environmental derating.

24 / SPECIFICATION SCHEDULE

The customer deployment at a glance.

Final platform compatibility and service levels are established by a customer-specific integration schedule.

ITEM	PROPOSED DESIGN / SERVICE BASIS
Platform zoning	Dedicated floor planning: B300, GB300 NVL72, MI355X, Vera Rubin NVL72 and MI455X / Helios; allocations demand-led.
Rack power envelopes	B300 60-70 kW; MI355X 80-150 kW; selected GB300 positions up to 200 kW; selected next-generation positions up to 300 kW.
Structural acceptance	Actual wet rack mass, local support loads, anchors, seismic restraint and all movement routes accepted individually.
Air-cooled capacity	10-30 kW/rack planning band; contained airflow and separate provision for residual air heat from liquid-cooled systems.
Network fabric	400/800 Gb/s-class compute links; segregated storage, front-end and management functions; physical fabrics matched to the reference design.
Campus connectivity	Carrier-neutral approach, dual meet-me rooms, diverse entries and routed paths; provider serviceability and route evidence required.
Security and operations	Layered access control, tenant isolation, OT separation, monitoring, 24×7 coordination and controlled change procedures.
Performance reporting	Agreed metering boundaries for energy, PUE, WUE, environmental conditions, network and storage service metrics.
Handover	As-built records, asset register, capacity schedules, test evidence, procedures, training and signed service acceptance.

All values on this page are proposed planning envelopes unless explicitly cited as a platform fact elsewhere. They are not uniform entitlements across all rack positions.

25 / DESIGN BASIS & INTERFACES

Clear specifications. Explicit acceptance boundaries.

Use this brochure to open a detailed technical discussion; use the approved design schedule to order capacity.

What comes from the project brief

The 100 MW IT allocation, site names, planned RFS dates, utility voltage levels, N+N UPS, N+1 generation, liquid/air mix, dedicated platform floors and 6 m height originate in Cyfuture’s instructions.

What has been proposed

Floor loads, rack envelopes, route payloads, UPS autonomy, fuel duration, cooling redundancy and setpoints, network options and operating procedures are proposed engineering requirements. They need design-owner approval.

How platform data is used

Published values describe specific reference implementations. Rack mass, power, coolant, footprint and network details can differ by integrator. Separate operating and transport schedules are required.

Capacity and resilience

IT MW excludes facility overheads. Rack allocations include supporting IT. A/B power and N+1 cooling/generation are tested within the declared block and failure boundary, including maintenance states.

Structure and delivery

UDL, local rack reactions and rolling loads are distinct checks. Use engineered heavy-load slabs or podiums. The 4,000 kg target requires accepted wet rack mass; the 6,000 kg route target includes rigging and packaging.

Standards and assurance

Detailed design establishes applicable building, electrical, fire, environmental and local statutory requirements. Any ISO, Tier, TIA or platform-validation scope must be evidenced separately for the relevant facility.

Platform names identify intended compatibility. No supplier selection or third-party endorsement is implied. Numerical targets are not guarantees of platform availability.

26 / PRIMARY REFERENCES

Grounded in published technical guidance.

Reference review: 7 October 2026. Linked publications support platform facts; proposed Cyfuture targets are identified separately.

01 **Cyfuture | About Us**
Established in 2001; service portfolio and existing regional presence.

02 **Cyfuture | Contact Us**
Published Noida office, Jaipur IT Park and sales telephone.

03 **NVIDIA | DGX B300 User Guide**
System size, weight, input power and network connections.

04 **NVIDIA | NVL72 AI Factory Components**
GB300 rack architecture and published up-to-142 kW demand.

05a **NVIDIA | NVL72 Network Topologies**
Dual-plane compute and converged services; failure bandwidth impact.

05b **NVIDIA | B300 SuperPOD Network Fabrics**
Platform-specific compute, storage and management patterns.

06 **NVIDIA | Vera Rubin POD Technical Overview**
16 March 2026; rack architecture, approximate mass and warm coolant.

07 **AMD | MI355X Platform Datasheet**
Module power, direct liquid cooling and coolant operating specification.

08 **AMD | Helios Rackscale Solution Brief**
July 2026; MI455X rack integration and scale-out architecture.

09 **NVIDIA | DSX AI Factory Platform**
Coordinated facility, compute, power, cooling and simulation approach.

10 **ASHRAE | Data Centres & Telecom Facilities**
Thermal environment, liquid loops and dew-point considerations.

P **Project sources | Supplied materials**
Cyfuture project brief and B300 proposal photograph. Internal source.

Click a reference title to open its source. Architectural diagrams are original schematic interpretations; the hall visual is conceptual and the building photo is supplied.

CYFUTURE

Let’s design your next AI deployment.

JAIPUR • DELHI NCR
100 MW PLANNED IT CAPACITY

Start the technical conversation

cyfuture.com

Sales: +91-8800766220

Cyfuture India Private Limited
Plot 197-198, NSEZ, Phase II
Noida, Uttar Pradesh 201305, India

TECHNICAL & SALES / OCTOBER 2026

Planned design basis. Final scope by customer agreement.

