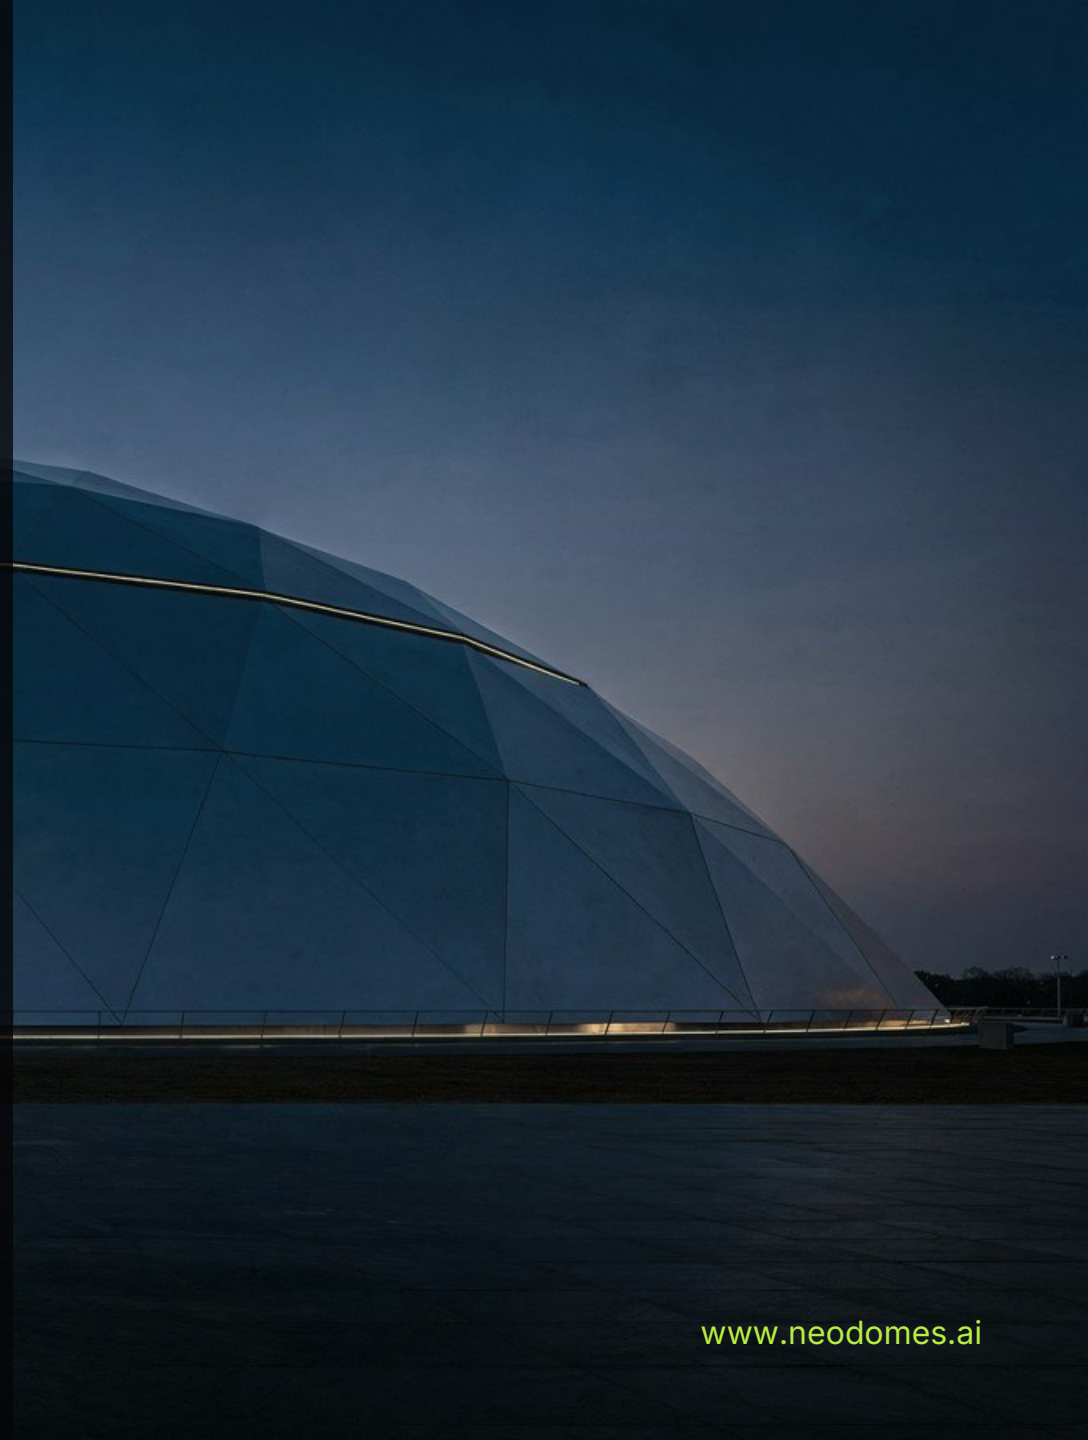
NeoDomes

THE FUTURE OF AI DATA CENTER INFRASTRUCTURE

NeoDome 3.0

Technical Overview

A next-generation AI infrastructure platform
— engineered, not assembled.



AI INFRASTRUCTURE · ENGINEERED, NOT ASSEMBLED

September 2026 · Confidential

NeoDomes Inc.

www.neodomes.ai

CONFIDENTIAL

The Platform at a Glance

A clustered, modular, AI-native dome platform.

NeoDome 3.0 is the production platform of NeoDomes Inc. — a next-generation AI infrastructure platform that co-locates high-density compute with on-site industrial energy generation, engineered as a single integrated system for the AI-era compute densities and delivery timelines that conventional data-center architectures were never designed to serve.

LAYER 1

The Dome

Engineered geodesic shell. Manufactured, repeatable, and permittable. Compact footprint, quiet operation, near-zero water draw — sitable in climates and communities conventional halls cannot enter.

LAYER 2

The Intelligence

AI-orchestrated dispatch, workload-aware optimization, predictive maintenance. Every site simulated end-to-end before commissioning. The model compounds with each deployment.

LAYER 3

The Network

Carrier-neutral interconnection, multi-cloud on-ramps, geo-redundant fiber. Discrete domes operate as one logically unified compute fabric across regions.

LAYER 4

The Power

Behind-the-meter generation engineered as part of the dome — natural gas today, with hydrogen-ready, solar + storage, and waste-heat optionality.

MODULE SIZING

Single dome 10–20 MW · Multi-dome campuses 50, 100 MW+ · Capacity is added by adding domes, not by re-engineering the platform.

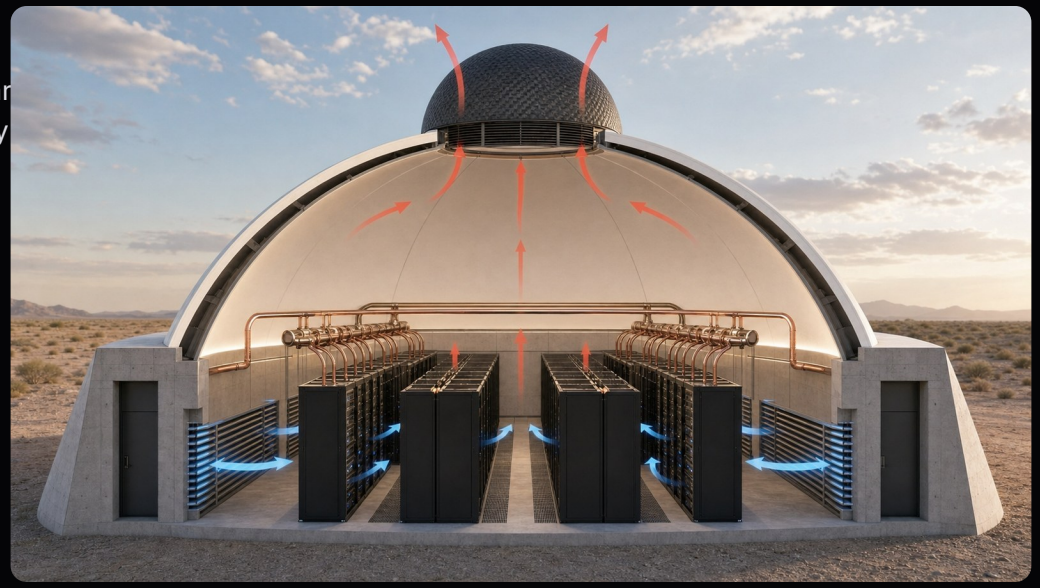
Physical Engineering

Inspired by nature. Engineered for AI-era density.

- **Geodesic shell.**
  Triangulation distributes structural load with roughly 30% less material than a comparable rectangular shell, opens a column-free interior, and assembles in modular phases. Shell erection in approximately 12 weeks — a decisive advantage on AI-scale delivery timelines.

- **Passive convection cooling.**
  The dome's curved interior creates natural convection currents that move heat without the chiller plants and water loops conventional designs depend on. Hybrid air + direct liquid + immersion supported out of the box. Built for NVIDIA GB200 / GB300-class density natively — not retrofitted to it.

- **Outside-in design.**
  Compact footprint, low-noise envelope, near-zero evaporative water draw, no industrial scar. Designed to be welcomed by host communities and to sit comfortably inside modern environmental and water-stewardship priorities.



Direct-to-chip cooling + passive thermal stratification through dome geometry.

SHELL ERECTION	SUSTAINED PUE	WATER DRAW	RACK DENSITY
~12 weeks	Sub-1.2	~90% less	GB200 / GB300

Operating Software and Interconnect

Two layers most operators sell separately. We engineer them in.

LAYER 2 · THE INTELLIGENCE

Software that compounds.

Pre-commissioning simulation.
Every dome is simulated end-to-end before construction — power, cooling, airflow, workload patterns. Performance is engineered, not discovered.

AI-orchestrated dispatch.
Continuous workload-aware optimization across compute, cooling, and power. Predictive maintenance keeps utilization high and OPEX low.

Cross-site learning.
Each new deployment feeds back into the model. The platform's operating intelligence gets sharper with every dome, every site, every megawatt — a compounding asset for the operator, not a static commissioning artifact.

Vendor-neutral.
Architecture accommodates NVIDIA GB200 / GB300 and like-type-class accelerators, custom AI silicon, and standard CPU compute under the same operating layer.

LAYER 3 · THE INTERCONNECT CLUSTER

Sites become a platform.

Regional fiber backbone.
Discrete dome locations operate as one geo-redundant, multi-tenant compute fabric across a region — not a collection of unrelated sites.

Carrier-neutral interconnection.
Multi-cloud on-ramps and owned route diversity at every dome cluster. AI-WAN class connectivity between domes, between sites, and to public-cloud and sovereign-AI endpoints under whatever data-residency regime the host jurisdiction requires.

Designed to interconnect with partner fiber.
The Network layer is built to interoperate with regional telecom and fiber carriers rather than compete with them.

Why it matters.
Without the Cluster, sites are valued as real estate. With it, they are valued as a platform.

Energy built for the load, not bolted on.

Behind-the-meter generation, engineered as part of the dome.

Global AI interconnection queues now stretch four to seven years in the markets where AI demand is heaviest. The NeoDome 3.0 platform was engineered to bypass that bottleneck by treating on-site generation as a first-class element of the design — not an afterthought strapped to a building already committed to a utility timeline.

Each dome is paired with on-site behind-the-meter generation, sized to the campus. Natural gas is the primary fuel today, with the architecture engineered to accept hydrogen blends, RNG, solar PV + battery storage, and waste-heat reuse as economics evolve. The system is dispatchable, islanded, and black-start capable.

Fuel-flexible by design. The power architecture is engineered to operate under whatever fuel and grid mix the host region actually has today — and to onboard next-generation fuels as they mature. One platform, multi-decade optionality. The integration is technical, not opportunistic.

Behind-the-Meter · Primary

On-site natural gas generation sized to the campus. Independent of utility queue timing. Dispatchable, islanded, black-start capable.

Grid-Parallel · Optional

Where regional utility capacity is available and economics support it, sites can interconnect for redundancy and surplus sales — without depending on the queue to energize.

Hybrid · Future-Fuel Ready

Architecture accepts hydrogen blends, RNG, solar PV, battery storage, and waste-heat reuse as economics evolve. One platform, multi-decade fuel optionality.

10, 20 or 25 MW NeoDomes
Sized for clustered campus buildout.

Modular sizes. Any campus scale. Grown to the customer plan.

NeoDome 3.0 campuses are delivered as manufactured, repeatable modules — 10, 20, or 25 MW per dome — and cluster into any campus scale the customer's plan requires: hundreds of megawatts and beyond. Capacity is added by adding domes, not by re-engineering the platform. Whether the target is a single dome or a multi-hundred-megawatt cluster, delivery uses the same architecture and the same build cadence.

INITIAL CAPACITY 10, 20 or 25 MW per dome	EXPANSION PATH 50, 100, 250, 500 MW+ by adding domes — scales to customer plan	SHELL ERECTION ~12 weeks per dome	SITE-TO-REVENUE TARGET Under 18 months
PRIMARY POWER Behind-the-meter natural gas	FUEL SOURCE Regionally sourced natural gas	POWER ARCHITECTURE Dispatchable, islanded, black-start	SUSTAINED PUE TARGET Sub-1.2
COOLING Hybrid air + direct liquid + immersion	RACK DENSITY NVIDIA GB200 / GB300 class	CONNECTIVITY Carrier-neutral, multi-cloud on-ramps	WATER DRAW Near-zero evaporative

Delivery structure (direct build, joint venture, or regional license) defined per engagement between NeoDomes Inc. and the counterparty.