



DUG Nomad puts AI where you need it.



dug
NOMAD

DUG Nomad delivers 80 kW of cooling in 80 sq ft.

Just add
power



DUG Nomad 10 puts high-density, AI-inference power wherever you need it — delivering 80 kW of IT heat rejection in a standard 10 ft Hi-Cube shipping container. Engineered for rapid deployment — with self-contained cooling that eliminates the need for external chillers, plumbing or site prep, DUG Nomad goes from delivery to operation in just hours. Designed by DUG's award-winning HPC engineers, Nomad is purpose-built for low-latency, high-throughput AI inference in edge, remote, or integrated environments.

For more information email info@dug.com

High-performance power

Small container. Huge capacity.

Up to 80 kW of IT heat rejection in a standard 10 ft Hi-Cube shipping container.

True mobility

Engineered for rapid deployment.

Fully self-contained in a standard shipping container for simple transportation. Operational within hours, not days — just add power.

Instant impact

Low latency. High throughput.

5-15 ms time-to-first-token for edge deployments. Build a redundant, distributed AI-inference network across your operational footprint. 3,600 to 5,400 tokens per second with NVIDIA H100 GPUs.

Scale with ease

Need more compute power?

Scale up to the DUG Nomad 40 for high-density computing and 1 MW of IT heat rejection.

