SPEC SHEET · MODULAR DATA CENTERS · DIRECT-TO-CHIP LIQUID

MDC-500 DLC

Liquid-cooled productized cluster · 2 to 5 MW · 150 kW per rack

LEAD TIME	SITE-READY	COOLING	NCP-ELIGIBLE
24–34 weeks	6 to 8 months	Direct-to-chip liquid	Quantum-2 / Spectrum-X

OVERVIEW

150 kW per rack on direct-to-chip liquid.

Same envelope, same UPS, same 6-to-8-month site-ready window.

The MDC-500 DLC is the liquid-cooled variant of the MDC-500 platform — same ISO-format container envelope, same 2N power and N+1 cooling topology, same operating-envelope rating. The differentiator is thermal architecture: closed-loop direct-to-chip liquid takes per-rack density to 150 kW for current-generation Blackwell-class AI training workloads, with headroom for the next GPU generation.

Buyers who deployed MDC-500 air for inference or current-density training can phase MDC-500 DLC modules into the same campus on a shared substation, shared SCADA, and shared weld-traceability standard. The cooling architecture is the only structural change between configurations. Site-ready in 6 to 8 months from PO; substation, generator, switchgear, and SCADA on the same schedule.

Core specifications

MECHANICAL

Form factor	ISO 668 standard container, 20-ft and 40-ft variants
Cluster size	2 to 5 MW productized; up to 50 MW phased
Operating envelope	–35°C to +52°C ambient
Ingress protection	IP54 sealed
Mounting	Concrete pad or pile foundation; anti-vibration optional
Transport	Standard road and rail; ISO corner castings

ELECTRICAL

Density	Up to 150 kW per rack direct-to-chip liquid
Power topology	2N power · N+1 cooling
Source	Bundled (substation + genset + ATS + SCADA) or your fence
UPS	Vertiv Liebert APM2 standard; Eaton/ Schneider available
Switchgear	Schneider Prisma Plus / Okken inside the box
Backup	Cummins generators to 3 MW
PDU	PDI/EnerSys floor-standing or in-row, high-density variant

Redeployment	Full disconnect and recommissioning per SAT package

THERMAL — LIQUID

Cooling type	Closed-loop direct-to-chip liquid
ASHRAE class	W4 (high-density liquid)
CDU	Vertiv liquid-cooling distribution unit, redundant pair
Coolant	Water-glycol (PG-25) standard; dielectric optional
Heat rejection	Dry cooler, evaporative, or hybrid — climate-specific
External water	Not required when paired with dry-cooler heat rejection
Air-cooled variant	Same envelope at 50 kW/rack — see MDC-500

FIRE & LIFE SAFETY

Suppression	Gas (FM-200 or Novec 1230) standard
Detection	VESDA aspiration + traditional smoke
Compliance	NFPA 75 · NFPA 110 · NFPA 70
Egress	Single or dual depending on layout

NVIDIA reference architecture compatibility

The fabric your AI tenants expect. Modules ship NCP-eligible — not bolted-on. Tier-2 GPUaaS buyers shortlist on whether the enclosure plays cleanly with the reference architectures their customers have already standardized on.

COMPUTE FABRIC

Quantum-2 InfiniBand

Spectrum-X Ethernet alternate. BlueField-3 DPUs validated.

STORAGE VALIDATED

VAST · WEKA · NetApp

Reference designs validated for AI training and inference workloads.

SERVER PLATFORMS

H100 · H200 · GB200 · Blackwell

AMD MI300 also supported. Vendor-agnostic IT — bring your own GPUs.

Configurations

Sized to your project. The productized 2 to 5 MW cluster is the Year-1 wedge; phased deployments use the same cluster as the building block.

EDGE / SINGLE MODULE

20 to 80 kW per module

ISO-format containers for ROBO, federal edge, cell-tower-adjacent, and bridge capacity below 100 kW. EMP-shielded variants available for federal and tactical-edge programs.

PRODUCTIZED CLUSTER — DLC

2 to 5 MW · 150 kW/rack

Direct-to-chip liquid cluster for current-generation Blackwell-class training. Substation, generator, switchgear, and SCADA on the same schedule. Site-ready in 6 to 8 months from PO.

PHASED DEPLOYMENT

Up to 50 MW phased

Pilot to full campus in three to five commissioning events. CapEx aligned to revenue. GPU generations change every 18–24 months — Phase 1 runs current hardware; Phase 3 specs the next generation.

Compliance and standards

ELECTRICAL

Switchgear (LV)	IEC 61439 TTA · UL 891 parallel
Switchgear (MV)	IEC 62271 · IEEE C37 parallel
Grounding	IEEE 142 · IEC 60364
Arc flash	NFPA 70E · IEEE 1584

MECHANICAL AND STRUCTURAL

Container	ISO 668 · ISO 1496
Seismic	ASCE 7 site-specific class
Wind / snow	Per local building code
Fire	NFPA 75 · NFPA 110 · NFPA 70

Lead time and delivery

Factory completion	24 to 34 weeks from PO and approved engineering
Site-ready (incl. delivery + commissioning)	6 to 8 months
Comparable US-incumbent equipment lead time	50 to 100+ weeks (Schneider, Eaton, ABB, Siemens direct)
Comparable stick-built data hall	18 to 24 months
Ships from	Pascagoula, MS (US final-mile assembly path available 2026.5)
FAT	Factory-tested at full load before shipment
SAT	On-site recommissioning per documented SAT package

REQUEST A LEAD-TIME QUOTE

teraplex.us/quote-lead-time

One business day to confirm available delivery window for your need-by date.

TALK TO ENGINEERING

teraplex.us/schedule-se-call

Multi-MW deployments, integrated power+compute scopes, federal/SLED programs.

Specifications subject to change. All claims are honest at time of publication; individual configurations may vary. Compatible with NVIDIA reference architectures noted; NCP eligibility verified per release. Equipment lead-times referenced are for typical Q2 2026 conditions and are tracked quarterly at teraplex.us/lead-time-tracker.