HERON

Heron Power Electronics
www.heronpower.com

A Blueprint for 800 VDC Datacenters

Standardize the Block, Accelerate the Build:
12 MW IT Block for the AI Era

May 2026

Executive Summary

AI data center construction is increasingly constrained not by GPU supply, but by the electrical equipment and labor required to deliver power to the GPU. The conventional path from medium-voltage (MV) to GPU was designed for a world of lower-density racks, requires four conversion steps, four separately procured equipment categories, and four sequential installation phases. Each adds cost, lead time, and energy losses. Together they add months to construction schedules and significant avoidable expense.

Heron Link simplifies the conversion architecture to two steps. A 4.2 MW solid-state transformer with integrated battery buffer replaces the MV transformer, UPS, switchgear, and PDU. The requirement to have an electrical room (gray space) disappears with them. This blueprint compares a traditional 480 VAC, 12 MW data hall with the Heron Link 800 VDC design. Hyperscale AI builders will see the following benefits:

FIGURE 1 800 VDC Blueprint performance vs. 480 VAC

65% $
lower MV-to-rack electrical cost

90%
lower MV-to-rack electrical installation labor

25%
faster construction schedule

50%
lower power delivery losses

© 2026. Heron Power Electronics Company.

1 Why Conventional Architecture Struggles at AI-Scale

As AI rack densities climb toward megawatt scale, the constraint is no longer compute, but how power is delivered to it. Rack power is rapidly increasing from tens of kilowatts to hundreds of kilowatts, with roadmaps pointing toward 1 MW+ per rack within this decade. Conventional 480 VAC architectures were not designed for this level of concentration. At these power levels, the current requirement becomes prohibitively high, driving rapid growth in conductor size, busway capacity, breaker ratings, and thermal management requirements. The result is a system that becomes physically larger, more complex, and increasingly difficult to deploy as power density rises.

> “
> *At AI-scale, the problem isn’t compute, it’s moving power.*

At the same time, the traditional multi-step AC power chain introduces structural inefficiencies that are amplified at AI scale. Power must pass through multiple conversion steps, from MV AC to low-voltage AC, to DC, back to AC, and finally to low-voltage DC at the rack, with each step adding loss, cost, and commissioning complexity.

At AI scale this translates into megawatts of lost capacity, longer construction schedules, and tightly coupled installation dependencies across specialized equipment. Meanwhile, highly dynamic AI training workloads introduce rapid load fluctuations that stress upstream infrastructure and complicate grid interconnection. Utilities are increasingly requiring AI campuses to provide fault ride-through (FRT) capabilities to support the grid rather than acting as passive loads. Together, these factors make the conventional 480 VAC architecture not just inefficient, but fundamentally misaligned with the performance, speed, and grid requirements of modern AI infrastructure.

FIGURE 2 Factors constraining traditional 480 VAC architecture[1]

5 discrete components

- MV Transformers
- LV Switchboards
- UPS Units
- Static Transfer Switches
- Power Distribution Units

46 week lead time

- **46-48 weeks — Switchgear**
- 20-28 weeks — UPS System
- 16-20 weeks — MV Transformer
- 14-18 weeks — PDU/Distribution

6.4% power wasted

- 1.0% — MV Transformer
- 3.0% — UPS
- 2.4% — Rack PSU

© 2026. Heron Power Electronics Company.

---

[1] Switchgear lead times from Cushman & Wakefield, Data Center Development Cost Guide, 2025.

2 The Blueprint’s Approach

NVIDIA’s Oberon and Kyber rack architectures both specify 800 VDC as the rack input voltage. Most facility utility feeds deliver MV AC (34.5 kV) from the substation next to the data hall. A single conversion stage from MV AC to 800 VDC is the most direct path to achieving this higher input voltage while minimizing losses and streamlining equipment architecture. This approach eliminates entire equipment categories with long lead times, high labor intensity, and significant inefficiency. Construction crews can begin commissioning from the first day equipment arrives on site with a simplified electrical chain.

FIGURE 3 Comparison of Traditional vs. Heron Link

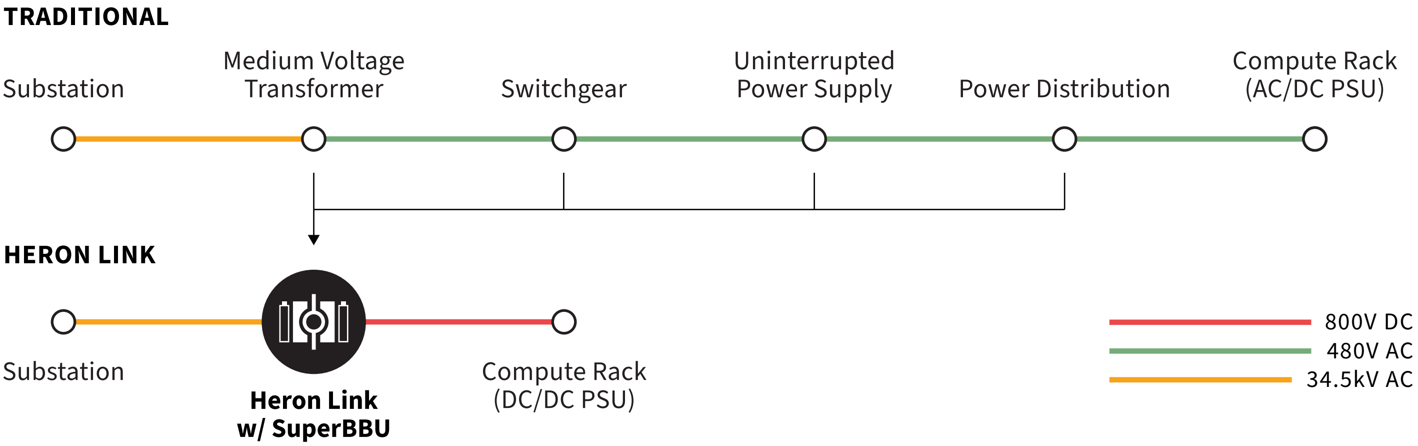

© 2026. Heron Power Electronics Company.

> “
> *A single Heron Link collapses a complex collection of distributed components into one elegant solution, delivering megawatt-scale power with fewer steps, higher efficiency, and no single point of failure*

→ Heron Link: Highly Efficient with Integrated Protection

The Heron Link is a 4.2 MW solid-state transformer (SST) that connects directly to 34.5 kV and delivers 800 VDC output with over 98.5% conversion efficiency. It integrates MV protection, an MV disconnect switch, and a 40-channel DC power distribution panel with per-channel overcurrent protection. Its modular, fault-tolerant architecture eliminates single points of failure, enabling higher system availability and uptime. The Link's power converter modules are field-replaceable via blind-mate connectors, allowing swaps to be completed in minutes rather than hours. A single Heron Link provides 2.5 times the power density per square foot when compared to the traditional equipment it replaces.

→ SuperBBU: Stable Grid Draw, Built-in Backup

A rack training a frontier model may swing between 30% and 150% of its thermally designed power within milliseconds, and because all racks in a cluster step through the same training iterations in lockstep, these swings are facility-wide. Left unmanaged, the entire campus load pulses in synchrony, creating severe flicker concerns and threatening utility interconnection agreements.

The SuperBBU addresses this at the source. A high C-rate battery, skid-integrated with each Heron Link, couples directly to the 800 VDC bus without an additional DC-to-DC converter. This design responds faster than any existing ripple mitigation solutions, absorbing rack-level power fluctuations and presenting a smooth, predictable load to the grid. In addition, it provides 30 seconds of full-load backup power during a grid disturbance, matched to the ~15-second start time of generators connected at the MV level. For applications requiring longer hold-up or utility-imposed ramp rate control, the SuperBBU can be substituted with long duration energy storage via external DC-DC coupling.

FIGURE 4 AI ripple mitigation with Heron Link and SuperBBU

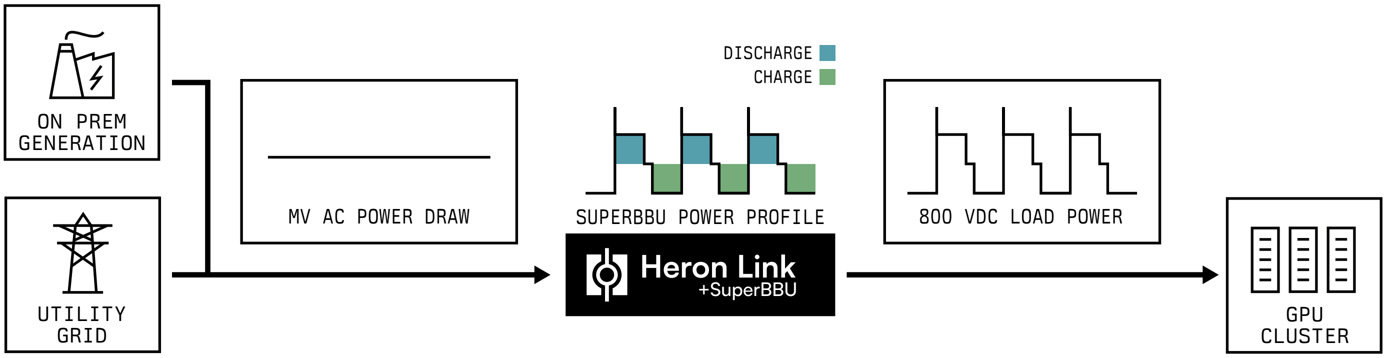

© 2026. Heron Power Electronics Company.

Together, a Heron Link with SuperBBU replaces the MV transformer, low voltage (LV) switchboard, UPS system, and power distribution unit. Four equipment categories with the longest lead times and highest installation costs are replaced by one with a single 24-week procurement lead time. The Heron Link will be built in the United States at a facility with 40 GW of annual production capacity, supporting hyperscale deployment without supply-chain risk.

3

12 MW Data Hall Building Block

For the sake of easy comparison, this blueprint proposes a 12 MW data hall using 4 to make 3 redundancy design as a repeatable building block; powered by four Heron Links. The block can be replicated to reach campus-scale capacity, eliminating custom engineering on each project and giving procurement, logistics, and installation teams a repeatable playbook.

Each of the 40 compute racks in the block receives power from all four Heron Links simultaneously through four independent power shelves. If any Heron Link fails, each rack continues operating on the remaining three shelves without interruption. Redundancy is achieved through topology, not by additional equipment layers. The following table contrasts the MV-to-rack equipment stack in a conventional 480 VAC[2] build against the Heron Link 800 VDC blueprint:

TABLE 1 MV-to-rack electrical equipment comparison

Equipment between MV feeder and IT rack	480 VAC (Traditional)	800 VDC (Heron Link)
Heron Link SST with SuperBBU (4.2 MW, 34.5 kV → 800 VDC)	0	**4**
MV Transformer	4	**Eliminated**
LV switchboard	4	
UPS system with battery cabinets	12	
PDU	4	
Remote Power Panel	20	

© 2026. Heron Power Electronics Company.

Each Heron Link includes an integrated power distribution panel with 40 individually overcurrent-protected outputs, meaning that no additional switchgear or PDU is required anywhere in the design. Units sit outdoors, adjacent to the data hall. The electrical room, along with the months of work required to build, wire, and commission it, is eliminated entirely. Power reaches each rack through a single direct cable per power shelf routed via overhead cable tray, using touch-safe 800 VDC whips. CDU and networking racks are also powered by 800 VDC. These loads are connected via a prefabricated trunk-bus harness; a distribution method with a proven track record in utility-scale solar deployments.

Together, whip-based rack feeds and trunk-bus harnesses eliminate the cost and complexity of traditional busway systems. No specialized labor is required, harnesses ship pre-terminated, and the combined approach reduces copper utilization by 30%, cutting both material and installation costs at scale.

---

[2] Schneider Electric, Eco Struxure Reference Design 113: 10.2 - 12.7 MW, Tier III, ANSI, Chilled Water, Liquid-Cooled AI Clusters (NVIDIA Vera Rubin).

FIGURE 5 12 MW Data Hall building block: 4 outdoor Heron Link SSTs serving 40 compute racks across 4 rows

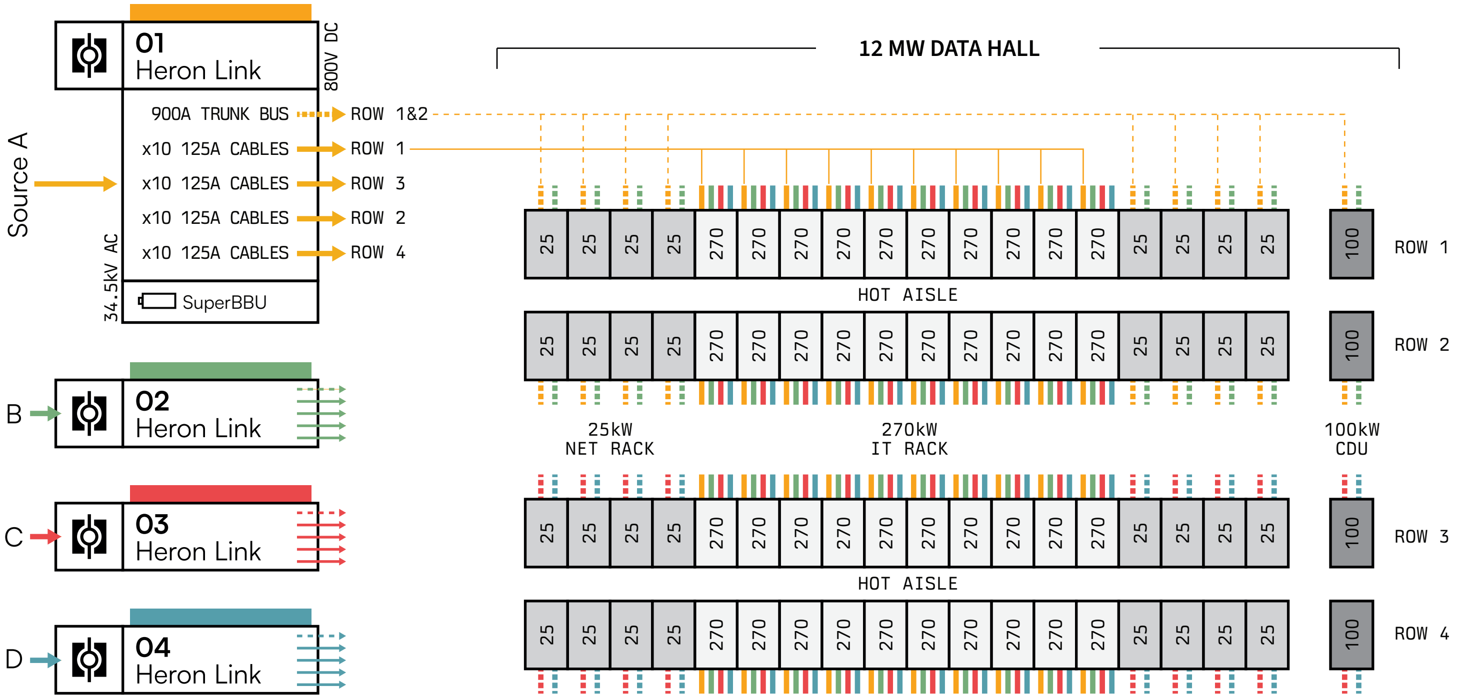

© 2026. Heron Power Electronics Company.

TABLE 2 Data hall block specifications

Specification	Value
Redundancy level	Compute: 4-to-make-3 CDU & Networking: 2N
Rated power of compute racks	up to 270 kW
Number of compute racks	40 (10 racks x 4 rows)
Total capacity of compute racks	10,800 kW
Total capacity of networking racks (800 VDC)	800 kW
Total CDU rated power (800 VDC)	400 kW (100 kW x 4)
Total Data Hall 800VDC load	12,000 kW

© 2026. Heron Power Electronics Company.

4 Cost, Schedule, and Efficiency Gains

75% Lower Electrical Cost

Reducing the complexity and footprint of electrical equipment lowers cost. With four equipment categories eliminated, total MV to rack electrical cost drops 75%, from roughly $1.77M/MW to $0.42M/MW, a saving of $1.35 million per MW installed.

The labor savings are even steeper than equipment savings: 90% versus 65%. Electrical labor falls from ~$0.73M/MW to ~$0.08M/MW. That asymmetry reflects the staggering complexity in a conventional build. UPS commissioning, switchgear installation, and PDU field wiring each require specialized trades working in strict order. Removing those steps saves not just the cost of the equipment but also collapses the commissioning dependency chain that adds months to the schedule.

TABLE 3 Cost breakdown in Millions per MW of compute capacity

$1.35M saved/MW of compute

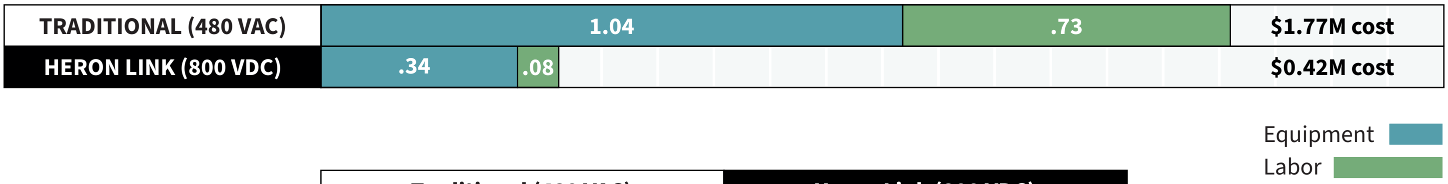

	Traditional (480 VAC)			Heron Link (800 VDC)			
Component	**Equipment**	**Labor**	**Total**	**Equipment**	**Labor**	**Total**	**Net Savings**
UPS System	0.60	0.40	1.00	*Eliminated*			1.00
MV Transformer	0.05	0.02	0.07				0.07
LV Switchgear	0.20	0.10	0.30				0.30
Grey Space Shell	0.04	0.03	0.07				0.07
Wiring (grey space & RPP)	0.02	0.03	0.05				0.05
Commissioning	—	0.05	0.05				0.05
Busway/Cables	0.13	0.10	0.23	0.09	0.07	0.16	0.07 (30% reduction in Copper)
Heron Link w/ SuperBBU	—	—	—	0.25	0.01	0.26	—
Total Cost	1.04	0.73	1.77	0.34	0.08	0.42	1.35 saved

All values in $Millions per MW of compute capacity
These figures are conservative. Based on an industry survey, cost totals are between $1.5–$3M/MW.

© 2026. Heron Power Electronics Company.

→ 3 Months Faster to 1st Compute

By removing four equipment categories, Heron's blueprint cuts $0.65M/MW in field labor associated with MV to rack power delivery. A conventional 480V, 12 MW data hall with grey space and electrical fit-out takes approximately 14 months to build, with warm-shell labor running approximately $2.50M/MW[3]. The $0.65M/MW saving represents 25% of that labor cost. Assuming labor reduction maps proportionally to schedule compression, removing 25% of field labor from a 14-month schedule recovers approximately three months.

> “
> *Heron Link and SuperBBU are NEMA 3R–rated for outdoor installation, eliminating the need for a dedicated electrical room.*

---

[3] Total facility labor ~$2.50/W, from McKinsey's 25% aggregate labor rate applied to a ~$10/W warm shell baseline
(refer to https://www.leanrs.com/insights/data-center-equipment-growth and https://www.mckinsey.com/industries/industrials/our-insights/beyond-compute-infrastructure-that-powers-and-cools-ai-data-centers)

→ 4MW Recovered at 100 MW Scale

Every conversion stage in the power chain loses some energy as heat. The conventional 480 VAC chain loses 6.4% of input power before it reaches the compute tray. The Heron blueprint's two-stage chain reduces losses by over 50%, to only 3.0%[4]. At 100 MW of grid input, this blueprint can deliver 3.4 MW more electricity directly to GPUs rather than dissipating in distribution equipment. That is 3.4 MW of additional GPU capacity from the same utility contract.

> “
>
> *Parallel commissioning: each SST is commissioned independently as it arrives, enabling earlier system readiness.*

5 | Reliability and Grid Interaction

Fail Operational Redundancy

Heron Link's power converter architecture is designed for fail-operational redundancy, ensuring continuous operation even in the presence of component failures. Each 4.2 MW Link consists of thirty 140 kW power cells, with ten cells connected in series per phase to achieve a 34.5 kV output. The system can tolerate up to two power cell failures per phase. Fast-acting bypass devices with sub-cycle response times automatically isolate and bypass faulted cells. Each cell failure reduces power handling capacity by 3.3%, but full output voltage and grid fault ride-through compliance is preserved.

This fail operational architecture allows flexibility in designing the redundancy architecture of the MV to LV power train. Depending upon the application and uptime requirements, N+1 redundancy can be reduced to N+0 where each compute rack is fed by only N+0 Heron Links. A power cell failure results in compute racks de-rating their performance proportionally. NVIDIA's Vera Rubin GPU is already capable of limiting its power draw via software controls (Max-P and Max-Q modes). This architecture provides a compelling alternative to the traditional approach of overbuilding MV to LV conversion infrastructure.

> “
>
> *Failures are inevitable; cascading failures are not. Heron Link contains faults and keeps the system running.*

→ Faults Stay Small and Get Cleared Fast

Each Heron Link includes a solid-state circuit breaker (SSCB) per rack output that limits fault let through current to less than 2 kA and clears faults in under 5 microseconds. A fault on any output is isolated before it can propagate, bounding failures to a single power shelf, a single panel output, or a single Heron Link unit. A mechanical contactor provides galvanic isolation after clearing.

On the MV side, UL-listed current-limiting fuses and an integrated disconnect rated for 40 kA at 34.5 kV coordinate with upstream utility protection without a dedicated switchgear enclosure. The 4-to-make-3 topology ensures that even a complete Heron Link failure leaves every rack running on its remaining three power shelves.

---

[4] Power delivery efficiency from medium-voltage input to compute load. Traditional chain: MV transformer (99%) × UPS (97%) × AC/DC PSU (97.5%) = 93.6%. Heron chain: SST + SuperBBU (98.5%) × DC/DC PSU (98.5%) = 97.0%. Net: 53% reduction in distribution losses.

A Stronger Interconnection Position

As AI campus loads grow to hundreds of megawatts on single feeders, grid operators are raising the bar for interconnection. ERCOT's Large Electrical Load framework, which is increasingly a reference point for other markets, requires large loads to demonstrate fault ride-through capability and stable, controllable power draw. Facilities built on conventional architecture cannot meet these requirements without significant additional equipment.

The Heron Link with integrated SuperBBU addresses these requirements natively. The SuperBBU absorbs AI training load transients directly on the 800 VDC bus, preventing ripple from propagating upstream and reducing thermal stress on MV cables, switchgear, and protection equipment. During grid disturbances, the Heron Link's grid-forming capability enables seamless fault ride-through while providing sufficient hold-up time for backup generator start-up or feeder reconfiguration.

TABLE 4 Comparison of grid support capabilities

Capability	Traditional 480 VAC + UPS	800 VDC: Heron Link + SuperBBU
AI training power fluctuations	Pass through to grid	<1% ripple on MV AC
Fault ride-through	None	IEEE 2800, ERCOT NOGRR 282
Grid Interaction	Passive Load	Active, grid-forming capable
Hold-up duration	5 minutes	40 seconds in 4-to-make-3 config 2 minutes controlled ramp-down

© 2026. Heron Power Electronics Company.

6 | Built for Today, Ready for Kyber: No Stranded Assets Tomorrow

One of the most significant risks in data center construction today is building power infrastructure for current rack densities only to find it cannot serve the next generation of rack density without costly rebuilds. This blueprint eliminates that risk by design.

A data hall built to this blueprint in 2027 will serve NVIDIA's 1 MW+ Kyber racks with no changes to the MV feeders, Heron Link units, or protection system. The Link's power whips can be reallocated to higher power density racks through row-level combiners or mechanical adapters. Looking further ahead, the same first principles that drove the industry from 48 to 800 VDC are pointing toward 1500 VDC as the next preferred distribution voltage for the ultra dense compute racks. Heron Link is natively designed to support 1500 VDC applications, thereby ensuring capital deployed today remains relevant through multiple generations of compute.

7 | Five Advantages From One Architectural Decision

Reducing the power delivery chain from four conversion stages to two produces compounding benefits across every dimension that matters for large-scale AI infrastructure:

- 65% lower electrical equipment cost and 90% lower installation labor, driven by eliminating four equipment categories with the longest lead times and highest field complexity
- Shaving three months off construction timelines with a single 24-week procurement replacing multiple sequential long-lead items
- Halving distribution losses and recovering 3.4 MW at 100 MW scale that would otherwise be dissipated as heat
- Stronger interconnection position, with built-in fault ride-through, grid-forming capability, and load smoothing via the SuperBBU
- No stranded assets: the same infrastructure serves today’s 270 kW racks and tomorrow’s 1 MW+ Kyber racks with minimal modification

The facilities that come online fastest, at lowest cost, with the strongest utility relationships will define who wins the AI infrastructure buildout. Heron’s proposed architecture is designed to be that facility. The 12 MW block is repeatable, the equipment list is fixed, and the schedule is three months shorter. The only question is how long a builder waits to design around it.

For detailed floor layouts, wiring diagrams and modeling contact Heron Power: product@heronpower.com