MARITIME FUSION

SHIELD

Superconducting High Integrity Energy Link & Distribution. A New Age of Lossless Power Delivery

From fusion magnets to power distribution

Maritime Fusion was founded to solve the hardest problem in energy — building a commercially viable fusion reactor. That mission meant engineering High Temperature Superconductor (HTS) cables operating at extreme current densities and magnetic fields.

SHIELD is the direct application of that expertise to the most urgent power infrastructure problem of the decade: moving gigawatts of clean, dense, reliable power into AI datacenters.

Where copper has hit hard physical limits, HTS opens a new design space. SHIELD carries 8,000 amps in a cross-section smaller than a tennis ball, with effectively zero resistive loss. It moves heat generation, power conversion stages, and failure modes out of the white space — allowing for substantially higher compute density.

SHIELD integrates with today's 800V DC reference architectures, and enables a fundamental redesign of AI rack power for tomorrow.

The case is economic as much as physical. SHIELD eliminates distribution losses, cuts white-space cooling load, and reclaims rack volume worth far more than the cable itself — and it pays for itself in well under three years, with returns that only sharpen as rack densities and energy costs climb.

> SHIELD carries 8,000 amps in a cross-section smaller than a tennis ball.

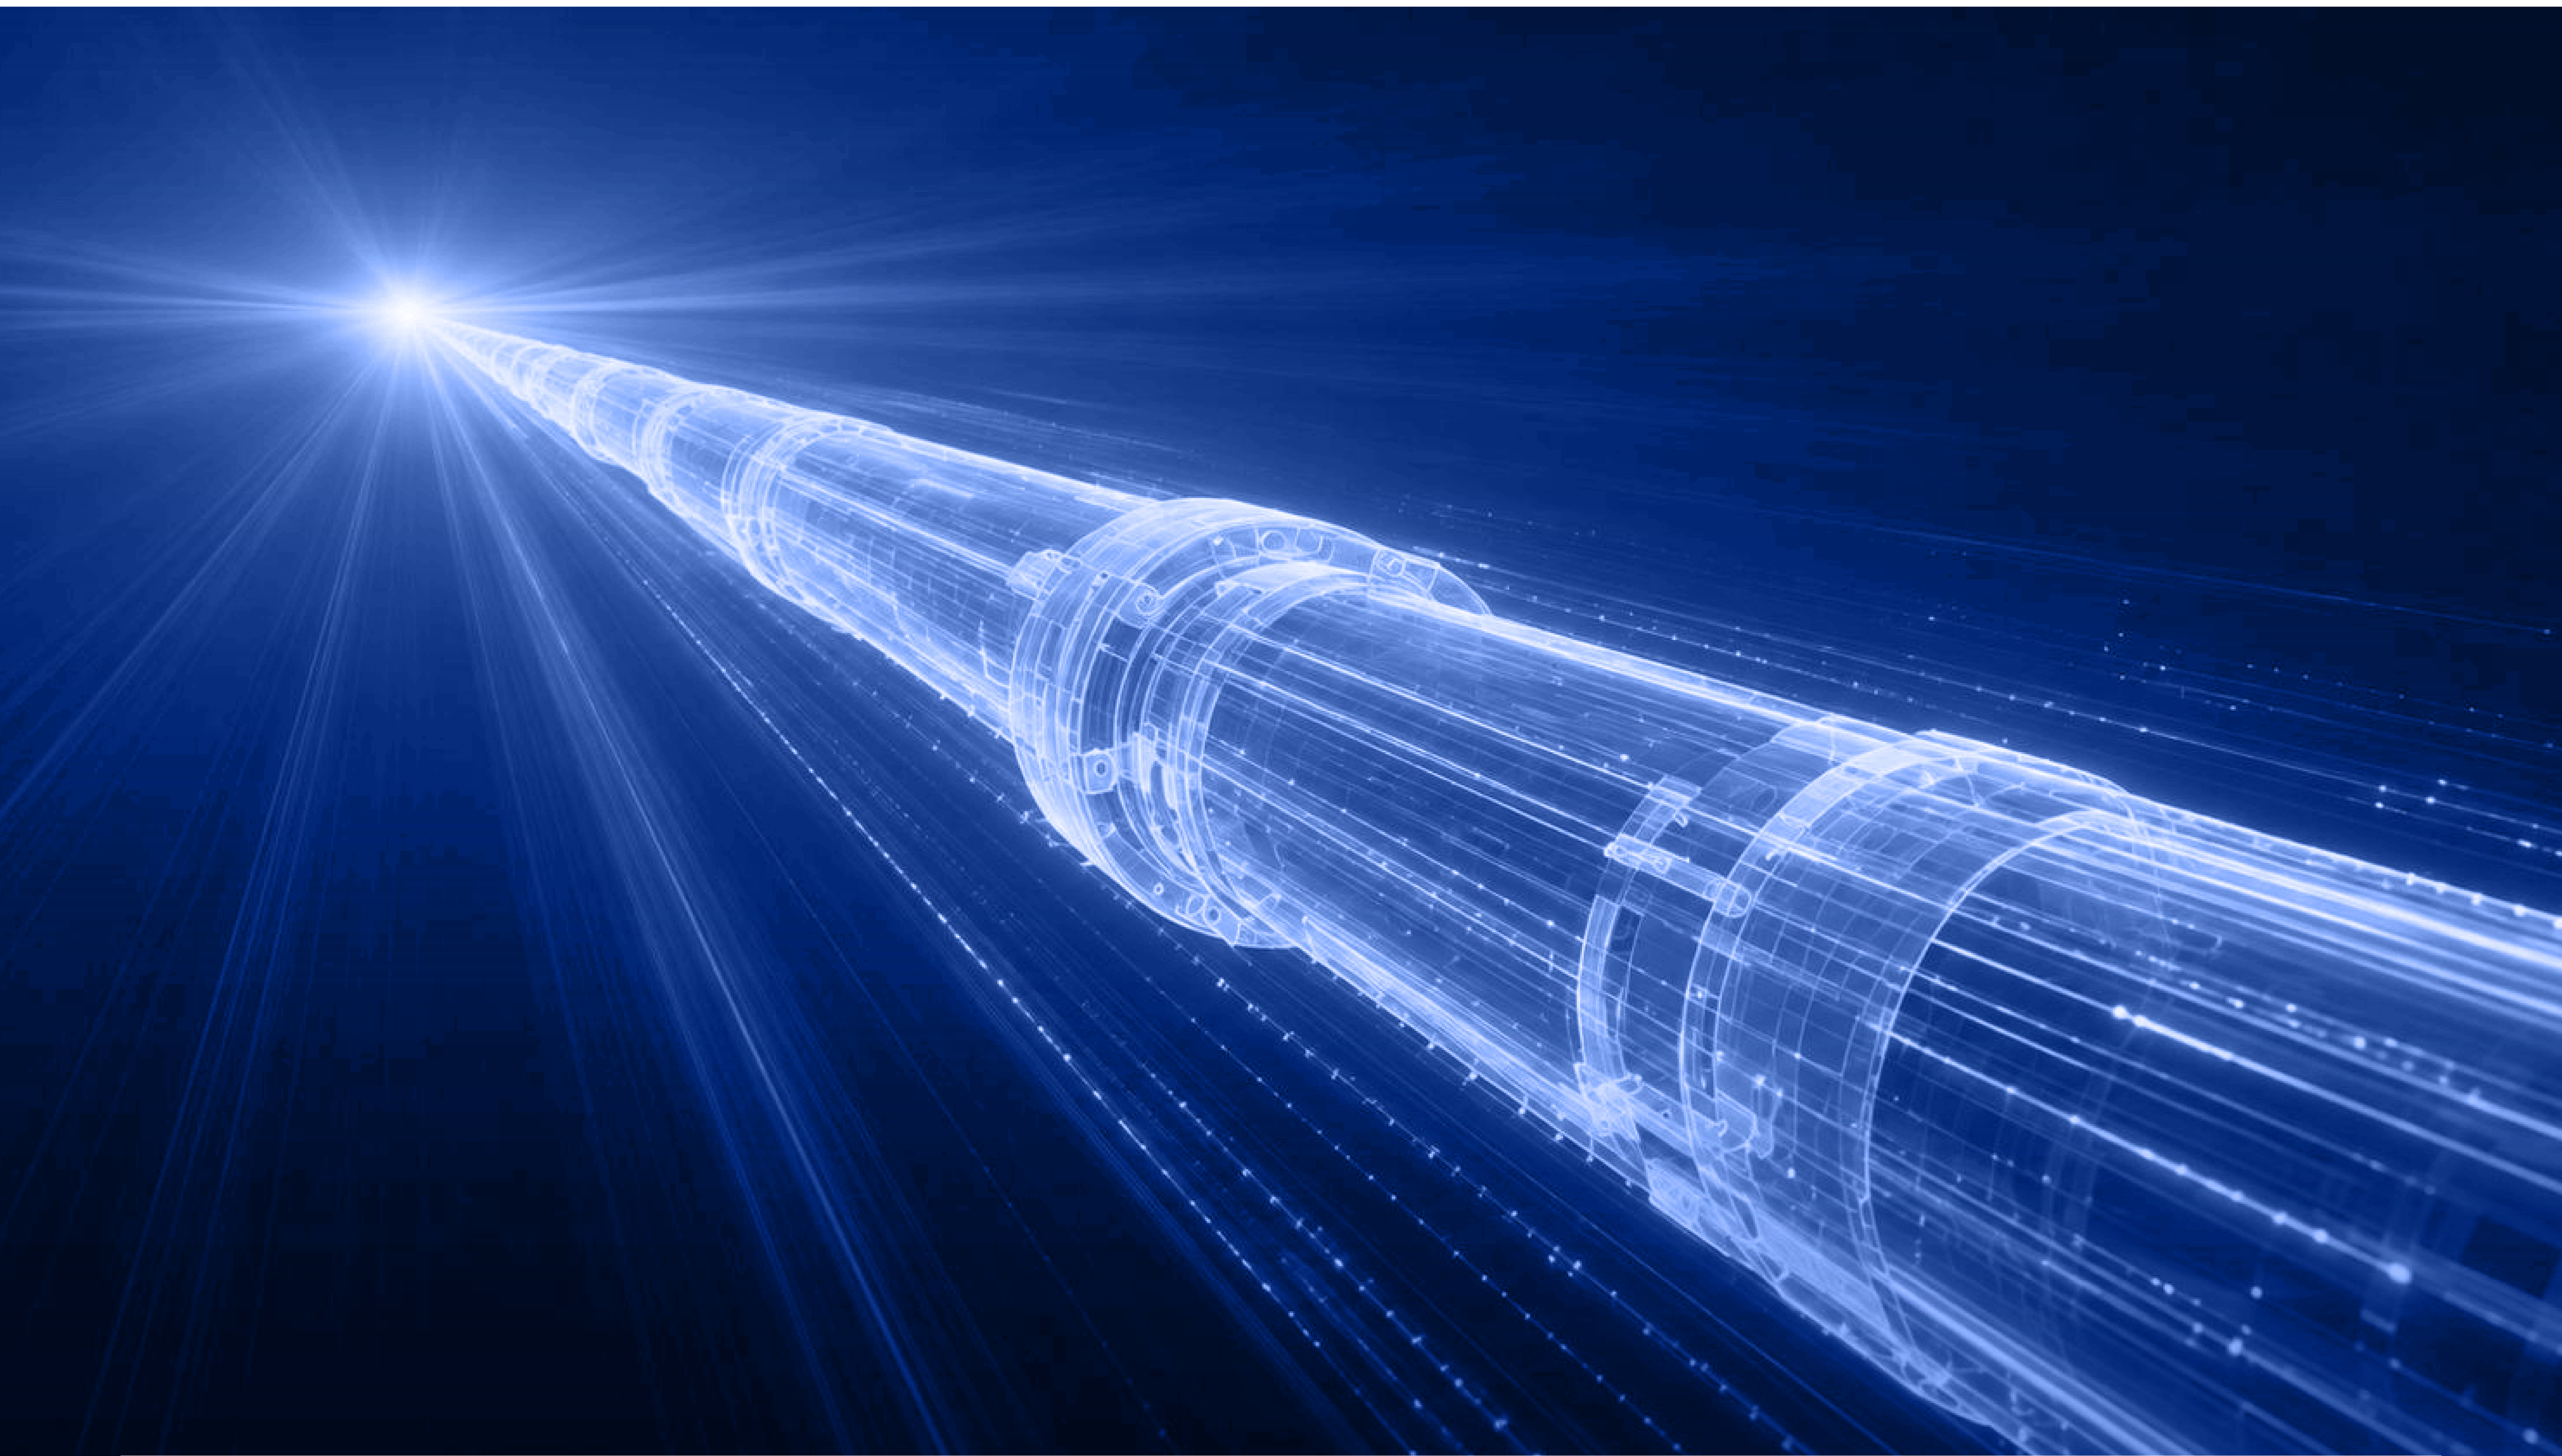

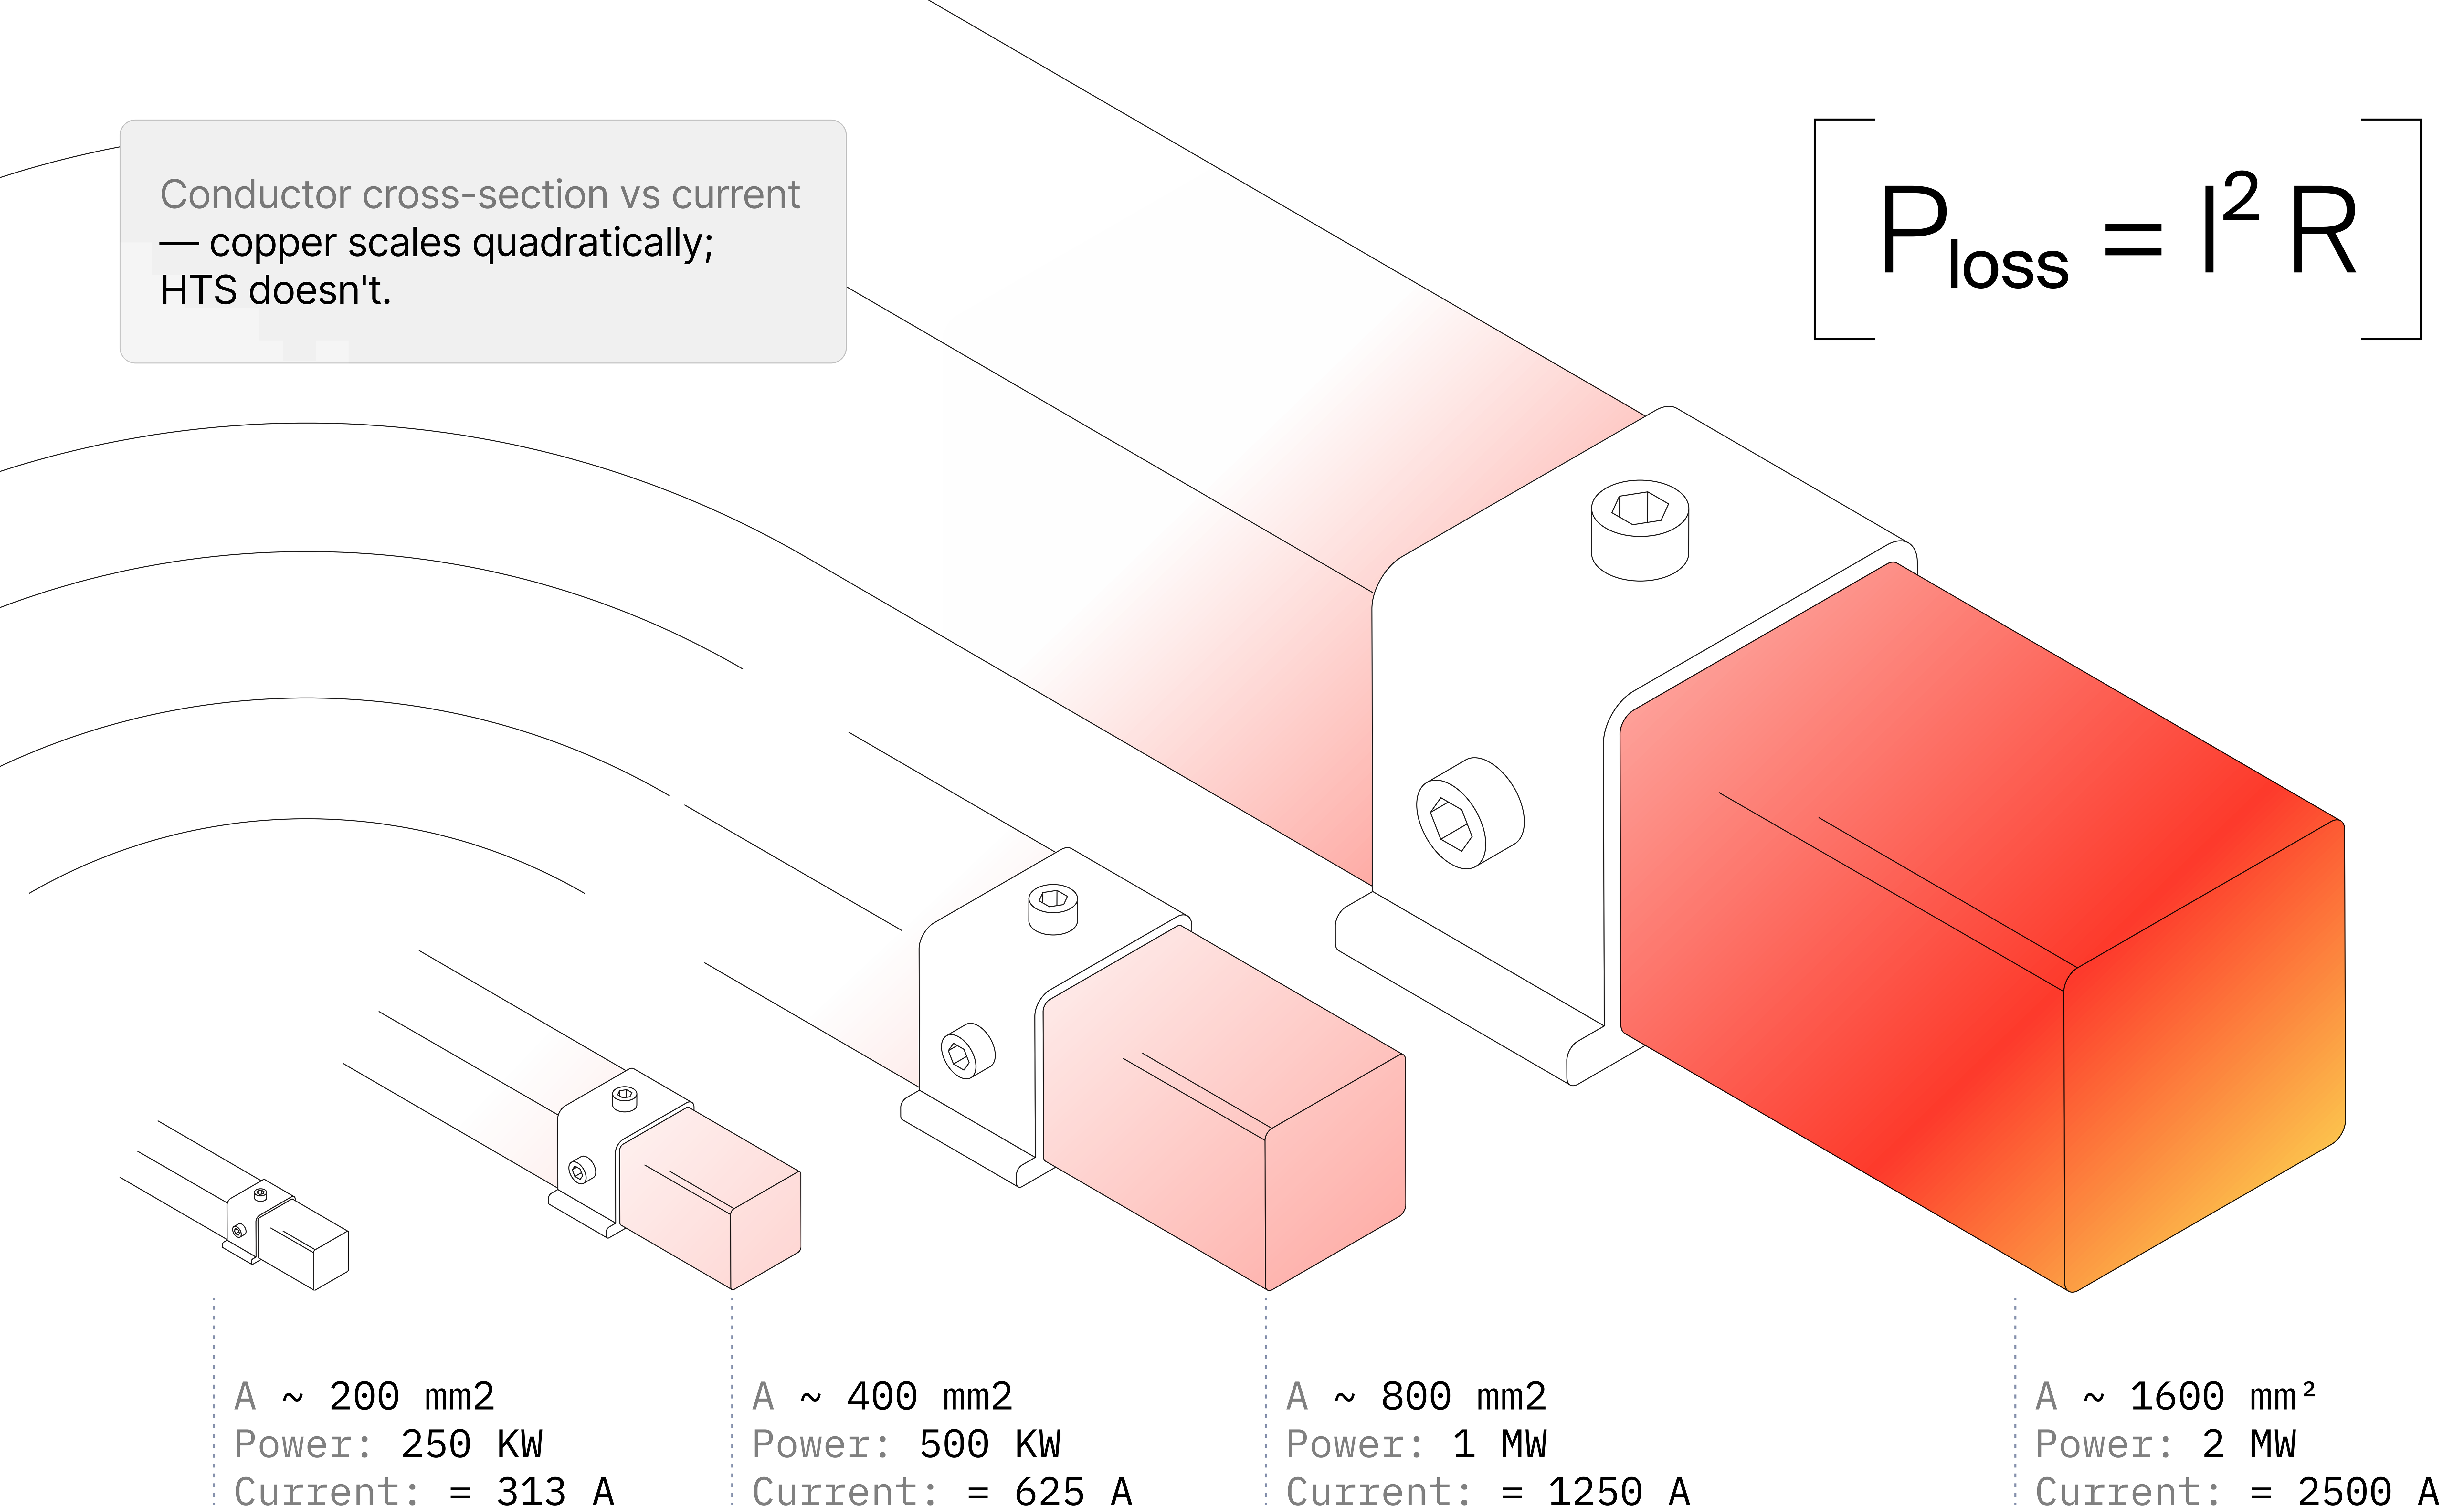

Conductor cross-section vs current — copper scales quadratically; HTS doesn't.

$$P_{loss} = I^2 R$$

Copper Hits a Wall

> Voltage delays the I^2R wall. It doesn't escape it.

Copper has carried the world's electricity for 140 years, but the AI era is pushing it past its limits. The problem is simple: conductor losses scale as I^2R. As rack currents rise from hundreds of amps to thousands, copper has to get dramatically larger just to keep losses from growing out of control.

A 1 MW rack fed at 800V DC draws 1,250A continuously; the busbar required to deliver that current at acceptable voltage drop is now thicker than a human wrist, weighs hundreds of pounds per rack, and dissipates kilowatts of waste heat directly into the white space — heat that must then be removed by the cooling plant at additional cost.

The industry is responding by climbing the voltage ladder — from the standard 240V rack to 415V AC, 800V DC, and eventually 1500V DC — and that's the right direction. But higher voltage only delays the I^2R wall; it doesn't escape it. At 1 MW per rack and rising, even 800V distribution demands copper geometries that strain what can be installed, cooled, and serviced inside a building.

Superconductors break the curve entirely. With zero DC resistance below their critical temperature, HTS conductors carry high current densities with no resistive loss. The trade-off is cryogenic cooling — and at gigascale, that trade is overwhelmingly favorable.

Cable size comparison —
8 kA Cu busbar vs SHIELD vs tennis ball (all measurements to scale)

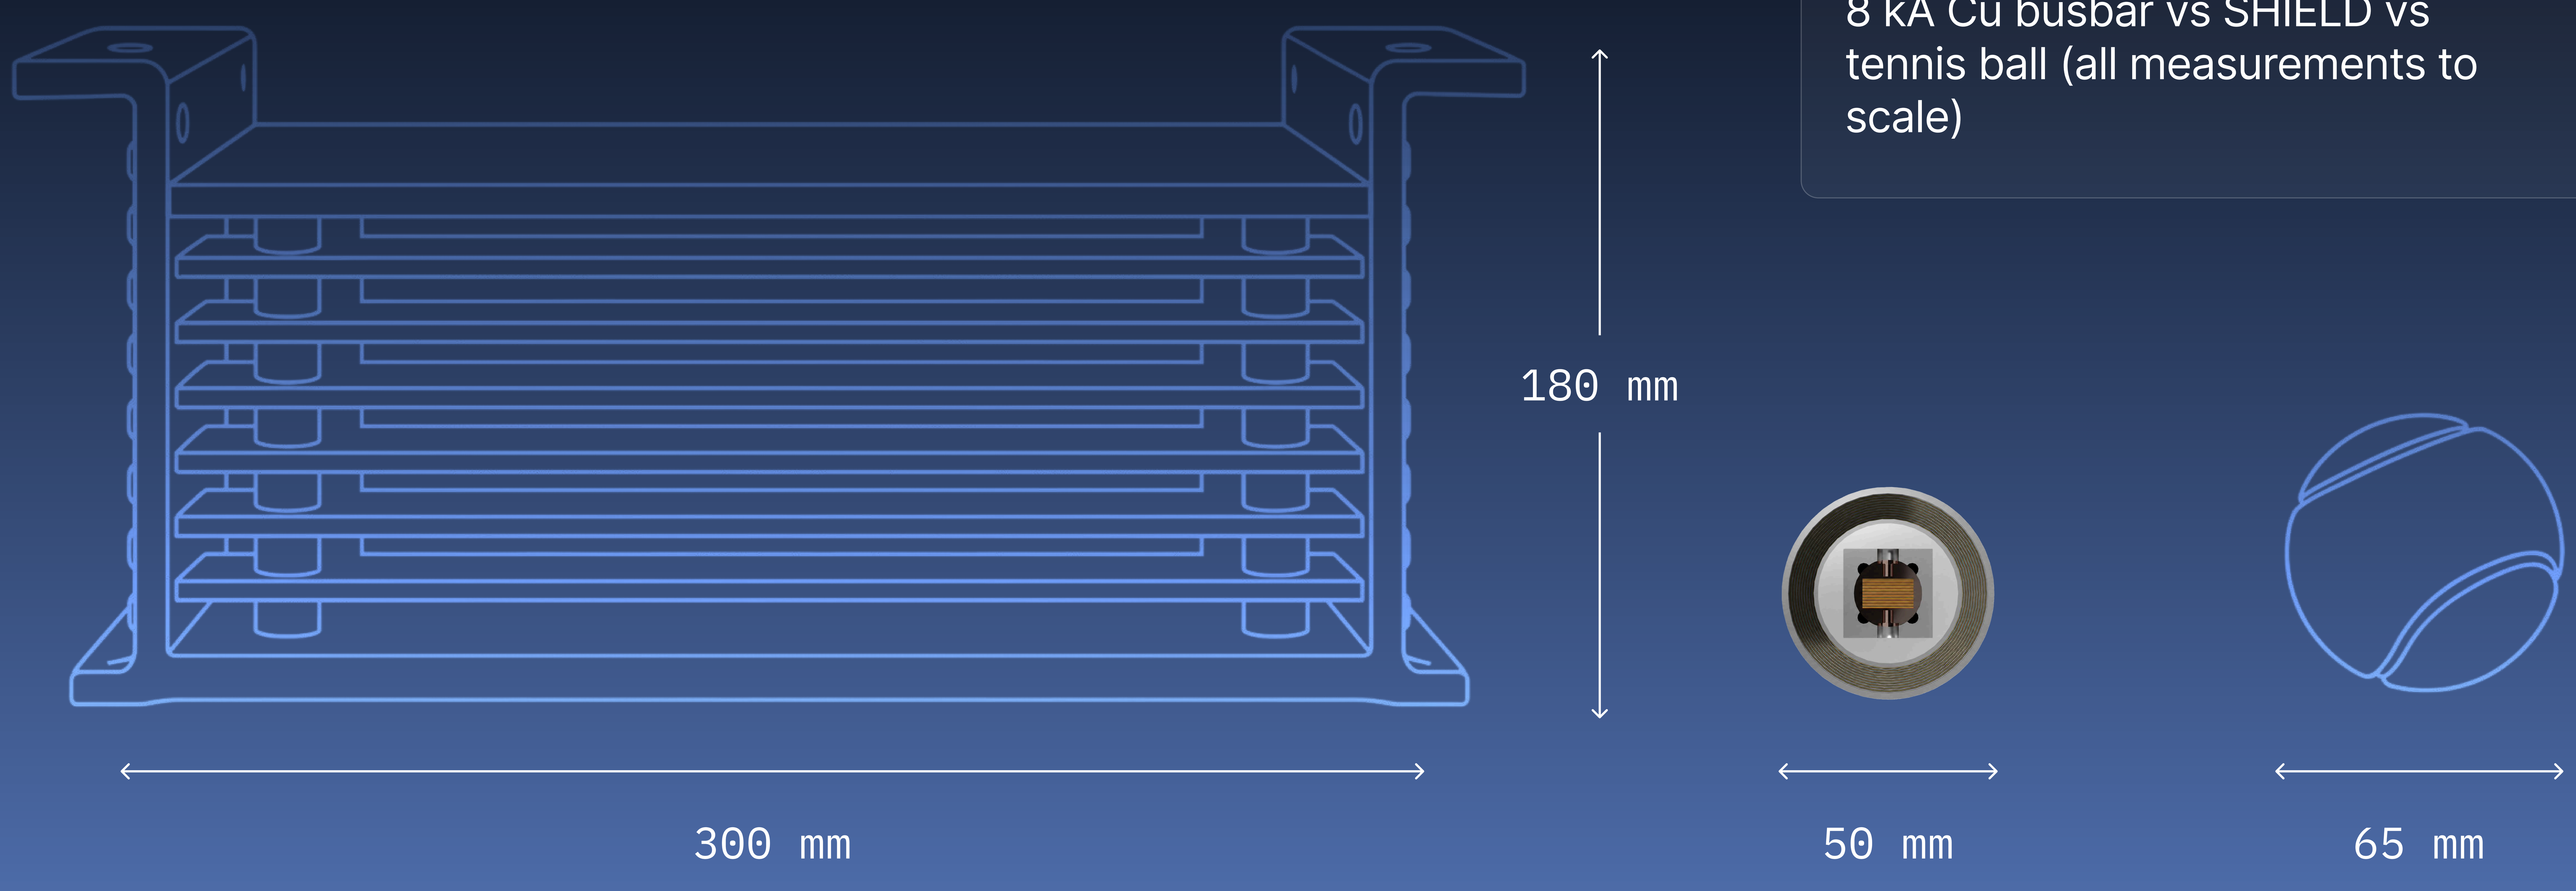

The SHIELD cable

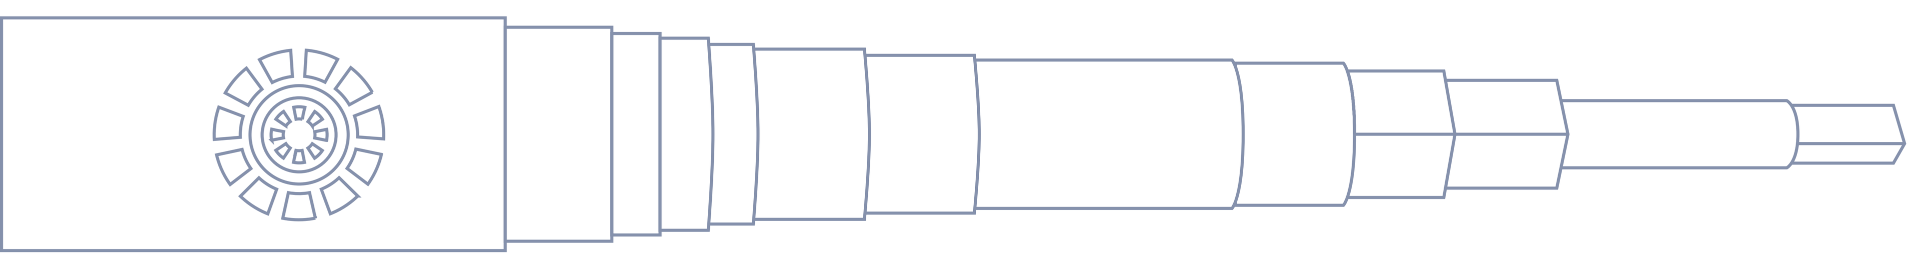

SHIELD — Superconducting High-Integrity Energy Link & Distribution — is a cryogenically cooled HTS power cable purpose-built for AI datacenters and DC power distribution.

Construction

SHIELD is built around a stack of REBCO HTS tapes soldered into a former, jacketed in stainless steel and a flexible cryostat that reduces cryogen to ambient heat leak. Touch-safe terminations land directly on standard HV DC busbars or 54V DC bus plates. The cable is field-routable at radii compatible with existing cable tray infrastructure.

System integration

A SHIELD link is delivered as a complete package: cable + cryogenic refrigeration plant + monitoring and protection electronics. The cryogenic load is small — less than 2 Watts/meter — and is offset many times over by eliminated I^2R losses and reduced HVAC load in the white space.

> “
> Current density of over 2.5kA/cm² without ohmic losses

Continuous current rating	up to 8,000 A
Outer cable diameter	≤ 67 mm (smaller than a tennis ball)
Effective current density	> 2.5 kA/cm² of cable cross-section (~50× copper)
Conductor resistive loss	effectively zero
Operating temperature	70-77 K (liquid nitrogen, closed loop)
End-to-end voltage drop, 100 m run	< 50 mV

Datacenter Layout

SHIELD enables two architectural moves, deployable in sequence.

Today: 800V DC distribution from SST to white space

In near-term deployments, SHIELD slots directly into NVIDIA's 800V DC reference architecture. Solid-state transformers, battery energy storage, and on-site solar or gas generation are co-located outdoors in a power yard. The SST converts medium-voltage AC to 800V DC at > 98 % efficiency. SHIELD then carries 800V DC from the outdoor power block, through the building envelope, to the white-space cabinets — at currents copper cannot match.

Tomorrow: 54V DC SHIELD distribution to the rack

The longer-term architecture goes further. Today, roughly half of every AI rack is occupied by power electronics that step 800V DC down to a 54V DC rack bus, then to 12V, then to
< 1V at the GPU. SHIELD makes it economical to distribute the 54V DC bus itself across the white space — at currents copper cannot sustain — and remove the rack-level conversion stage entirely.

The rack reclaims over 8U per chassis for compute, memory, or thermal hardware, while power conversion moves to a centralized, serviceable plant outside the white space. The freed volume and lower in-rack heat load also make conventional GPU liquid cooling easier to deploy.

- Copper busways and their associated losses are eliminated; a single SHIELD link replaces parallel runs of 4/0 cable
- I^2R losses drop to zero across the run, recovering 1–3 % of total facility power
- Battery storage and solar tie directly to the 800V DC bus that feeds the racks

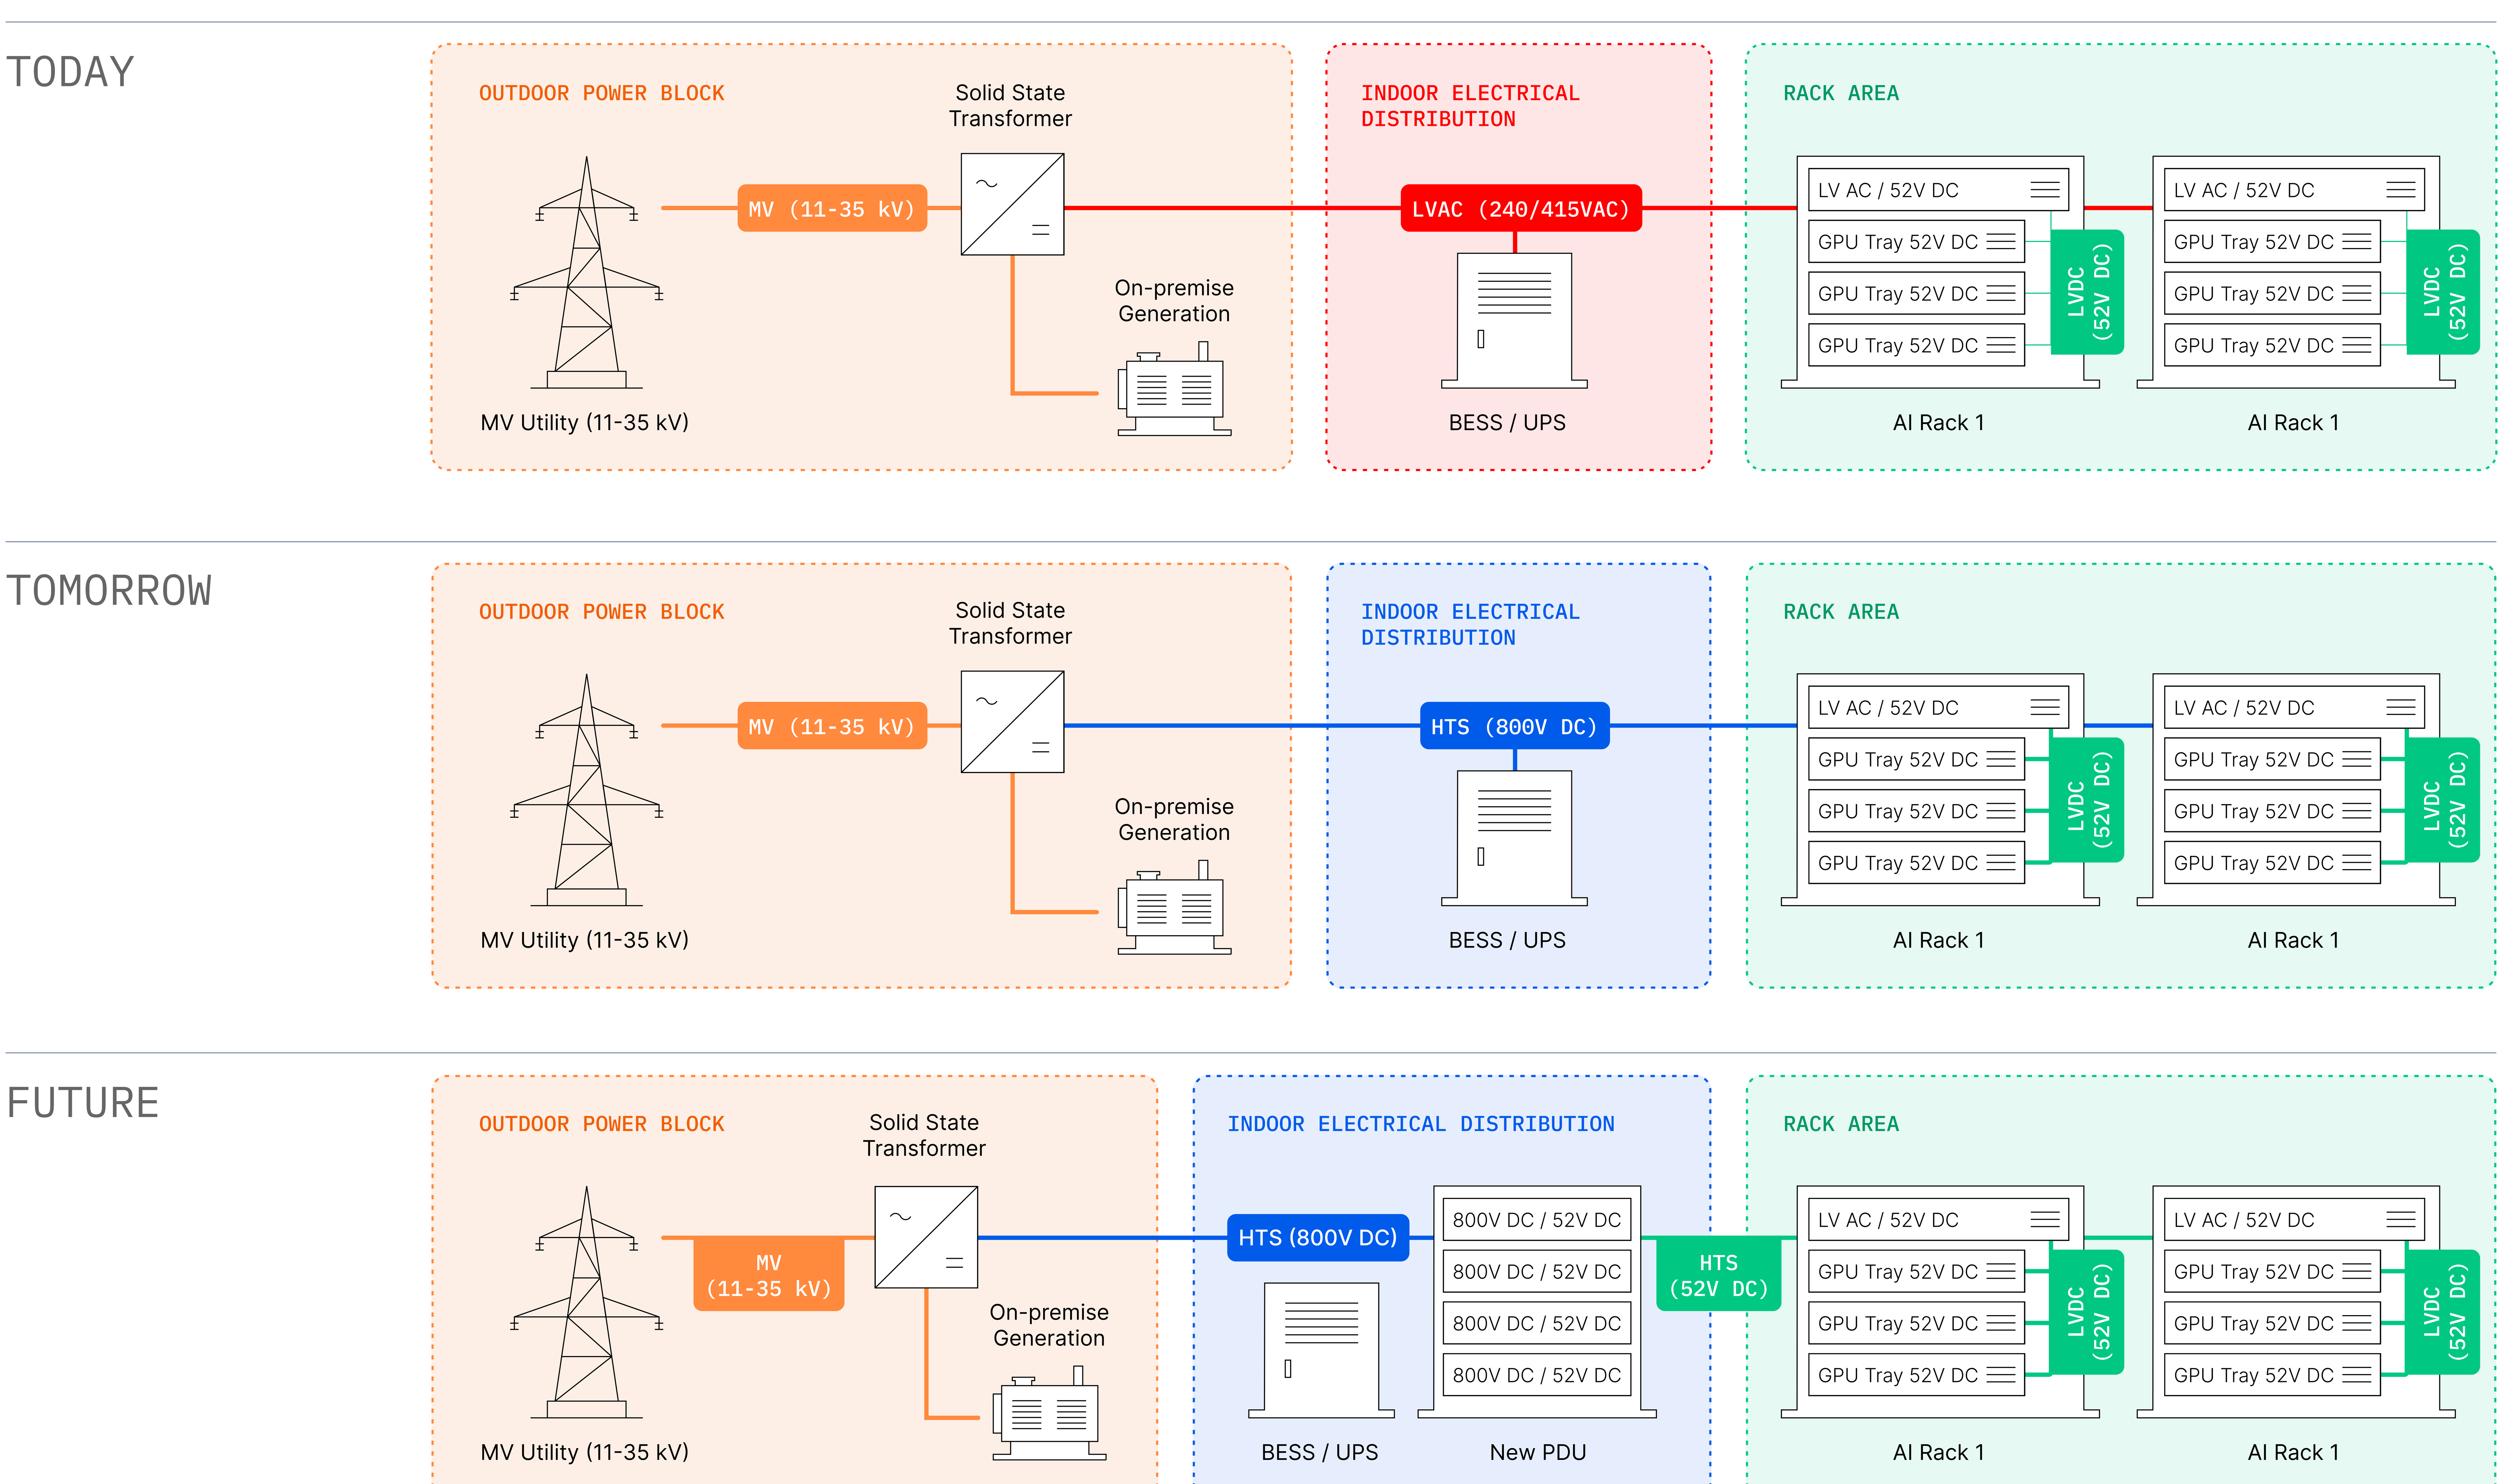

Rack architectures — left, NVL72 reference with 4+4 PSU shelves consuming 8 U at 33 kW each; right, SHIELD-fed with the 52V DC bus arriving directly from a remote conversion cabinet. The recovered 8 U is available for additional compute trays, in-rack CDU capacity, or accelerator memory.

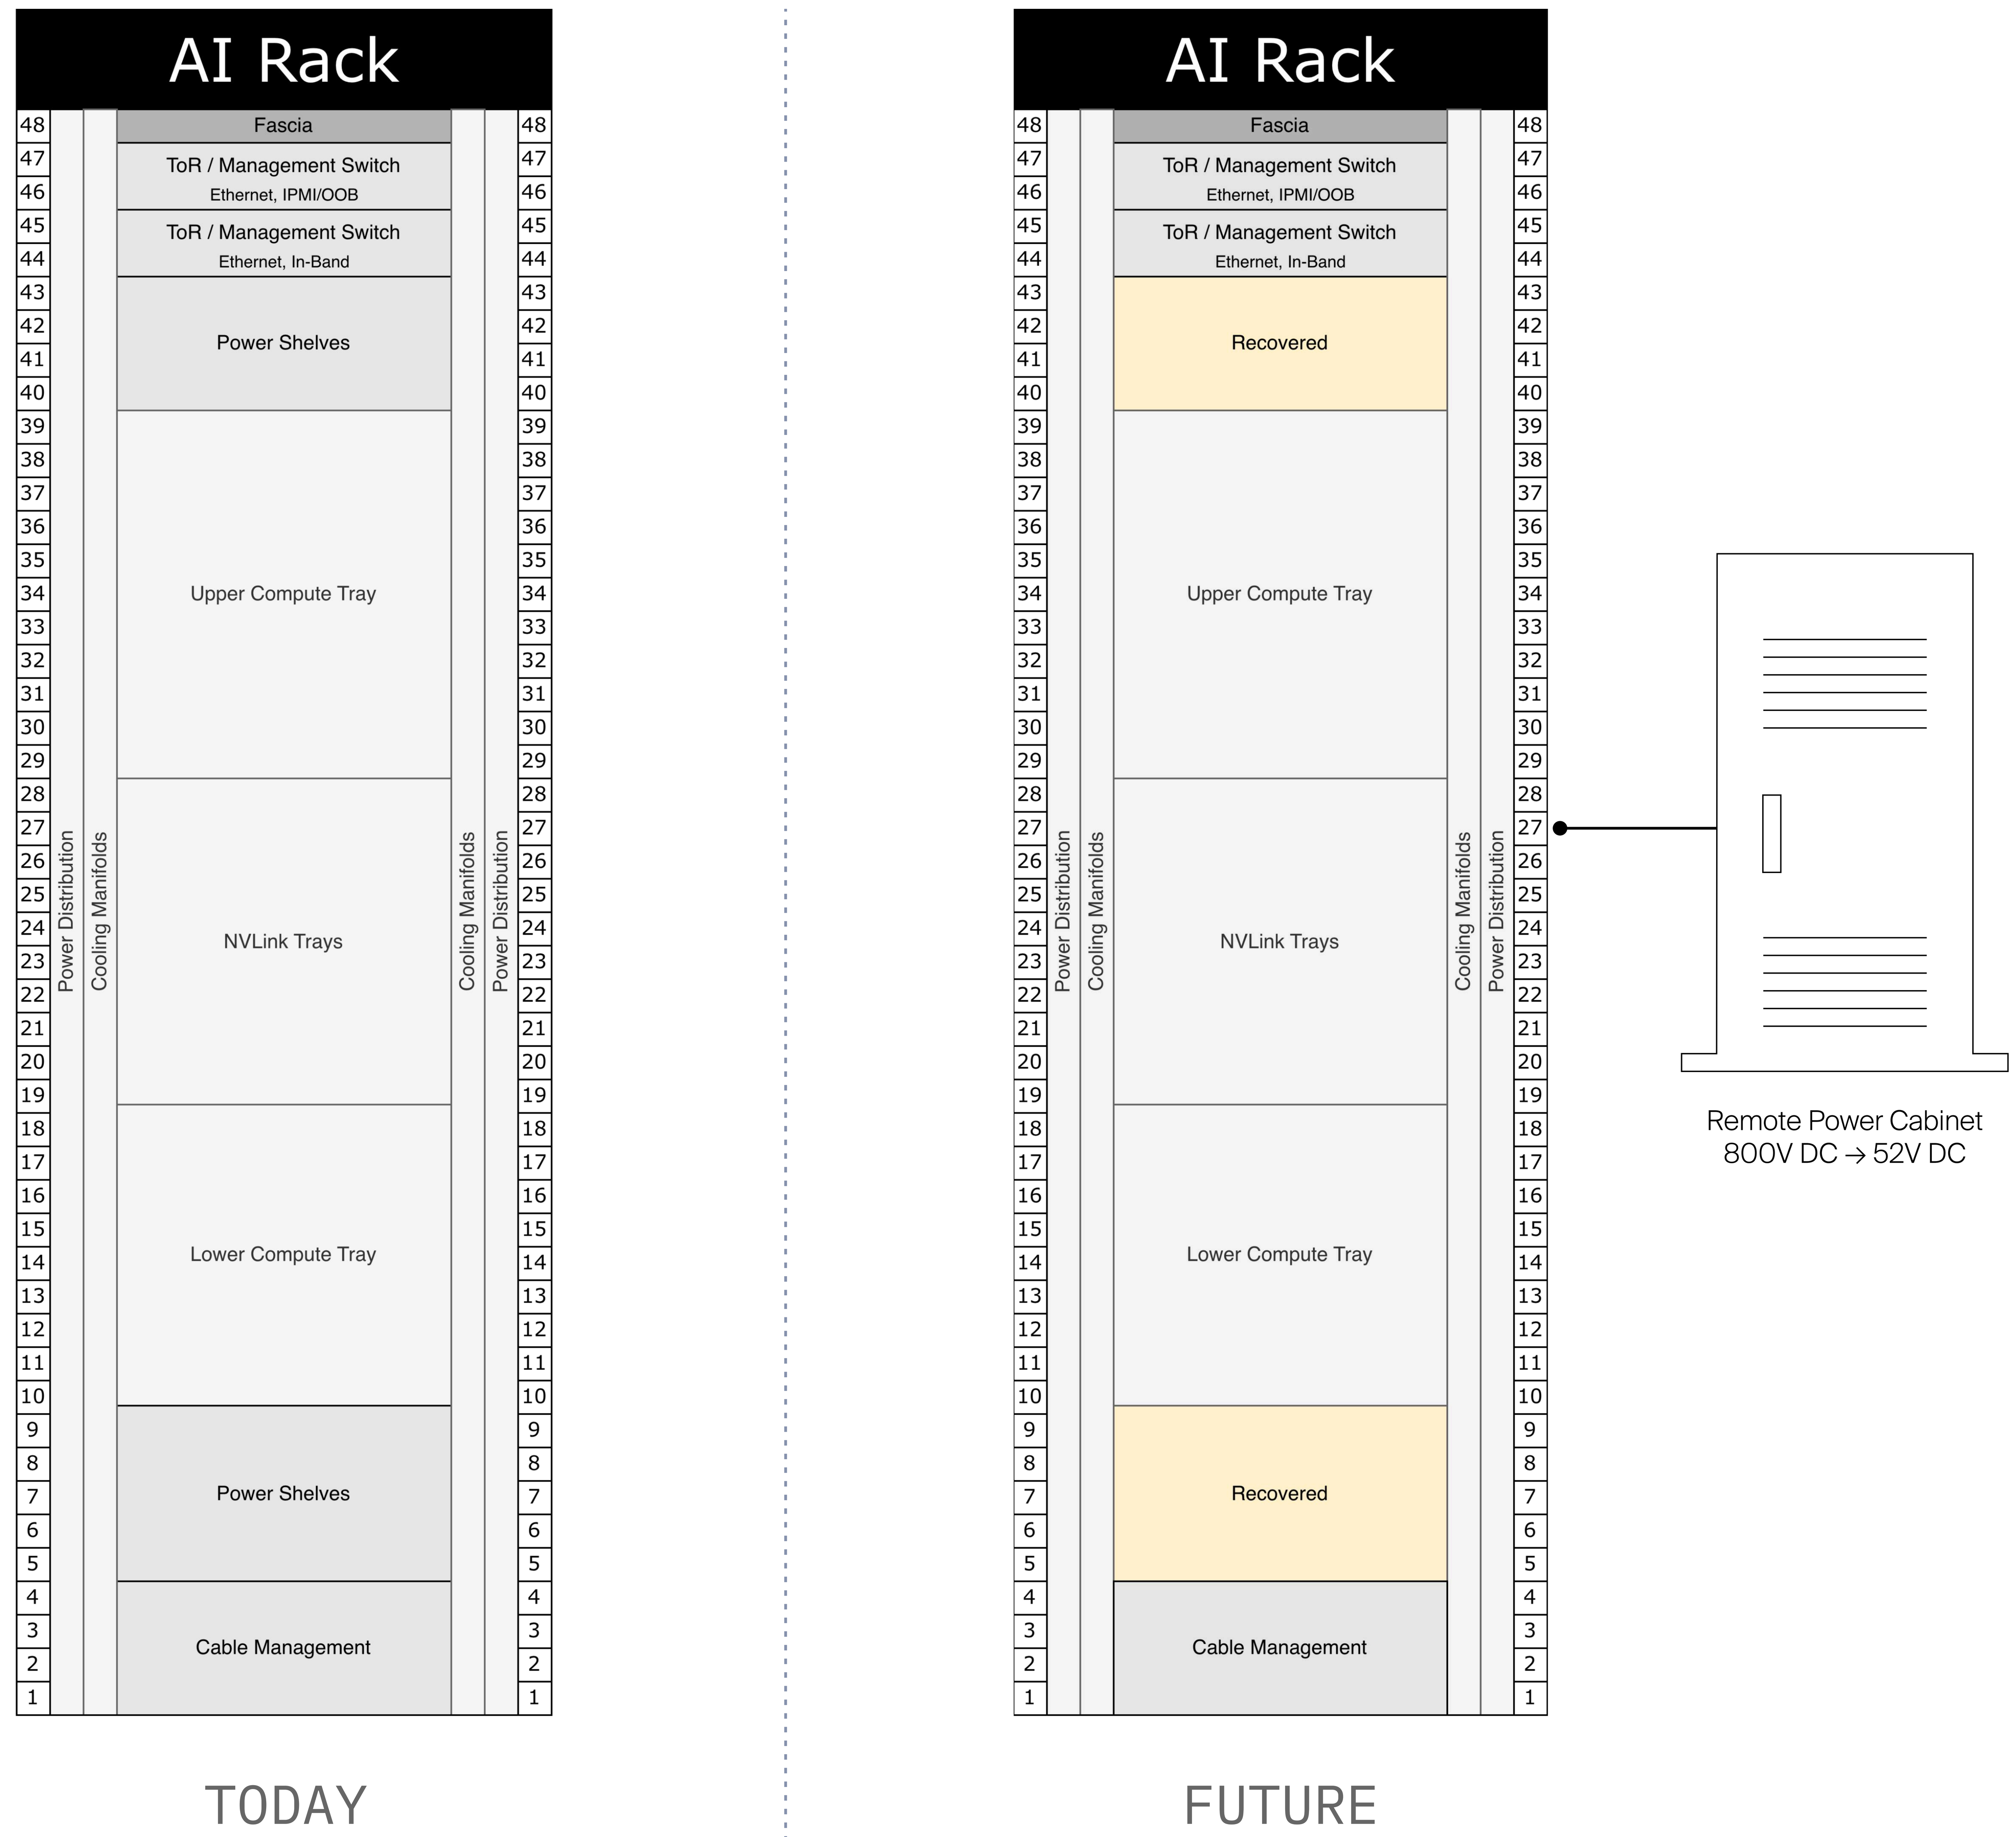

Rack Improvements

8 Units
Recovered per rack

~17%
Rack volume returned

~3.6 kW
In-rack PSU heat removed

+11%
GPU density uplift

Today's AI racks still inherit a power architecture built around copper. Power is distributed through the white space typically at 240V/415V AC and delivered to the rack, where rack-mounted power supplies convert it to 52V DC. From there, additional stages take power to 12V at the tray and finally to sub-volt levels at the GPU package. These stages consume rack volume, add heat, and create more hardware that must be cooled, serviced, and protected.

The next step architecture is already emerging: 800V DC distribution to the rack, enabled by solid-state transformers that convert utility medium voltage directly to a regulated DC bus. SHIELD is designed to fit naturally into that transition, moving 800V DC from outdoor or edge-of-building SSTs into the white space with very low electrical loss and minimal heat dumped inside the data hall. It can support today's 415V AC facilities, transitional 800V DC rack architectures, and future low-voltage DC networks, but its value becomes strongest as the data center moves toward distributed DC power.

At 52V DC, SHIELD can carry thousands of amps with effectively negligible I^2R loss, making it possible to move bulk conversion out of the rack and feed rows or clusters from centralized, serviceable power modules. PSU shelves are deleted from the rack, heat is removed from the white space, and operators recover volume for GPUs instead of power electronics. The result is not just lower electrical loss; it is higher rack density, cleaner serviceability, and a power architecture that scales with AI rather than fighting it.

Cumulative savings
vs conventional copper —
100 MW campus, 10-year window

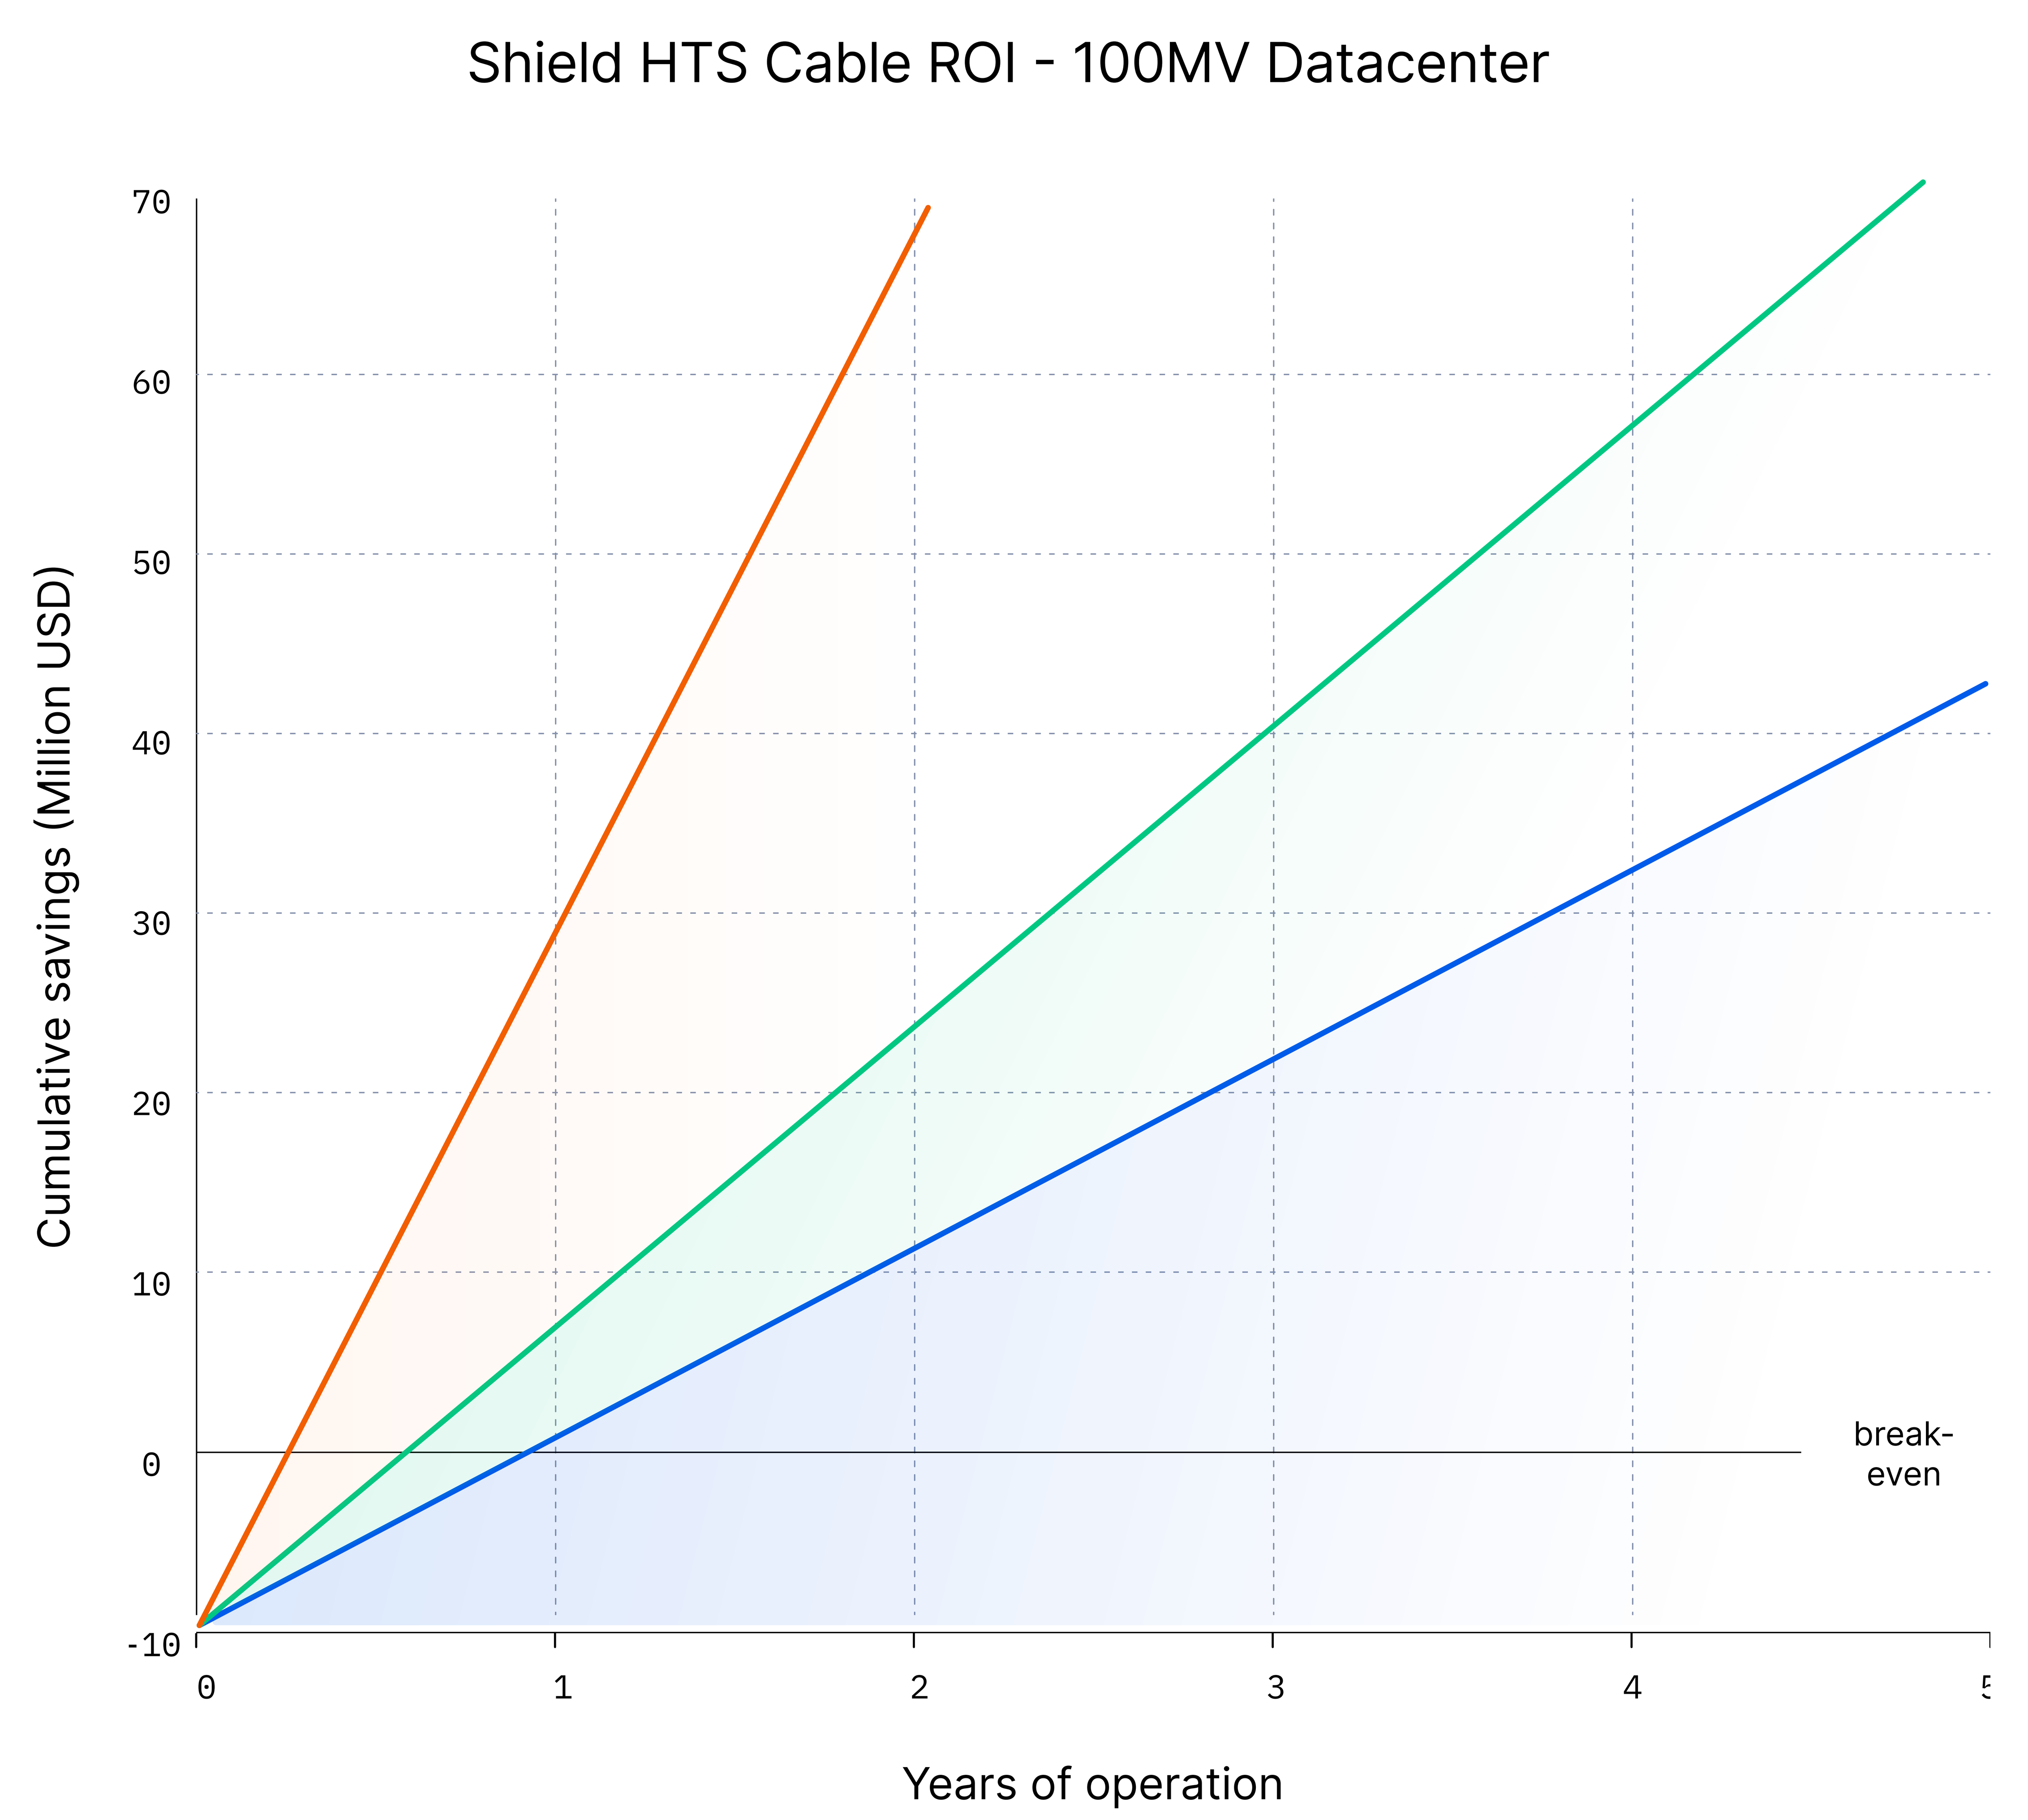

- LV AC / 52V (11 month payback)
- SST 800V (7 month payback)
- SST 52V (3 month payback)

Economics

6-8 Units
Rack volume recovered

5-8%
White-space cooling offset

$1–2 M
Savings per year per 100 MW

Energy savings A 100 MW AI campus on conventional copper distribution loses 2–4 MW continuously to I^2R heating and conversion overhead. SHIELD recovers approximately 1.5–3 MW of that loss at the distribution layer alone. At industrial rates of $0.08/kWh and 90 % utilization, this corresponds to $1–2 M per year per 100 MW in direct energy savings — before considering avoided whitespace HVAC cooling load.

HVAC offset Every kilowatt of copper I^2R heat dumped into the white space requires roughly 0.3 kW of additional chiller capacity to remove. By moving conversion outdoors and eliminating in-building I^2R heating, SHIELD reduces whitespace cooling load by 5–8 %: a meaningful CapEx avoidance on the cooling plant and an ongoing OpEx reduction.

Compute density For operators measuring revenue in dollars per rack-U, the 8U recovered per AI rack is huge. With current AI infrastructure constraints, that volume is worth 10x more than the cost of the HTS cable that enabled it.

Payback Across realistic datacenter scenarios, SHIELD reaches payback in 18–36 months and delivers 3–6x lifetime ROI over a 10-year operating window. The crossover happens earlier at higher rack densities, higher utilization, and higher local electricity costs — exactly the conditions defining the AI buildout.

Conclusion

SHIELD is the convergence of two trajectories: hyperscale AI's relentless demand for dense, reliable power — and Maritime Fusion's mastery of high-temperature superconductors designed for mass manufacturability and reliable operations.

By moving 8,000 amps in a tennis-ball-sized cross-section with effectively zero loss, SHIELD eliminates the I^2R wall copper cannot escape, and frees the white space, the cooling plant, and the rack itself for compute.

The technology integrates with NVIDIA's 800 DC reference architecture today, and enables a fundamental redesign of AI-rack power tomorrow. The economics provide substantial ROI today and only improve as densities grow.

Partner with Maritime Fusion to define the next standard of AI infrastructure.

> High temperature superconductors enable economic lossless power distribution

Get in touch

product@maritimefusion.com

maritimefusion.com/shield

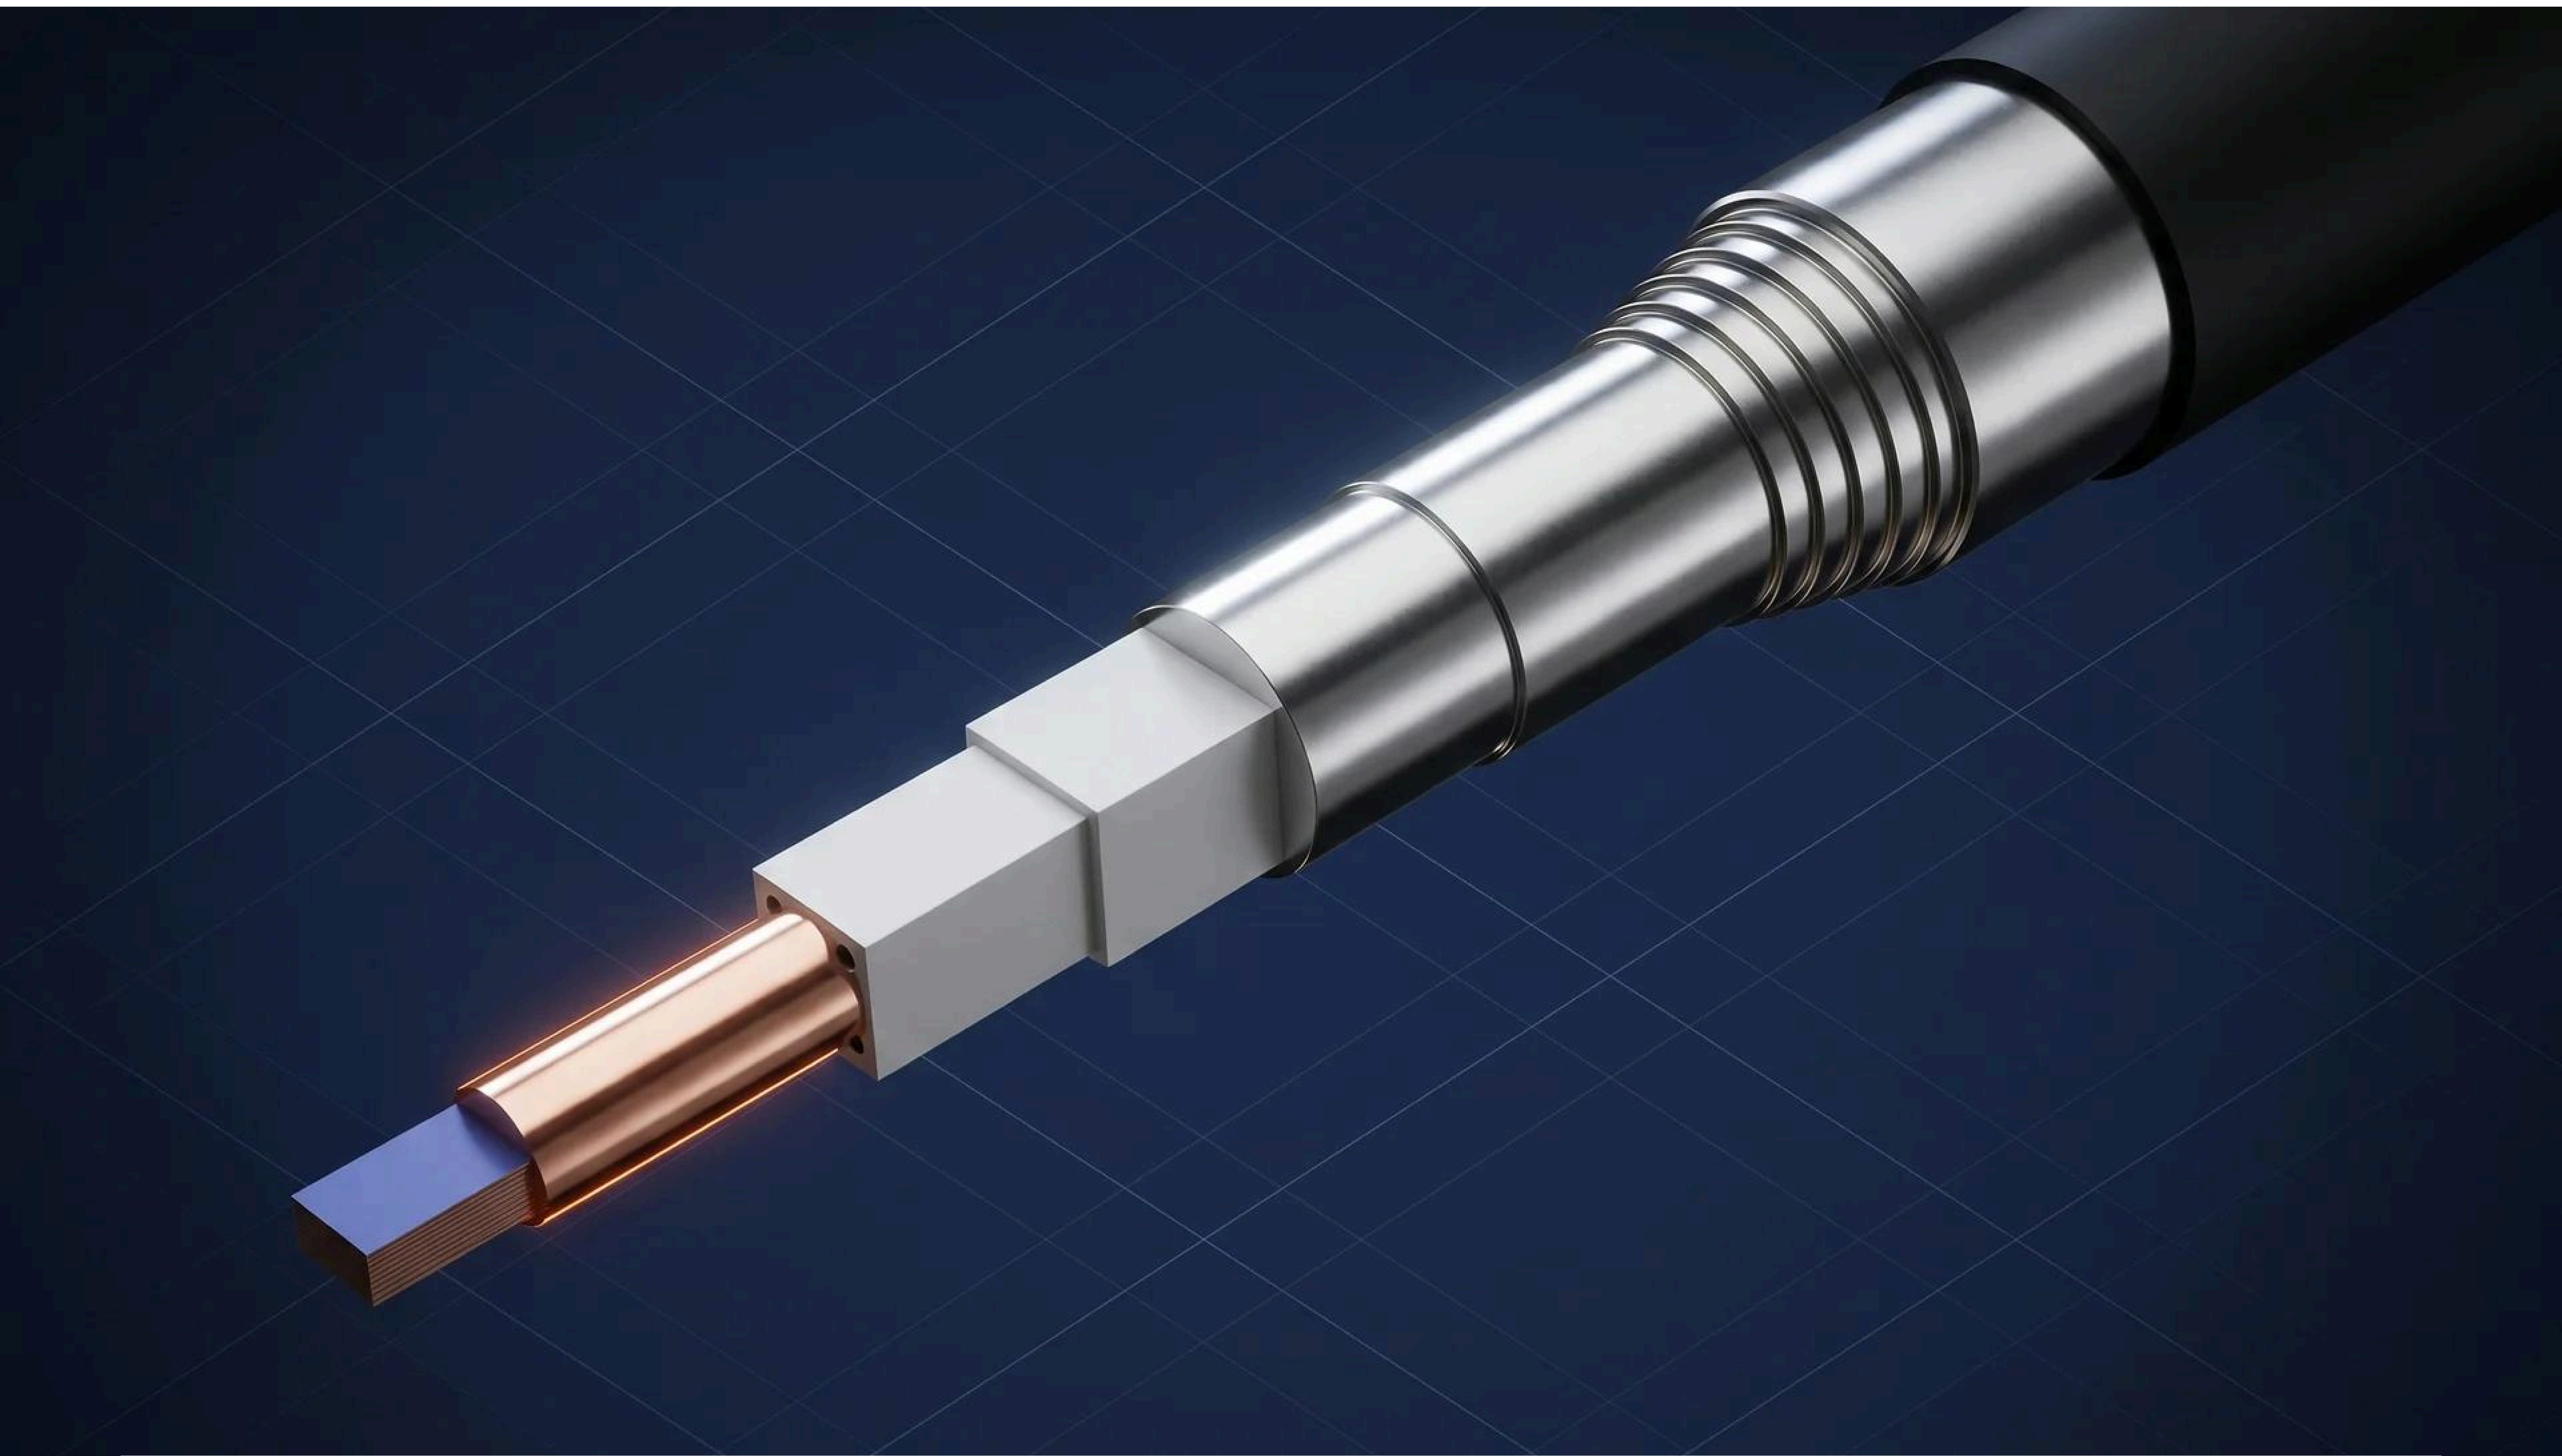

MARITIME FUSION