EMERGING SOLUTIONS FOR GRID EFFICIENCY

Practical and emerging approaches to a more efficient electric grid

Andre Canelhas |

International Rethinking Electricity Forum (IREF) · Le Méridien, Paris · 25 August 2026

The architectures presented are open and multi-vendor — nothing here is proprietary to, or dependent on, any single supplier.

FROM DETERMINISTIC TO TRIPLE-STOCHASTIC

Generation, load and flows now move at once.
GIGA-scale AI loads add **extreme density** and five-nines continuity requirements.
The CAISO curve makes the daily reality visible: solar creates a deep midday trough; the evening ramp must be supplied by storage and flexible resources.

Engineering response: co-located compute absorbs surplus; storage and flexible load shape the ramp.

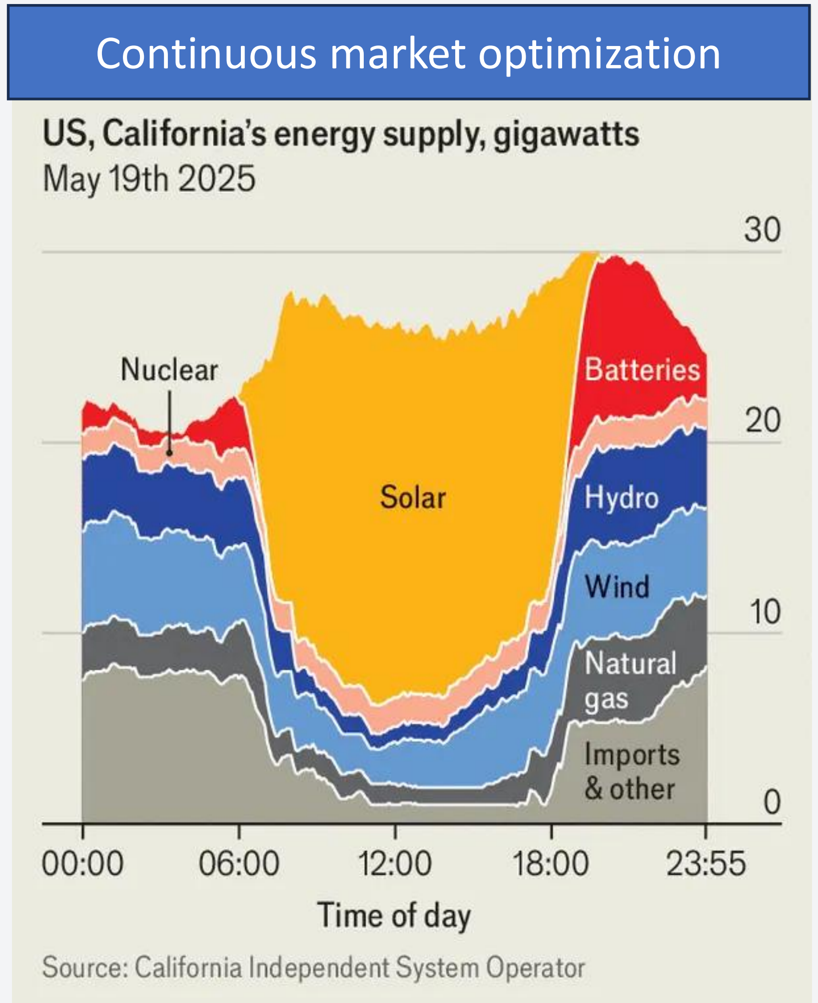

CAISO daily net-load profile: midday surplus shifts the flexibility challenge into the evening ramp.

Source: California Independent System Operator (CAISO)

THREE ACTORS · ONE STRUCTURAL DEADLOCK

UTILITIES	HYPERSCALERS	MANUFACTURERS
System reliability for all consumers	Speed to Power at five-nines	Firm order book before expansion
Timescale	**Timescale**	**Timescale**
Years	**Months**	**Years**
Trip risk · stranded upgrades	Losing the AI race waiting for power	Build for a bubble or become the bottleneck

Three rational actors · Three clocks that do not synchronise · Risk allocation is the real problem.

No single actor can break the loop alone.

"Speed to Power is fundamentally a risk-allocation problem." — Arushi Sharma Frank, CSIS

HVDC TECH

GRID EFFICIENCY IS A COLLABORATIVE BUILD

Grid Efficiency is the system-level ability to deliver **firm, reliable, scalable megawatts** to dense, time-critical loads faster, at lower societal cost, with stronger reliability and sustainability.

FRAME · AND solutions

DEMAND · GIGA-scale compute

ENERGY · Lowest-cost supply

METRIC · Speed to Power

THE BUILD DELIVERS

Speed, affordability, resilience and sustainability — simultaneously, not as trade-offs.

Source: Emerging Solutions for Grid Efficiency, Sections 1 and 7. Definition and component set adapted from the supplied paper.

First power in months → progressive firming over time

1. Site for deliverability
2. Build modular
3. Bridge power first
4. Operate as VPP day 1
5. Overlay & reinforce
6. Firm progressively

COMMERCIAL & RISK PRINCIPLE

Each party underwrites what it is best placed to carry.

Private firming by the developer · Milestone-based co-development · Standardised contracts · Productised hardware

Nothing quietly socialised onto ordinary consumers.

HOW THE LOAD BECOMES A GRID ASSET

Under the hood — native DC, VPP control and the 800 VDC architecture

Before I jump into compute-rack engineering —
power engineering and compute-rack engineering collaborate not for a compromise, but so the compute asset becomes a grid asset too.

A DC HUB FROM STANDARD EQUIPMENT — TODAY'S BENCHMARK

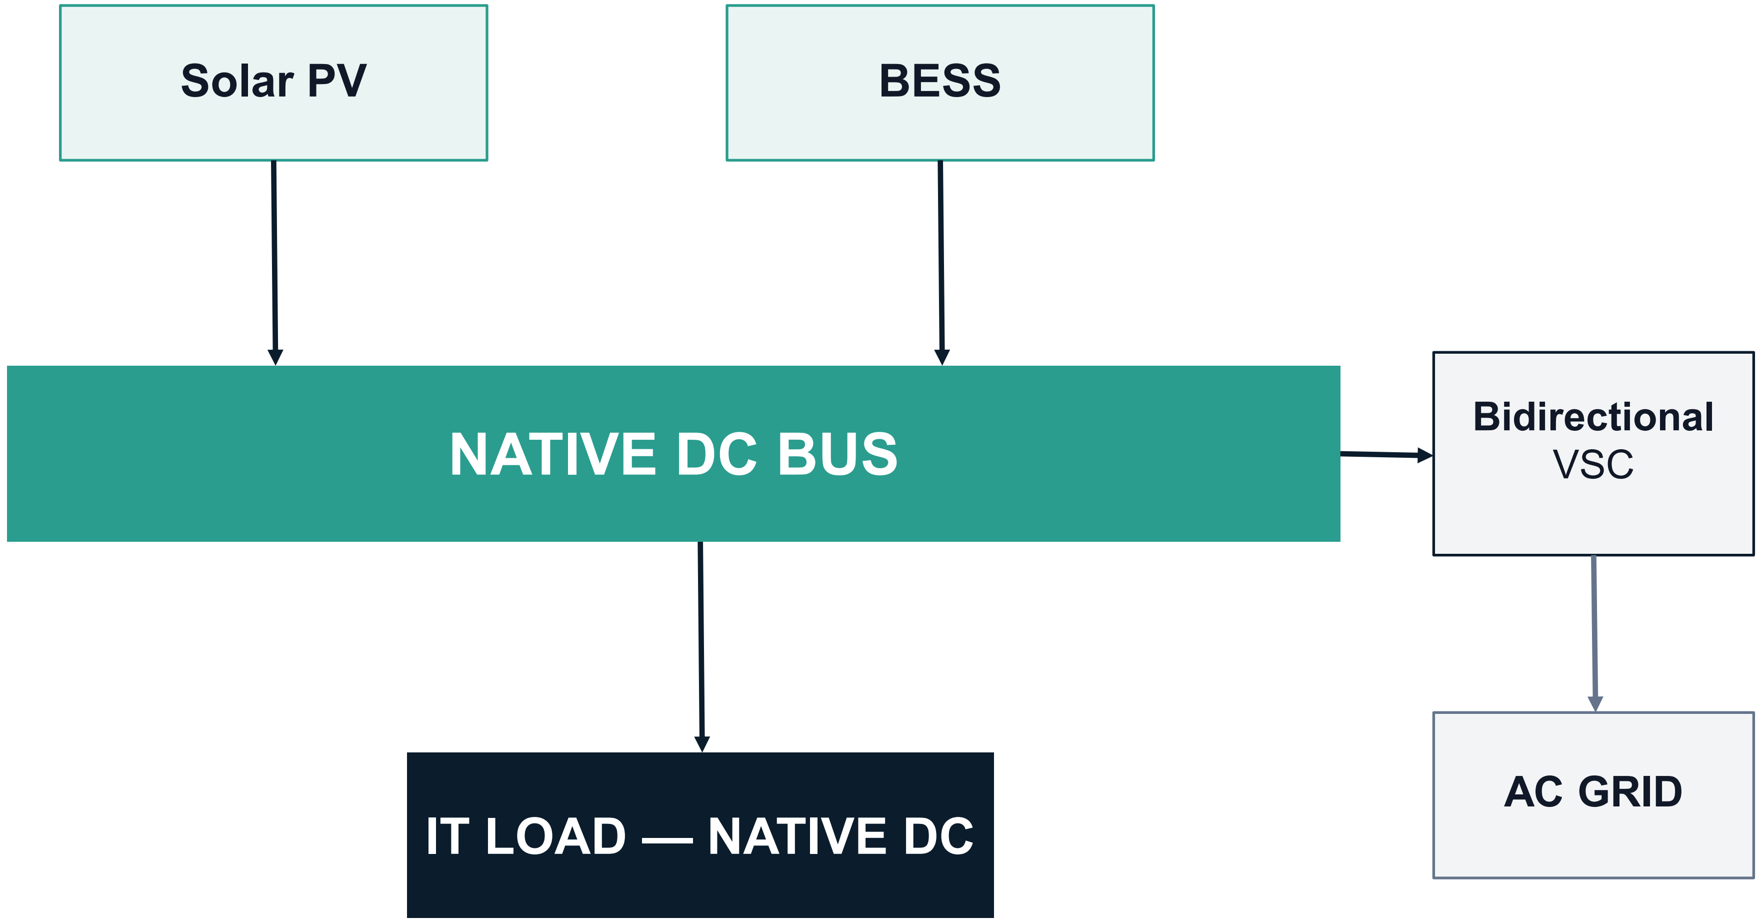

Heavy auxiliaries — cooling and other AC loads — remain on AC; never convert them twice.

WHY IT WINS

- Fewest conversion stages — IT stays on native DC.
- Field-proven MPPT, DC/DC, VSC, transformer and GIS equipment.
- Modular designs deliver speed to power today.
- VSC is rated for net exchange only — a fraction of facility power.

Indicative design target: generation/storage → IT ≈ 1–1.5% loss at rated load — HVDC Technologies estimate; boundary and assumptions as defined in the paper.

Source: Emerging Solutions for Grid Efficiency, 2026. Typical-loss estimate HVDC Technologies Ltd; site-specific studies required.

800 VDC: MOVE CONVERSION OUT, MOVE DENSITY IN

Facility-to-rack power path — direct 800 VDC delivery with near-GPU conversion

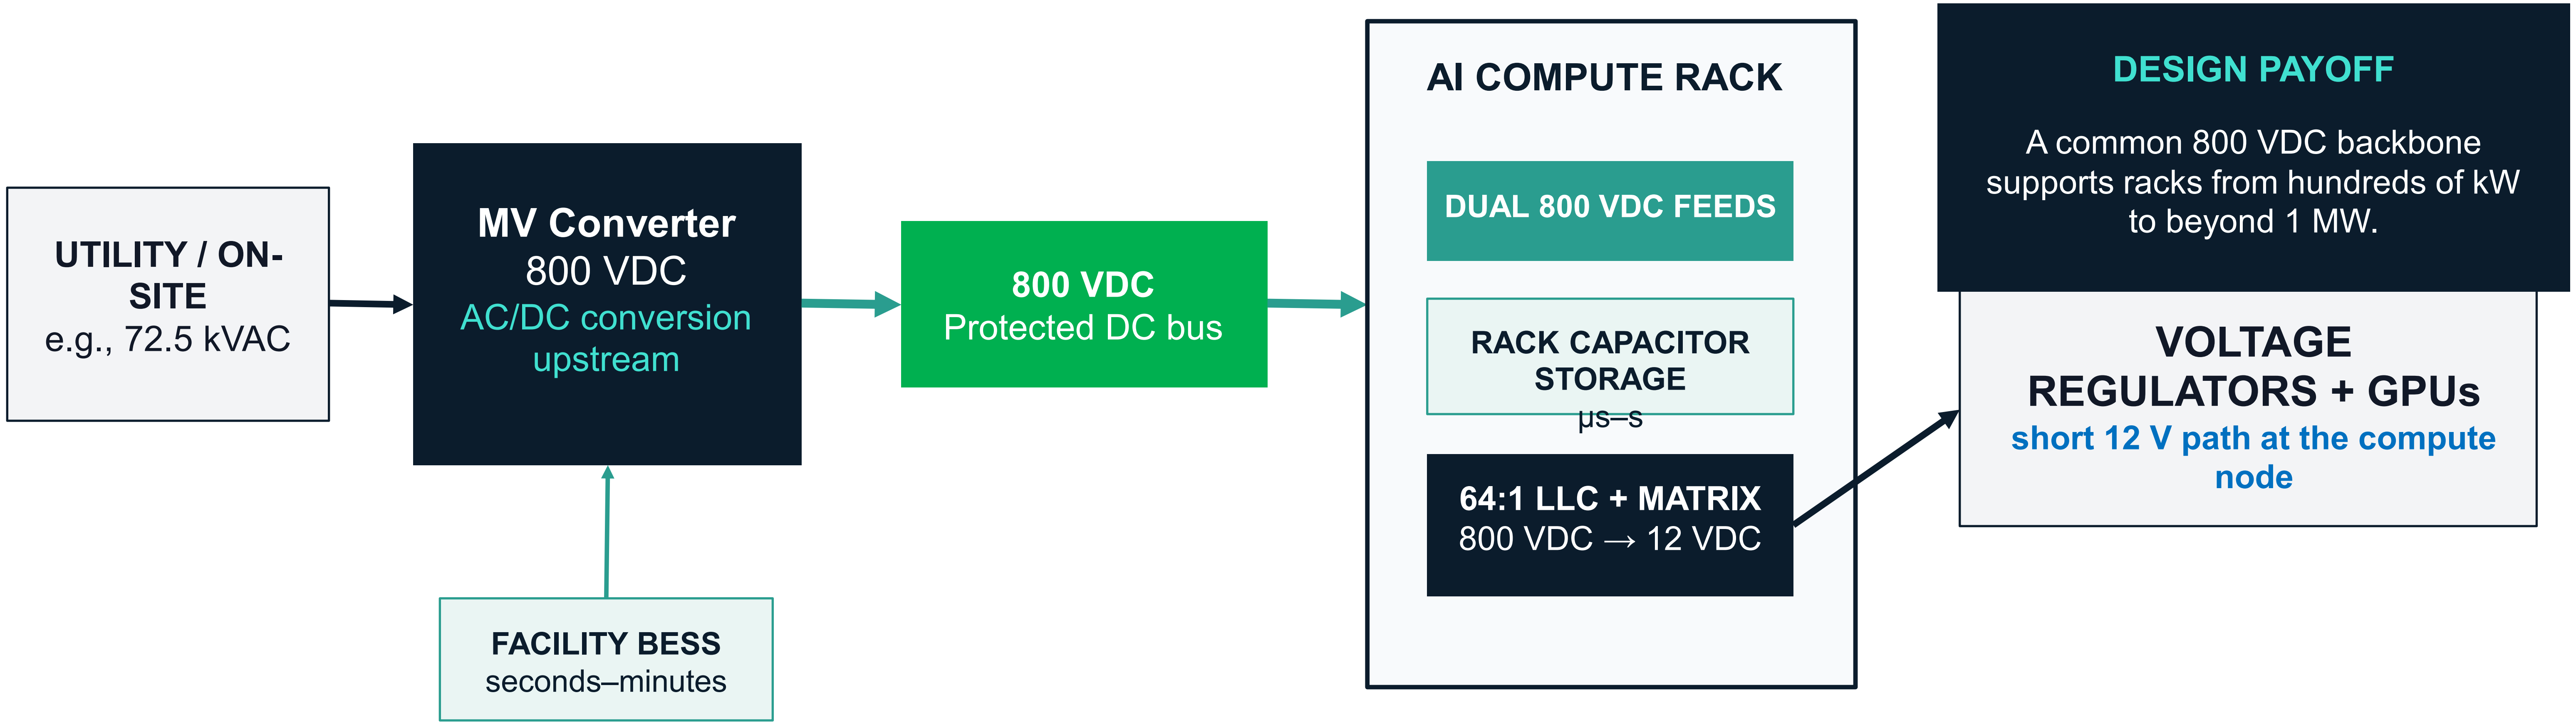

Safety architecture: touch-safe, mechanically interlocked connectors; protected DC zones; clear isolation points for maintenance.

Sources: NVIDIA, 800 VDC Architecture for Next-Generation AI Infrastructure; NVIDIA 800 VDC technical blogs (2025) and HVDC Technologies Ltd

THE HUB IS A VPP — CONTROL ACROSS TIMESCALES

THE MAIN SUBJECT IS CONTROL

• Compute is schedulable, curtailable and migratable between clusters.
• The compute mix can be designed for elasticity: training versus inference.
• Thermal load is a storage resource: over-run and release cooling.
• Generation, storage, elastic compute and thermal storage behind one VSC become a virtual power plant offered to the grid.

GRID CONTROLS & MARKETS	hours
WORKLOAD ORCHESTRATION & VPP DISPATCH	min–h
CAMPUS EMS · STORAGE & THERMAL CONTROL	s–min
CONVERTER & SEMICONDUCTOR CONTROL	µs–ms

must interact correctly

Recommendation: proportionate performance requirements for converter-interfaced mega-loads — ride-through, power quality, model provision, grid support — referenced to generation of comparable rating; jurisdiction- and connection-specific.

NVIDIA AI LOAD: POWER ENVELOPE OVER A FULL COMPUTE CYCLE

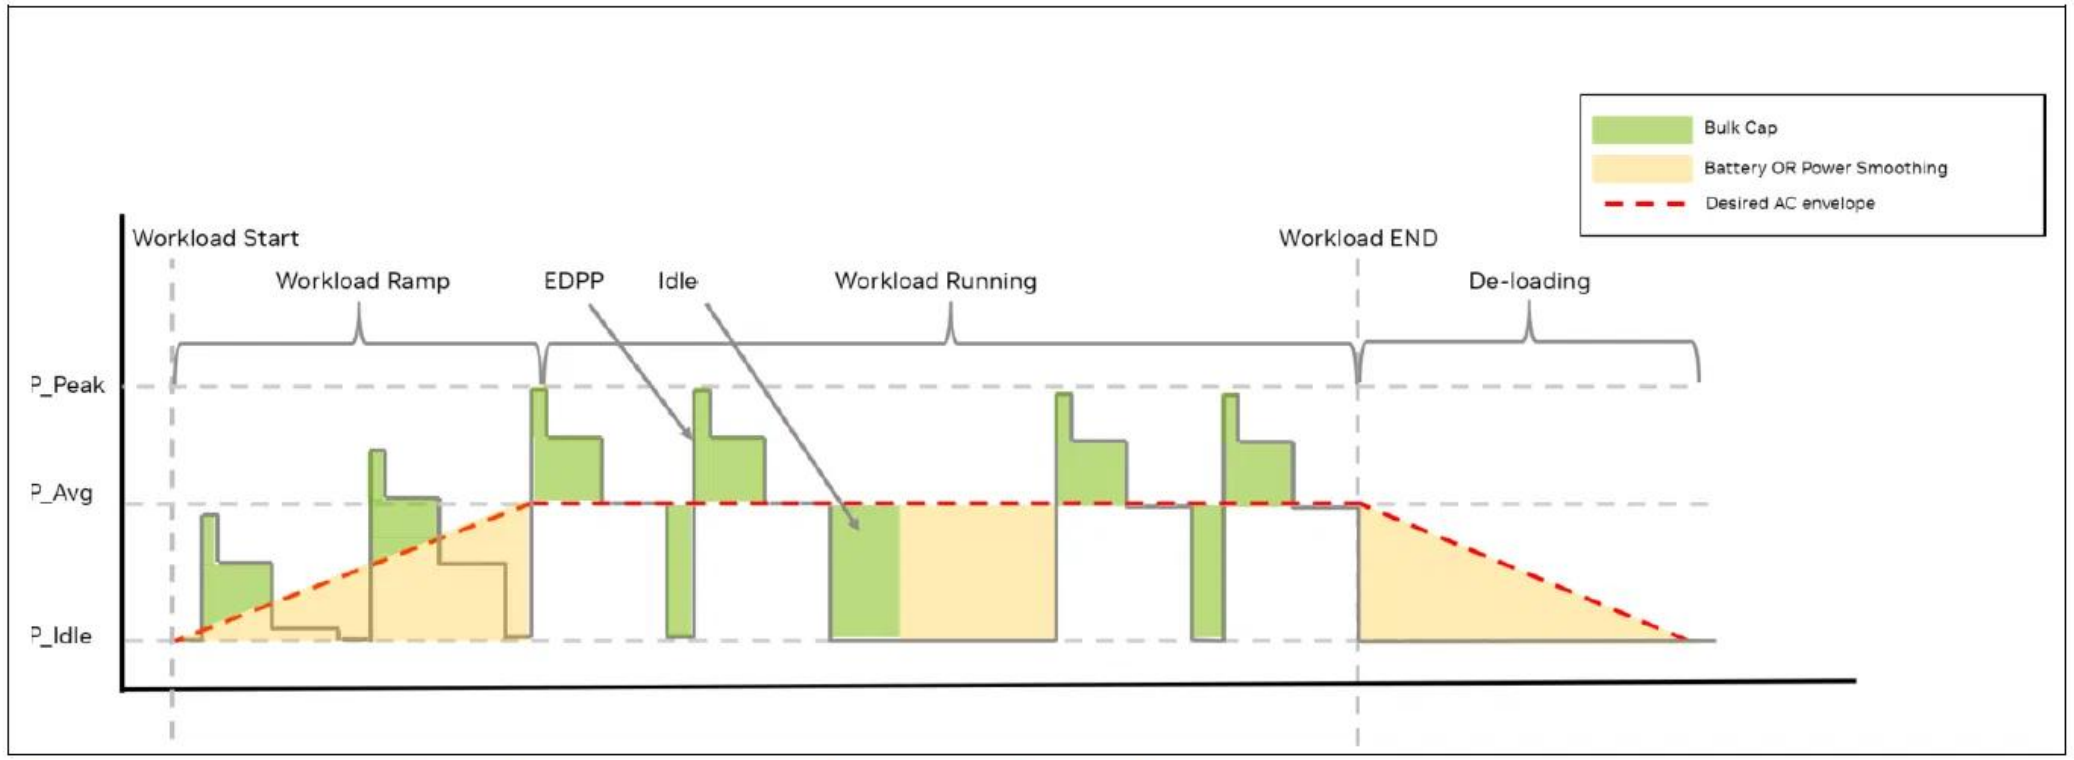

Figure 4: Power Envelope over a full compute cycle

HOW NVIDIA SHAPES THE LOAD

GREEN — Bulk cap
Power caps raise the GPU demand progressively during workload ramp-up.

YELLOW — Storage / smoothing
Energy storage charges in valleys and discharges into peaks to flatten input demand.

RED — Desired AC envelope
The grid sees a controlled, predictable power request rather than raw GPU pulses.

1 · WORKLOAD RAMP

A power-cap driver increases permitted GPU draw at a rate the grid can tolerate.

2 · WORKLOAD RUNNING

Power-shelf storage absorbs fast peaks and fills troughs, acting as a rack-level low-pass filter.

3 · DE-LOADING

If work does not resume, GPU burn maintains a floor then ramps demand down at a controlled rate.

Source: NVIDIA, 800 VDC Architecture for Next-Generation AI Infrastructure, Figure 4; NVIDIA, "How New GB300 NVL72 Features Provide Steady Power for AI" (2025).

INTEGRATED ENGINEERED OUTCOME

Native-DC efficiency + control across every timescale + 800 VDC density.

The data centre becomes a generation-like, dispatchable, grid-supporting asset.

One VSC — one integrated Power–Compute system.

The threat of the dual shock becomes part of the solution — the load joins the platform.

CALL TO ACTION

Utilities	From rationing to enabling — performance-based standards and flexible access.
Hyperscalers	Bring your own firming; operate as transparent, dispatchable assets.
Manufacturers	Productise and standardise adaptable, grid-forming equipment.
Policymakers	Write the shared-risk framework; shield ordinary consumers.
All parties	Open standards and shared high-fidelity models before energisation.

From high-entropy rationing → low-entropy platform

Let there be power.

BACK UP

LEGEND

- HVDC backbone
- AC backbone
- DC Data-centre anchors

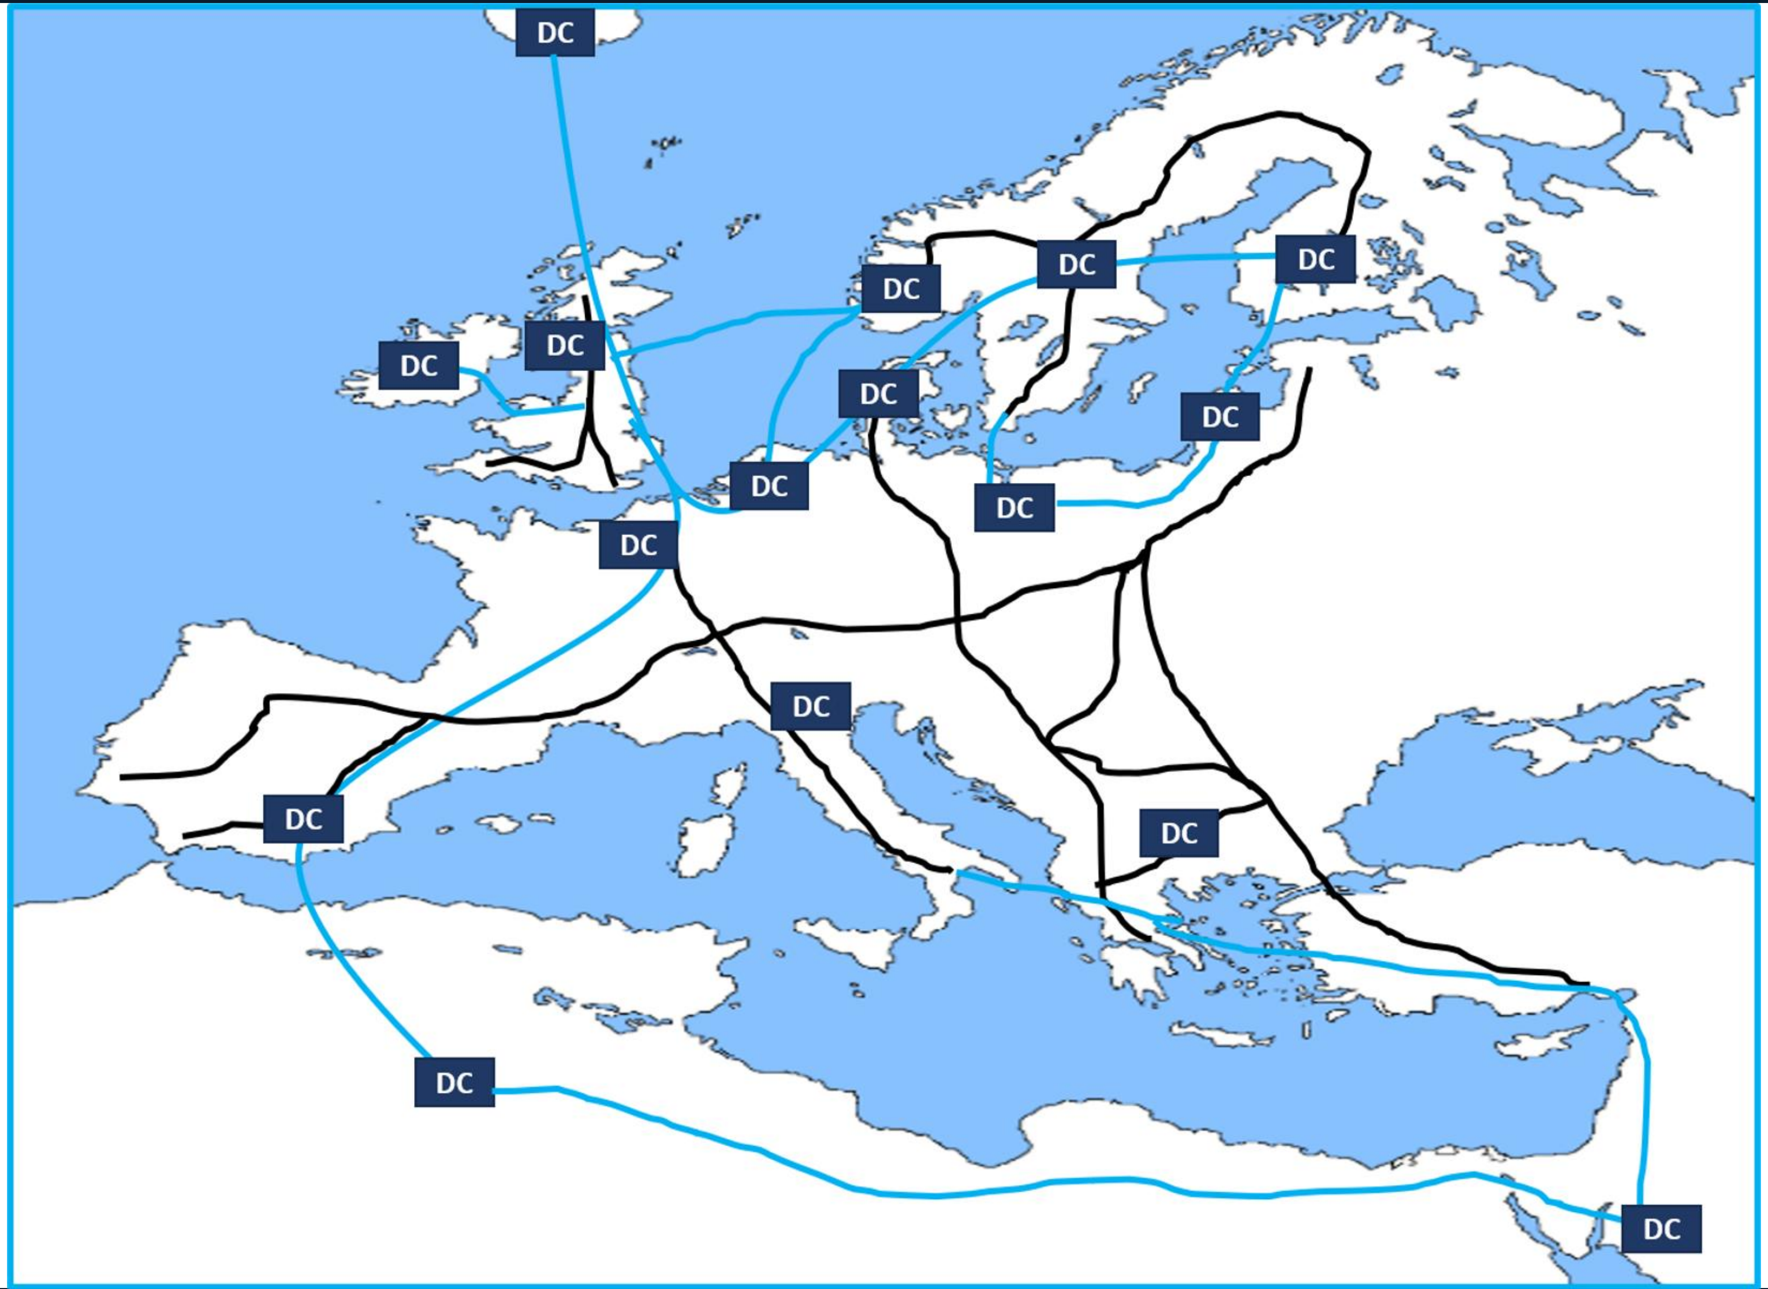

THE EUROPEAN SYSTEM

An interconnected canvas for controlled HVDC, FACTS, storage and active-load coordination.

Engineering does not replace markets or policy; it gives every actor a physical platform on which they can align.

THE LOAD, SHAPED AT SOURCE

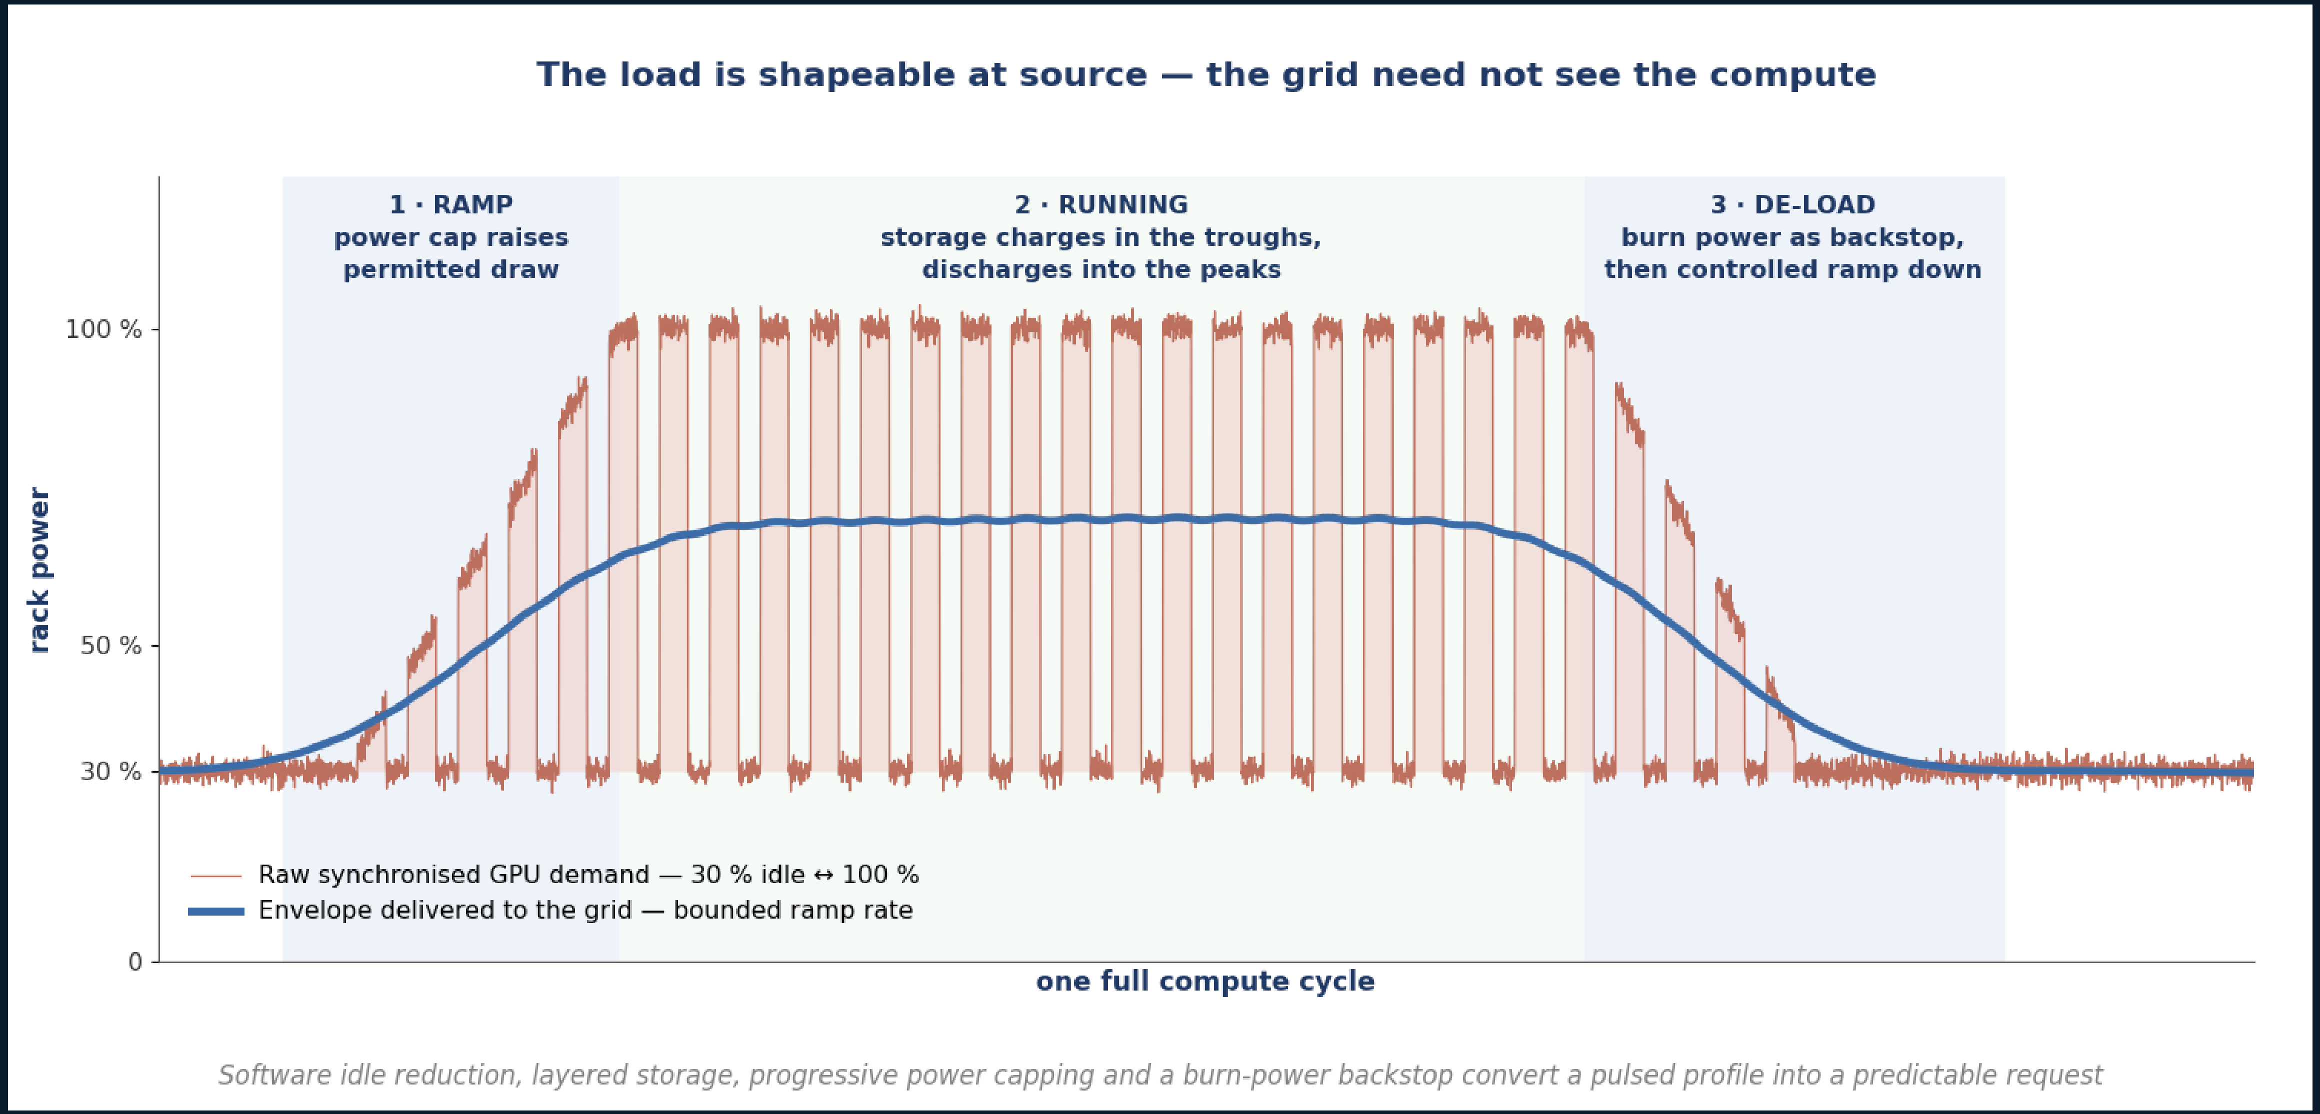

Raw synchronised GPU demand vs the envelope delivered to the grid — layered storage, progressive power capping and a burn-power backstop bound the ramp rate at the point of connection. Illustrative; residual spectrum to be shown by study.