



ESS in High-Density AI Infrastructure: Systems, Safety, and Standards

Sandia Energy Storage Safety & Reliability Forum 2026 · Session 5

Josh Gerber · Founder & CEO, Symphonic Energy

josh@symphonicenergy.com

Outline

1. **The scale of AI capex and revenues**
2. **GW-scale facilities under construction**
3. **Stepping up voltage, increasing power density**
4. **AI training workloads**
5. **Standards and suppression**

Presenter Background

- Led one of the world's earliest utility-scale battery projects at SDG&E
- 25+ years in energy and information& communications technology
- Co-author, industry reference document for batteries in data centers
- Advisor to battery technology and project developers from utility scale to residential

Josh Gerber

Founder & CEO,
Symphonic Energy



S E C T I O N

1

The scale of AI capex and revenues

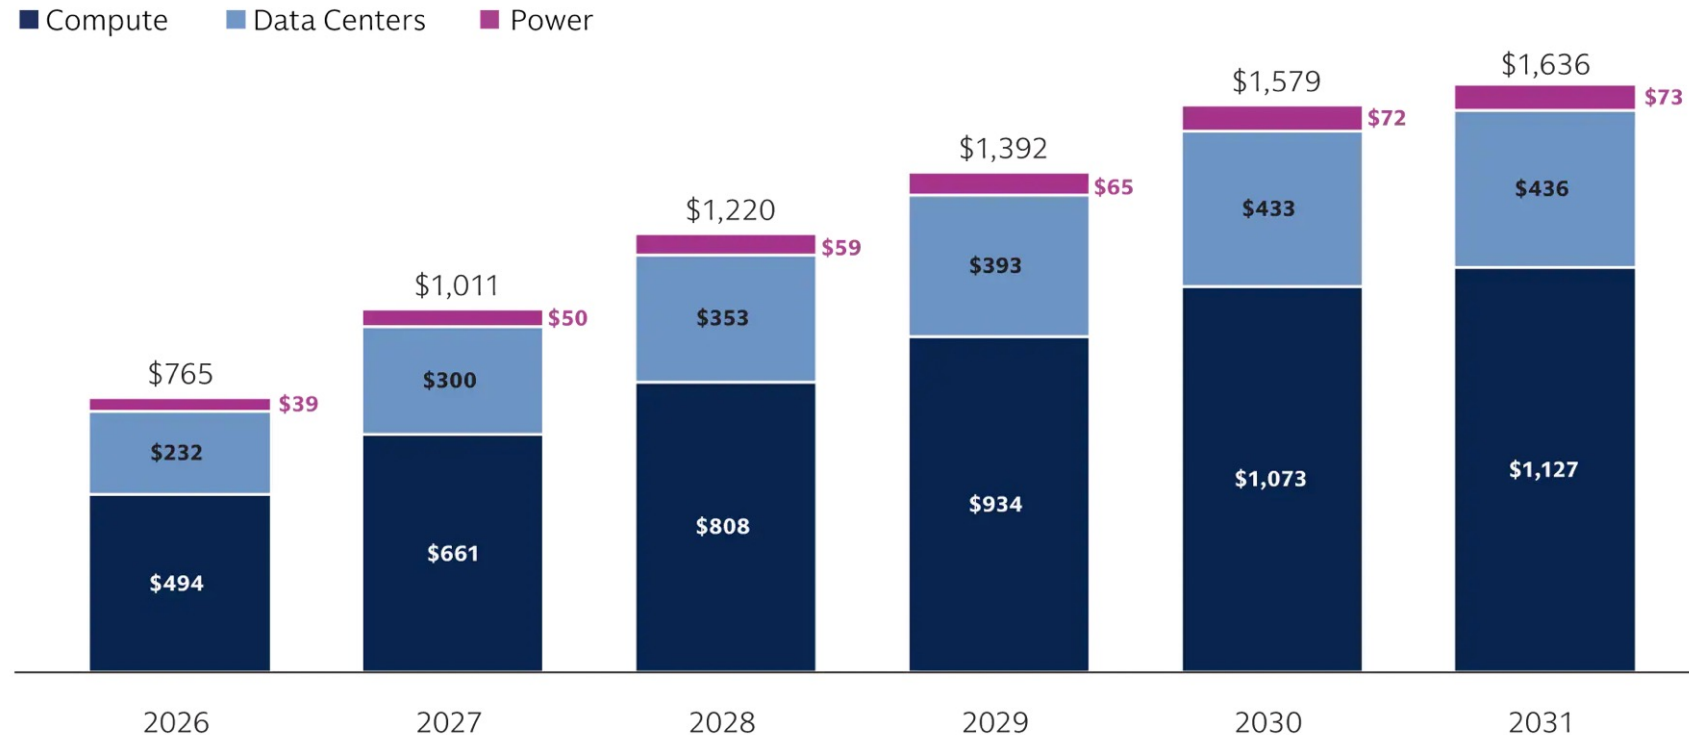
AI data center economics

Unprecedented Capital Investment



Baseline aggregate AI CapEx estimates (bn)

~\$7.6tr of capital between 2026 and 2031 across compute, data centers, and power



Source: Goldman Sachs Global Institute, Goldman Sachs Global Investment Research NVIDIA projections (as of March 3, 2026) Note: Forecasts and expectations are based on material assumptions subject to change. Assumes NVIDIA accounts for 75% of total compute spend in each period. Assumes 5% YoY compute growth past the projection period (2031). Uses VR200 (Rubin) chip as baseline spec (\$80.5K per GPU [incl. node costs] and 3,000 W per package) across all years. Assumes 1.2 PUE, \$15mn per MW for data centers, and \$2,500 per kW for new power. Assumes 15% of required data center space is brownfield (i.e., excluded from calculation) in 2026, growing to 30% in 2031. Totals may not sum due to rounding.

Comparisons:

- ~12x Interstate highway system (\$634B in 2024\$ [Wikipedia])
- ~30x Apollo Program (\$257B in 2023\$ [Wikipedia])

The Claimed Economics

\$59M / MW

upfront DC + IT capex

DC + power plant \$19M · GPUs
\$30M · other AI infra \$10M

~\$1.1M / MW

annual operating cost

Revenue is 15-30× operating
cost

\$15–30M / MW

annual revenue

IaaS baseline → managed-AI
services at 2× premium

- This implies**

•

Operating margin
(cash)

Revenue – Opex

\$1,587–
\$3,300/MW/hr

Net cash flow after
debt service

Operating margin –
Debt service

\$796–
\$2,509/MW/hr

Source: Crusoe CEO Chase Lochmiller, Stanford MS&E 435 "Economics of the AI Supercycle," Spring 2026, Class #3, Symphonic Energy calculations with 100% debt financing at 10% for 20 years

S E C T I O N

2

GW-scale facilities under construction

Two examples from the Stargate program

Facility scale, example 1 – Stargate Abilene, Texas

- 1.2GW shown, 2.1GW planned for campus
- Tenants: Microsoft, Oracle
- Each of ~40 data halls (4 per building shown) is ~100,000 sq ft, 25MW each



Source: Crusoe CEO presentation, Stanford MS&E 435, Spring 2026, Google Maps measurements

Images: left: DPR Construction; right: Image © Airbus DS 2026, captured 2026-03-24

Facility scale, example 2 – Project Jupiter, Doña Ana County, New Mexico

- Tenants - Oracle (anchor), OpenAI (customer)
- Footprint - 1,400 acres, 4 buildings + Campus Ultra
- IT capacity > 1 GW
- Primary power - Bloom Energy SOFC fuel cells, 2.45 GW
- Other infra - Dedicated battery storage center + on-site desalination
- Operations start before end of 2026
- Investment \$50B / 5 yr → \$165B / 30 yr



“To ensure Project Jupiter does not impact the cost of electricity for other customers in the region, the project will construct a microgrid to independently power the site.”

SECTION

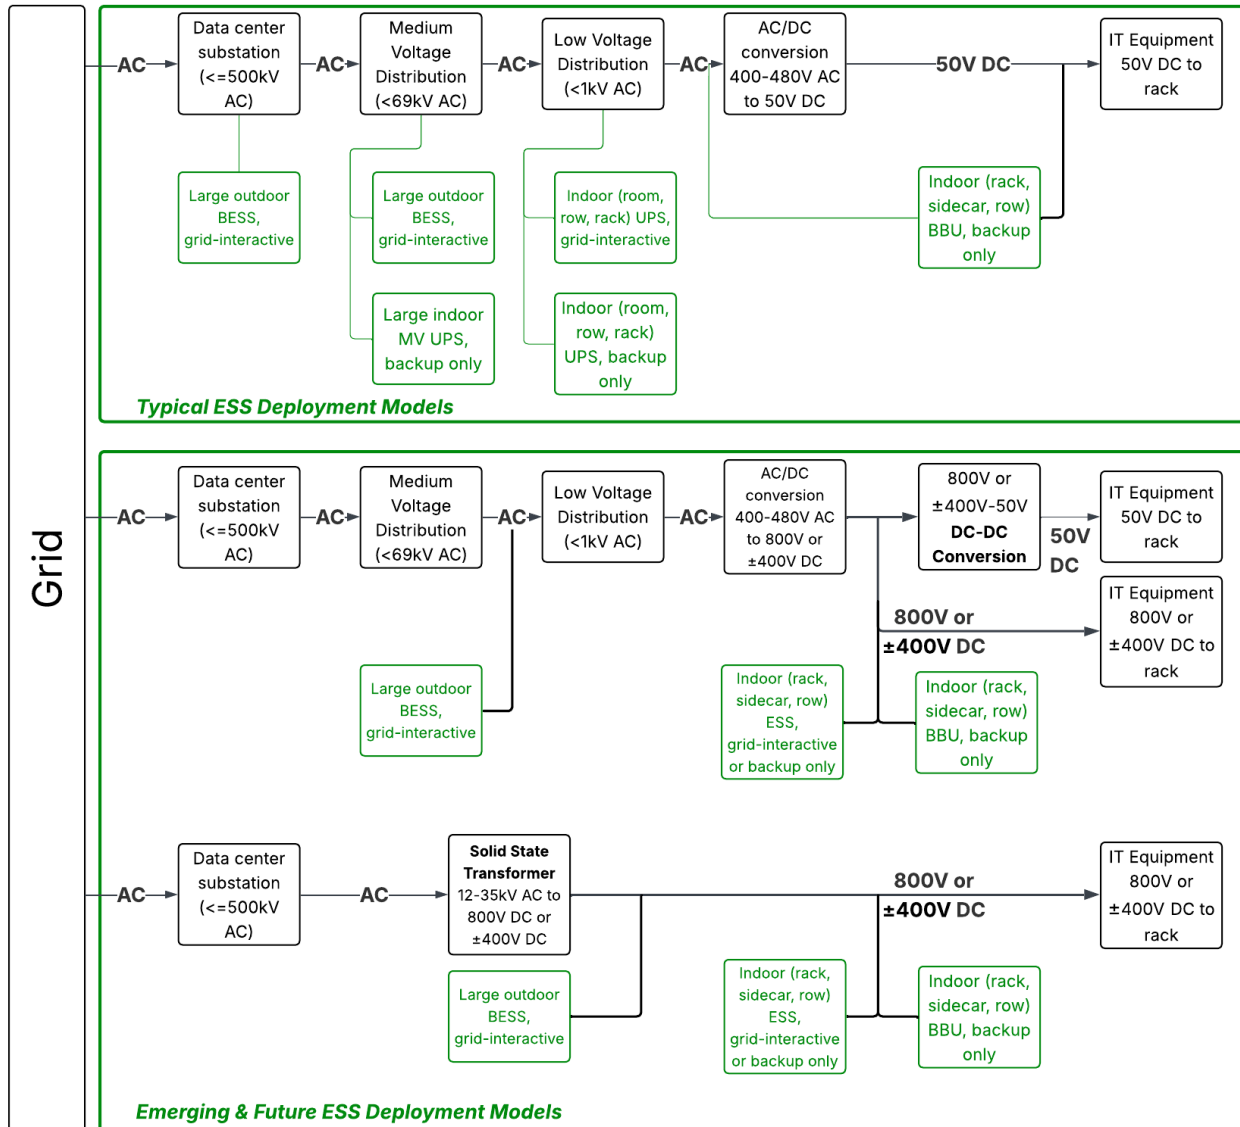
3

Stepping up voltage, increasing power
density

48 V → ±400 V → 800 V

Where ESS lives in the data center

Illustrative Data Center Power Systems with ESS



Today

ESS distributed across the AC chain – outdoor BESS (grid-interactive), indoor MV/LV UPS rooms, in-rack/sidecar Battery Backup Units (BBU, backup-only)

Emerging

Higher-voltage DC (800 V or ± 400 V) shifts where AC/DC conversion happens; in-rack/sidecar ESS can now be grid-interactive, not just backup

Future

Solid-state transformers collapse multiple conversion stages – MV AC direct to high-voltage DC, with ESS retained at the rack/sidecar/row level

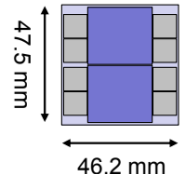
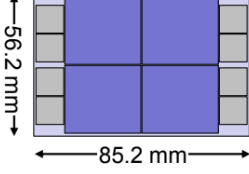
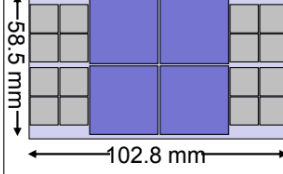
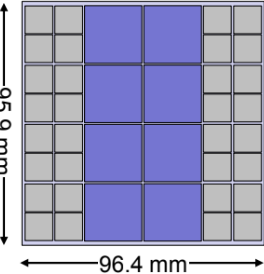
Source: [OCP ESS White Paper, "Requirements for Energy Storage Systems Used in Data Centers" \(Rev 1.0, January 2026\), Figure 2](#)

Increasing Chip Size and Power Reqs

Future NVIDIA chip generations enable higher rack power density


2.2kW->4.4kW->6kW->15kW
per chip-memory module
(800W->1200W/GPU while
reducing die size ~20%)

High Bandwidth Memory
(HBM) on the module
increasing from 0.4TB to 6TB


Next-Generation GPU-HBM Roadmap : More GPU & HBM Integrated Above Interposer				
GPU Architecture	Rubin (2026)	Feynman (2029)	Post Feynman (2032)	Next-Gen Architecture (2035)
GPU Die Size	728 mm ²	750 mm ²	700 mm ²	600 mm ²
GPU Power	800 W	900 W	1,000 W	1,200 W
GPU-HBM Module	R200	F400	Post Feynman GPU-HBM Module	Next-Gen GPU-HBM Module
Interposer Size				
# of GPU Dies	×2	×4	×4	×8
# of HBM Stack	HBM4×8	HBM5×8	HBM6×16	HBM7×32
Interposer Die Size	2,194 mm ² (46.2 mm x 48.5 mm)	4,788 mm ² (85.2 mm x 56.2 mm)	6,014 mm ² (102.8 mm x 58.5 mm)	9,245 mm ² (96.4 mm x 95.9 mm)
Total Bandwidth	16 / 32 TB/s	40 TB/s	120/256 TB/s	1,024 TB/s
Total HBM Capacity	288/384 GB	400/500 GB	1,536/1,920 GB	5,120/6,144 GB
Total Power	2,200 W	4,400 W	5,920 W	15,360 W

Source: Korea Advanced Institute of Science & Technology (KAIST) TeraByte Interconnection and Package Laboratory
<https://tera.kaist.ac.kr/researches/teralab-hbm-milestone-and-roadmap>

Google Ironwood – in-rack BBU is current practice



Google: “our Li-ion BBUs benefit from our distributed UPS architecture that offers significant availability and TCO benefits compared to traditional monolithic UPS systems. The distributed UPS architecture **improves machine availability by: 1) reducing the failure-domain blast radius to a single rack, and 2) locating the batteries in the rack** to eliminate intermediate points of failure between the UPS and machines.”



- “Pods” with up to 9,216 Google Tensor Processing Units (TPUs) and 10MW load
- Liquid cooling to the chips
- Battery Backup Unit ← bottom of every rack

Sources: Google - <https://blog.google/innovation-and-ai/infrastructure-and-cloud/google-cloud/ironwood-tpu-age-of-inference> and <https://cloud.google.com/blog/products/compute/ironwood-tpus-and-new-axion-based-vms-for-your-ai-workloads> and <https://cloud.google.com/blog/topics/systems/100-million-li-ion-cells-in-google-data-centers> 13

Delta 800 VDC In-Row 660 kW Rack

- BBU's with "NMC, NCA, LFP, etc." and "Unwavering adherence to UL 9540A/UL 9540 is maintained across the cell, module, and unit levels, incorporating robust thermal and gas mitigation designs critical for the safety of Lithium-Ion technology in data center environments."
- CBU's "optimized to stabilize the rapid power fluctuations characteristic of high-performance GPU clusters ...with "Electric Double Layer Capacitors (EDLCs), or supercapacitors, Lithium-ion capacitors (LIC), and more."




Fig. 2 In-Row Power (HVDC) System

SECTION


4

Training workloads

Duty profile, heat, and cell wear

Training – high-amplitude, coordinated, synchronous

Per-rack swings can synchronize across tens of thousands of GPUs



idle <-> TDP

per-GPU power range each iteration

0.2–3 Hz

dominant frequency band (sub-synchronous)

~1 sec

typical iteration period

10s–100s
MW/s

aggregate datacenter swings (&GW in future?)

Frontier training – duration, power, cost

Each frontier training run has grown ~2× per year; durations now 90-180 (++) days

Model	Year	Duration	Cluster power
Year	Cluster (rated)	Energy per training run*	Energy Cost @ \$100/MWh
2025	100–150 MW	324–486 GWh (\approx 0.3–0.5 TWh)	\$32M – \$49M
2028	1–2 GW	3.24–6.48 TWh	\$324M – \$648M
2030	4–16 GW	13–52 TWh	\$1.3B – \$5.2B

Source: EPRI / Epoch AI, "Scaling Intelligence: The Exponential Growth of AI's Power Needs" (August 2025)

*Illustrative estimate: * Energy = cluster × 75% avg utilization × 180 days × 24 hr · Cost @ \$100/MWh delivered energy (nominal)*

What does that duty cost in cycle life?



Illustrative effective cycle life accumulation

Use case scenario	EFC/day	180-day run	1 yr*	5 yr*
A. Outage-only (grid sags, gen-set transitions)	0.04	7	10	51
B. Throttle + checkpoint / stall bridging	1.5	270	383	1,916
C. Active rack-throttle bridging (120-sec window)	7.5	1,350	1,916	9,581
D. Bidirectional smoothing residual after capacitors	30	5,400	7,665	38,325

*Symphonic Energy estimates: * 1 yr and 5 yr assume 70% training-active calendar duty; within a run, BBU engagement varies with compute/comm/checkpoint phases. Methodology: Choukse-derived engagement bracketing – presenter’s estimate –obviously, this could be wrong.*

SECTION

5

Standards & suppression

Where deployments outrun the codes

Standards applicability – potential gaps

Each ESS tier sits under different (or no) normative guidance

Standard, Test Method, Insurance or Industry Guidance	Covers	Falls short for AI DC ESS
NFPA 855 (2026)	Indoor BESS in dedicated rooms; mandatory HMA	Multi-kWh (increasing in size) BBUs and sidecars distributed across thousands of racks [attached directly to the served equipment]
UL 9540A 5th Ed (Mar 2025)	Cell / module / unit / installation fire-propagation tests	Fresh cells at benign rate; not aged cells at high duty cycles
FM Global DS 5-33	Stationary Li-ion ESS	Explicitly excludes Li-ion BBU for distributed power systems in server racks
OCP Diablo 400 v0.7.0 §11.1	Hyperscaler rack-power spec	<i>Fire propagation tests per UL 9540A shall be performed at the cell, module, and rack levels. No external flames are permitted during rack-level testing and shall comply with the criteria not to trigger the installation base fire tests.</i> Stricter test envelope; silent on chemistry, charging, BMS, aging

IT protection vs. ESS protection – same room, competing requirements

Distributed in-rack / sidecar BBU forces both to coexist

Need	IT protection paradigm	ESS protection paradigm
Agent	Clean agent (FM-200, Novec, inert) or agent/water mist hybrid (NFPA 770)	Allow to consume fuel / burn out, possible water application for cooling, sprinklers if indoors
Envelope	Sealed sub-room for clean-agent hold time	Actively vented for off-gas evacuation with emergency power source
Detection	Smoke / heat	Off-gas (VOC, H ₂), VESDA before thermal
Hold time	10 min typical	“where designers want to distribute lithium-ion batteries throughout the data servers, however, AHJs will often require high-density sprinkler systems
Discharge	Single discharge OK	ESS can reignite hours later

Recent public-record incidents

BESS fires inside or adjacent to occupied data-center space

Seoul, South Korea

September 2025

- Lithium-ion battery fire at data-center facility
- Permanent data loss for government services reported



Sources: <https://www.koreaherald.com/article/10585058>, screengrab of WRIC ABC News video from <https://www.youtube.com/watch?v=pDq5nRDod08>

Three observations to close

- **The standards stack hasn't caught up to the deployment pattern.**

NFPA 855 covers indoor BESS in dedicated rooms; FM 5-33 excludes distributed in-rack; OCP Diablo 400 §11.1 imposes a stricter rack-level test but is silent on chemistry, sizing, charging, and aging.

- **The field duty profile of high voltage BBUs and other emerging BESS architectures isn't yet known.**

Operational C rates, depth of discharge distribution, aging trajectories, even cell chemistry — all unknown as high density, high voltage systems emerge. The safety envelope at certification changes over time, and we don't know how long end-of-life actually is in this duty regime.

- **The IT-protection and ESS-protection paradigms haven't been reconciled.**

Clean agent needs a sealed envelope; Li-ion needs ventilation. Distributed in-rack BBU, sidecars, have both in the same room. There isn't a settled engineering practice yet



Symphonic
Energy

*Energy Storage Safety & Orchestration for
Power-Dense Data Centers*

Thank you

Josh Gerber · Founder & CEO, Symphonic Energy

josh@symphonicenergy.com · 858-344-3443 · [linkedin.com/in/joshgerber](https://www.linkedin.com/in/joshgerber)