Inside the 800VDC Revolution – Part 1

深入探究 800VDC 变革 – 第 1 部分

Four-Phase 800VDC Transition, Power Rack Economics, SST, Equipment Content/MW Build, Supplier Implications

800VDC 转型四个阶段、功率机架经济学、固态变压器（SST）、设备含量/每兆瓦建设成本、供应商影响

NICOLAS BONTIGUI, JEREMIE ELIAHOU ONTIVEROS, KONRAD WANG, AND 3 OTHERS

NICOLAS BONTIGUI, JEREMIE ELIAHOU ONTIVEROS, KONRAD WANG, 以及其他 3 位作者

MAY 26, 2026 2026 年 5 月 26 日 · PAID · 付费内容

21 2 Share 分享

We'd like to thank [DG Matrix](), [Novos Power](), and [Aran Industries]()for their contributions and insights during the preparation of this deep dive.

我们衷心感谢 DG Matrix、Novos Power 和 Aran Industries 在本深度报告准备过程中所提供的贡献与见解。

Introduction: Welcome to the Power Chain Roller Coaster

引言：欢迎来到动力链过山车

Across every major industry conference in the first half of 2026, our research team kept walking past the same scene: a booth ten or fifteen people deep, leaning in to catch every word from another datacenter equipment messiah preaching the gospel of 800VDC. The pitch was the same every time. 800VDC is about to change the electrical infrastructure of the datacenter.

在 2026 年上半年的每一场重大行业会议上，我们的研究团队总会路过同样的场景：一个展位前挤满了十到十五个人，他们身体前倾，全神贯注地聆听着又一位数据中心设备“先知”宣讲 800VDC 的福音。每一次的推介内容都如出一辙：800VDC 即将彻底改变数据中心的电力基础设施。

Every architectural shift looked excessive at first. Operators spent decades keeping water and leaks out of the data hall, then GPU thermal density made running coolant right up against the precious silicon unavoidable. Each shift happened anyway, because physics and the economics of compute do not negotiate. 800VDC is next, and the logic is the same. Tokens per watt are what matters.

每一次架构变革起初看起来都显得激进。运营商花费了数十年时间防止水和泄漏进入数据大厅，随后 GPU 的热密度使得将冷却液直接引向珍贵的硅片变得不可避免。无论如何，每一次变革终究都会发生，因为物理规律和计算经济学是不容商量的。800VDC 就是下一个趋势，其逻辑如出一辙：每瓦代币（Tokens per watt）才是核心所在。

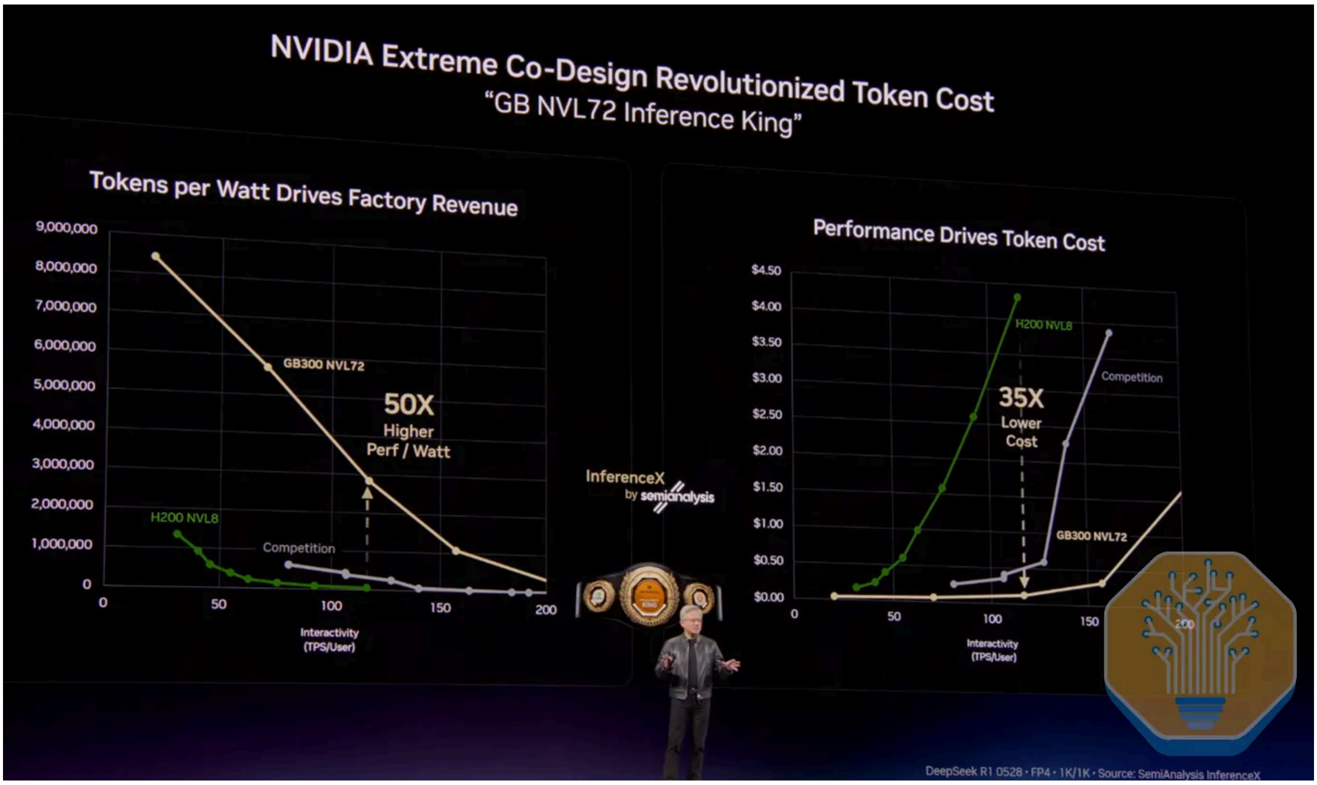

Source: Nvidia, InferenceX

来源：Nvidia, InferenceX

As GPU clusters become increasingly dense, with Kyber Ultra approaching 660kW per rack, the physics start to break down. Resistive losses scale with current squared, and at these power levels copper mass and thermal envelope exceed what fits inside a rack. Moving to 800VDC eliminates conversion stages, reduces resistive losses, and cuts facility-level power consumption by ~5%. At 1GW of IT load, that is over 50MW of continuous savings, tens of millions in annual electricity costs, or new compute capacity unlocked. For all the inference-king proponents out there, 800VDC is a transition forced by physics and motivated by system economics.

随着 GPU 集群的密度不断增加，例如 Kyber Ultra 已接近每机架 660kW，物理极限开始显现。电阻损耗随电流的平方成比例增长，在如此高的功率水平下，铜缆质量和散热范围已经超出了机架的承载能力。转向 800VDC 架构可以消除多级转换，降低电阻损耗，并将设施层面的功耗降低约 5%。在 1GW 的 IT 负载下，这意味着超过 50MW 的持续节能，每年可节省数千万美元的电费，或者释放出全新的计算容量。对于所有“推理为王”的拥趸来说，800VDC 是一场由物理定律驱动、受系统经济效益促成的必然转型。

We have been tracking this transition through our InferenceX and Industrials Models, which provide a bottom-up view of where efficiency gains materialize and which equipment categories absorb the disruption. The Industrials Model includes a dedicated 800VDC module, building up from individual accelerator architectures to a top-down view of 800VDC penetration, MW adoption, and market sizing for equipment like the power sidecar and Solid-State Transformers (SSTs).

我们一直通过 InferenceX 和工业模型（Industrials Models）跟踪这一转型，这些模型自下而上地展示了效率提升的具体环节，以及哪些设备类别正在消化这种颠覆性影响。工业模型包含一个专门的 800VDC 模块，从单个加速器架构出发，逐步构建出 800VDC 渗透率、兆瓦（MW）采用量，以及电源侧柜（power sidecar）和固态变压器（SST）等设备的市场规模等自上而下的全景视图。

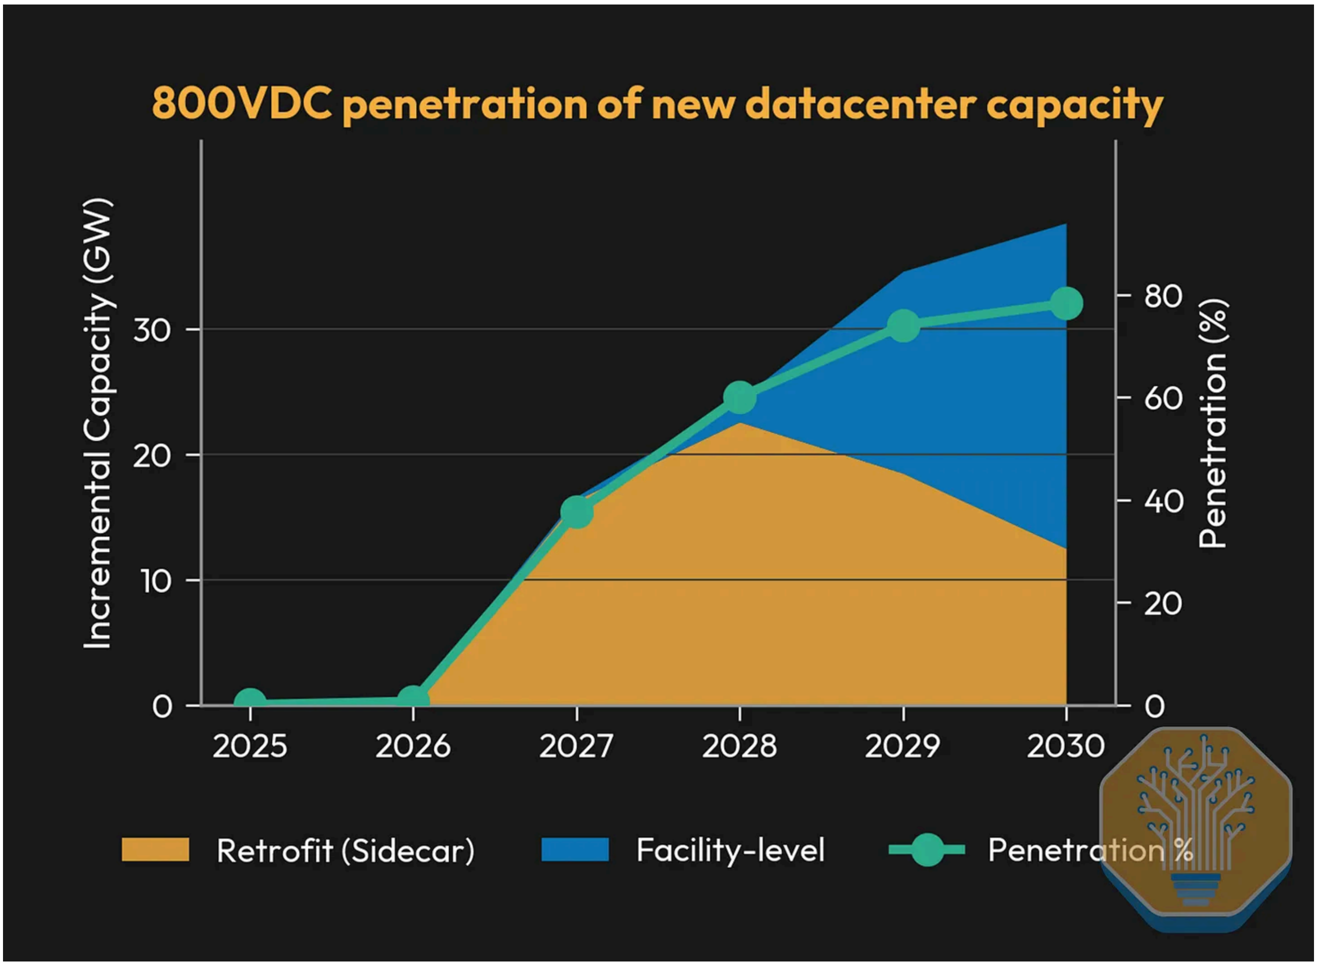

Source: SemiAnalysis Industrials Model

来源：SemiAnalysis 工业模型

This deep dive traces the transition phase by phase: from the sidecar retrofit, through faciliy-level DC distribution, to the SST endgame. For each phase, we analyze the BoM and map the changes in equipment content/MW, what survives, what gets redesigned, and what gets eliminated.

本次深度探讨将分阶段追溯这一转型过程：从侧柜改造，到设施级直流配电，再到固态变压器（SST）的终局。针对每个阶段，我们都会分析物料清单（BoM），并绘制设备内容/每兆瓦的变化图谱，明确哪些设备将继续存在、哪些需要重新设计，以及哪些将被淘汰。

The 800VDC revolution is set to dramatically change the revenue trajectory of certain suppliers. We've been tracking winners and losers for over a year in Industrials Model, which estimates the BoM for 20+ different datacenter designs broken down into 70+ equipment types and lays out the impact for 500+ suppliers. It is built on our industry-leading Datacenter Model which forecasts quarter-by-quarter MWs for 6000+ datacenters and anticipates design changes.

800VDC 革命注定将戏剧性地改变某些供应商的营收轨迹。一年多来，我们一直在工业模型中跟踪其中的赢家和输家。该模型估算了 20 多种不同数据中心设计的物料清单（细分为 70 多种设备类型），并阐述了对 500 多家供应商的影响。该模型建立在我们行业领先的数据中心模型之上，后者预测了 6000 多个数据中心逐季度的兆瓦数，并预判了设计变更。

This has enabled us to successfully call out both winners, and companies inaccurately pictured as losers by the market, before anyone else. If you are wondering whether UPS systems have a place in upcoming 800VDC distribution, what is the market opportunity for SSTs, or which suppliers are leading this transition, stick with us.

这使我们能够先于他人成功识别出真正的赢家，以及被市场误判为输家的公司。如果你想知道 UPS 系统在即将到来的 800VDC 配电中是否占有一席之地，SST（固态变压器）的市场机会如何，或者哪些供应商正在引领这一转型，请继续关注我们。

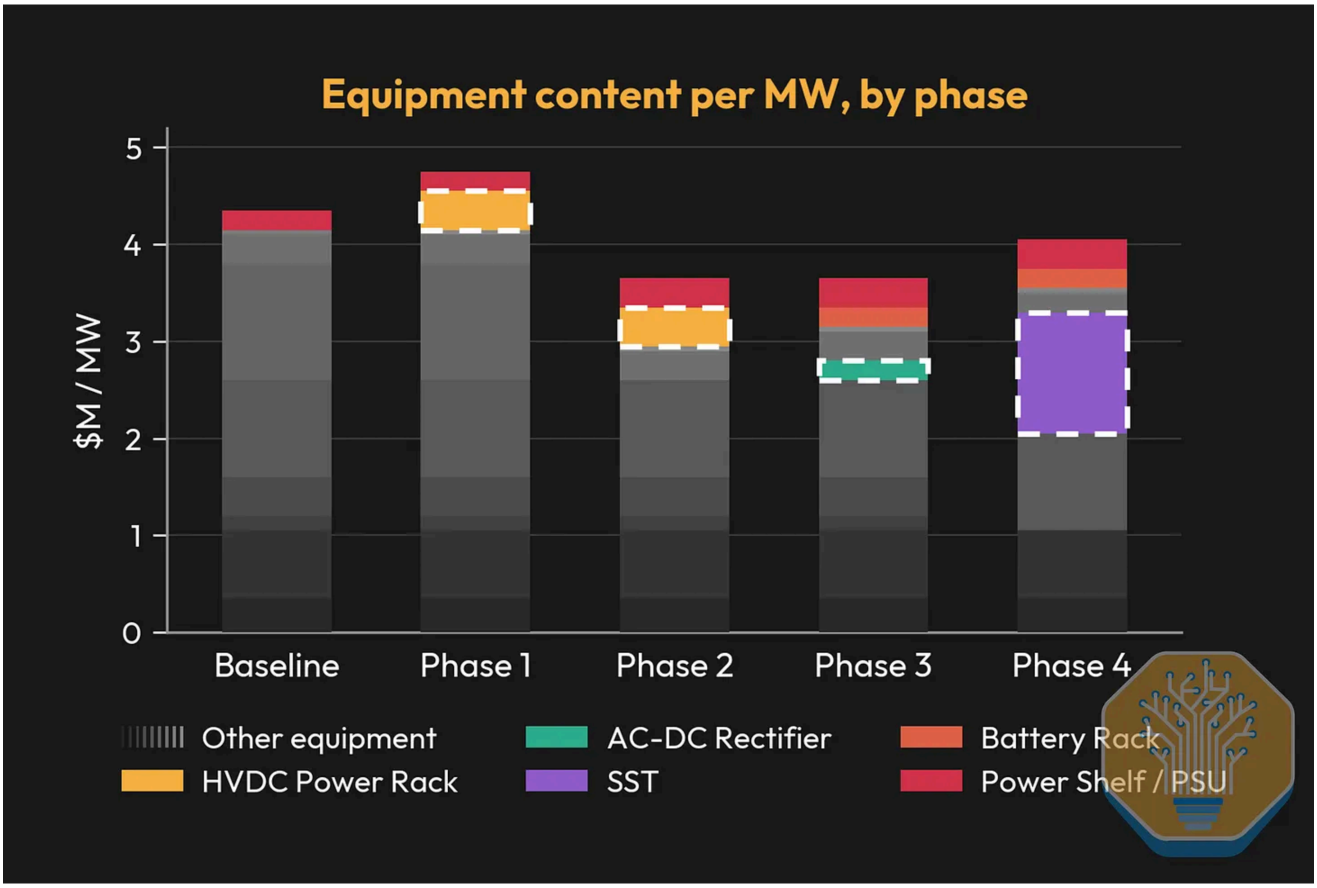

[Source: SemiAnalysis Industrials Model]

来源：SemiAnalysis 工业模型

Part 1 of this 800VDC Revolution series covers datacenter layout and equipment implications. Part 2 will focus on power electronics and the semiconductor revolution underneath it.

《800VDC 革命》系列的第一部分涵盖了数据中心布局和设备影响。第二部分将重点关注电力电子技术及其底层的半导体革命。

Understanding The Basics: What is 800VDC and Why It's Inevitable

了解基础知识：什么是 800VDC，以及为什么它是必然趋势

At its simplest, 800VDC in this context means distributing power at ~800 volts direct current through the data hall or row and into the rack, then stepping it down near the compute. The number 800 is not arbitrary, but a voltage high enough to materially reduce current (and therefore copper loss and thermal burden) while remaining within the broad regulatory and product-safety classification of "low-voltage DC" in many jurisdictions. For context, EU rules around the Low Voltage Directive scope reference DC equipment ratings up to 1,500 V DC (and AC up to 1,000 V).

简单来说，本文语境下的 800VDC 是指在数据大厅或机柜排中以约 800 伏直流电进行配电并引入机架，然后在靠近计算设备的地方进行降压。800 这个数字并非随意选定，而是因为该电压既高到足以显著降低电流（从而减少铜损和热负荷），又在许多司法管辖区内仍处于"低压直流电"的广泛监管和产品安全分类范围内。作为参考，欧盟《低电压指令》范围内的相关规则参考的直流设备额定值最高可达 1,500 V DC（交流电最高可达 1,000 V）。

Current datacenter electrical architectures generally rely on AC distribution at the facility level. Datacenters today use three-phase AC at 415V or 480V, and the topology relies on conventional UPS architectures before distributing 48-54V DC within the rack.

目前的数据中心电气架构通常依赖于设施层面的交流配电。如今的数据中心使用 415V 或 480V 的三相交流电，其拓扑结构依赖于传统的 UPS 架构，然后在机架内分配 48-54V 直流电。

This works at today's rack power levels, but starts to fail as rack densities in the next two years approach ~600 kW+, for several reasons:

这种模式在目前的机架功率水平下行之有效，但随着未来两年机架密度接近 600 kW+，由于以下几个原因，该模式开始失效：

- **Copper becomes unmanageable at 48–54 V.** A 1 MW rack at 48–54 VDC needs ~200 kg of copper busbars. At 1 GW scale, that's hundreds of tons of copper — brutal on cost, weight, installation complexity, and routing space.

  在 48–54 V 电压下，铜的使用变得难以掌控。一个功率为 1 MW 的机架在 48–54 VDC 下需要约 200 公斤的铜母线。在 1 GW 的规模下，这意味着数百吨的铜——在成本、重量、安装复杂性和布线空间方面都是沉重的负担。

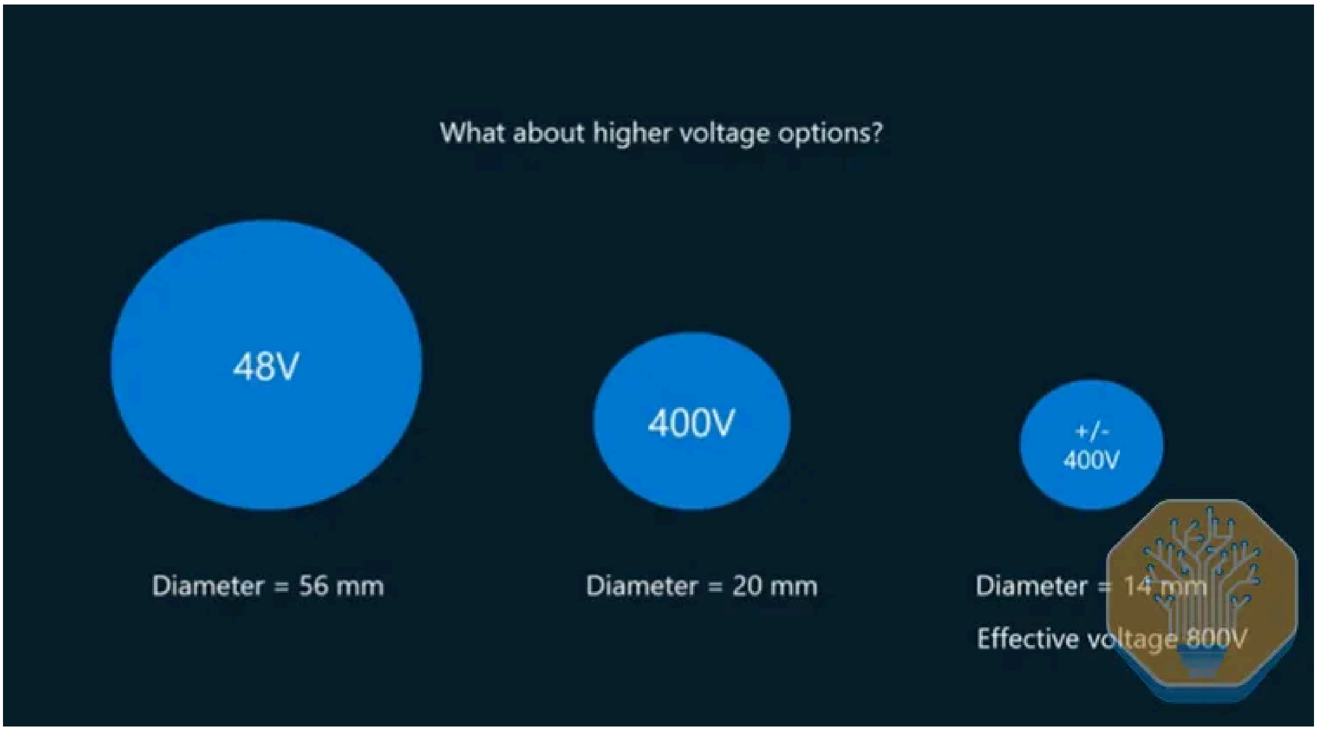

Source: Microsoft Microsoft

- **Power shelves crowd out compute.** Today's NVL72 racks already use up to 8 power shelves. At Kyber-class rack power, a 48–54V approach would require ~64U-equivalent of power hardware, effectiviely an entire rack, leaving no volume for compute.

  电源架挤占了计算空间。如今的 NVL72 机架已经使用了多达 8 个电源架。在 Kyber 级机架功率下，采用 48–54V 方案将需要约 64U 等效的电源硬件，实际上占据了整个机架，导致没有空间留给计算设备。

- **Current becomes the real limiter.** Delivering 600 kW at 48–54 V implies ~12,500A. At 800 V, that drops to ~750 A (~16.7× less), enabling dramatically smaller conductors/busbars and far lower thermal stress. If conductor resistance were held constant, I^2R losses fall ~278×, so in practice you shrink copper and "buy" size/weight reductions.

  电流成为了真正的限制因素。在 48–54V 电压下输送 600 kW 功率意味着约 12,500A 的电流。而在 800V 下，电流降至约 750A（减少约 16.7 倍），这使得导线/母线排可以大幅缩小，并显著降低热应力。如果保持导线电阻不变，I^2R 损耗将下降约 278 倍，因此在实践中，你可以通过缩小铜缆尺寸来"换取"体积和重量的减少。

- **Conversion losses compound and hurt reliability.** Stacked AC-to-DC and DC-to-DC stages reduce end-to-end efficiency, increase heat, and introduce failure points, raising cooling loads, downtime risk, and maintenance costs.

  转换损耗会叠加并损害可靠性。堆叠的 AC-to-DC 和 DC-to-DC 阶段降低了端到端效率，增加了发热，并引入了故障点，从而提高了冷却负荷、停机风险和维护成本。

At the end of the day, 800VDC is the physics enabler for 2,300W TDP chips and 600kW racks, and those 600kW racks are the direct consequence of the push for density, because density is what drives cost per token down. Cost per token is dictated by the size of the scale-up world you can build at full NVLink bandwidth: bigger domains mean wider Expert Parallelism (EP) / Tensor Parallelism (TP), MoE routing on NVLink rather than scale-out, and less serialization across decode. As we laid out in our Vera Rubin Deep Dive and GTC 2026 pieces, Nvidia's design rule is to pack compute tightly enough that copper reaches everything in the rack. Reiner Pope made this point cleanly on our friend Dwarkesh's podcast a few weeks ago, indicating that a single rack bounds the size of the expert layer you can build, because the moment an all-to-all crosses a rack boundary, it falls onto a scale-out fabric that is roughly eight times slower than NVLink.

归根结底，800VDC 是实现 2,300W TDP 芯片和 600kW 机架的物理基础，而这些 600kW 机架是追求密度的直接结果，因为密度是降低每代币成本（cost per token）的驱动力。每代币成本取决于你在全 NVLink 带宽下能构建的多大规模的纵向扩展（scale-up）领域：更大的域意味着更宽的专家并行（EP）/ 张量并行（TP）、在 NVLink 而非横向扩展（scale-out）网络上进行 MoE 路由，以及更少的解码串行化。正如我们在《Vera Rubin 深度解析》和《GTC 2026》文章中所阐述的，Nvidia 的设计原则是将计算单元打包得足够紧密，以便铜缆能覆盖机架内的所有设备。Reiner Pope 几周前在我们朋友 Dwarkesh 的播客中清晰地阐述了这一点，他指出单个机架限制了你可以构建的专家层规模，因为一旦全对全（all-to-all）通信跨越机架边界，它就会落入横向扩展网络，而其速度比 NVLink 慢大约八倍。

Bigger scale-up worlds mean denser racks, denser racks mean 600kW envelopes, and 800VDC is what makes those envelopes possible.

更大规模的扩展世界意味着更高密度的机架，更高密度的机架意味着 600kW 的功率包络，而 800VDC 正是实现这些功率包络的关键。

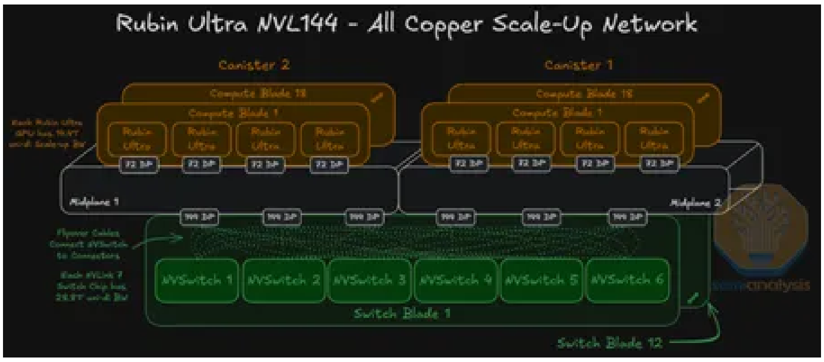

[Source: SemiAnalysis AI Networking Model]

来源：SemiAnalysis AI 网络模型

The Four Chapters of the HVDC Transition

高压直流（HVDC）转型的四个篇章

The move to 800VDC is a complex metamorphosis that rewrites the entire electrical architecture, introduces new safety standards, requires new regulatory frameworks, and, most importantly, forces operators to make very different strategic choices about when to walk away from their legacy AC distribution.

向 800VDC 的转变是一场复杂的蜕变，它重写了整个电气架构，引入了新的安全标准，需要新的监管框架，并且最重要的是，迫使运营商在何时放弃其传统的交流（AC）配电系统这一问题上，做出截然不同的战略选择。

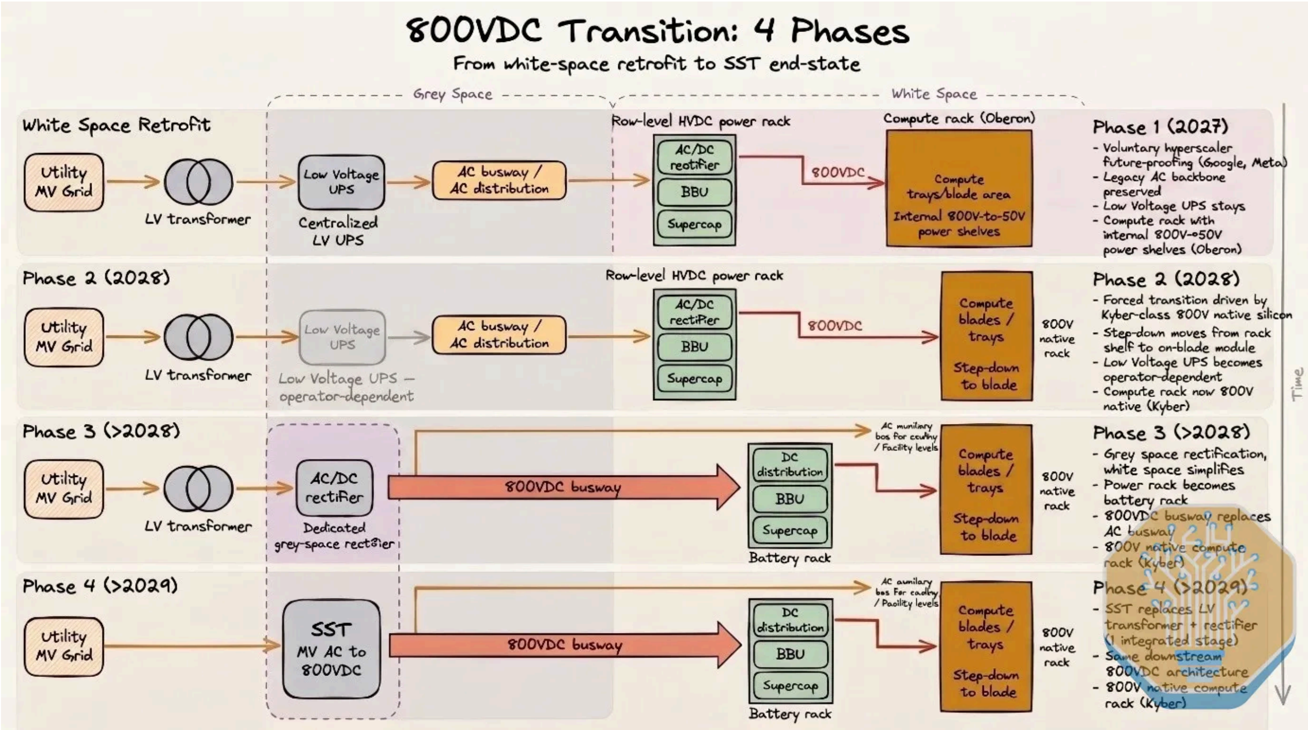

Source: SemiAnalysis SemiAnalysis

We frame the 800VDC transition as progressing through four distinct phases. Phases 1 and 2, starting in late 2026 / early 2027, retrofit the existing AC distribution into 800VDC at the rack level via the power rack. Phase 1 is the early-mover stage, driven by hyperscalers willing to pay up for future-proofing and efficiency gains. Phase 2 kicks in once 800VDC-native systems begin shipping at volume. Phase 3 rewrites the electrical architecture itself, taking 800VDC distribution facility-wide. Phase 4 is the end state, built around new pieces of equipment that promise to render much of today's electrical stack obsolete.

我们将 800VDC 的转型划分为四个不同的阶段。第 1 和第 2 阶段将于 2026 年底或 2027 年初开始，通过电力机架将现有的交流配电系统在机架层面改造为 800VDC。第 1 阶段是早期采用者阶段，由愿意为前瞻性设计和效率提升支付溢价的超大规模云服务商驱动。一旦 800VDC 原生系统开始批量出货，第 2 阶段便会启动。第 3 阶段将重写电气架构本身，在整个设施范围内实现 800VDC 配电。第 4 阶段是最终状态，围绕新型设备构建，有望使当今的大部分电气堆栈过时。

Equipment	Baseline	Phase 1	Phase 2	Phase 3	Phase 4
Timings	Today	2027	2028	>2028	>2029
GREY SPACE					
MV Transformer	0.35	0.35	0.35	0.35	0.35
MV Switchgear	0.70	0.70	0.70	0.70	0.70
LV Transformer	0.15	0.15	0.15	0.15	-
LV Switchgear	0.40	0.40	0.40	0.40	-
Generator	1.00	1.00	1.00	1.00	1.00
Central UPS	1.20	1.20	-	-	-
AC-DC Rectifier	-	-	-	0.20	-
SST	-	-	-	-	1.25
WHITE SPACE					
Busway / PDU	0.30	0.30	0.30	0.30	0.20
Rack PDU	0.05	-	-	-	-
Busbar	-	0.05	0.05	0.05	0.05
HVDC Power Rack	-	0.40	0.40	-	-
Battery Rack	-	-	-	0.20	0.20
Power Shelf / PSU	0.20	0.20	0.30	0.30	0.30
TOTAL	**4.35**	**4.75**	**3.65**	**3.65**	**4.05**
Delta vs Baseline		*0.40*	*(0.70)*	*(0.70)*	*(0.30)*

Source: SemiAnalysis Industrials Model

来源：SemiAnalysis 工业模型

The result is a progressive adoption curve for 800VDC. We expect total incremental capacity powered by 800VDC to reach ~39GW by 2030. Through Phases 1 and 2, all addressable capacity is served by sidecars, since the underlying facility is still AC-distributed and the conversion happens at the power rack. The mix inflects in 2029 as facility-level HVDC distribution becomes viable and the first 800VDC-native sites come online, shifting the conversion stage upstream from the rack to the SST or MV rectifier.

其结果是 800VDC 的采用曲线呈现渐进式增长。我们预计到 2030 年，由 800VDC 供电的总增量容量将达到约 39GW。在第一和第二阶段，所有可寻址容量都由侧挂式设备（sidecars）提供，因为底层设施仍采用交流配电，转换发生在电源机架处。这种构成比例将在 2029 年发生转折，届时设施级高压直流（HVDC）配电将变得可行，首批原生支持 800VDC 的站点将上线，从而将转换环节从机架端移至上游的固态变压器（SST）或中压（MV）整流器。

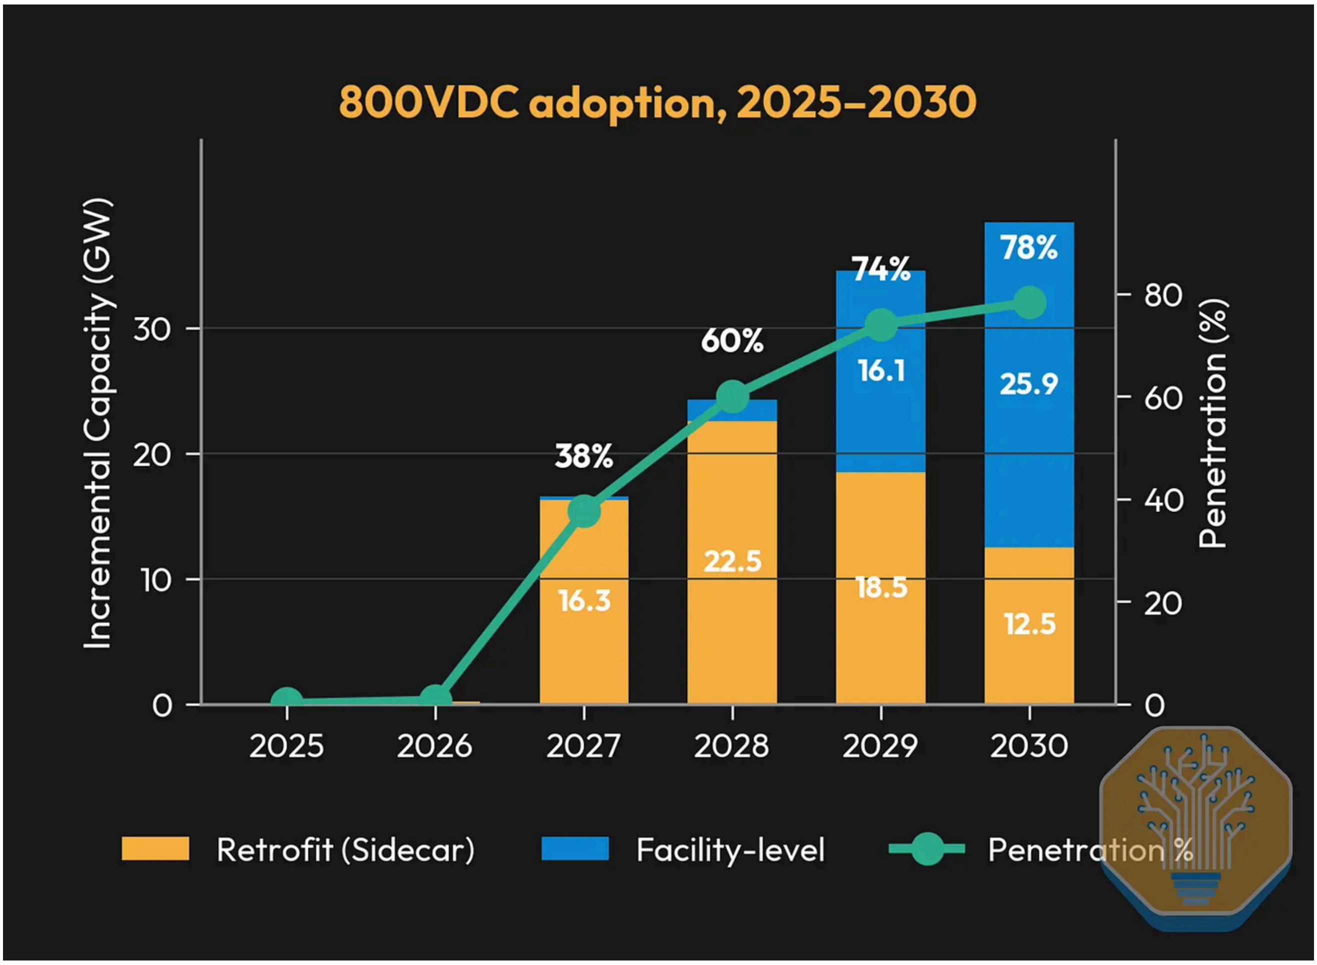

[Source: SemiAnalysis Industrials Model]

来源：SemiAnalysis 工业模型

Before diving into how the datacenter layout changes, we encourage readers to revisit [Part 1 of our datacenter anatomy series], which explains many of the core concepts behind datacenter electrical equipment.

在深入探讨数据中心布局如何演变之前，我们建议读者先回顾我们数据中心解构系列的第 1 部分，其中解释了数据中心电气设备背后的许多核心概念。

Phase 1 (2026/2027): The White Space Retrofit

第 1 阶段（2026/2027 年）：白区改造

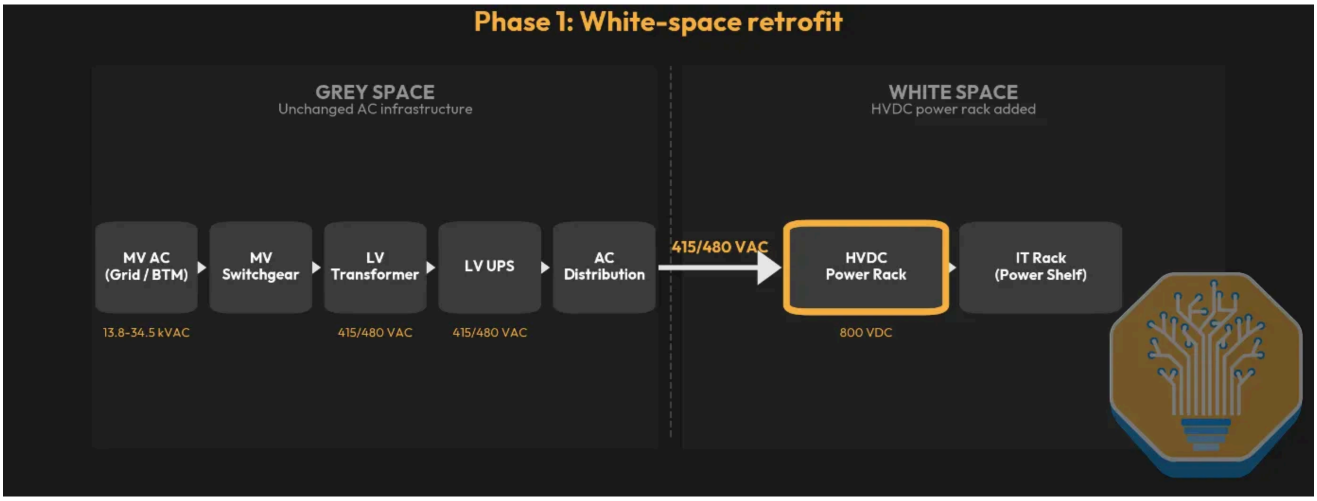

[Source: SemiAnalysis Industrials Model]

来源：SemiAnalysis 工业模型

The HVDC journey begins primarily with two operators, Google and Meta. Both have been pushing their 800VDC architectures through the OCP working groups for over 18 months, most visibly with the Mt. Diablo reference design, first announced in October 2024 and published as an open specification in May 2025. Neither is being forced into the transition, but they are doing it to take a leading position in the upcoming shift and because they want to squeeze every megawatt and every point of efficiency out of their existing power chain before the rest of the market is forced to catch up.

高压直流（HVDC）的演进历程主要始于两大运营商：Google 和 Meta。18 个月以来，双方一直通过 OCP 工作组推动其 800VDC 架构，其中最引人注目的是 Mt. Diablo 参考设计——该设计于 2024 年 10 月首次公布，并于 2025 年 5 月作为开放规范发布。这两家公司并非被迫转型，而是为了在即将到来的变革中占据主导地位，并希望在市场其他参与者被迫跟进之前，从其现有的电力链中榨取每一兆瓦的容量和每一个百分点的效率。

This matters because 800VDC is not yet a hard requirement. The chip generations ramping in late 2026 and 2027, like Vera Rubin NVL72, top out at rack densities of 180-220kW. Three-phase AC can still deliver that without hitting the physical limits of conductor sizing or distribution losses. Phase 1 is therefore voluntary future-proofing, not a forced response to a hardware constraint.

这一点至关重要，因为 800VDC 目前还不是一项硬性要求。即将在 2026 年底和 2027 年量产的芯片世代（如 Vera Rubin NVL72），其机架功率密度峰值在 180-220kW 之间。三相交流电（AC）仍能满足这一需求，而不会触及导体尺寸或传输损耗的物理极限。因此，第一阶段是自发性的前瞻性布局，而非针对硬件瓶颈的被迫响应。

This initial phase kicks off the “White Space Retrofit” era. New HVDC hardware, primarily a row-level cabinet called the HVDC power rack, layers on top of existing white space infrastructure rather than replacing it. The datacenter’s electrical backbone stays intact. Same transformers, same UPS, same switchgear, same ATS.

这一初始阶段开启了“白区改造”时代。新型高压直流（HVDC）硬件——主要是被称为 HVDC 电源机柜的列级机柜——叠加在现有的白区基础设施之上，而非将其替换。数据中心的电气骨干保持不变。同样的变压器、同样的 UPS、同样的开关柜、同样的 ATS。

Power Flow Overview with HVDC Power Rack

带有高压直流（HVDC）电源机架的功率流概览

At the facility level, Medium-voltage AC enters the grey space and is stepped down via transformer to 415V or 480V three-phase AC. That feeds into a UPS, which performs double conversion (AC-DC-AC), then outputs 415V AC. AC is then distributed through the data hall via busway. So far, this is the traditional power flow [we have extensively covered in previous articles.]

在设施层面，中压交流电进入灰色空间，并通过变压器降压至 415V 或 480V 三相交流电。该电流进入 UPS，进行双变换（交流-直流-交流），然后输出 415V 交流电。随后，交流电通过母线槽在数据大厅内进行分配。到目前为止，这仍是我们之前文章中广泛涵盖的传统功率流。

The change occurs when we get closer to the IT racks. Instead of feeding 415V directly into in-rack power supply units, the AC feed now terminates at a standalone 42U cabinet named the HVDC power rack deployed at the row level.

当电流接近 IT 机架时，变化发生了。交流馈线不再直接进入机架内的电源单元（PSU），而是终止于一个部署在列级的、名为 HVDC 电源机架的独立 42U 机柜。

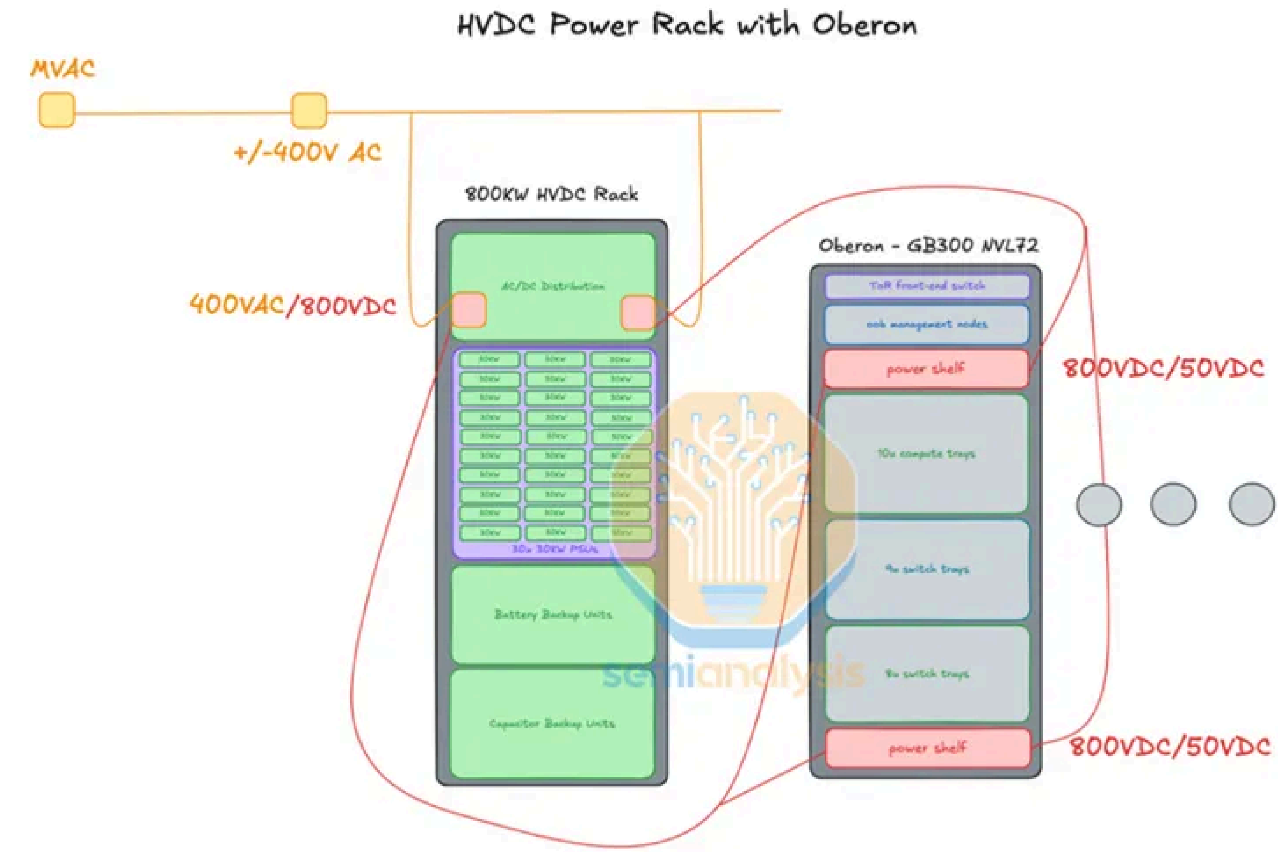

Source: SemiAnalysis SemiAnalysis

The rack receives AC from the overhead busway and outputs 800VDC through cable to adjacent IT racks. Inside, it performs three jobs: rectification of 415V AC to 800VDC, BBU modules for ride-through during outages, and optionally, capacitor shelves for transient buffering during GPU load spikes.

机架从架空母线槽接收交流电，并通过电缆向相邻的 IT 机架输出 800V 直流电。在其内部承担着三项任务：将 415V 交流电整流为 800V 直流电；配备用于停电期间维持运行的 BBU 模块；以及可选配的电容架，用于在 GPU 负载激增时提供瞬态缓冲。

In a Nutshell: The Power Rack

简而言之：电源机架

It is worth looking in more detail into the building block that underpins Phases 1 and 2 of the 800VDC transition: the disaggregated power rack. This is a dedicated rack that consolidates AC-to-DC rectification, energy storage (BBU and/or capacitor bank), and power management into a single unit, freeing the compute rack to be entirely dedicated to GPUs, networking, and cooling. Microsoft’s Mt Diablo project originated the concept; the OCP Diablo 400 specification, co-authored by Google, Meta, and Microsoft, standardizes it.

深入研究支撑 800VDC 转型第一和第二阶段的构建模块——解耦式电源机架（disaggregated power rack）是非常有意义的。这是一个专用机架，它将交流到直流的整流、能量存储（BBU 和/或电容器组）以及电源管理整合到一个单元中，从而使计算机架能够完全专注于 GPU、网络和冷却。微软的 Mt Diablo 项目发起了这一概念；由 Google、Meta 和微软共同编写的 OCP Diablo 400 规范则将其标准化。

Key components that are commonly found in a sidecar power rack:

边车电源机架中常见的关键组件：

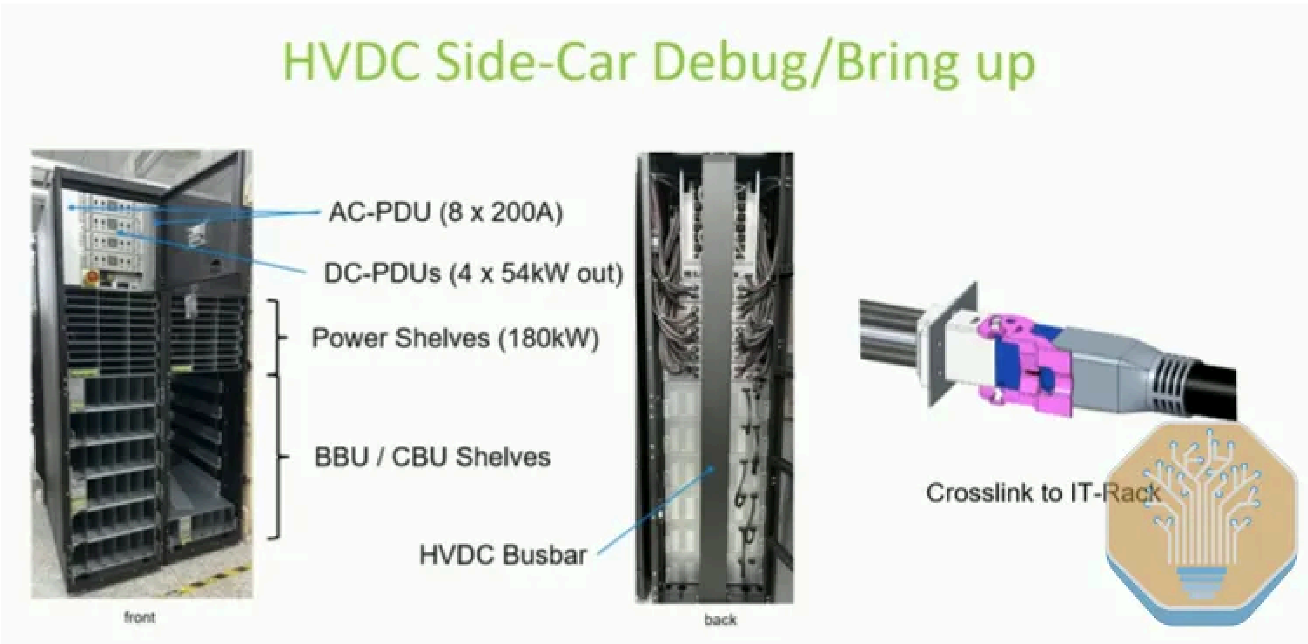

Source: Rittal 威图

Component	Function
AC-PDU	Receives 200A AC whips from the tap box, distributes AC to the power shelves
AC-DC Power Shelves	Rectify 415-480V AC to 800VDC (or +/-400VDC bipolar per Mt Diablo)
BBU Shelves	Lithium-ion battery backup for ride-through during utility outage
CBU / Supercapacitor Shelves (optional)	Capacitor Backup Units. EDLC or HSC cells for peak buffering during GPU load transients
HVDC Busbar	Distributes 800VDC from the rectifier shelves to the DC-PDU output stage.
DC-PDU	Connects the HVDC busbar to multiple 50kW cables going to adjacent IT racks
Power Management / BMS	Rack-level controller managing cell health, charge cycles, thermal monitoring, and CAN/I2C communications

Source: SemiAnalysis SemiAnalysis

But the sidecar concept did not emerge fully formed. It evolved through several OCP rack and power specification versions. The earlier iterations (ORv2 at 12V, ORv3 at 48V, and the HPR V1/V2 variants that pushed single-rack 48V designs up to ~190 kW with liquid-cooled busbars and upgraded 72 kW power shelves) are covered in our Datacenter Anatomy series. Here we focus on the versions directly relevant to 800VDC: the disaggregated sidecar designs where the voltage transition occurs.

但侧车（sidecar）概念并非一蹴而就，而是通过多个 OCP 机架和电源规范版本演进而来。早期的迭代版本（如 12V 的 ORv2、48V 的 ORv3，以及通过液冷母线和升级后的 72 kW 电源机架将单机架 48V 设计推升至约 190 kW 的 HPR V1/V2 变体）已在我们的《数据中心解剖》系列中有所涵盖。本文我们将重点关注与 800VDC 直接相关的版本：即发生电压转换的解耦式侧车设计。

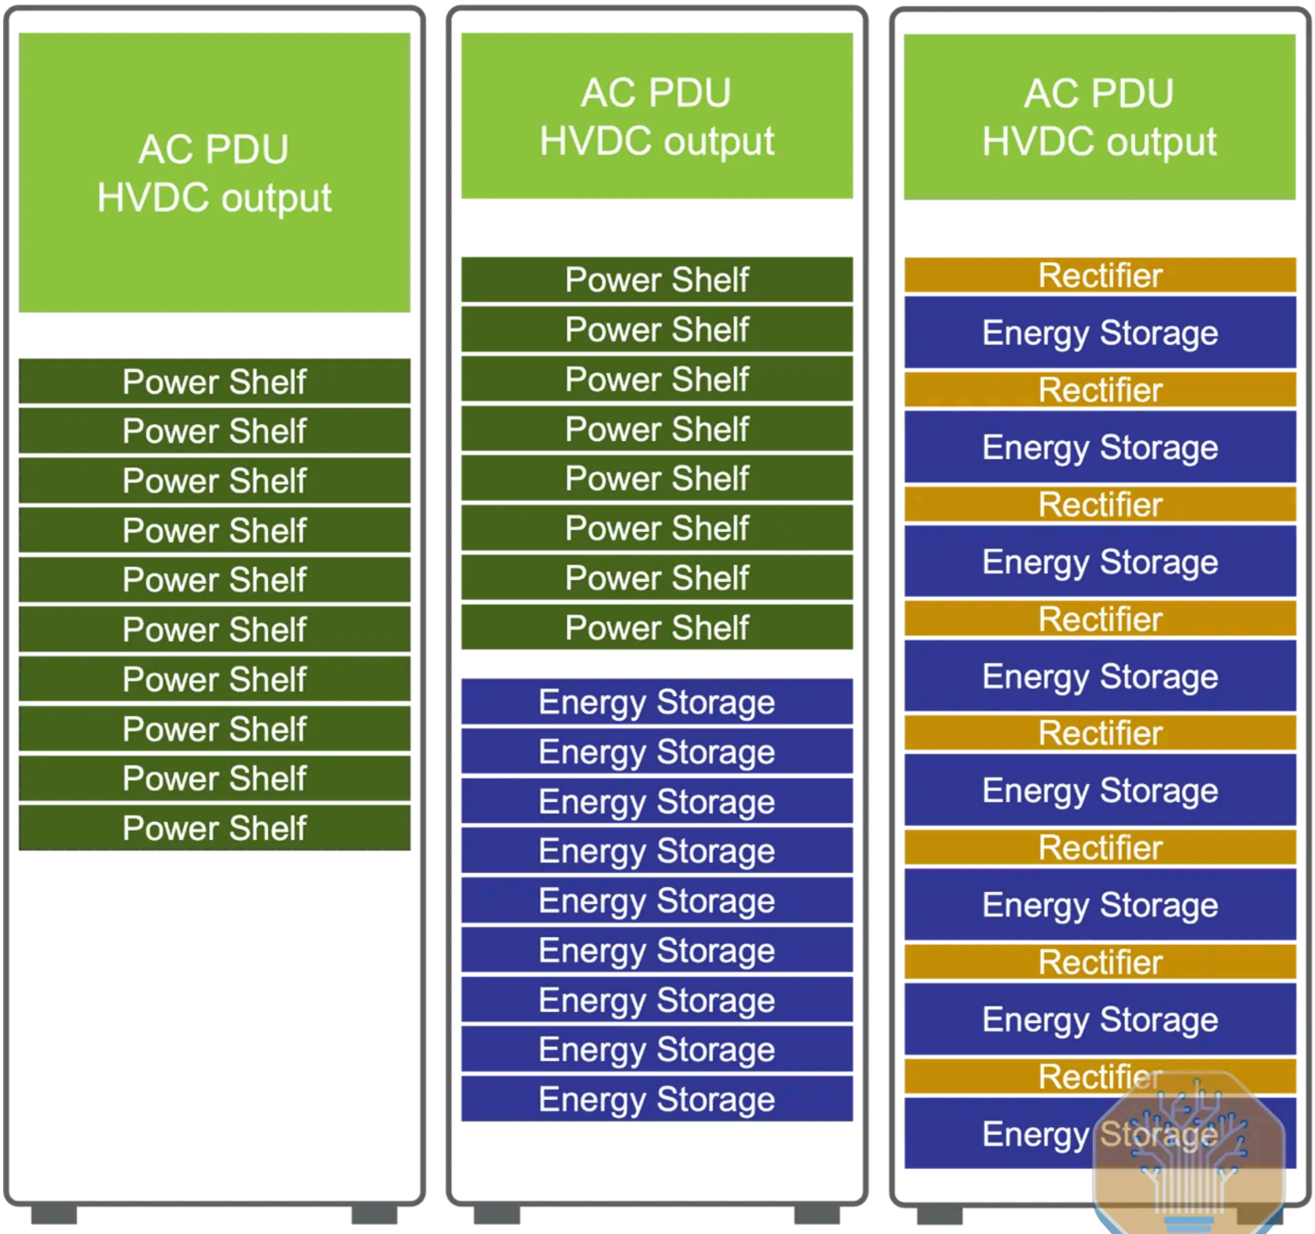

Three examples of power rack configurations

Source: OCP OCPp

ORv3 HPR V3: The Disaggregation Threshold (50V Sidecar, up to 300 kW)

ORv3 HPR V3：解耦临界点（50V Sidecar，最高 300 kW）

HPR V3 is really where power and compute separate into distinct racks, the genesis of the "sidecar" concept. PSU and BBU shelves move into a dedicated 50VDC side power rack connected to the IT rack through horizontal busbars at the top and bottom of both. Both remain ORv3 HPR standard form factor. Power capacity tops out at 300 kW, limited by the horizontal crosslinks and the air-cooled vertical busbar inside the power rack.

HPR V3 真正实现了电源与计算资源在不同机架上的分离，这也是“Sidecar（侧柜）”概念的起源。电源模块（PSU）和电池备用模块（BBU）机架移至专用的 50VDC 侧置电源机架中，通过顶部和底部的水平母线与 IT 机架相连。两者均保持 ORv3 HPR 标准外形尺寸。受限于水平互连链路和电源机架内部的风冷垂直母线，其功率容量最高可达 300 kW。

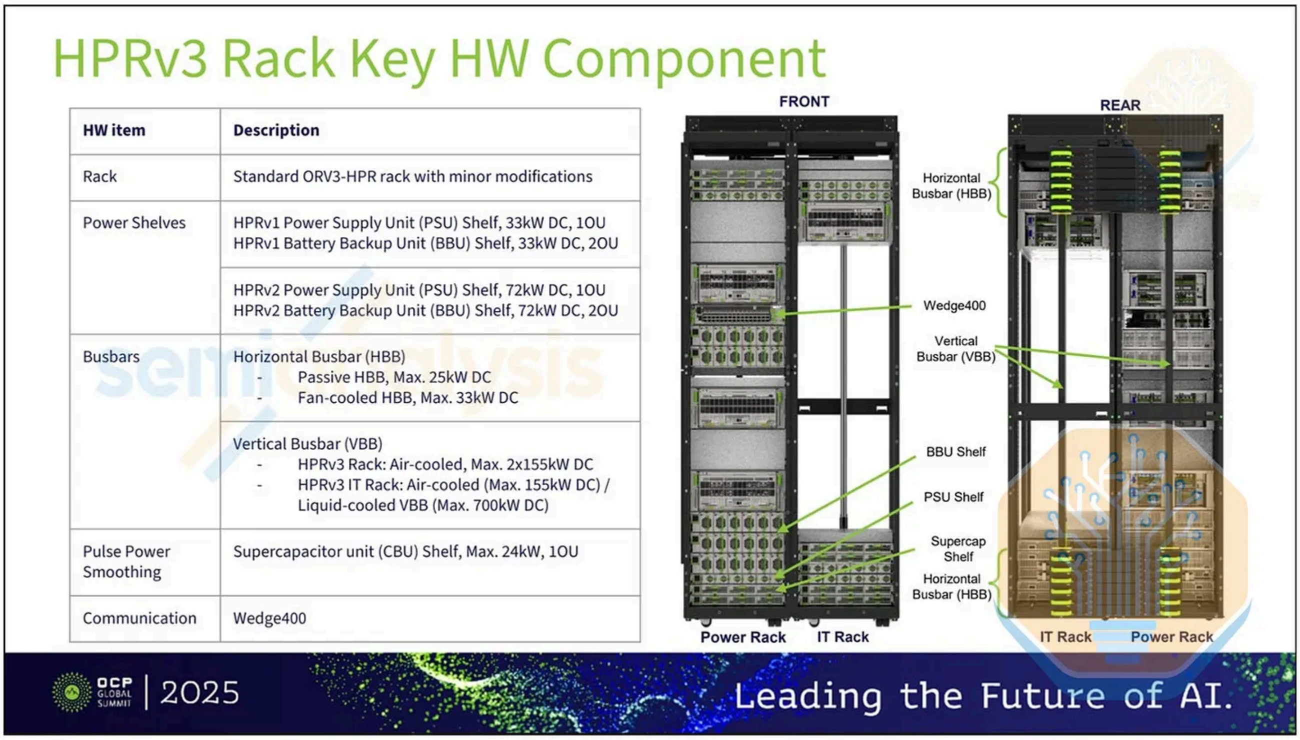

Source: OCP OCPp

The insight is putting power conversion hardware in a rack optimized for power, with appropriate cooling, safety, and serviceability, rather than cramming it into a rack optimized for compute. The V3 power rack can be serviced independently, shrinking the blast radius of power-side failures. But V3 still distributes at 50VDC, which means busbar currents remain high (6,000A at 300 kW) and the crosslinks become the bottleneck.

核心见解在于将功率转换硬件置于针对电力优化的机架中，并配备适当的冷却、安全和可维护性设计，而不是将其强行塞入针对计算优化的机架内。V3 电力机架可以独立进行维护，从而缩小了电力侧故障的影响范围。但 V3 仍采用 50VDC 进行配电，这意味着母线电流依然很高（300 kW 时为 6,000A），且跨接链路成为了瓶颈。

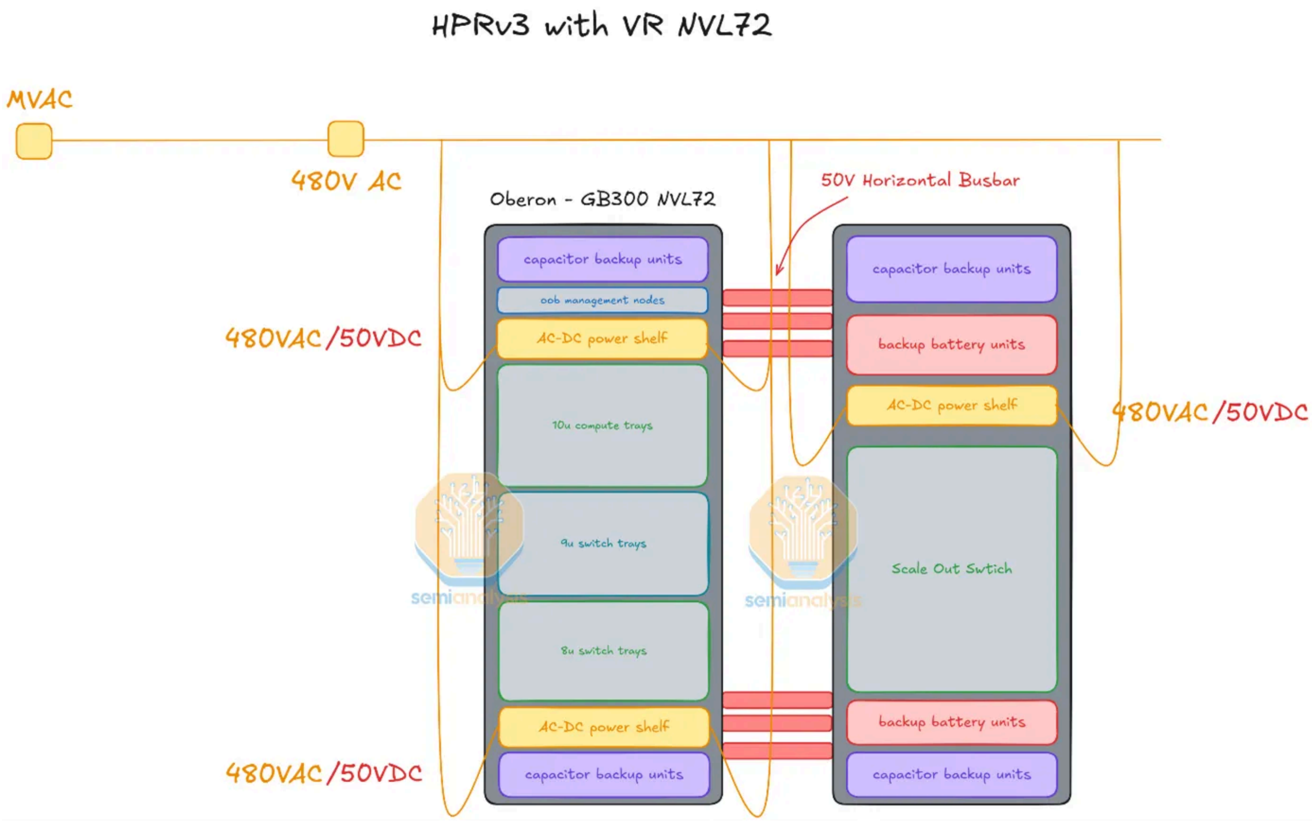

Source: SemiAnalysis SemiAnalysis

This persists today. Even the VR NVL72 rack, when fed by an HVDC power rack at 800VDC (Nvidia spec) or ±400VDC (OCP spec), still distributes internally over a 50V busbar. A DC-DC power shelf inside the rack steps the high-voltage DC down to 50VDC before it reaches the compute trays. At the far end, VRMs on the GPU board convert from 50V to sub-1V.

这种情况一直延续至今。即使是 VR NVL72 机架，在由 800VDC（Nvidia 规范）或 ±400VDC（OCP 规范）的高压直流电源机架供电时，其内部仍通过 50V 母线进行分配。机架内部的 DC-DC 电源层会在高压直流电到达计算托盘之前将其降压至 50VDC。在末端，GPU 板上的电压调节模块（VRM）再将 50V 转换为低于 1V 的电压。

We have more detailed power and architecture details in our [VR NVL72 Component BoM and Power Budget Model].

我们在 VR NVL72 组件物料清单 (BoM) 和功耗预算模型中提供了更详尽的功率及架构细节。

ORv3 HPR V4: HVDC Sidecar at +/-400VDC (up to 800 kW)

ORv3 HPR V4：+/-400VDC 高压直流 (HVDC) 侧柜（功率高达 800 kW）

HPR V4 is the version that bridges the OCP HPR lineage into the HVDC era. It makes two critical changes: the voltage steps up from 50VDC to +/-400VDC (800V total), and the busbar-based crosslink is replaced with discrete power cables.

HPR V4 是将 OCP HPR 系列带入高压直流 (HVDC) 时代的版本。它进行了两项关键改进：电压从 50VDC 提升至 +/-400VDC（总计 800V），并且原有的基于母线槽 (busbar) 的互连被离散电源线缆所取代。

- **Architecture**: PSU and BBU shelves move into a +/-400VDC side power rack, which also houses AC input and DC output PDUs

  架构：电源单元 (PSU) 和电池备份单元 (BBU) 机架移至 +/-400VDC 侧边电源柜中，该机柜同时容纳交流 (AC) 输入和直流 (DC) 输出配电单元 (PDU)

- **Power delivery**: The power rack connects to the IT rack through 16x 50 kW HVDC cables (replacing the horizontal busbars of V3), each carrying +/-400VDC

  电力传输：电源机架通过 16 条 50 kW HVDC 电缆（取代了 V3 的水平母线）与 IT 机架连接，每条电缆承载 +/-400VDC

- **Power capacity**: Up to 800 kW maximum. If capacitor-based energy storage (CBUs) occupies half the BBU slots, effective capacity drops to ~400 kW

  功率容量：最高可达 800 kW。如果电容式储能单元 (CBU) 占据了一半的 BBU 插槽，有效容量将降至约 400 kW

- **AC input**: 200A single conductor wire from tap boxes

  交流输入：来自分线箱的 200A 单芯导线

- **Form factor**: Same ORv3 HPR rack dimensions as V3

  外形尺寸：与 V3 相同的 ORv3 HPR 机架尺寸

- **Why cables instead of busbars**: At the power levels V4 targets (400-800 kW), the horizontal busbar crosslinks from V3 become current-limited. Replacing them with discrete cables allows each cable to be independently routed, fused, and managed, and eliminates the single-point busbar as a thermal and mechanical constraint

为什么使用电缆而非母线：在 V4 目标功率水平（400-800 kW）下，V3 的水平母线互连会出现电流限制。使用独立电缆替代母线，可以让每根电缆独立布线、熔断和管理，并消除了单点母线在散热和机械方面的限制。

V4 effectively represents the “pre-Diablo” state of HVDC sidecar design, developed primarily by Meta’s rack and power team. It proved the concept of disaggregated HVDC power delivery but was not yet a multi-vendor, multi-hyperscaler specification.

V4 实际上代表了高压直流（HVDC）侧置式设计的“前 Diablo”状态，主要由 Meta 的机架与电源团队开发。它证明了分布式 HVDC 供电的概念，但尚未成为多供应商、多超大规模云厂商通用的规范。

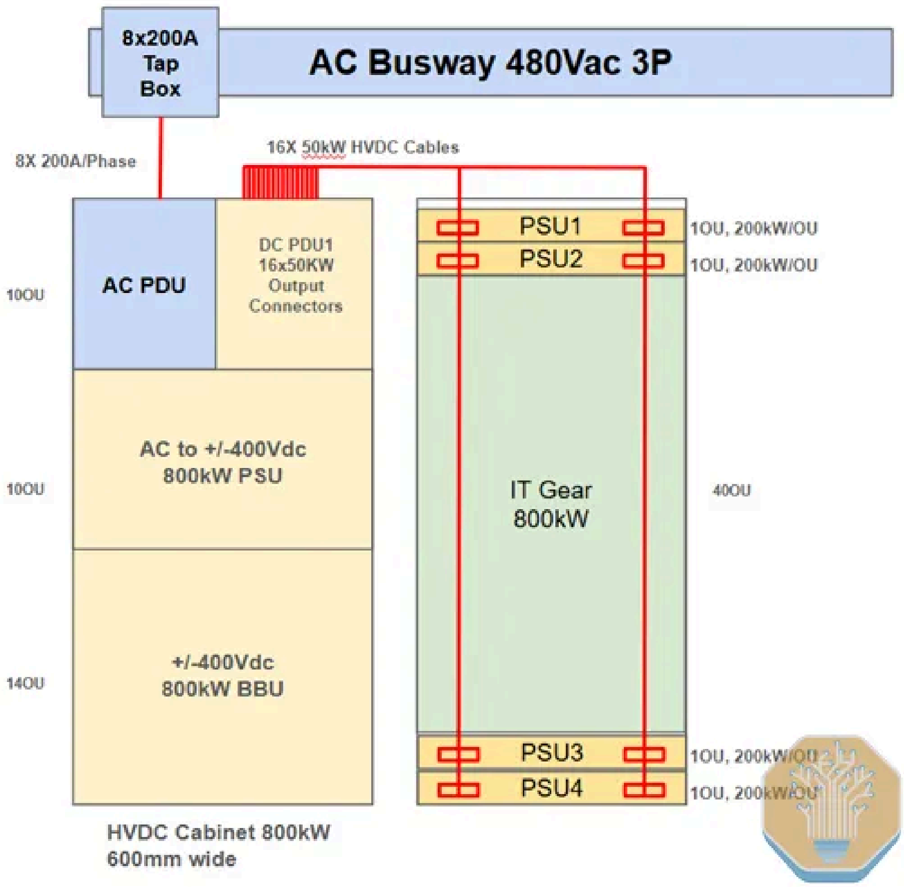

Source: Meta 元数据

The Diablo 400 Specification: Standardizing the HVDC Sidecar

Diablo 400 规范：高压直流（HVDC）侧挂式电源标准化

The Diablo 400 specification (named after Mt Diablo, Microsoft's original internal project name) formalizes and standardizes the HVDC sidecar concept that HPR V4 pioneered. Co-authored by Google, Meta, and Microsoft, Diablo 400 was released as a draft specification (v0.5.2) in May 2025, with a subsequent v0.7.0 revision following industry feedback.

Diablo 400 规范（以微软最初的内部项目名称 Mt Diablo 命名）将 HPR V4 首创的 HVDC 侧挂式电源概念正式化并标准化。Diablo 400 由 Google、Meta 和微软共同起草，于 2025 年 5 月发布了规范草案（v0.5.2），随后根据行业反馈修订至 v0.7.0 版本。

What Diablo 400 standardizes that HPR V4 did not:

Diablo 400 实现了标准化，而 HPR V4 未能做到的方面：

- **Multi-vendor interoperability**: Standardized electrical and mechanical interfaces so that PSU shelves from Delta, power management from Advanced Energy, busbars from TE Connectivity, and BBUs from multiple suppliers can all work together in a single rack

  多厂商互操作性：标准化的电气和机械接口，使得来自 Delta 的 PSU 机架、来自 Advanced Energy 的电源管理系统、来自 TE Connectivity 的母线以及来自多个供应商的 BBU 都能在同一个机架中协同工作

- **Dual voltage support**: The base specification defines +/-400VDC bipolar as the standard configuration (3-wire: +400V, -400V, and Common/Midpoint/Return at the rectifier shelf output), with 800VDC monopolar as an explicit design option (2-wire: 800VDC and Return, safety-isolated from PE ground)

  双电压支持：基础规范将 +/-400VDC 双极定义为标准配置（3 线制：整流机架输出端为 +400V、-400V 以及公共端/中点/回流线），并将 800VDC 单极作为明确的设计选项（2 线制：800VDC 和回流线，与 PE 地线安全隔离）

- **Power range**: 100 kW to 1 MW per IT rack

  功率范围：每个 IT 机架 100 kW 至 1 MW

- **PSU design**: 3-phase AC input, +/-400VDC output. PSU modules are front-of-rack accessible, hot-swappable, and hot-pluggable, with droop and active current sharing between PSUs and power shelves

  PSU 设计：三相交流输入，+/-400VDC 输出。PSU 模块支持机架前部操作、热插拔和热拔插，且在 PSU 与电源机架之间具备下垂控制和主动均流功能。

- **Cable spec**: Voltage drop budget of 0.1% at 5m cable length for output cables between power rack and IT rack

  电缆规格：电源机架与 IT 机架之间的输出电缆，在 5 米电缆长度下的电压降预算为 0.1%

- **Holdup time**: Minimum 20 ms without energy storage at 100% loading; distributed holdup acceptable between the AC/DC PSU in the Diablo 400 rack and downstream DC/DC converters located outside the rack

  保持时间：100% 负载下，无储能设备的最小保持时间为 20 ms；允许在 Diablo 400 机架内的 AC/DC PSU 与机架外下游的 DC/DC 转换器之间进行分布式保持

- **Mechanical**: Sliding shelves for push-in/pull-out of large building blocks (e.g., 4OU BBU), blind-mate connectors with static rail/sliding rails for PSU/BBU/CBU hot-swap

  机械结构：用于大尺寸构建模块（如 4OU BBU）推入/拉出的滑动托架；带有静态导轨/滑动导轨的盲插连接器，支持 PSU/BBU/CBU 热插拔

- **Seven standardization areas**: Connectivity, power rack form factor, AC-DC PSU topology, DC-DC modules, redundancy architecture (single/dual feed, N+x), safety standards for HVDC and liquid-cooled systems, and data/power management backplane

  七个标准化领域：连接性、电源机架外形尺寸、AC-DC PSU 拓扑结构、DC-DC 模块、冗余架构（单/双路供电，N+x）、HVDC 和液冷系统的安全标准，以及数据/电源管理背板

The choice of 400VDC as the nominal voltage was deliberate. As Google's engineers stated at OCP EMEA 2025: "selecting 400 VDC as the nominal voltage allows us to leverage the supply chain established by electric vehicles, for greater economies of scale, more efficient manufacturing, and improved quality and scale." In the bipolar configuration, each individual rail sits only 400V from the grounded midpoint, keeping the system within the voltage range where mature EV-grade power electronics (650V GaN FETs, 400V-class capacitors, connectors, and fuses) can be used directly.

选择 400VDC 作为标称电压是经过深思熟虑的。正如 Google 工程师在 OCP EMEA 2025 大会上所言："选择 400 VDC 作为标称电压，使我们能够利用电动汽车建立的供应链，从而获得更大的规模经济效应、更高效的制造工艺，以及更高的质量和规模。" 在双极配置中，每个单独的母线距离接地中点仅 400V，这使系统保持在成熟的电动汽车级功率电子器件（如 650V GaN FET、400V 级电容器、连接器和保险丝）可以直接使用的电压范围内。

No One-Size-Fits-All 没有万能的解决方案

There is no one-size-fits-all 800VDC power rack. Yes, Diablo 400 provides a shared base specs, but the reality on the ground is fragmented. Nvidia sits entirely outside it and is developing a monopolar 800V reference design at 660kW, with air-cooled samples and production in mid-2026, and a liquid-cooled VR Ultra variant sampling in late-2026.

目前并没有放之四海而皆准的 800VDC 电源机架。诚然，Diablo 400 提供了一套共享的基础规格，但实际情况却是碎片化的。Nvidia 完全独立于该标准之外，正在开发一种功率为 660kW 的单极 800V 参考设计，其风冷样品和量产预计在 2026 年中期实现，而液冷版 VR Ultra 变体则将在 2026 年底提供样品。

Even within Diablo 400, the three co-authors diverge meaningfully. Meta runs 600-800kW with 50kW HVDC output cables and 8x 200A AC input whips. Google push to 900Kw by reallocating rack space from BBU and supercap slots to PSUs, run 100kW output cables, and need 12 AC whips at the 1.1MW roofline. Amazon's design lands at 800kW on ±400V. Microsoft co-authored the spec but we believe they are making slower progress.

即使在 Diablo 400 规范内部，三位共同作者的设计也存在显著差异。Meta 运行功率为 600-800kW，采用 50kW 高压直流（HVDC）输出电缆和 8 条 200A 交流（AC）输入电缆。Google 通过将机架空间从备用电池单元（BBU）和超级电容插槽重新分配给电源单元（PSU），将功率推高至 900kW，运行 100kW 输出电缆，并在 1.1MW 的峰值功率下需要 12 条交流输入电缆。Amazon 的设计则定格在 ±400V 电压下的 800kW。Microsoft 虽然共同起草了该规范，但我们认为其进展相对较慢。

Besides, an alternative sidecar topology uses an LV-input SST in place of the conventional rectifier-plus-PSU stack, like DG Matrix's Interport Cell Series.

此外，另一种旁置式拓扑结构使用低压输入固态变压器（SST）来取代传统的整流器加电源单元（PSU）堆栈，例如 DG Matrix 的 Interport Cell 系列。

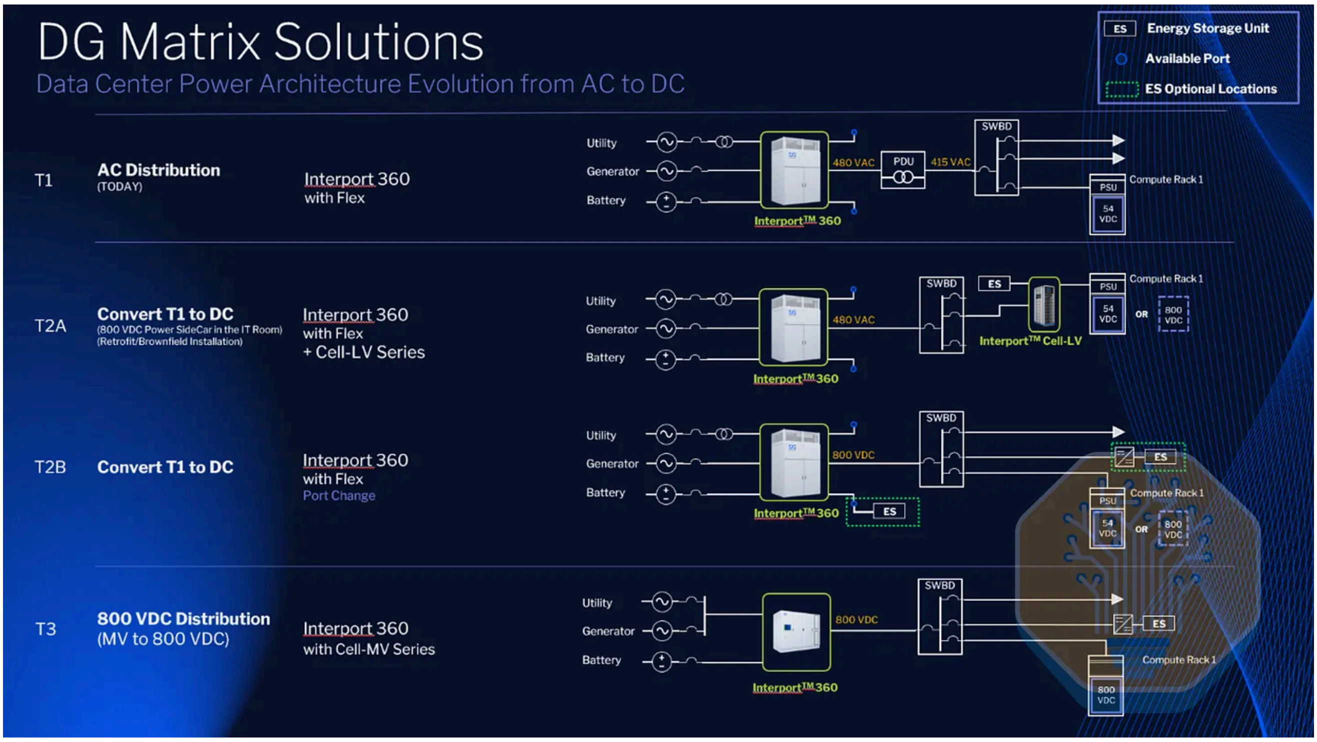

Source: DG Matrix DG Matrix

The cost of the power rack

电源机架的成本

The HVDC power rack is the headline new-equipment cost in the early retrofit phases. We estimate the ASP for the Power Rack to reach $400-500k per unit, roughly 10x the ~$40k ASP of standard AC power-rack equipment. On a deployed-MW basis, that lands near $500k/MW.

在高压直流（HVDC）改造的早期阶段，电源机架是新增设备成本中的核心部分。我们预计电源机架的平均售价（ASP）将达到每台 40-50 万美元，约为标准交流电源机架设备（ASP 约 4 万美元）的 10 倍。按部署的兆瓦（MW）计算，其成本接近 50 万美元/MW。

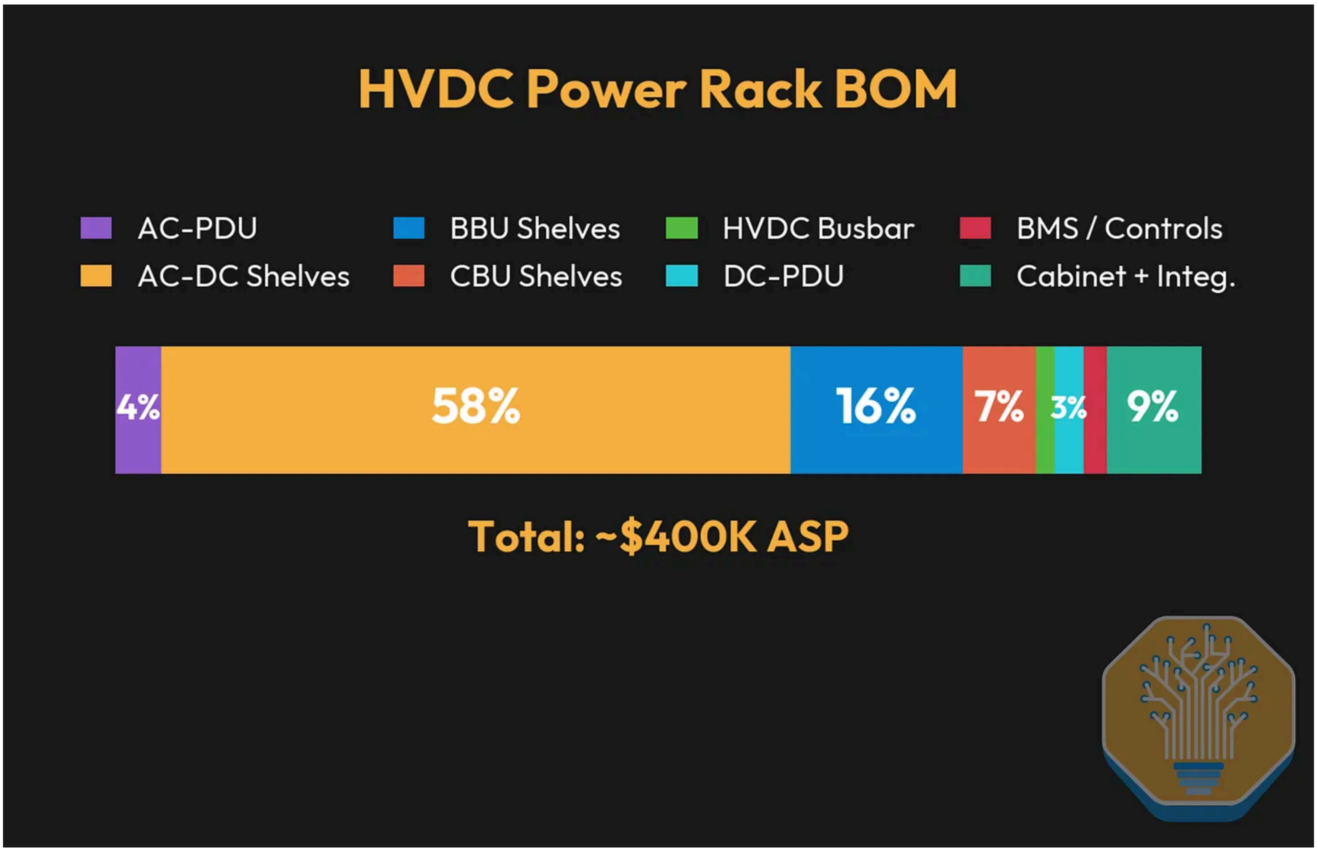

Source: SemiAnalysis Industrials Model

来源：SemiAnalysis 工业模型

The Sidecar Market Opportunity and TAM sizing

Sidecar 市场机遇与 TAM 规模估算

In our SemiAnalysis Industrials Model, we size the 800VDC equipment TAM, specifically for the sidecar (power rack) and Solid State Transformer (SST), by applying this phase-by-phase adoption timelines to incremental datacenter capacity build and doing a chip-by-chip SKU calculation.

在我们的 SemiAnalysis 工业模型中，我们通过将这种分阶段的采用时间表应用于增量数据中心建设容量，并进行逐芯片的 SKU 计算，从而估算了 800VDC 设备（特别是 Sidecar 电源机架和固态变压器 SST）的潜在市场总量（TAM）。

We expect sidecar TAM to peak at ~$11B in 2028 before declining as facility-level 800VDC takes share in Phase 3. We assume a power rack content of $0.5M/MW.

我们预计旁置式（sidecar）解决方案的可服务市场总量（TAM）将在 2028 年达到约 110 亿美元的峰值，随后随着设施级 800VDC 在第三阶段占据市场份额而下降。我们假设电源机柜的价值量为 50 万美元/兆瓦。

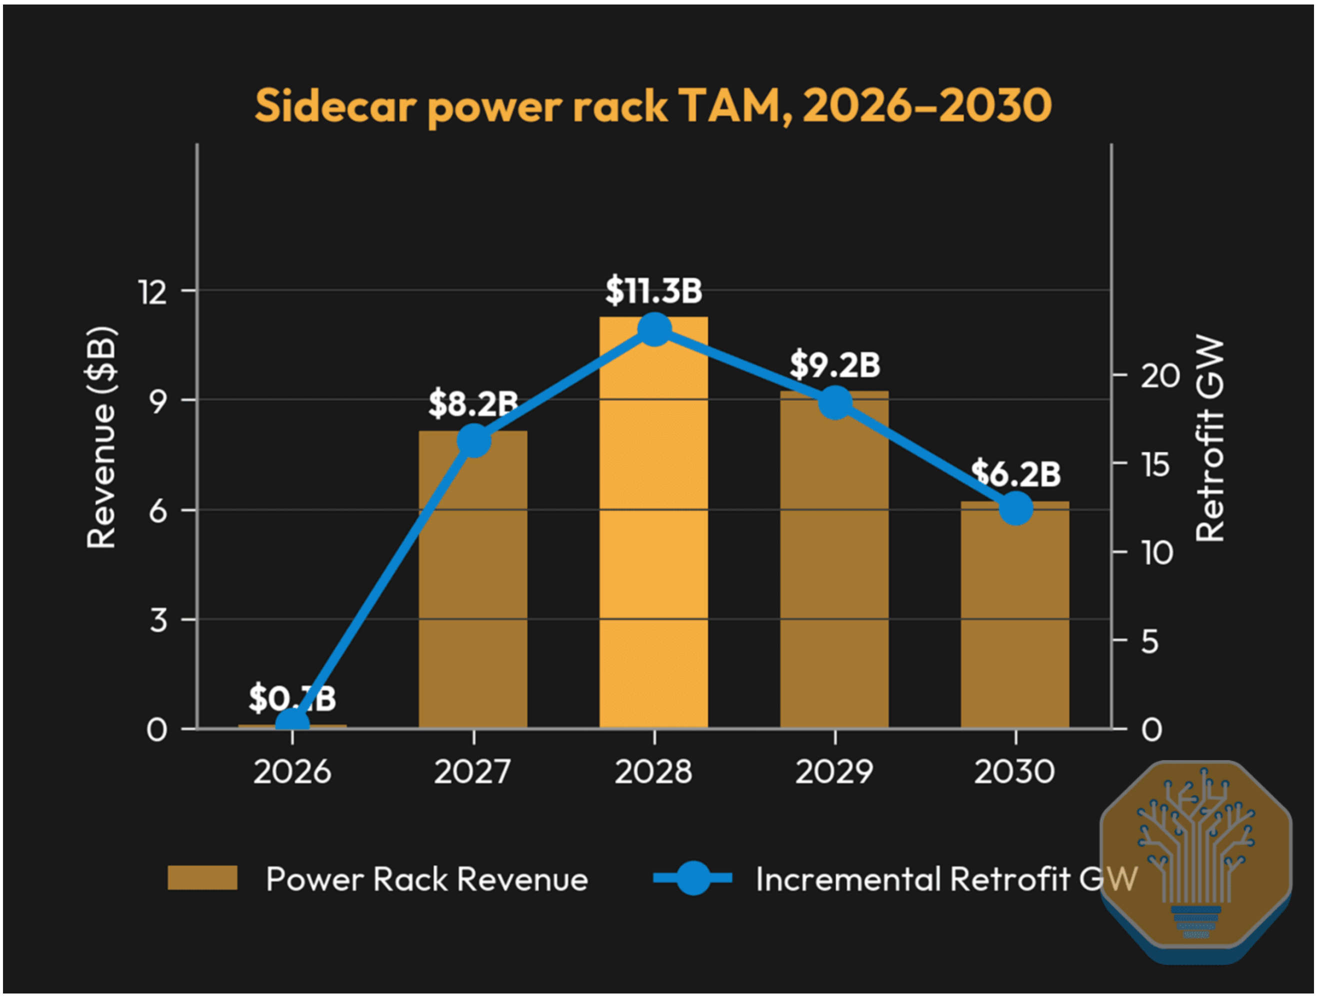

Source: SemiAnalysis Industrials Model

来源：SemiAnalysis 工业模型

Phase 1 Summary 第一阶段总结

The white-space retrofit represents a clear cost uplift in electrical content/MW versus current architectures, because Phase 1 essentially deletes nothing. We estimate the delta at roughly +$400-500k/MW, with the HVDC power rack accounting for the large majority.

白区改造代表了与当前架构相比，每兆瓦（MW）电气内容的成本明显提升，因为第一阶段基本上没有删除任何原有设施。我们估计这一差额约为每兆瓦 +40-50 万美元，其中高压直流（HVDC）电源机架占了绝大部分。

Phase 1 power chain cost delta vs baseline

Equipment	Baseline $M/MW	Phase 1 $M/MW	Delta $M/MW
MV Transformer	$0.35M	$0.35M	—
MV Switchgear	$0.70M	$0.70M	—
LV Transformer	$0.15M	$0.15M	—
LV Switchgear	$0.40M	$0.40M	—
Generator	$1.00M	$1.00M	—
Central UPS	$1.20M	$1.20M	—
Busway / PDU	$0.30M	$0.30M	—
Rack PDU	$0.05M	$0.00M	-$0.05M
Busbar	$0.00M	$0.05M	+$0.05M
HVDC Power Rack	$0.00M	$0.40M	+$0.40M
Power Shelf / PSU	$0.20M	$0.20M	—
TOTAL	**$4.35M**	**$4.75M**	**+$0.40M**

Source: SemiAnalysis Industrials Model

来源：SemiAnalysis 工业模型

Phase 2 (2027/2028): The Turning Point Comes with 800VDC-Native Compute

第二阶段（2027/2028 年）：800VDC 原生计算带来的转折点

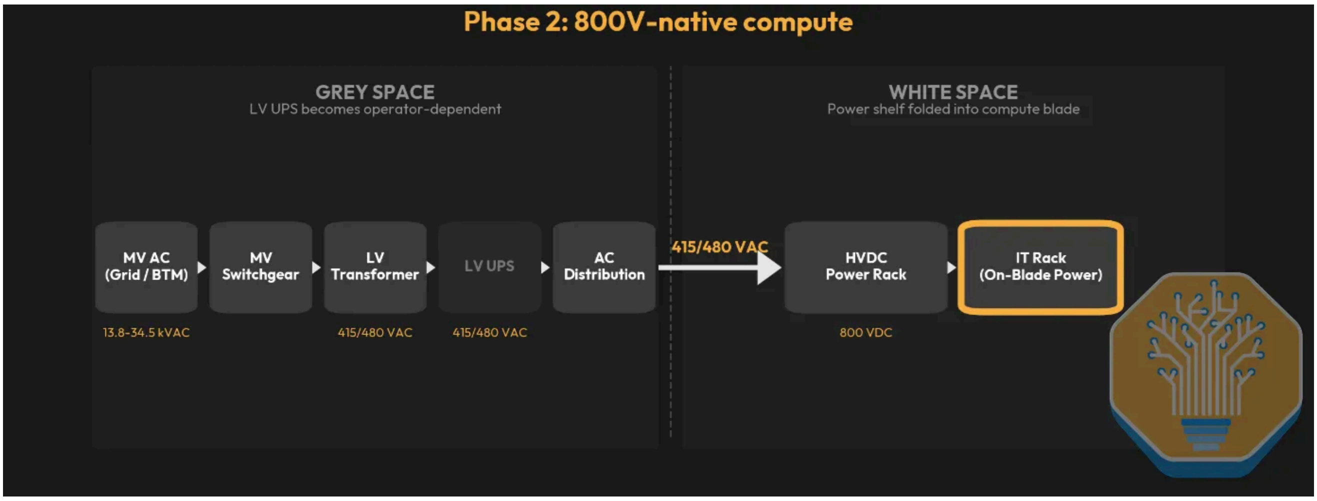

Source: SemiAnalysis Industrials Model

来源：SemiAnalysis 工业模型

Phase 1 was the start of the retrofit era. The real inflection point comes with the arrival of 800VDC-native systems. At that point, 800VDC stops being a future-proofing pilot and becomes a mandatory transition forced by physics and rack density. Operators electrifying the Kyber Rack have no AC fallback at the rack inlet, and we expect 800VDC penetration to spike sharply in this window. Because 800VDC-native silicon will land before facility-level 800VDC distribution is ready, the retrofit phase persists.

第一阶段是改造时代的开始。真正的拐点随着 800VDC 原生系统的到来而出现。届时，800VDC 将不再是面向未来的试点项目，而是由物理特性和机架密度强制推动的必然转型。为 Kyber 机架供电的运营商在机架入口处没有交流电（AC）备选方案，我们预计 800VDC 的渗透率将在这一窗口期急剧飙升。由于 800VDC 原生芯片将在设施级 800VDC 配电准备就绪之前落地，因此改造阶段仍将持续。

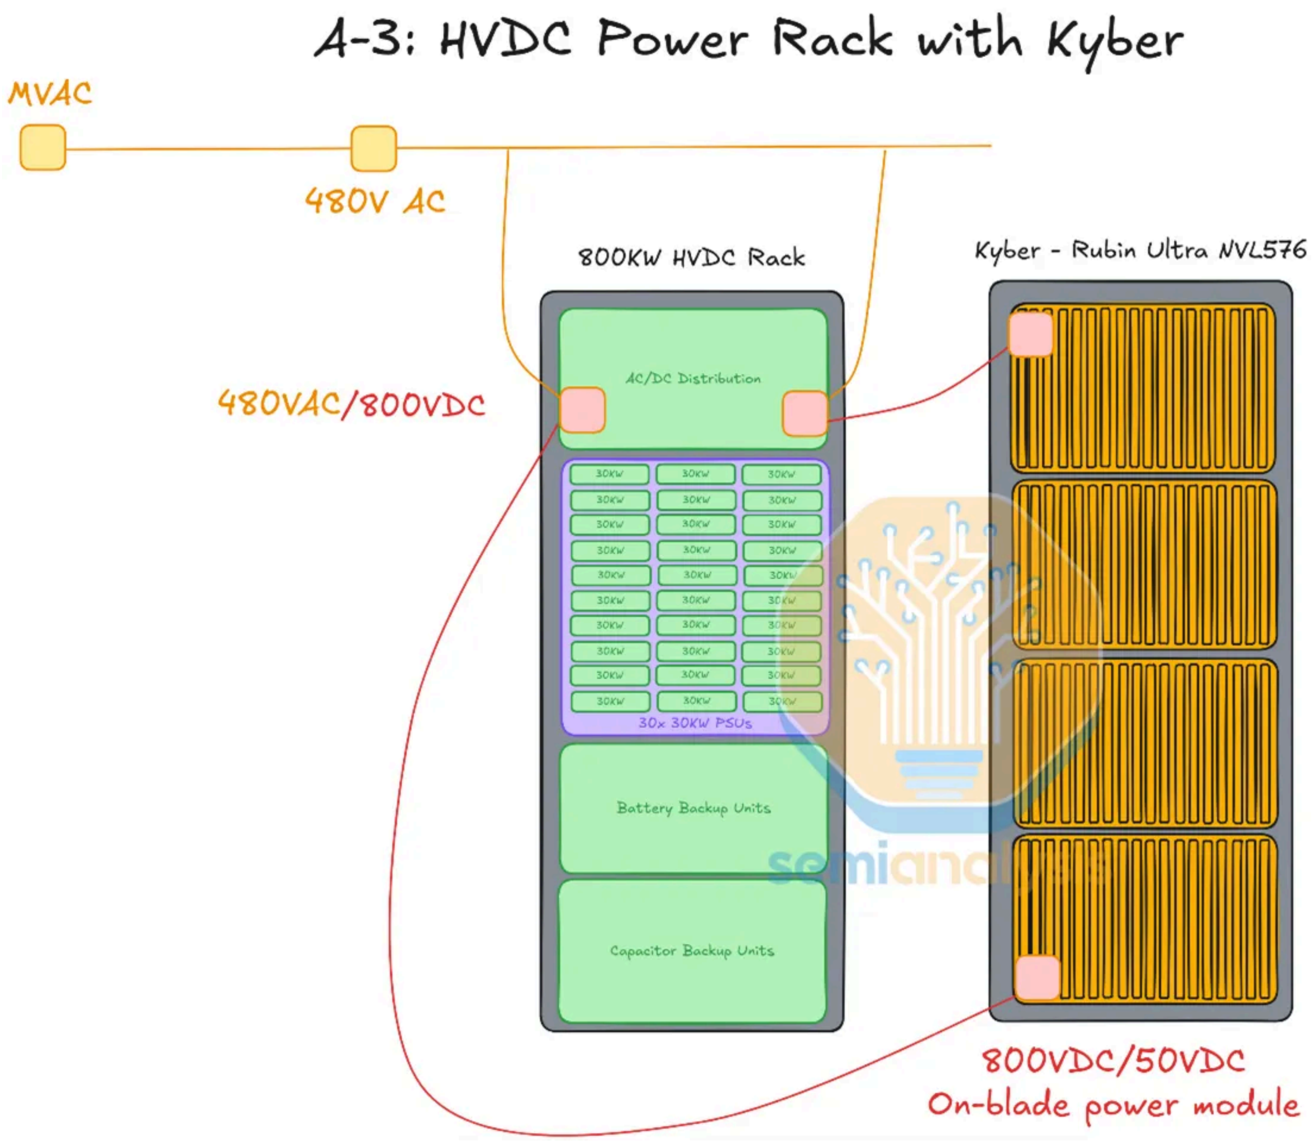

Source: SemiAnalysis SemiAnalysis

Architecturally, Phase 2 looks very similar to Phase 1. Both retrofit the white space with the HVDC power rack, both leave the grey space intact, and both rectify AC to DC in the row-level power rack. The key difference is where the voltage steps down to chip-usable levels. In Phase 1 (Oberon rack), a power shelf inside the IT rack converts 800VDC to ~50VDC before it reaches the compute trays. In Phase 2 (Kyber rack), the 800VDC bus runs directly to the compute blade, and an on-blade power module handles the final step-down to 50V.

从架构上看，第二阶段与第一阶段非常相似。两者都使用高压直流（HVDC）电源机架对白色空间（机房设备区）进行改造，都保持灰色空间（动力设施区）不变，且都在列级电源机架中将交流电（AC）整流为直流电（DC）。关键区别在于电压降至芯片可用水平的位置。在第一阶段（Oberon 机架）中，IT 机架内部的电源层在电流到达计算托盘之前将 800VDC 转换为约 50VDC。而在第二阶段（Kyber 机架）中，800VDC 母线直接通向计算刀片，由刀片上的电源模块负责最后的降压至 50V。

Earlier Kyber designs shown at OCP depicted a DC-DC PSU sidecar adjacent to the compute rack, but we now believe this approach is unlikely to be adopted at scale. A standalone sidecar consumes more aggregate floor and rack space than integrating the conversion stage into the blade itself, and the power module form factor has proven feasible within the compute tray's volume constraints.

早先在 OCP 上展示的 Kyber 设计描绘了一个紧邻计算机架的 DC-DC PSU 侧柜，但我们现在认为这种方案不太可能被大规模采用。与将转换级集成到刀片本身相比，独立的侧柜会消耗更多的总地面和机架空间，而且事实证明，在计算托盘的体积限制内，电源模块的形态是可行的。

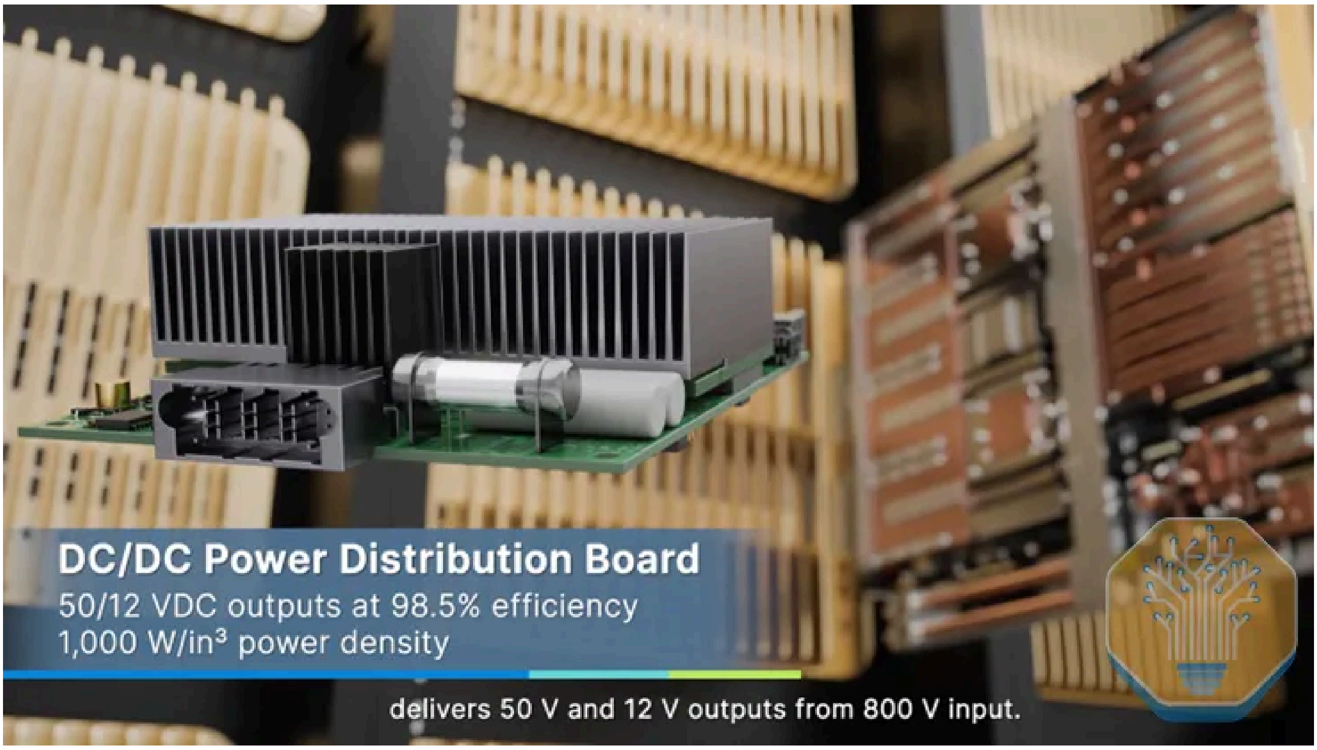

Source: Delta 台达

Because most servers and trays still take roughly ~50V input, both architectures retain a high-power 800V-to-~50V DC-DC conversion stage. The difference is where that conversion happens.

由于大多数服务器和托盘仍采用约 50V 的输入电压，因此这两种架构都保留了大功率 800V 转约 50V 的 DC-DC 转换阶段。区别在于该转换发生的具体位置。

Some discussions have explored delivering 800VDC directly into the compute tray and stepping it down to an intermediate bus voltage (IBV) before further conversion to point-of-load rails. While Kyber's on-blade power module does accept 800V input, it converts to the established ~50V bus level rather than an IBV scheme. A full 800V-to-IBV-to-PoL architecture within the tray remains extremely challenging given the limited space and safety constraints involved.

一些讨论探讨了将 800VDC 直接输送到计算托盘，并将其降压至中间总线电压（IBV），然后再进一步转换为负载点（PoL）电压轨。虽然 Kyber 的板载电源模块确实接受 800V 输入，但它是将其转换为已确立的约 50V 总线电平，而非采用 IBV 方案。考虑到有限的空间和安全约束，在托盘内实现完整的 800V-转-IBV-转-PoL 架构仍然极具挑战性。

What Happens With UPS and Battery Storage

UPS 和电池储能系统的变化

Traditional central UPS systems are probably the most contested piece of infrastructure in the 800VDC transition. In the 800VDC architecture, we expect centralized Low Voltage UPS systems to progressively lose their role and eventually become obsolete. In the retrofit era, the power rack sits directly on the 800VDC bus and houses BBU modules and supercapacitors, which we cover shortly. Both are natively DC-coupled. BBUs bridge seconds-to-minutes during outages and supercapacitors absorb millisecond-scale GPU load transients. Together, they replace the centralized short-term battery storage and UPS ride-through function without the

2-3% conversion loss of an AC-DC-AC UPS pair.

在 800VDC 转型过程中，传统的中央 UPS 系统可能是最具争议的基础设施。在 800VDC 架构中，我们预计集中式低压 UPS 系统将逐渐失去其地位并最终被淘汰。在改造时代，电源机架直接连接在 800VDC 总线上，并容纳 BBU 模块和超级电容器（我们稍后会详细介绍）。两者都是原生直流耦合的。BBU 在停电期间提供秒级到分钟级的过渡，而超级电容器则吸收毫秒级的 GPU 负载瞬变。它们共同取代了集中式短期电池储能和 UPS 跨越功能，且没有交流-直流-交流（AC-DC-AC）UPS 组合所产生的 2-3% 的转换损耗。

As we covered in our electricals deep dive note, Google and Meta already took this aggressive approach years ago, bypassing the central monolithic UPS with “distributed UPS” architectures. In their architecture, AC power is distributed directly to the rack, the in-rack PSU handles AC-DC conversion, and rack-level Li-Ion Battery Backup Units (BBUs) provide the short-duration bridge power. This removes the central UPS’s AC-DC-AC conversion pair and improves efficiency, while also cutting in half the total battery capacity needed for the datacenter, since there is no longer a need for both an A-side and a B-side UPS.

正如我们在电气系统深度解析报告中所提到的，谷歌和 Meta 早在几年前就采取了这种激进的做法，通过“分布式 UPS”架构绕过了中央单体式 UPS。在他们的架构中，交流电直接输送到机架，机架内的电源单元（PSU）负责交流-直流（AC-DC）转换，而机架级的锂离子电池备份单元（BBU）则提供短时间的桥接电力。这消除了中央 UPS 的“交流-直流-交流”双重转换环节，提高了效率，同时也使数据中心所需的总电池容量减少了一半，因为不再需要同时配备 A 路和 B 路 UPS。

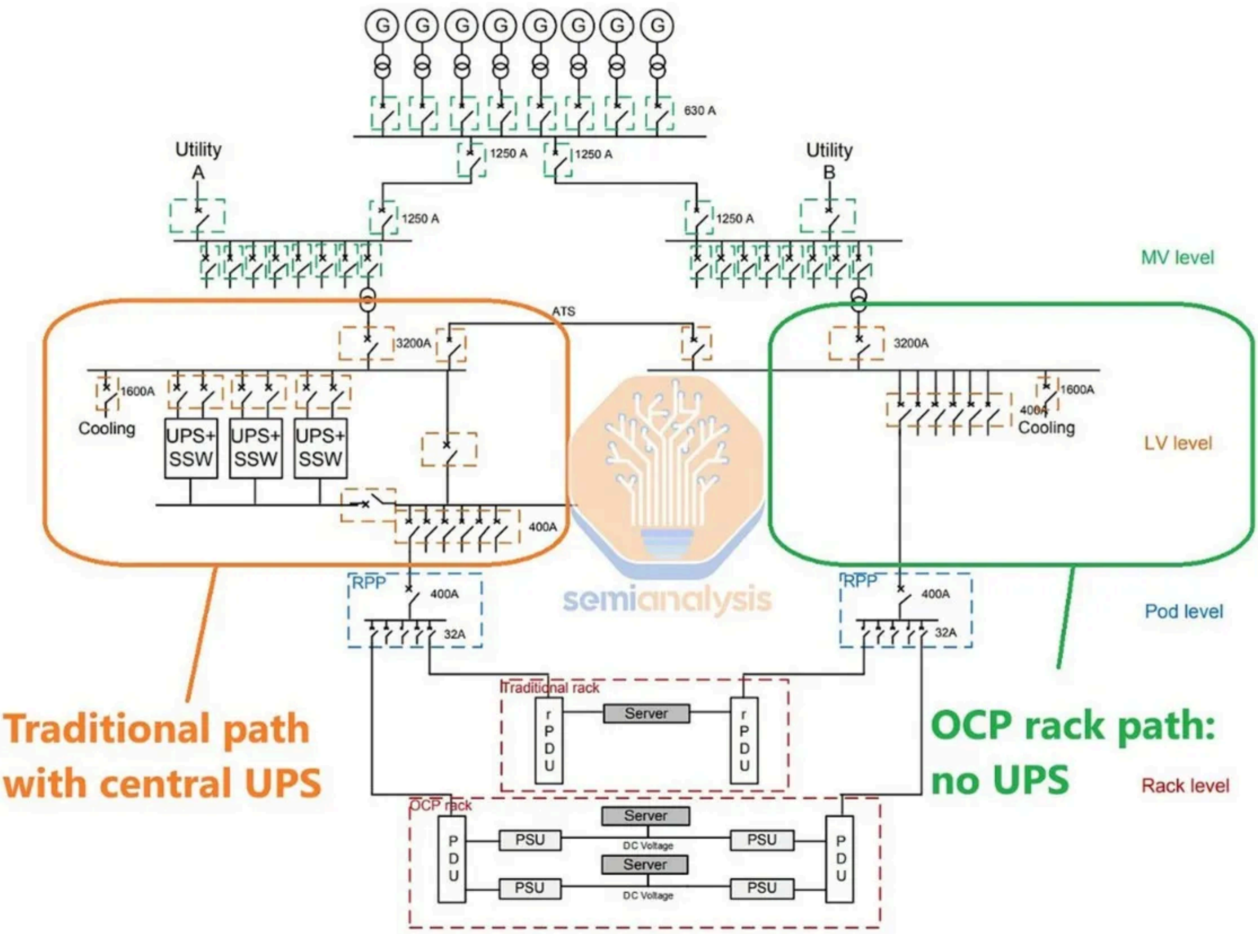

Source: SemiAnalysis SemiAnalysis

That said, managing distributed UPS or battery backup is more operationally challenging than running a traditional central UPS. We expect operators other than vertically integrated hyperscalers like Google and Meta to keep the Low Voltage UPS in place for redundancy and load fluctuation management, at least in the medium term.

即便如此，管理分布式 UPS 或电池备份在运维上比运行传统的中央 UPS 更具挑战性。我们预计，除了像 Google 和 Meta 这样垂直整合的超大规模运营商外，其他运营商至少在中期内仍将保留低压 UPS，以用于冗余和负载波动管理。

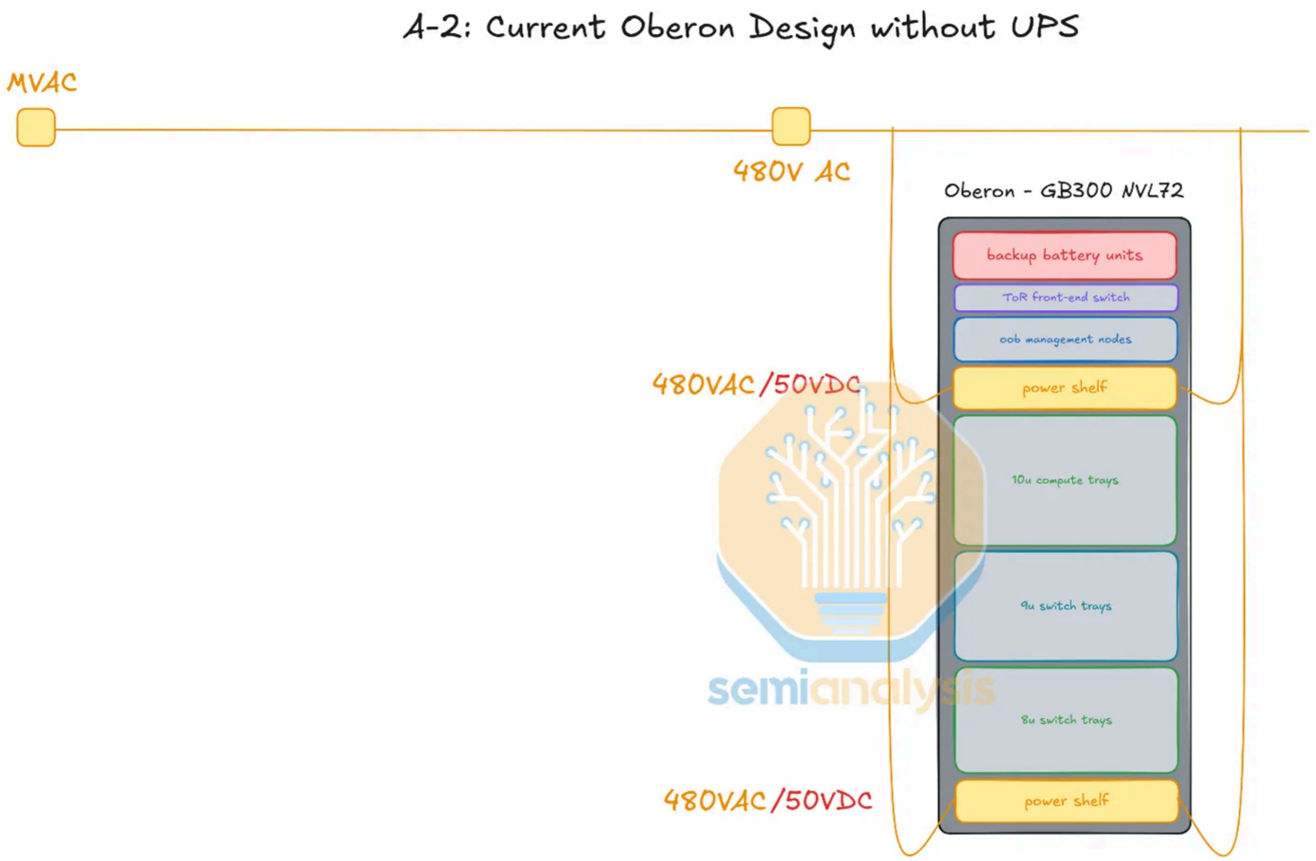

Source: SemiAnalysis SemiAnalysis

This is especially true for colocation providers, which prioritize flexibility and need to support mixed workloads: CPU racks, storage arrays, networking equipment, and older GPU racks that still run on AC. Keeping the grey-space AC infrastructure intact lets these operators deploy 800VDC for their densest AI racks while running standard AC distribution for everything else.

对于托管服务提供商而言尤其如此，他们优先考虑灵活性，并需要支持混合工作负载：CPU 机架、存储阵列、网络设备，以及仍运行在交流电（AC）上的旧款 GPU 机架。保持灰色空间（grey-space）交流基础设施的完整性，让这些运营商能够在其密度最高的 AI 机架上部署 800VDC，同时为其他所有设备运行标准的交流配电。

We expect different operators to adopt different architectural approaches to backup, and new alternatives are emerging. Medium Voltage UPS, operating at 4.16-34.5 kV directly at the grid connection point, is functionally similar to the rack-level Battery Rack but centralized at the grid interface rather than distributed across the data hall. ABB’s HiPerGuard runs at 98% efficiency and is already deployed at Applied Digital’s 400MW North Dakota AI campus. ON.energy was awarded few weeks ago a US patent that protects their MV double-conversion UPS architecture. The second alternative is facility-level BESS, which as we covered in our deep dive operates at megawatt-to-hundreds-of-megawatts scale, provides 1-4 hour duration backup, and increasingly replaces or shrinks the diesel generator.

我们预计不同的运营商将对备份采取不同的架构方案，而新的替代方案也正在涌现。中压 UPS（MV UPS）直接在电网连接点以 4.16-34.5 kV 的电压运行，其功能类似于机架级电池架（Battery Rack），但它是集中在电网接口处，而非分布在整个数据大厅。ABB 的 HiPerGuard 运行效率高达 98%，并已部署在 Applied Digital 位于北达科他州的 400MW AI 园区。几周前，ON.energy 获得了一项美国专利，该专利保护了他们的中压双变换 UPS 架构。第二种替代方案是设施级电池储能系统（BESS），正如我们在深度探讨中所介绍的，其运行规模在兆瓦到数百兆瓦之间，提供 1-4 小时的备份时长，并正日益取代或缩小柴油发电机的规模。

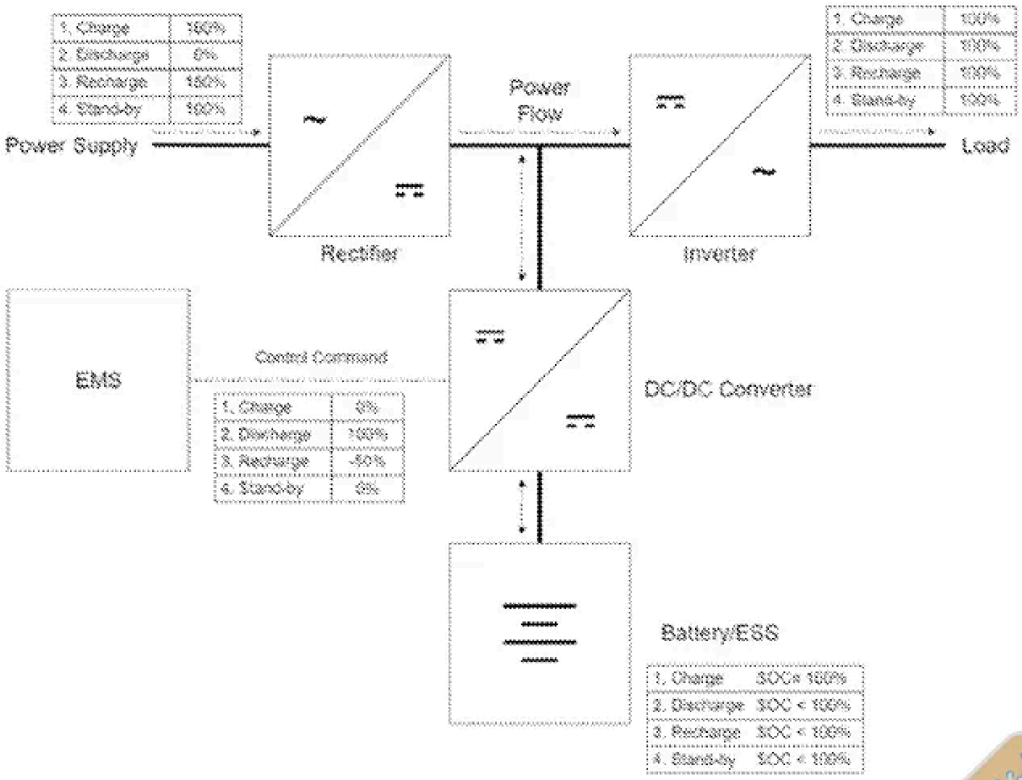

FIG. 14

Source: United States Patent and Trademark Office

美国专利及商标局

Phase 3 (late 2028/2029): Redesigning the Electrical Architecture With a Centralized Rectifier

第 3 阶段（2028 年底/2029 年）：通过集中式整流器重新设计电气架构

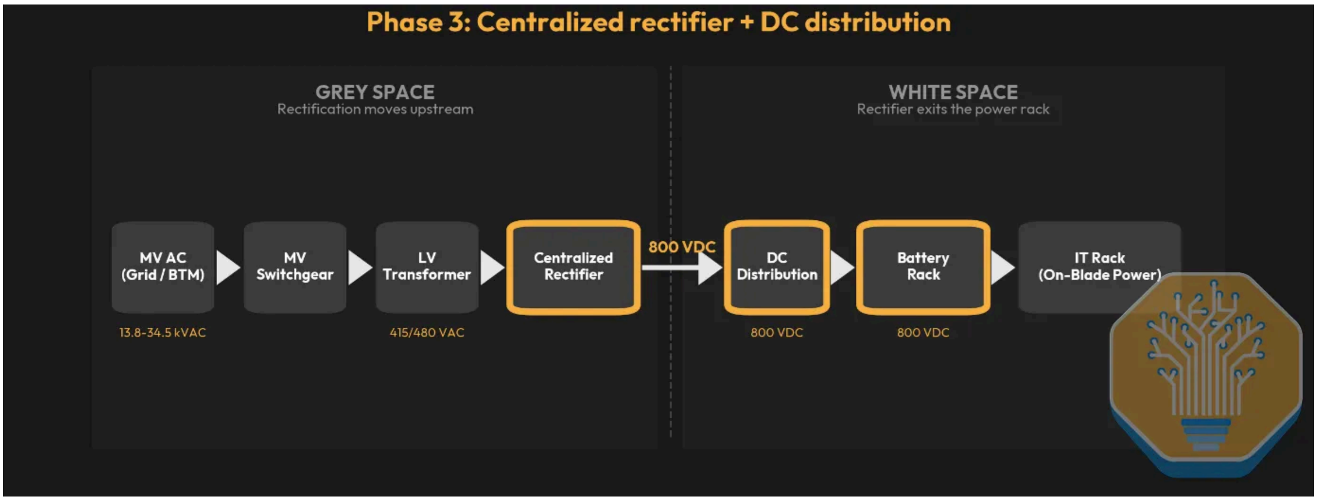

[Source: SemiAnalysis Industrials Model]

来源：SemiAnalysis 工业模型

In Phases 1 and 2, the AC-DC conversion happens close to the rack, inside the row-level HVDC power rack. Phase 3 changes the datacenter layout itself, and 800VDC becomes the building's electrical core. This is the true inflection point, where things start to become interesting. Let's unpack what happens in each area of the datacenter.

在第一和第二阶段，交流-直流（AC-DC）转换发生在机架附近，即列级高压直流（HVDC）电源机架内部。而第三阶段改变了数据中心本身的布局，使 800VDC 成为建筑的电气核心。这是一个真正的转折点，事情开始变得有趣起来。让我们来拆解数据中心每个区域所发生的变化。

What Happens in the Grey Space: Power Distribution Goes DC

灰色空间发生了什么：配电转向直流化

In Phase 3 a dedicated upstream rectifier that sits in the grey space or outdoors converts 415V AC to 800VDC, distributing DC across the entire hall. These are mature units using silicon IGBTs or thyristors rated at 1200-1700V.

在第三阶段，位于灰色空间或室外的专用上游整流器将 415V 交流电转换为 800VDC，并在整个机房内分配直流电。这些是成熟的设备，采用额定电压为 1200-1700V 的硅基 IGBT 或晶闸管。

The grey space splits in two. MV transformers connecting the datacenter to the grid are unchanged. MV switchgear stays because the utility feed is still AC, and the upstream MV infrastructure (11-34 kV) is expected to grow more complex as facilities scale to gigawatt clusters. LV transformers remain, stepping MV down to 415V AC for the upstream rectifier. The 480V AC switchgear between LV transformers and PDUs has no role once 800VDC flows through the busway, and AC floor PDUs are eliminated along with it, since the DC busway feeds the battery rack directly with no AC distribution PDU in between. In summary, everything above the AC-DC conversion point stays, while everything below it, designed for AC distribution, goes.

灰色空间一分为二。将数据中心连接到电网的中压（MV）变压器保持不变。中压配电装置得以保留，因为公用事业供电仍为交流电，且随着设施规模扩大至吉瓦级集群，上游中压基础设施（11-34 kV）预计将变得更加复杂。低压（LV）变压器依然存在，将中压降至 415V 交流电以供给上游整流器。一旦 800V 直流电通过母线槽，位于低压变压器和配电单元（PDU）之间的 480V 交流配电装置将不再发挥作用，交流地板 PDU 也随之被取消，因为直流母线槽直接为电池架供电，中间无需交流分配 PDU。总而言之，交流-直流转换点以上的所有设施均予以保留，而其下方专为交流分配设计的设施则全部移除。

Understanding DC Distribution: Switchboards, Busway, and Protection

深入了解直流配电：配电柜、母线槽与保护

In Phase 3, the AC switchboard's function of splitting one feed into multiple protected outputs has to land somewhere. Three product categories are positioned to absorb it: (i) MW-scale rectifiers built with multiple outputs and integrated SSCB protection per output, turning the rectifier into its own distribution device; (ii) DC busway with breaker-equipped tap-off boxes that hold protection in the distribution medium, once DC-rated tap-offs with adequate arc interruption mature; and (iii) prefabricated grey-space pods that bundle rectifier, switchboard, and busway into a factory-built skid, particularly for hyperscaler procurement.

在第三阶段，交流配电柜将一路输入拆分为多路受保护输出的功能必须有所承接。目前有三类产品定位用于吸收这一功能：(i) 兆瓦（MW）级整流器，采用多路输出设计并为每路输出集成 SSCB（固态断路器）保护，使整流器本身成为配电装置；(ii) 直流母线槽，配有带断路器的插接箱，将保护功能置于配电介质中（前提是具备充足灭弧能力的直流额定插接组件技术成熟）；以及 (iii) 预制化灰区模块（Pods），将整流器、配电柜和母线槽整合到工厂制造的撬块中，特别适用于超大规模数据中心（Hyperscaler）的采购模式。

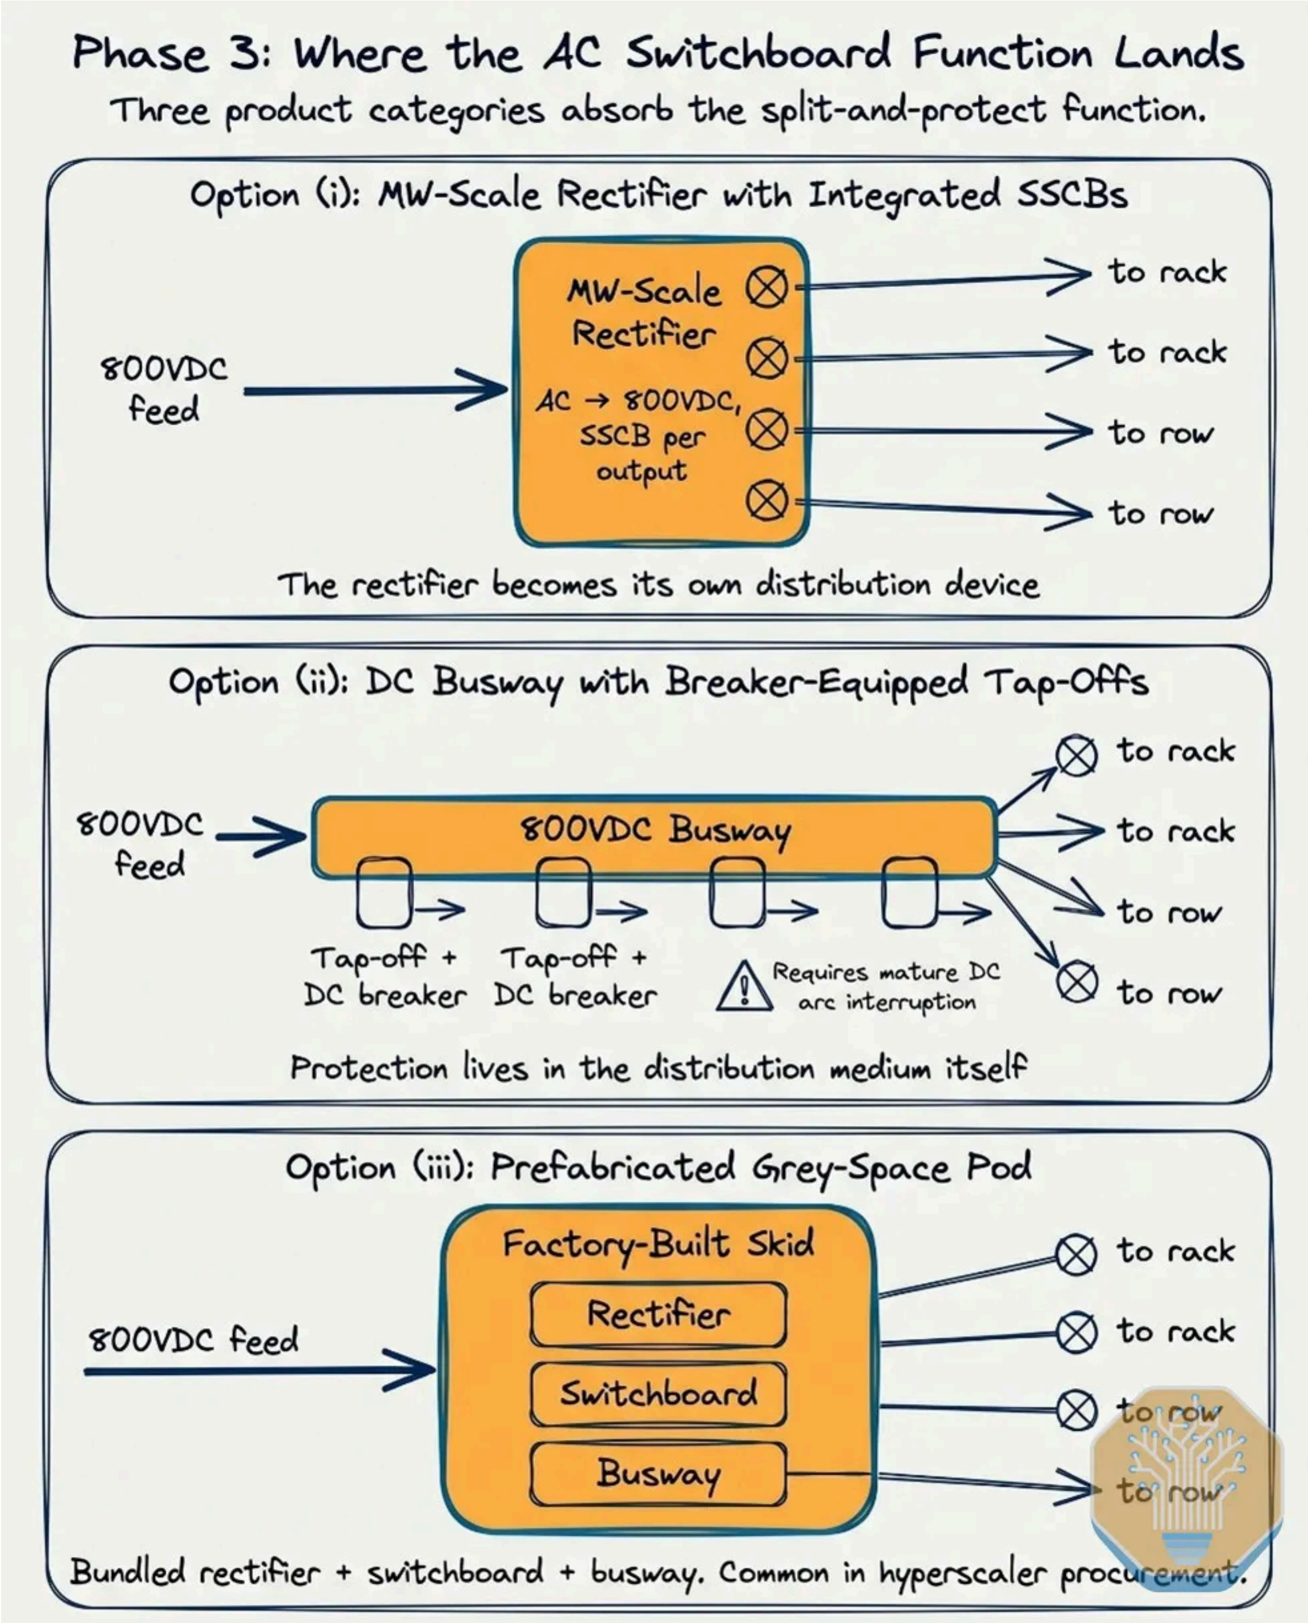

Source: SemiAnalysis SemiAnalysis

Major AC switchboard incumbents (Schneider Electric, ABB, Eaton, Vertiv) have not named discrete 800VDC switchboard products. ABB's October 2025 Nvidia partnership covers distribution inside its "modular power block" rather than as a standalone switchboard. EPEC Solutions sells a publicly marketed 800VDC LV switchboard with high-interrupting-capacity DC breakers. We expect the discrete switchboard to retain a niche in retrofits with existing single-output rectifiers, and where operators want vendor-neutrality across the rectifier and protection layers.

主要的交流配电柜现有厂商（施耐德电气、ABB、伊顿、Vertiv）尚未命名独立的 800VDC 配电柜产品。ABB 在 2025 年 10 月与 Nvidia 建立的合作伙伴关系涵盖了其“模块化电源块”内部的配电，而非作为独立的配电柜。EPEC Solutions 销售一种公开市场化的 800VDC 低压配电柜，配备了高分断能力的直流断路器。我们预计，在现有单输出整流器的改造项目中，以及运营商希望在整流器和保护层之间保持供应商中立性的情况下，独立配电柜将保留一定的利基市场。

Once the power is rectified, a DC busway replaces AC busway for hall-level 800VDC distribution. In traditional AC datacenters, busway systems have modular plug-in connections called tap-offs that branch power to individual racks or rows, similar to outlets on a power strip. You can add or remove these while the busway is energized. A feeder-only busway, by contrast, has no intermediate openings or tap-offs. Power enters at one end and exits at the other end or at predefined termination points.

一旦电源完成整流，直流母线槽将取代交流母线槽，用于机房级的 800VDC 配电。在传统的交流数据中心中，母线槽系统具有被称为分接箱（tap-offs）的模块化插接连接，可将电力分支到各个机架或机柜列，类似于电源插座板上的插口。您可以在母线槽带电的情况下添加或移除这些分接箱。相比之下，仅馈电式母线槽没有中间开口或分接箱。电力从一端进入，从另一端或预定义的端接点输出。

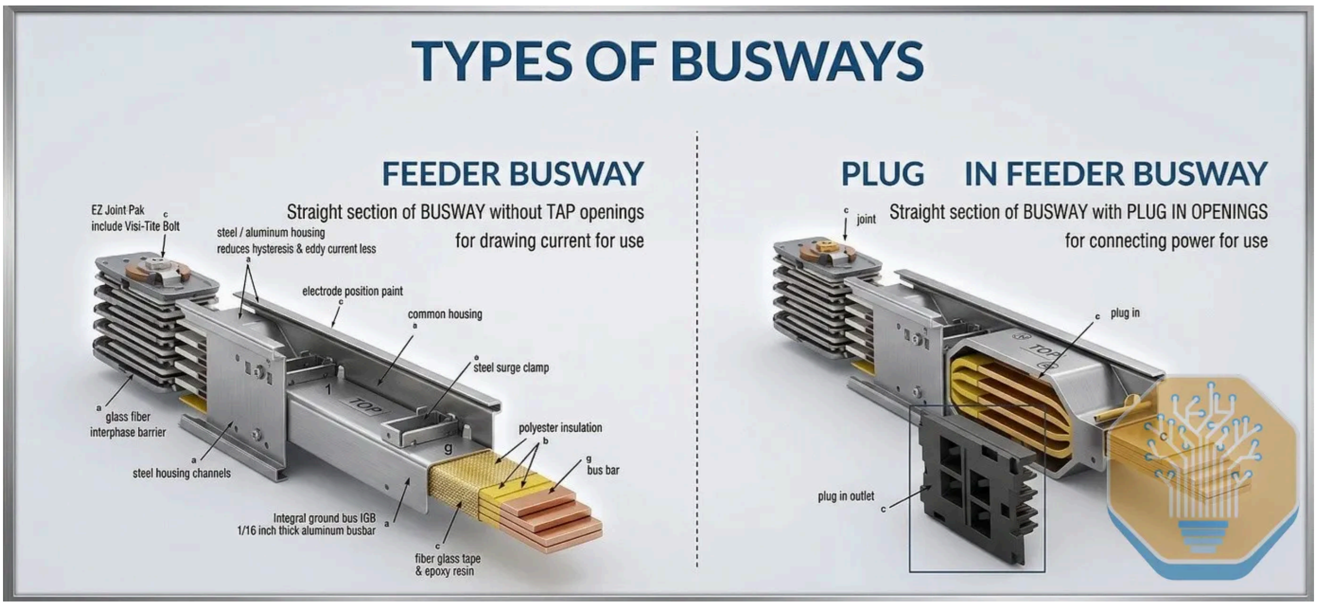

Source: SemiAnalysis SemiAnalysis

We expect early 800VDC deployments to use feeder-only busway because, essentially, tap-offs become more complex. At 800VDC, interrupting current under load creates a sustained arc (a plasma discharge producing extreme heat) that does not self-extinguish because DC has no zero-crossing point, while AC arcs naturally extinguish 100-120 times per second as the waveform crosses zero. Besides, DC-rated tap-off units with adequate arc interruption are physically larger, making them impractical today. Delta and ABB have publicly disclosed 800VDC busway programs, and we expect other major busway vendors like Legrand and EAE to follow in 2026.

我们预计早期的 800VDC 部署将采用仅限馈线的母线槽，因为从本质上讲，插接箱（tap-offs）会变得更加复杂。在 800VDC 电压下，带载切断电流会产生持续的电弧（一种产生极高热量的等离子体放电），且由于直流电没有过零点，电弧不会自动熄灭；相比之下，交流电弧在波形过零时自然熄灭，频率为每秒 100-120 次。此外，具备充分灭弧能力的直流额定插接单元体积庞大，在目前并不实用。台达（Delta）和 ABB 已公开披露了 800VDC 母线槽项目，我们预计罗格朗（Legrand）和 EAE 等其他主要母线槽供应商也将在 2026 年跟进。

To address these challenges, multiple proven protection paradigms exist at this voltage class from adjacent industries. The likely implementation combines multiple approaches, one being new generations of circuit breakers. More specifically, following the same solid-state trend already underway with Solid State Transformers, Solid State Circuit Breakers (SSCBs) are now being adopted. SSCBs use SiC or GaN to interrupt fault current in microseconds. Because semiconductor switches can simply stop conducting with no physical contact separation, there is no arc to extinguish in the first place.

为了应对这些挑战，相关相邻行业在这一电压等级上已存在多种成熟的保护范式。可能的实现方式是结合多种方法，其中之一便是新一代断路器。更具体地说，顺应固态变压器（Solid State Transformers）领域已经出现的固态化趋势，固态断路器（SSCBs）目前正被采用。固态断路器利用碳化硅（SiC）或氮化镓（GaN）在微秒内切断故障电流。由于半导体开关只需停止导通而无需物理触点分离，因此从根本上就不会产生需要熄灭的电弧。

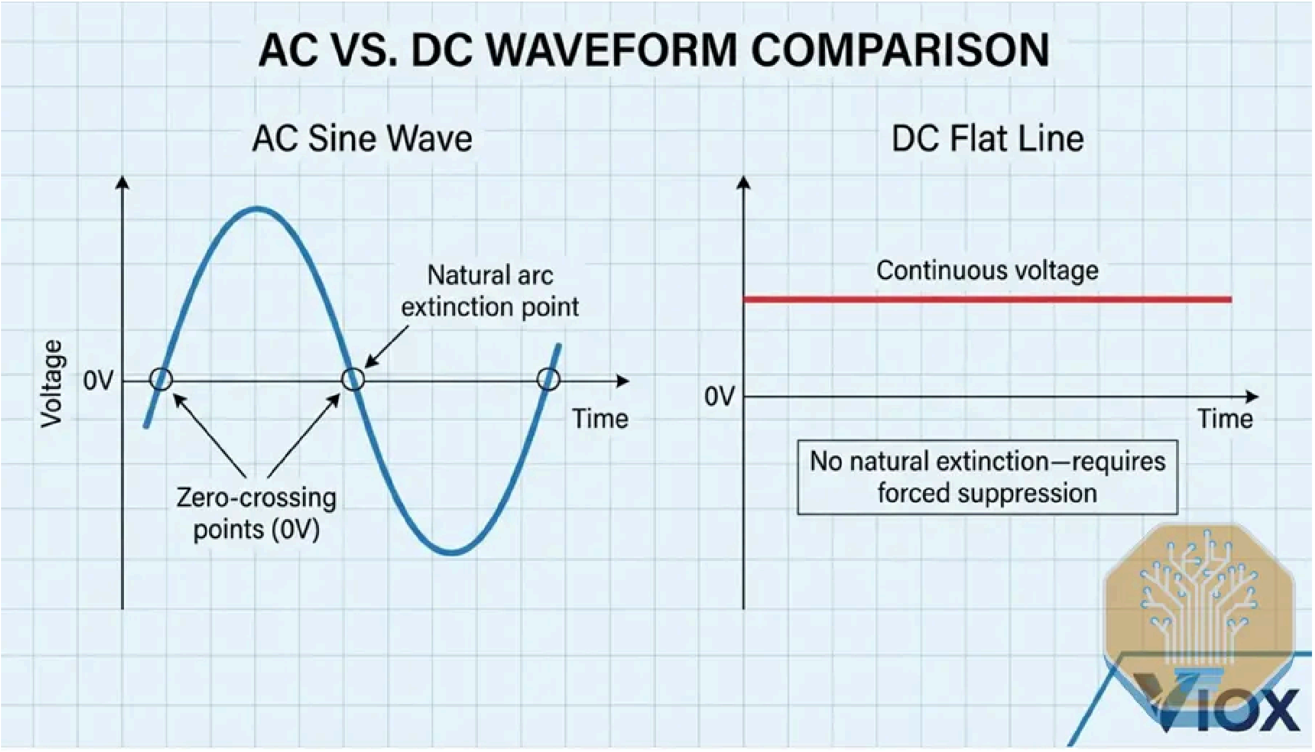

Source: VIOX VIOX

The new-generation circuit breakers are already commercialized today. ABB has the Emax 2 (1500V DC) used in solar, energy storage or marine, as well as the SACE Infinitus (solid-state, 1000V/2500A, datacenter adaptation with Nvidia announced October 2025). LS Electric has the first UL-certified DC molded case circuit breaker at 1500V, listed for datacenter applications.

新一代断路器现已实现商业化。ABB 推出了用于太阳能、储能或船舶领域的 Emax 2（1500V DC），以及 SACE Infinitus（固态断路器，1000V/2500A，并于 2025 年 10 月宣布与 Nvidia 合作进行数据中心适配）。LS Electric 则拥有首款通过 UL 认证的 1500V 直流塑壳断路器，并被列入数据中心应用清单。

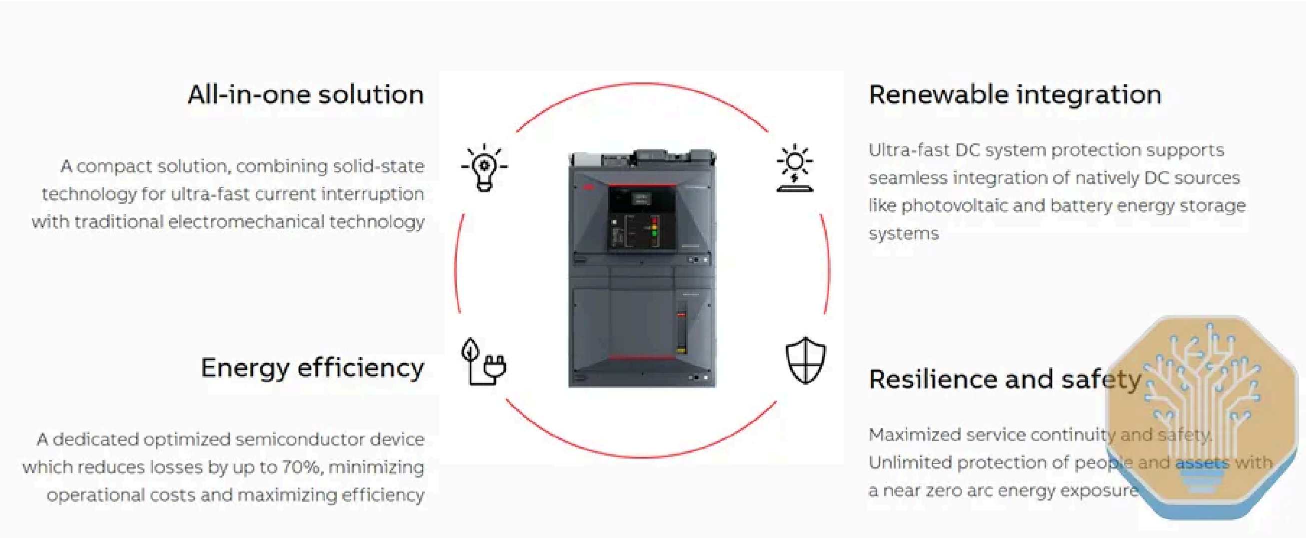

Source: ABB ABB

Alternative Path Using LV Solid State Transformers

使用低压固态变压器（LV SST）的替代路径

An emerging alternative to the centralized AC/DC rectifier is using LV SSTs. It performs the same conversion, 415V AC to 800VDC in the grey space or outdoors, but in a more compact and programmable form factor. The LV-SST sidesteps the 3,300V-class SiC supply constraint that gates MV-input SSTs, making it the earlier-to-market SST variant.

一种新兴的替代集中式交流/直流（AC/DC）整流器的方案是使用低压固态变压器（LV SST）。它在灰色空间或室外执行相同的转换（将 415V AC 转换为 800VDC），但具有更紧凑和可编程的形态。LV-SST 避开了限制中压输入（MV-input）SST 发展的 3,300V 级碳化硅（SiC）供应链约束，使其成为更早推向市场的 SST 变体。

What Happens in the White Space: From the Power Rack to the Battery Rack

白色空间内发生了什么：从功率机架到电池机架

As you can imagine, in Phase 3 we no longer need the Power Rack doing the 800VDC conversion. Instead, we salute a new friend, the battery rack.

正如你可以想象的那样，在第三阶段，我们不再需要功率机架（Power Rack）来进行 800VDC 转换。取而代之的是我们的新朋友：电池机架（Battery Rack）。

The battery rack shares most of the power rack's components and functions. The main difference is that it no longer performs AC-DC rectification, because it receives 800VDC directly from the grey space. Three main components remain:

电池柜共享了电源柜的大部分组件和功能。主要区别在于它不再执行交流-直流（AC-DC）整流，因为它直接从灰色空间接收 800VDC。保留了三个主要组件：

- **DC/DC distribution units:** manage power distribution, switching, and monitoring across the 800VDC bus. They do not step-down voltage. The full 800VDC travels from the battery rack to the compute blade.

  DC/DC 分配单元：管理 800VDC 总线的功率分配、切换和监控。它们不进行降压。完整的 800VDC 从电池机架传输到计算刀片。

- **BBU shelves:** provide ride-through power during supply interruptions.

  BBU 机架：在供电中断期间提供跨越电力。

- **Supercapacitors (optional):** absorb microsecond-to-millisecond transients that batteries are too slow to catch. They sit between the DC bus and the BBU, handling fast voltage excursions.

  超级电容器（可选）：吸收电池反应过慢而无法捕捉的微秒至毫秒级瞬态。它们位于 DC 总线和 BBU 之间，处理快速的电压偏移。

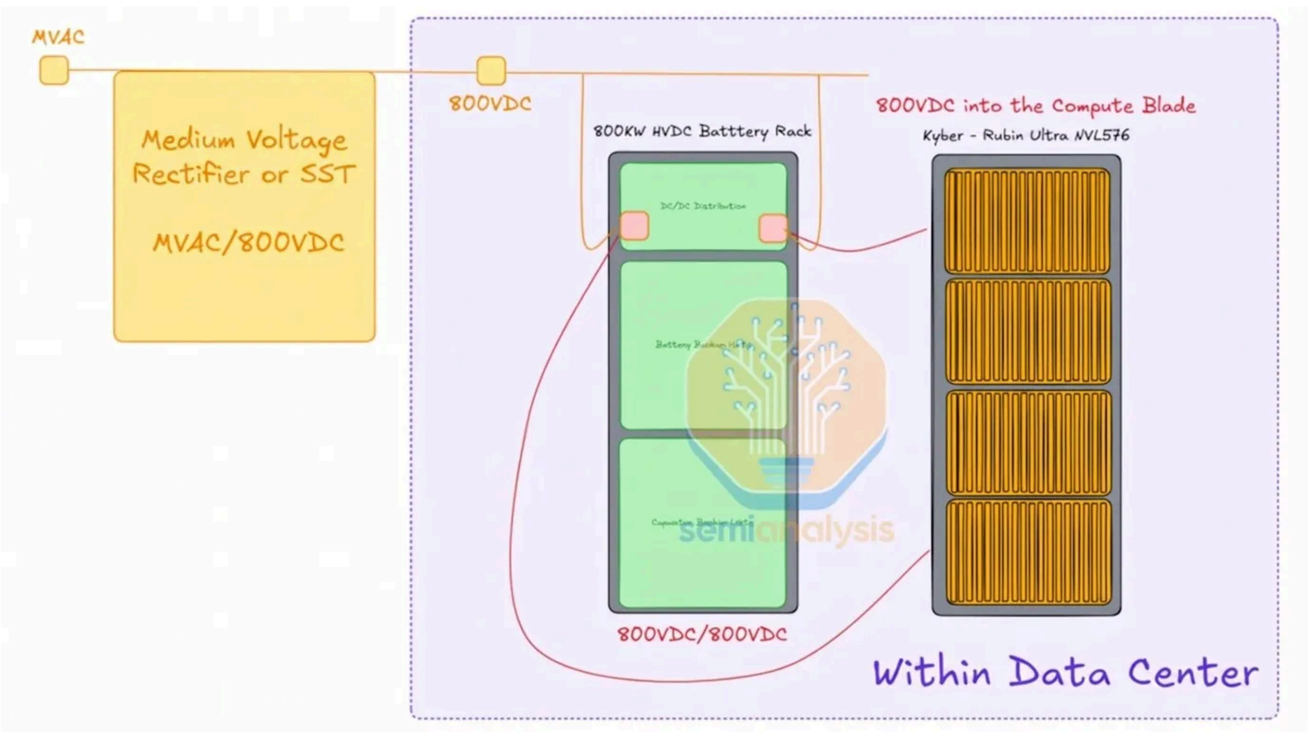

Source: SemiAnalysis SemiAnalysis

The battery rack sits generally at the same row level as the power rack it replaces, although some operators are deploying these in the adjacent grey space or in outdoor enclosures. The trade-off is simple: rectifiers go away, BBU and supercapacitor content goes up. We expect content per MW for the battery rack to reach around $200k/MW.

电池机架通常与它所取代的功率机架处于同一排水平，尽管一些运营商正将其部署在相邻的灰色空间或室外机柜中。这种权衡很简单：整流器消失了，而 BBU 和超级电容器的内容量增加了。我们预计电池机架每兆瓦（MW）的价值量将达到约 20 万美元/MW。

We covered supercapacitor chemistry and technical specifications in our deep dive on AI training load fluctuations. Part 2 of this 800VDC series will go deeper on supercapacitor economics, cell chemistry, the vendor landscape, and the practical tradeoffs of deploying them in production.

我们在关于 AI 训练负载波动的深度解析中已经介绍了超级电容器的化学成分和技术规格。本 800VDC 系列的第二部分将深入探讨超级电容器的经济性、电芯化学、供应商格局，以及在生产中部署它们的实际权衡。

BBU modules scale up BBU 模块扩展

Current modules are rated at roughly 5.5kW. With Rubin Ultra and 800VDC architectures, individual module wattage rises to 8-12kW. Infineon's BBU roadmap, announced in March 2025, uses modular 4kW Partial Power Converter cards that parallel to 12kW per unit at up to 99.5% peak efficiency.

目前的模块额定功率约为 5.5kW。随着 Rubin Ultra 和 800VDC 架构的出现，单个模块的功率将提升至 8-12kW。英飞凌在 2025 年 3 月公布的 BBU 路线图中，采用了模块化的 4kW 部分功率转换器（Partial Power Converter）卡，通过并联使单机功率达到 12kW，峰值效率高达 99.5%。

Delta, at GTC 2026, went further at the shelf level: its new 110kW power shelves embed 80kW of BBU capacity each, totaling 480kW across a six-shelf rack. Higher rack power demands proportionally more backup energy per rack, and higher-wattage modules deliver that energy with fewer physical modules, preserving space in the power rack.

台达在 GTC 2026 上进一步展示了机架级技术：其新型 110kW 电源机架每个都嵌入了 80kW 的 BBU 容量，在六机架式机柜中总计达到 480kW。更高的机架功率需求意味着每个机架需要成比例增加的备用能量，而更高瓦数的模块能以更少的物理模块提供这些能量，从而节省了电源机架的空间。

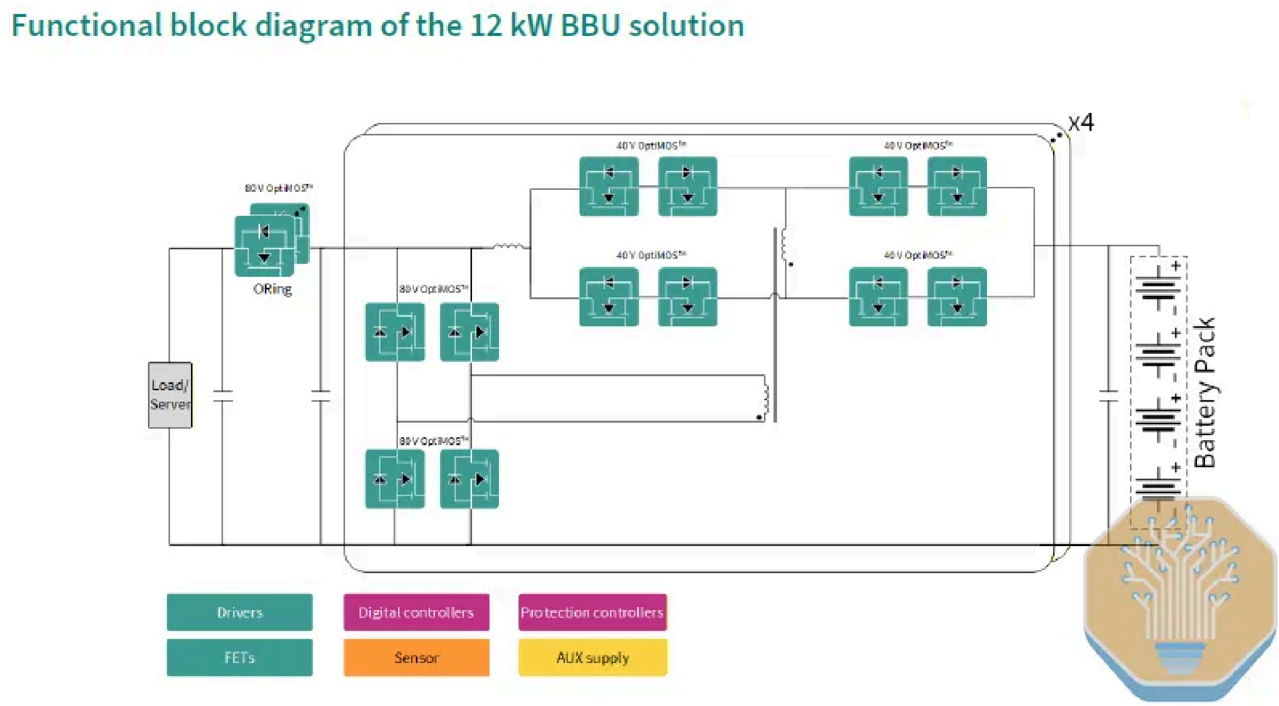

Source: Infineon 英飞凌

What Happens at the Facility Level

设施层面发生了什么

After analyzing the complete transformation of the grey space and the white space, the facility level is the part that changes least.

在分析了灰区和白区的完整转型后，设施层面是变化最小的部分。

Here, cooling stays on AC. Chillers, pumps, and fans still run on AC motors, requiring DC-to-AC inverters. Delta unveiled a 2.4MW In-Row CDU supporting 800VDC at GTC 2026, the first major cooling component engineered for native DC. But the full stack (chillers, compressors, pumps, building controls) remains AC-dependent, and no vendor sells an integrated DC-native cooling system.

在此阶段，冷却系统仍采用交流电（AC）。冷水机组、水泵和风扇仍由交流电机驱动，因此需要直流-交流（DC-to-AC）逆变器。台达（Delta）在 GTC 2026 上展示了一款支持 800VDC 的 2.4MW 列间级 CDU，这是首个专为原生直流电设计的重大冷却组件。然而，整个系统栈（冷水机组、压缩机、水泵、建筑控制系统）仍然依赖交流电，且目前尚无供应商销售集成的直流原生冷却系统。

Generator architecture is already loosening at some hyperscalers independently of 800VDC. Meta is likely bypassing generators at new sites entirely, and Microsoft's new designs use partial generator coverage. 800VDC could accelerate that direction, as supercapacitors, BBUs, and BESS form a distributed backup hierarchy that absorbs the functions generators used to own.

在一些超大规模运营商中，发电机架构已经开始脱离 800VDC 独立演进。Meta 可能会在新站点完全跳过发电机，而微软的新设计则采用了部分发电机覆盖。800VDC 可能会加速这一趋势，因为超级电容器、电池备用单元（BBU）和电池储能系统（BESS）形成了一个分布式备份层级，从而接管了发电机过去承担的功能。

Medium Voltage Rectifiers: Is There Room for Everyone?

中压整流器：每个人都有立足之地吗？

One reasonable question to ask is why power is rectified at the LV level and not directly from MV? The answer comes from semiconductor ratings. Rectifying from 13.8kV or 34.5kV requires devices rated above 10kV, which barely exist in commercial form today. That said, the gap is closing, and Wolfspeed's 10kV SiC MOSFET has been commercially available as bare die since March 2026.

一个合理的问题是，为什么要在低压（LV）级别进行整流，而不是直接从中压（MV）级别整流？答案在于半导体的额定电压。从 13.8kV 或 34.5kV 进行整流需要额定电压超过 10kV 的器件，而这类器件在目前的商业形态中几乎不存在。尽管如此，这一差距正在缩小，Wolfspeed 的 10kV SiC MOSFET 自 2026 年 3 月起已作为裸芯片投入商用。

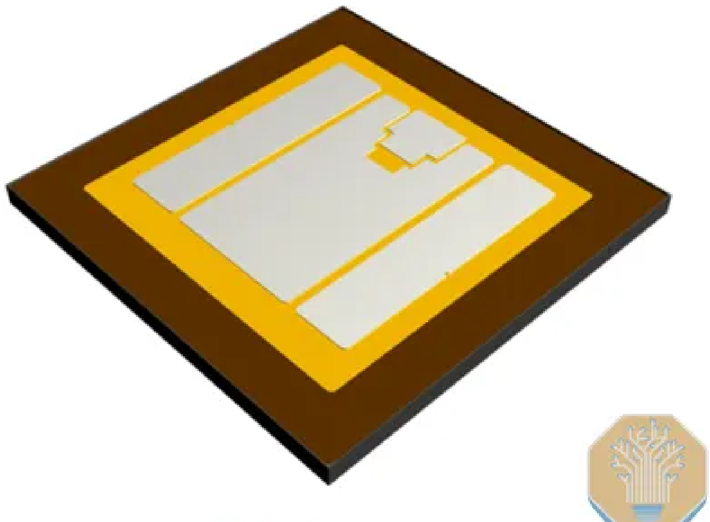

Wolfspeed's 10 kV Bare Die MOSFET

Source: Wolfspeed Wolfspeed

The development of SiC MOSFETs above 10kV opens the door to a second evolution of Phase 3, where even the LV equipment drops out of the main power bus. Continuing the trend, this collapses additional conversion steps and brings new efficiency gains.

10kV 以上 SiC MOSFET 的开发为第三阶段的第二次演进开启了大门，届时甚至连低压设备也将从主功率总线中移除。顺应这一趋势，这将减少额外的转换步骤，并带来新的效率提升。

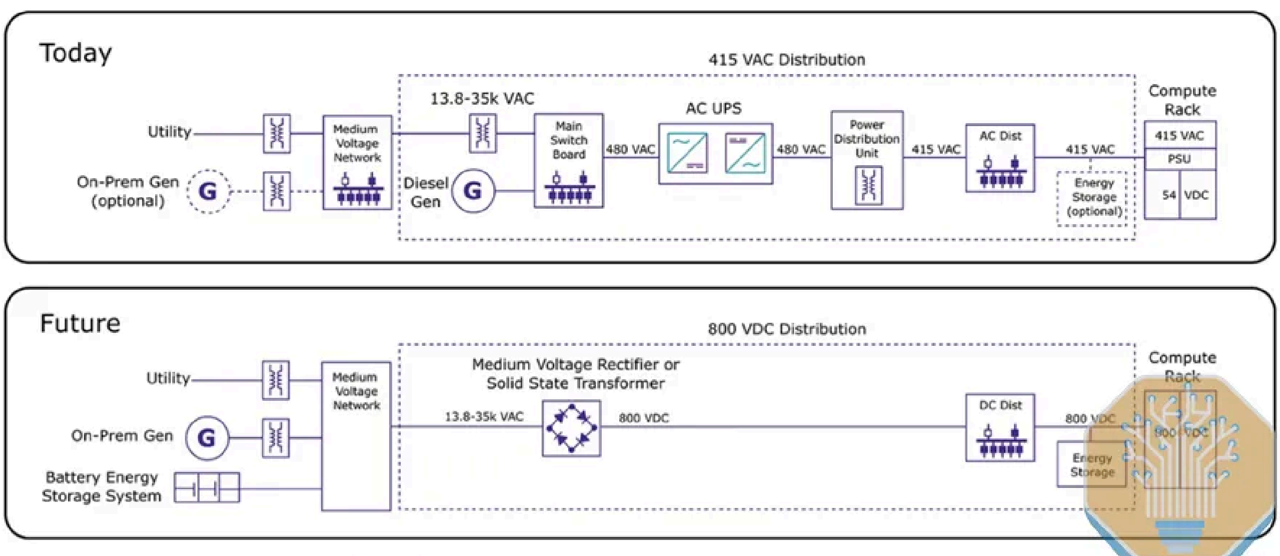

Figure 1: NVIDIA's 800 VDC power distribution in next-generation AI data centers[1]

Source: Wolfspeed Wolfspeed

The end state of our HVDC timeline will push even further. Even though conventional rectifiers with series-stacked silicon devices can handle MV rectification, an emerging technology promises to do it in a much more efficient, compact, and faster way. That technology is our protagonist for the next chapter of our journey: Solid State Transformers.

我们的高压直流（HVDC）时间线的终点将推向更远。尽管采用串联堆叠硅器件的传统整流器可以处理中压（MV）整流，但一项新兴技术有望以更高效、更紧凑、更快速的方式实现这一目标。这项技术正是我们下一章旅程的主角：固态变压器。

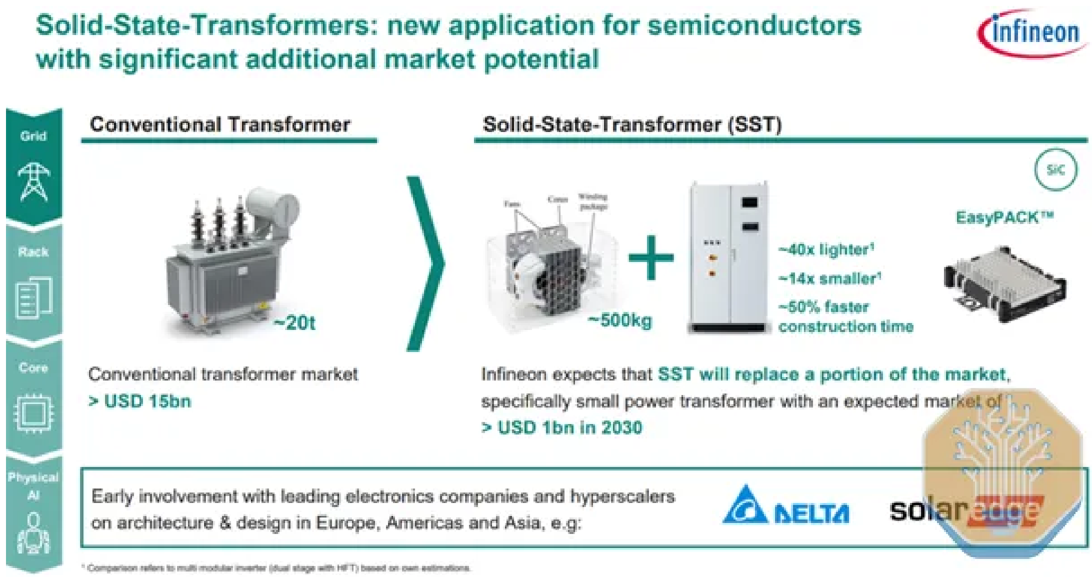

Source: Infineon 英飞凌

Phase 4 (>2029): SSTs, the End-State

第 4 阶段（2029 年以后）：SST，最终形态

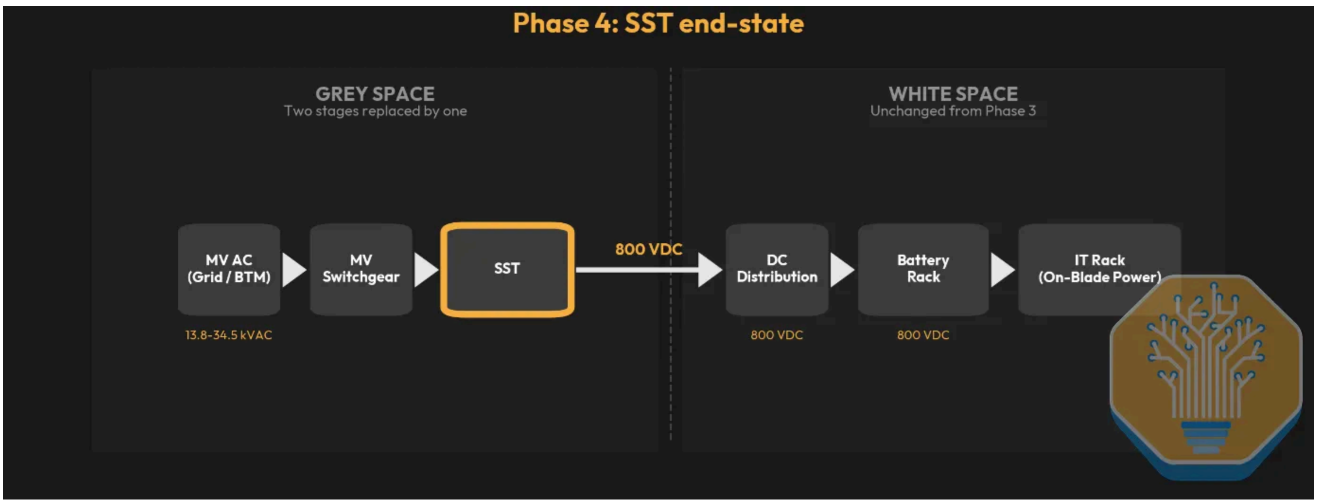

Source: SemiAnalysis Industrials Model

来源：SemiAnalysis 工业模型

Finally, we get to the holy grail of DC power distribution: Solid State Transformers, or SSTs. These are a new category of power electronic devices that replace conventional iron-core transformers with high-frequency, semiconductor-based converters.

最后，我们来到了直流配电的终极目标：固态变压器（SST）。这是一种新型电力电子设备，它利用基于半导体的高频转换器取代了传统的铁芯变压器。

Source: DGMatrix DGMatrix

Phase 4 and its datacenter layout is very similar to Phase 3. The major change is that the SST replaces the LV AC-DC rectifier and low-voltage transformer with a single piece of equipment that converts directly from medium voltage to 800VDC. If we consider the ending of the previous section, the possibility of using a MV rectifies that rectifies directly from MV AC, the architecture is essentially identical.

第四阶段及其数据中心布局与第三阶段非常相似。主要的改进在于，固态变压器（SST）取代了低压交流-直流整流器和低压变压器，仅通过单一设备即可将中压直接转换为 800VDC。如果我们考虑到上一节末尾提到的可能性，即使用直接从中压交流电进行整流的中压整流器，那么这两种架构在本质上是完全相同的。

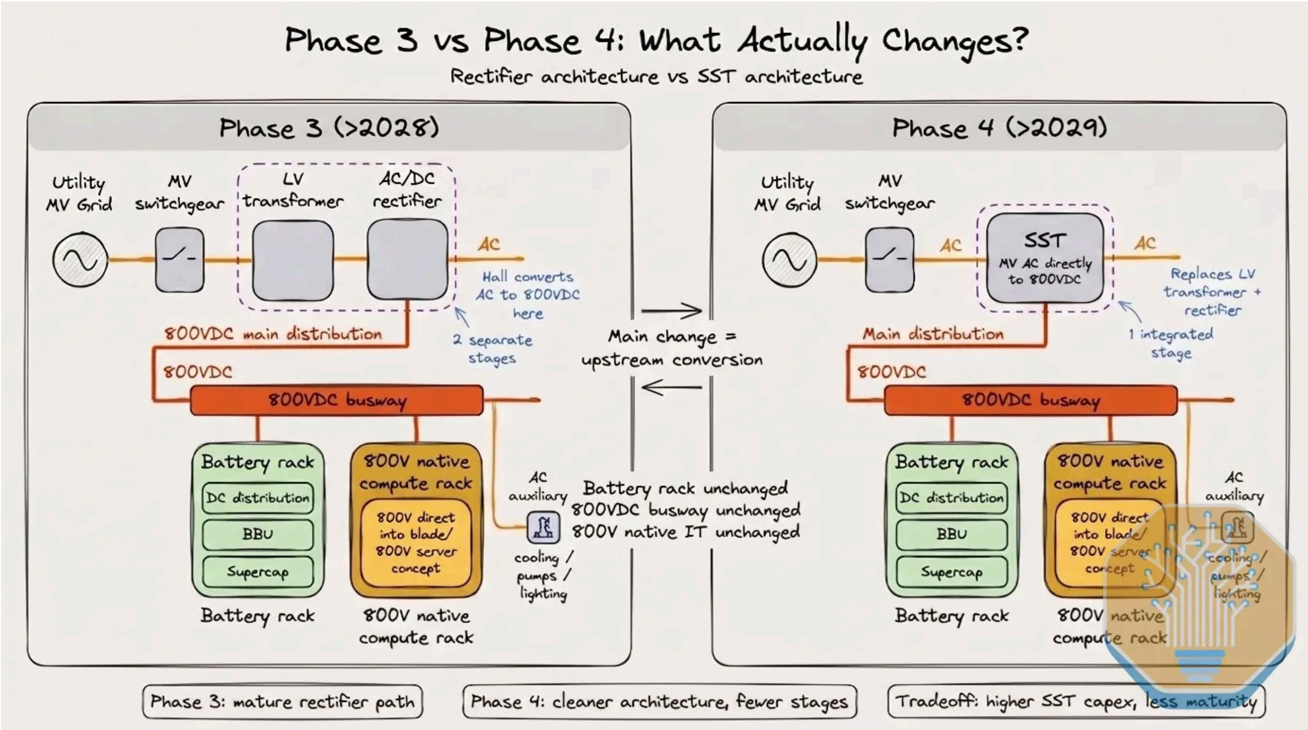

Source: SemiAnalysis SemiAnalysis

In a Nutshell: Solid-State Transformers

简而言之：固态变压器

Introduction to SSTs SST 简介

An SST does the same job as the massive iron-and-copper transformers in every datacenter’s grey space: step voltage from utility-level medium voltage to a level IT equipment can use. A conventional transformer uses magnetic induction at grid frequency. An SST uses semiconductor switching stages to achieve the same conversion in a fraction of the volume.

SST 的作用与每个数据中心灰色空间中那些庞大的铁铜变压器相同：将电压从公用事业级的中压降低到 IT 设备可以使用的水平。传统变压器在电网频率下利用磁感应原理工作。而 SST 则利用半导体开关级，以极小的体积实现相同的转换。

The datacenter SST is a three-stage device. The input stage converts AC to DC, handling the dangerous medium-voltage level (13.8 to 45kV) using SiC MOSFETs rated at 3,300V or higher. The isolation stage is where the size reduction happens. A high-frequency transformer steps the voltage down while providing galvanic isolation between the utility / power-source and the datacenter. The output stage produces the final 800VDC that the distribution system needs, with no inverter required.

数据中心 SST 是一种三级设备。输入级将交流电（AC）转换为直流电（DC），使用额定电压为 3,300V 或更高的 SiC MOSFET 来处理危险的中压水平（13.8 至 45kV）。隔离级是实现体积缩减的关键所在。高频变压器在降低电压的同时，在公用事业/电源与数据中心之间提供电隔离。输出级则产生配电系统所需的最终 800VDC，且无需逆变器。

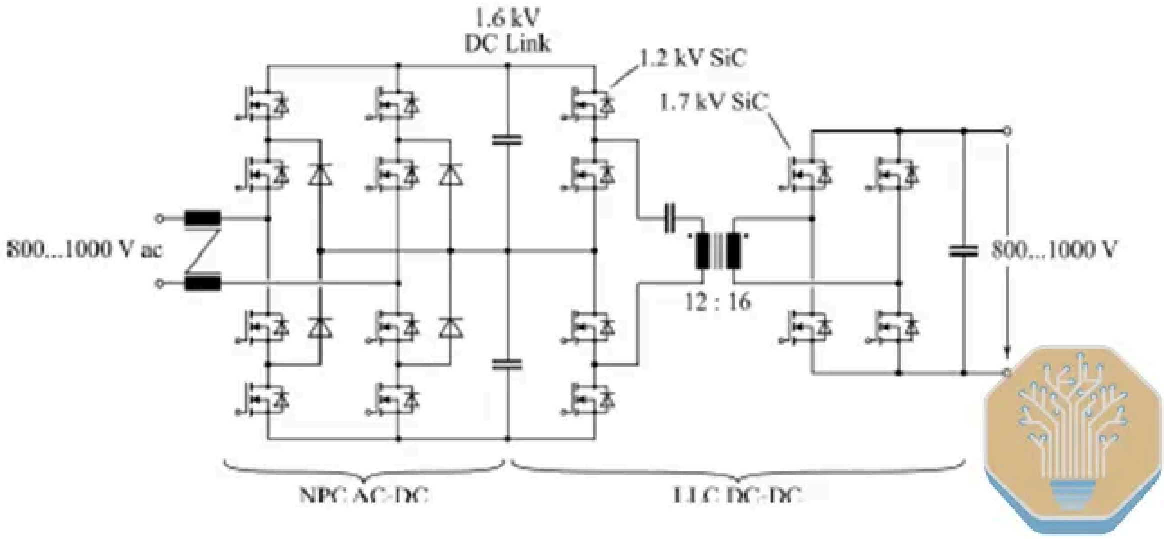

Source: ETH Zurich 苏黎世联邦理工学院

Pros and Cons of SSTs

固态变压器（SST）的优缺点

SSTs' core value proposition is energy efficiency, which translates directly to OPEX savings or unlocked compute capacity. By collapsing the medium-voltage transformer and rectifier into a single power-electronic stage, SSTs eliminate two conversion steps from the electrical chain. Vendors target up to 15% total system efficiency improvement, claiming the path rises from around 82-85% towards over 97%.

SST 的核心价值主张在于能源效率，这可以直接转化为运营成本（OPEX）的节省或计算能力的释放。通过将中压变压器和整流器合并为单个电力电子级，SST 消除了电气链中的两个转换步骤。厂商的目标是将系统总效率提高多达 15%，声称效率路径将从 82-85% 左右提升至 97% 以上。

SSTs are also dramatically smaller. A conventional transformer operates at 50 or 60 Hz and needs a massive iron core. An SST switches at 20,000 Hz or higher, shrinking the core by roughly 90%. That is where Infineon's claimed 40x weight reduction and 14x size reduction (!) come from.

SST 的体积也显著缩小。传统变压器在 50 或 60 Hz 下运行，需要巨大的铁芯。而 SST 的开关频率高达 20,000 Hz 或更高，使铁芯体积缩小了约 90%。这就是 Infineon 所声称的重量减轻 40 倍、体积缩小 14 倍（！）的由来。

Source: EENews EENews

In addition, SSTs are programmable. A conventional transformer steps voltage at a fixed ratio. An SST actively regulates output, adjusting under load. It also supports bidirectional power flow (pushing power to the grid during demand response, or charging a BESS). That said, SSTs with bidirectional capability and integrated BESS may trigger DER reclassification by the interconnecting utility, requiring IEEE 1547/2800 compliance.

此外，固态变压器（SST）具有可编程性。传统变压器按固定比例进行电压变换，而 SST 则能主动调节输出，根据负载情况进行调整。它还支持双向功率流（在需求响应期间向电网输送电力，或为电池储能系统 BESS 充电）。尽管如此，具备双向能力并集成 BESS 的 SST 可能会触发互连公用事业公司对分布式能源（DER）的重新分类，从而要求符合 IEEE 1547/2800 标准。

One additional major value proposition from SSTs is input flexibility. Some SST architectures extend this flexibility into multi-port topology, where a single device aggregates several inputs (utility AC, on-site generation, DC sources) and routes power across multiple outputs in software, including bidirectionally. The case for multi-port is that it reduces stranded power between zones and lets operators orchestrate flows across the site.

固态变压器（SST）的另一个主要价值主张是输入的灵活性。一些 SST 架构将这种灵活性扩展到多端口拓扑结构中，单个设备可以聚合多个输入（公用事业交流电、现场发电、直流电源），并通过软件在多个输出之间路由电力，包括双向路由。多端口方案的优势在于它减少了区域间的闲置功率，并允许运营商协调整个场地的潮流。

Reliability 可靠性

Conventional transformers last 30-40 years as passive devices. No SST vendor has published field reliability data at datacenter scale, as the longest deployment is the Hitachi-ABB PETT on Swiss Federal Railways, running since 2011. SSTs concentrate heat in semiconductor junctions and require active cooling, with DG Matrix using integrated liquid cooling and Novos Power using air cooling through proprietary insulation.

作为无源器件，传统变压器的寿命可达 30-40 年。目前还没有 SST 供应商发布过数据中心规模的现场可靠性数据，因为运行时间最长的部署是瑞士联邦铁路上的 Hitachi-ABB PETT，自 2011 年开始运行。SST 的热量集中在半导体结中，需要主动冷却，其中 DG Matrix 采用集成液冷技术，而 Novos Power 则通过专利绝缘技术采用风冷。

ETH Zurich's comparative evaluation found that a line-frequency transformer paired with a SiC rectifier can match SST efficiency and functionality. Datacenter-scale SSTs depend on SiC MOSFETs at 3,300V+ for the MV input stage, still in limited production. GaN, capped at roughly 650V, serves only downstream stages converting 800VDC to rack-level voltages.

苏黎世联邦理工学院（ETH Zurich）的对比评估发现，工频变压器与碳化硅（SiC）整流器的组合在效率和功能上可与固态变压器（SST）相媲美。数据中心级 SST 的中压（MV）输入级依赖于 3,300V 以上的 SiC MOSFET，而此类器件目前产量有限。氮化镓（GaN）的耐压上限约为 650V，仅适用于将 800VDC 转换为机架级电压的下游功率级。

Current Efficiency State 效率现状

The best public SST benchmark comes from ETH Zurich: 98% efficiency at 400kW in a 13.2kVAC-to-800VDC prototype presented at INTELEC 2025. Johann Kolar frames 98.0-98.5% as today's state of the art for full-scale SSTs, with 99% as the next engineering target for datacenter units.

目前最出色的公开固态变压器（SST）基准数据来自苏黎世联邦理工学院（ETH Zurich）：在 INTELEC 2025 上展示的一台 13.2kVAC 转 800VDC 原型机，在 400kW 功率下实现了 98% 的效率。Johann Kolar 指出，98.0-98.5% 是当今全规模 SST 的最高技术水平，而 99% 是数据中心单元的下一个工程目标。

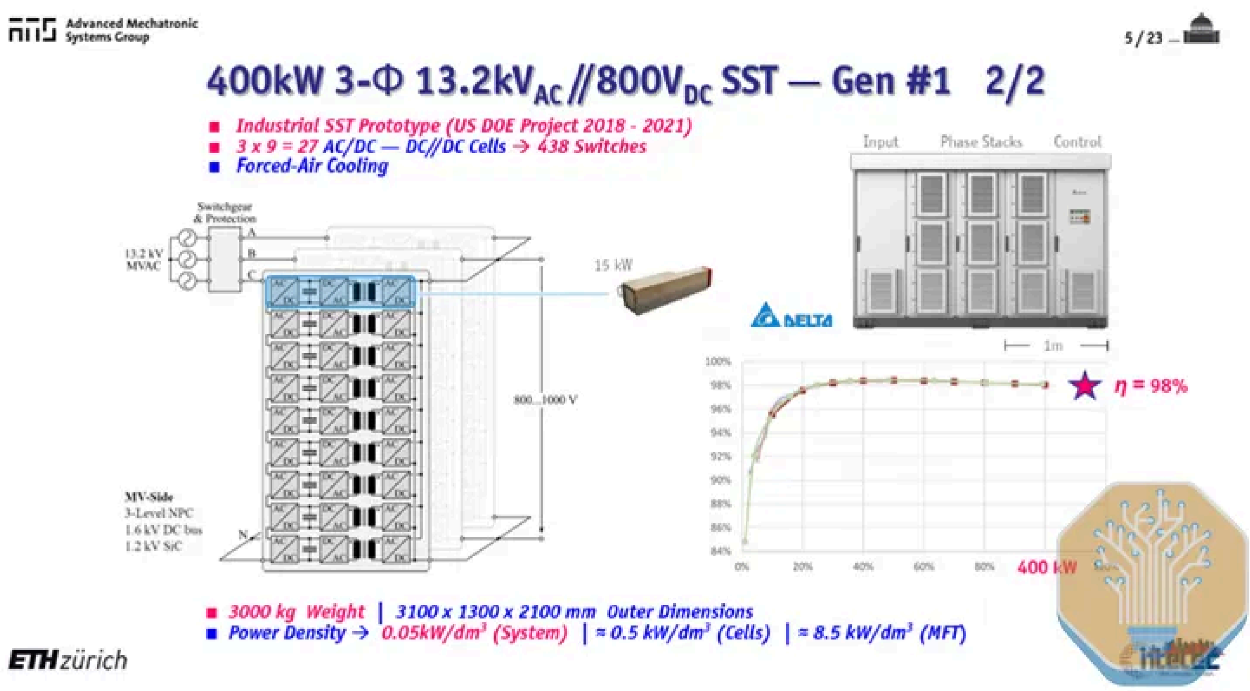

Source: ETH Zurich 苏黎世联邦理工学院

Different vendors now converge on that 98.5% ceiling: DG Matrix's Interport platform claims up to 98.5%, Amperesand's third-generation system claims greater than 98.5%, and Heron Power's Heron Link targets 98.5% MV-to-rack efficiency. Novos Power reports peak efficiencies over 98%. These are encouraging, but datacenters will need 3-

6MW units sustaining 99%+ efficiency under continuous load.

各大厂商目前正趋于 98.5% 这一效率上限：DG Matrix 的 Interport 平台声称效率高达 98.5%，Amperesand 的第三代系统声称效率超过 98.5%，而 Heron Power 的 Heron Link 则以 98.5% 的中压到机架（MV-to-rack）效率为目标。Novos Power 报告的峰值效率超过 98%。这些数据令人鼓舞，但数据中心仍需要 3-6MW 的机组在持续负载下保持 99% 以上的效率。

Two data points suggest the scale-up is underway. Chinese trade press reports that China XD Electric has deployed 2.4MW datacenter SSTs under the “East Data West Compute” program. NC State’s FREEDM Systems Center, the academic foundation which marked the origin of DG Matrix, has demonstrated 210 kHz switching at 3.3kV SiC with a 99% efficiency target for modular DC-DC SST variants.

有两个数据点表明规模化应用正在进行中。据中国行业媒体报道，中国西电已在“东数西算”工程中部署了 2.4MW 的数据中心电力电子变压器（SST）。作为 DG Matrix 技术起源的学术基础，北卡罗来纳州立大学的 FREEDM 系统中心已展示了在 3.3kV 碳化硅（SiC）下实现 210 kHz 的开关频率，其模块化直流-直流 SST 变体的效率目标为 99%。

Vendor Landscape 供应商概况

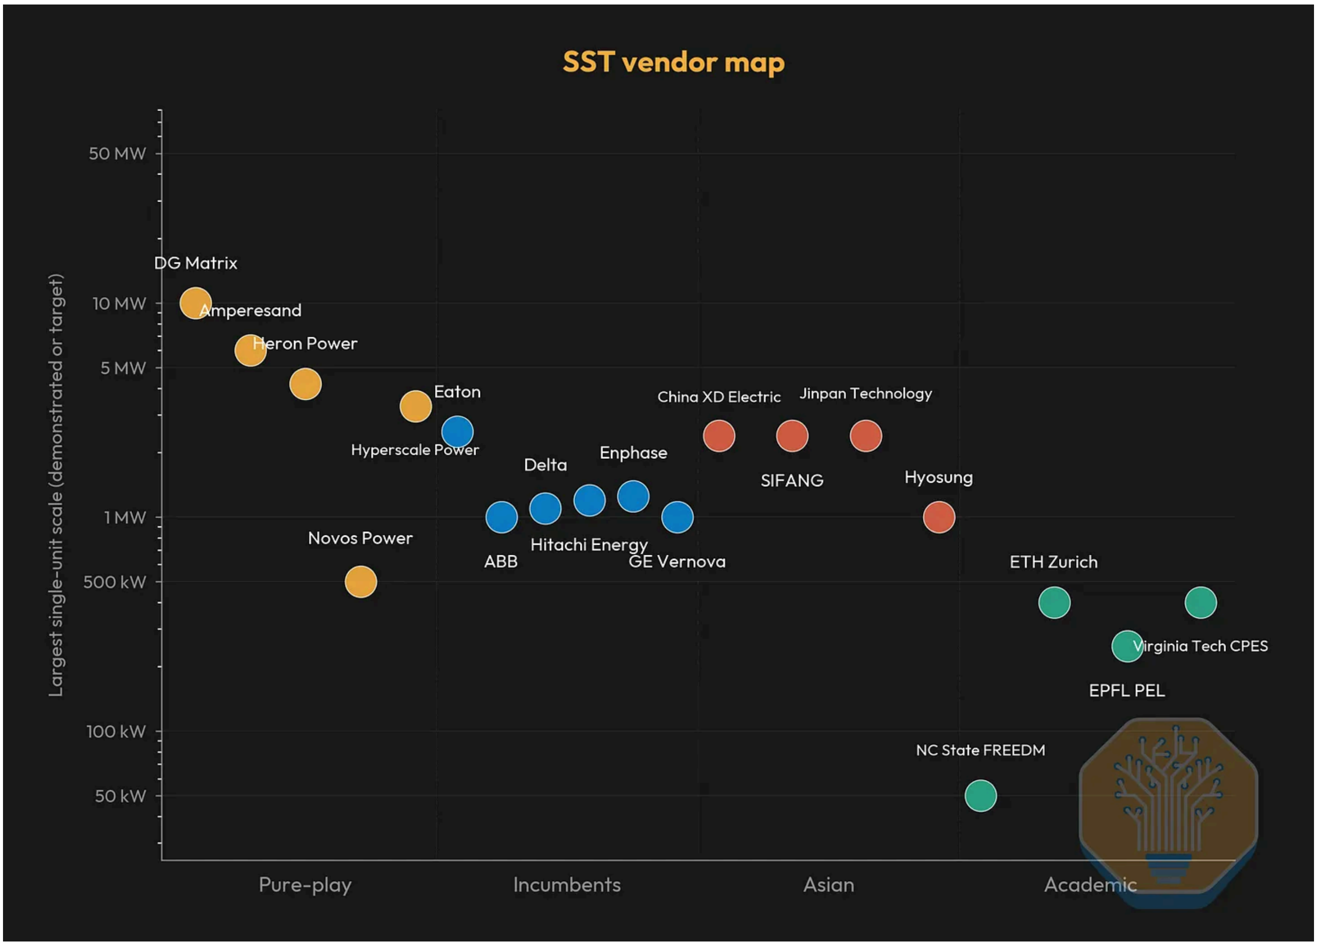

[Source: SemiAnalysis Industrials Model]

来源：SemiAnalysis 工业模型

The vendor landscape is moving fast. DG Matrix (ABB-backed, Infineon SiC supply deal) is shipping pre-certification units and targeting UL certification by end of Q2 2026. It is the only SST included in Nvidia's MGX reference architecture. Amperesand targets 30MW of commercial deployments in 2026. Heron Power is building a 40GW US manufacturing facility for its 4.2MW Heron Link units.

供应商格局正在快速演变。DG Matrix（由 ABB 支持，并与英飞凌签署了 SiC 供应协议）正在交付预认证单元，并计划在 2026 年第二季度末获得 UL 认证。它是唯一被纳入 Nvidia MGX 参考架构的固态变压器（SST）。Amperesand 的目标是在 2026 年实现 30MW 的商业部署。Heron Power 正在为其 4.2MW 的 Heron Link 单元建设一座产能达 40GW 的美国制造工厂。

Within the SST category, products are bifurcating along LV and MV input. DG Matrix and Amperesand are pursuing both, starting with LV-input SST skids (3.2-4.8 MW) that can be deployed today alongside existing AC distribution, and following with MV-input units as 3,300V-class SiC matures. Heron Power and Novos Power are concentrating on direct MV-input units that collapse the LV transformer and rectifier into a single device. Both paths converge on 800VDC at the output, but the LV path offers a shorter time-to-deployment at the cost of retaining the upstream MV-to-LV transformer.

在固态变压器（SST）类别中，产品正沿着低压（LV）和中压（MV）输入两个方向分化。DG Matrix 和 Amperesand 正在双管齐下，首先推出可与现有交流配电系统共同部署的低压输入 SST 撬装设备（3.2-4.8 MW），并随着 3,300V 级碳化硅（SiC）技术的成熟，后续推出中压输入单元。Heron Power 和 Novos Power 则专注于直接中压输入单元，将低压变压器和整流器整合为单一设备。这两条路径最终都在输出端汇聚为 800VDC，但低压路径以保留上游中压转低压变压器为代价，换取了更短的部署周期。

Source: DG Matrix DG Matrix

Novos Power claims a direct MV-to-800VDC SST with 50% smaller footprint and air cooling. On the incumbent side, Eaton acquired Resilient Power Systems in August 2025 for SST expertise. More than $320M flowed into SST startups in the twelve months ending March 2026.

Novos Power 声称其直接中压转 800VDC 的固态变压器（SST）占地面积减小了 50%，并采用风冷技术。在现有巨头方面，伊顿（Eaton）于 2025 年 8 月收购了 Resilient Power Systems，以获取 SST 领域的专业技术。在截至 2026 年 3 月的十二个月内，超过 3.2 亿美元的资金流向了 SST 初创企业。

Source: Novos Power Novos Power

Datacenter Layout Implications

数据中心布局影响

The SST eliminates the LV equipment at ~\$0.55M/MW and the Phase 2 rectifier at ~\$0.20M/MW. At an estimated SST cost of ~\$1.0-1.5M/MW, we expect the first instances of SSTs to come at an upfront Capex premium over directly replaced equipment.

SST（固态变压器）省去了约 0.55 美元/瓦的低压设备和约 0.20 美元/瓦的第二阶段整流器。按照 SST 约 1.0-1.5 美元/瓦的估算成本，我们预计首批 SST 的应用在前期资本支出（Capex）上会高于其直接替换的设备。

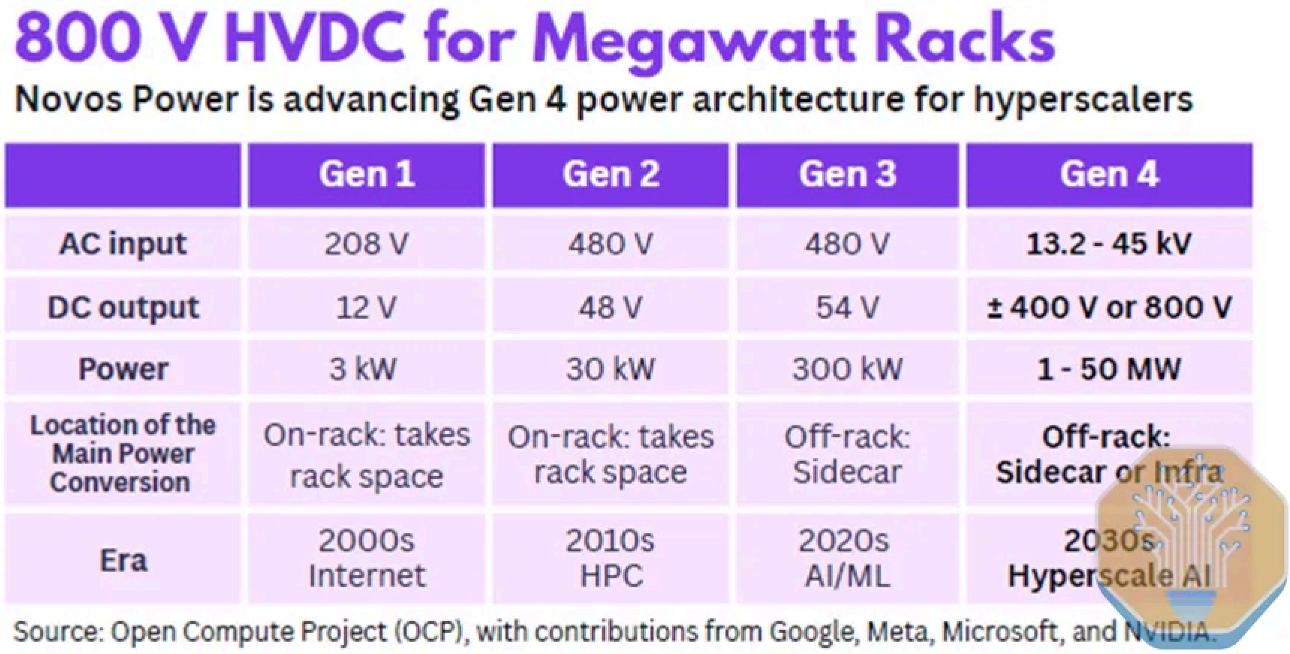

800 V HVDC for Megawatt Racks

Novos Power is advancing Gen 4 power architecture for hyperscalers

	Gen 1	Gen 2	Gen 3	Gen 4
AC input	208 V	480 V	480 V	13.2 - 45 kV
DC output	12 V	48 V	54 V	± 400 V or 800 V
Power	3 kW	30 kW	300 kW	1 - 50 MW
Location of the Main Power Conversion	On-rack: takes rack space	On-rack: takes rack space	Off-rack: Sidecar	Off-rack: Sidecar or Infra
Era	2000s Internet	2010s HPC	2020s AI/ML	2030s Hyperscale AI

Source: Open Compute Project (OCP), with contributions from Google, Meta, Microsoft, and NVIDIA.

Source: Novos Power Novos Power

The rest of the electrical architecture remains the same as in Phase 3. The 480V AC auxiliary bus for cooling, lighting, and facility systems carries over unchanged. On the IT rack-side, we expect that by the time SSTs are deployed, compute trays are already 800VDC native. However, we could see deployments of SST adoption with 800V microgrid and IT racks using a DC-DC power shelve converter, which could accelerate the adoption.

其余的电气架构保持与第 3 阶段相同。用于冷却、照明和设施系统的 480V 交流辅助总线保持不变。在 IT 机架端，我们预计到固态变压器（SST）部署时，计算托盘已经原生支持 800VDC。然而，我们也可能会看到 SST 的部署采用 800V 微网，而 IT 机架则使用 DC-DC 电源机架转换器，这可能会加速其普及。

On Phase 4 timings, this emerging technology is still in design phase, and we don’t expect major SST adoption at scale until early 2029. That said, we are aware that all major hyperscalers are running pilot and testing with main SST vendors, with commercial contracts already in place. As we cover in the following section, technology development itself will not be the only factor determining the adoption curve here. Regulatory framework and standards is a big one. In the SST space, no vendor has completed UL certification for datacenter SST deployment as of May 2026.

关于第 4 阶段的时间节点，这项新兴技术仍处于设计阶段，我们预计直到 2029 年初才会有大规模的 SST 应用。尽管如此，我们了解到所有主要的超大规模云服务商（hyperscalers）都在与主要的 SST 供应商进行试点和测试，并且已经签署了商业合同。正如我们在下一节中所讨论的，技术开发本身并不是决定此处普及曲线的唯一因素。监管框架和标准是一个重要因素。在 SST 领域，截至 2026 年 5 月，尚无供应商完成用于数据中心 SST 部署的 UL 认证。

The SST Market Opportunity and TAM sizing

SST 市场机遇与市场总量（TAM）估算

By 2030, we expect SST TAM to reach ~$13B, capturing the demand displaced from the sidecar layer plus the incremental MV-to-800VDC conversion. We consider a content of $1.25M/MW. A portion of this opportunity is contested by MV rectifiers, but we expect SSTs to capture the majority share.

到 2030 年，我们预计 SST 的市场总量（TAM）将达到约 130 亿美元，涵盖了从侧置层（sidecar layer）转移的需求以及从中压（MV）到 800VDC 转换的增量需求。我们估算的价值量为 125 万美元/兆瓦。这一机遇中的一部分正受到中压整流器的竞争，但我们预计 SST 将占据大部分份额。

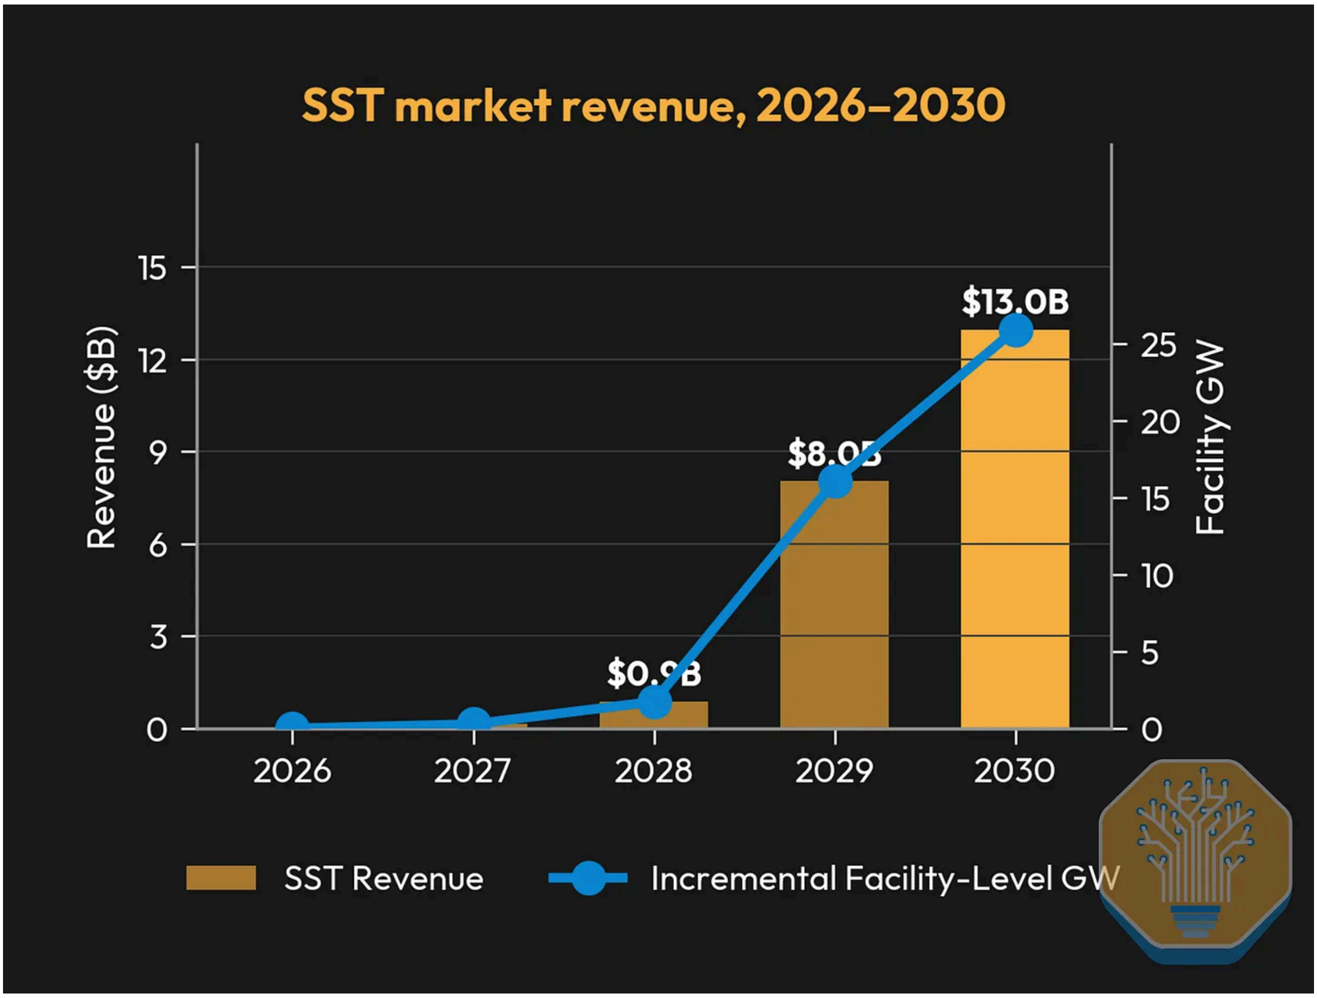

[Source: SemiAnalysis Industrials Model]

来源：SemiAnalysis 工业模型

Datacenter Layout Summary: Total Cost Barely Moves, Content Shifts, Efficiency Climbs

数据中心布局摘要：总成本几乎持平，构成发生转移，效率显著提升

Electrical System Cost 电气系统成本

Total electrical content per MW stays in a $3.6-4.8M band across four of the five architectures we model. The main headline is a content migration from grey space to white space, and the resulting change in equipment mix.

在我们模拟的五种架构中，有四种架构的每兆瓦（MW）总电气成本保持在 360 万至 480 万美元之间。核心要点是成本构成从灰色空间（grey space）向白色空间（white space）转移，并由此导致了设备组合的变化。

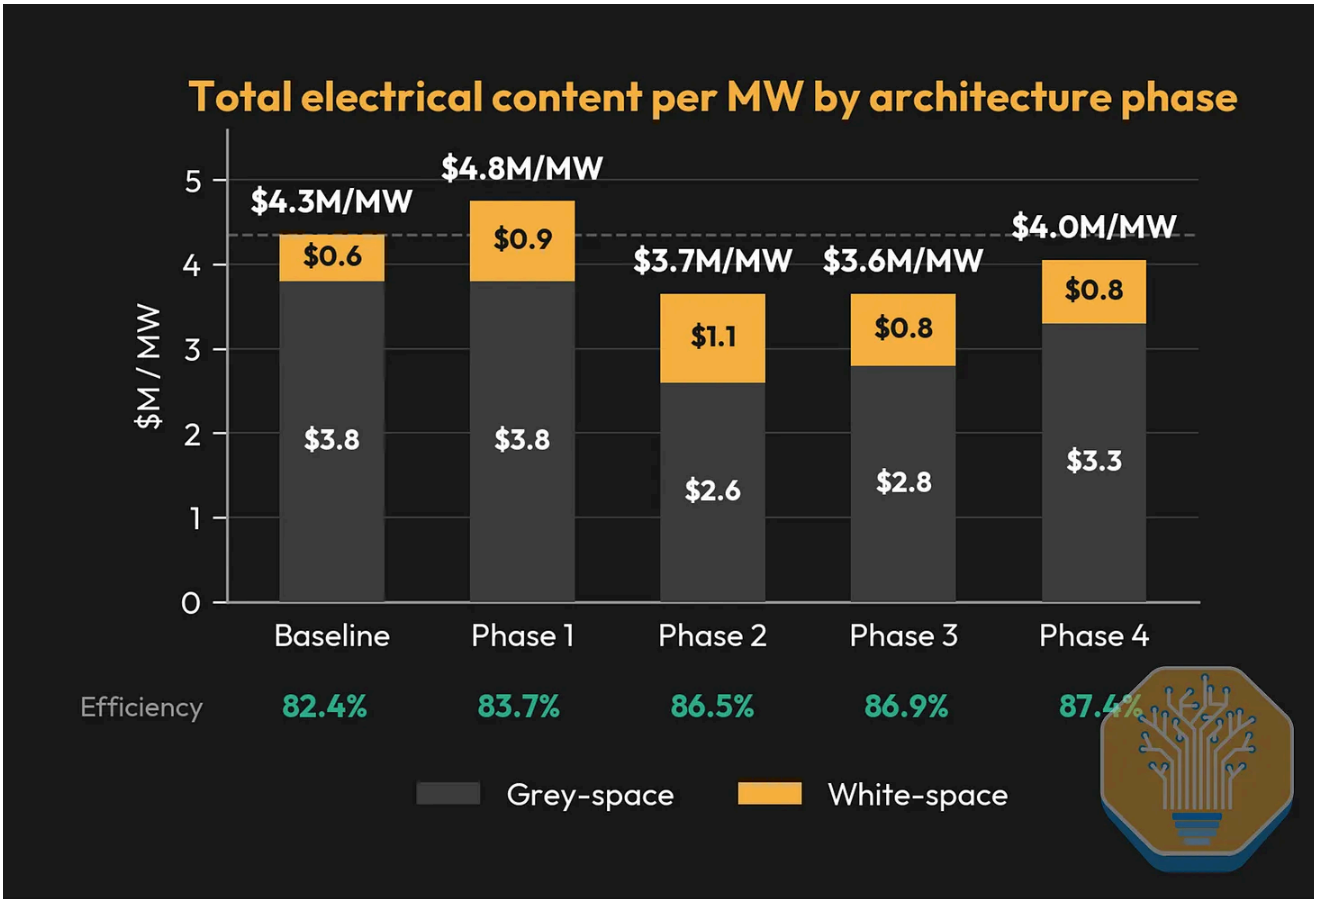

[Source: SemiAnalysis Industrials Model]

来源：SemiAnalysis 工业模型

Grey space content shrinks in Phase 2 as the centralized UPS ($1.2M) exits. White space peaks in Phase 1 because the HVDC power rack arrives. By Phase 4, total content climbs to $4.0M as the SST replaces the LV transformer and rectifier.

在第二阶段，随着集中式 UPS（120 万美元）的退出，灰色空间（设备区）内容有所缩减。由于高压直流（HVDC）电源机架的引入，白色空间（机房区）内容在第一阶段达到峰值。到第四阶段，随着固态变压器（SST）取代低压变压器和整流器，总价值攀升至 400 万美元。

Electrical System Efficiency

电气系统效率

We calculate the baseline AC power path at 82.0% cumulative efficiency across seven conversion stages. The VRM (92%) and PSU (94%) are the two largest single-stage losses. The VRM stays in every architecture, but the PSU’s loss is the largest penalty the 800VDC transition can eliminate. Phase 1 barely improves to an estimated 83.7%. The UPS double-conversion loop still eats 3 percentage points, and the new power rack rectifier (97.5%) plus DC-DC stage (97.0%) only marginally outperform the old single-stage PSU.

我们计算出基准交流电源路径在七个转换阶段的累计效率为 82.0%。电压调节模块（VRM，92%）和电源单元（PSU，94%）是两个最大的单级损耗点。VRM 存在于每种架构中，但 PSU 的损耗是 800VDC 转型可以消除的最大不利因素。第一阶段的效率仅略微提升至约 83.7%。UPS 的双变换回路仍消耗 3 个百分点，而新的电源机架整流器（97.5%）加上直流-直流（DC-DC）阶段（97.0%）的性能仅略优于旧的单级 PSU。

The real jump comes in Phase 2 (86.5%) when UPS elimination cuts the chain from seven stages to five. Phase 3 pushes to 86.9% because the centralized grey-space rectifier operates at MW scale (higher efficiency than modular rack-mounted units) and 800VDC hall-level distribution eliminates AC skin effect and reactive power losses. We estimate Phase 4 to reach 87.4% as the SST replaces two stages with a single device.

真正的飞跃发生在第二阶段（86.5%），此时取消 UPS 将链路从七个阶段缩减至五个。第三阶段效率提升至 86.9%，这是因为集中式灰区整流器以兆瓦级规模运行（效率高于模块化机架式单元），且 800VDC 大厅级配电消除了交流趋肤效应和无功功率损耗。我们预计第四阶段效率将达到 87.4%，因为 SST（固态变压器）用单个设备取代了两个阶段。

At 1GW of IT load, the Phase 2 gain translates to roughly 58MW of continuous grid power savings. Phase 3 extends that to 63MW and Phase 4 to 69MW. Nvidia cites up to 5% efficiency improvement, implying roughly 50MW at 1GW. Our Phase 4 efficiency delta calculations of 5% vs baseline matches Nvidia’s reported figures.

在 1GW 的 IT 负载下，第二阶段的增益相当于节省了约 58MW 的持续电网功率。第三阶段将这一数字扩大到 63MW，第四阶段则达到 69MW。Nvidia 提到的效率提升高达 5%，这意味着在 1GW 负载下约为 50MW。我们计算出的第四阶段与基准线之间 5% 的效率差值与 Nvidia 报告的数据相吻合。

The Other Side of the 800VDC Transition: Challenges and Limitations

800VDC 转型的另一面：挑战与局限

So far, we've mapped out a promising path, but as always, various challenges will arise along the way. We now unpack four main obstacles that will determine how fast 800VDC moves from small-case pilots to broader adoption.

到目前为止，我们已经规划出了一条充满希望的路径，但一如既往，沿途会出现各种挑战。我们现在将剖析四个主要障碍，它们将决定 800VDC 从小规模试点走向广泛应用的速度。

Challenge 1: Regulation, Safety and Grounding

挑战 1：监管、安全与接地

Regulation 法规

The National Electrical Code (NEC), published by NFPA on a three-year cycle, governs electrical installation in the United States. Adopted by nearly every state and municipality as binding law, it determines whether an operator can build to a standard design or must negotiate site-by-site with the local Authority Having Jurisdiction (AHJ). Full 800VDC code support targets NEC 2029. Pre-2029 deployments therefore require custom AHJ approvals and OEM-level UL certification for each site. This is workable for hyperscalers with in-house code engineering teams but could represent a meaningful barrier for colocation operators and smaller builders.

由美国国家消防协会（NFPA）每三年发布一次的《国家电气规范》（NEC）监管着美国的电气安装。由于几乎每个州和市政府都将其采纳为具有约束力的法律，它决定了运营商是能够按照标准设计进行建设，还是必须逐个站点与当地管辖机构（AHJ）进行协商。800VDC 的全面规范支持目标定于 NEC 2029。因此，2029 年之前的部署需要针对每个站点获得定制化的 AHJ 批准和 OEM 级的 UL 认证。这对于拥有内部规范工程团队的超大规模云服务商（Hyperscalers）来说是可行的，但对于托管运营商和小型建设者而言，可能构成实质性的障碍。

We see a useful parallel in the early days of the EV industry, where Tesla designed and approved its own internal safety frameworks because industry-wide standards had not yet arrived. Hyperscalers deploying 800VDC pre-2029 will be in a similar position.

我们在电动汽车（EV）行业早期看到了一个有益的类比：当时由于行业范围内的标准尚未出台，特斯拉设计并批准了自己内部的安全框架。在 2029 年之前部署 800VDC 的超大规模云服务商将处于类似的境地。

NEC 2029 would already be fast by historical standards, considering prior DC power standardization timelines in marine, telecom, and EV. However, timelines could benefit from extreme buyer concentration with five hyperscalers and Nvidia as both demand creator and solution architect and an EV 800V component supply chain.

考虑到船舶、电信和电动汽车领域以往的直流电（DC）标准化时间表，NEC 2029 的进度按历史标准衡量已经算很快了。然而，受益于极高的买家集中度（五大超大规模云服务商和作为需求创造者兼方案架构师的 Nvidia），以及现成的电动汽车 800V 组件供应链，这一时间表可能会进一步缩短。

We think NEC 2029 will achieve partial provision, while full code maturity probably lands at NEC 2032 or 2035. Partial means the basic framework exists (voltage classification, conductor sizing, overcurrent protection) but DC-specific arc flash PPE tables, busway standards, and stored energy maintenance protocols will likely be absent.

我们认为 NEC 2029 将实现部分条款的制定，而规范的全面成熟可能要等到 NEC 2032 或 2035 版本。所谓“部分”是指基本框架已经确立（电压分类、导体选型、过流保护），但针对直流特有的电弧闪络个人防护装备（PPE）表格、母线槽标准以及储能维护协议届时可能仍会缺失。

Safety 安全

The biggest safety risk is arc flash. IEEE 1584 does not cover DC, and NFPA 70E has no PPE table for 600-1000VDC. UL Solutions has launched a Direct Current Safety Research Consortium to build the missing hazard models, explicitly citing 800V DC datacenter architectures among the target applications.

最大的安全风险是电弧闪络。IEEE 1584 标准并不涵盖直流电，而 NFPA 70E 也没有针对 600-1000VDC 的个人防护装备（PPE）对照表。UL Solutions 已启动直流安全研究联盟（Direct Current Safety Research Consortium），旨在建立缺失的危险模型，并明确将 800V 直流数据中心架构列为目标应用之一。

Even beyond the code gap, daily reality may be harder. At 48V, a technician can hot-swap a server tray with minimal PPE. At 800V, many rack-adjacent tasks that were routine at 48V likely require a qualified person under NFPA 70E, with arc-rated clothing, insulated gloves rated to 1000V, and a face shield. Capacitor banks and BBU modules retain dangerous charge after power-down, and standard lockout-tagout procedures for AC do not account for stored DC energy. Multiple sources must each be verified de-energized before maintenance.

即便撇开规范缺失不谈，日常操作的难度也会增加。在 48V 电压下，技术人员只需穿戴极少的 PPE 即可热插拔服务器托盘。但在 800V 电压下，许多在 48V 时属于常规的机架旁操作，根据 NFPA 70E 的规定，可能需要由具备资质的人员执行，并配备防电弧服、耐压 1000V 的绝缘手套以及面罩。电容器组和电池备用单元（BBU）模块在断电后仍保留危险电荷，且标准的交流电挂牌锁死（LOTO）程序并未考虑到存储的直流能量。在维护之前，必须分别验证多个电源均已去耦。

Flex, a major Nvidia manufacturing partner, has publicly advocated for in-depth hazard identification and safety training at 800VDC facilities.

Nvidia 的主要制造合作伙伴 Flex 公开倡导在 800VDC 设施中进行深入的危险识别和安全培训。

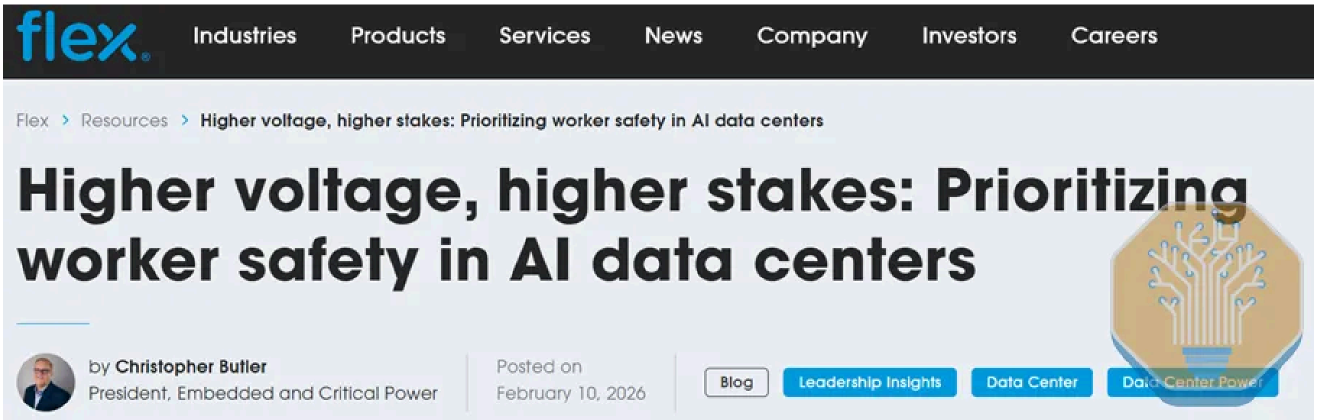

Source: Flex Flex

Grounding 接地

Grounding cascades into protection-device count, fault behavior, insulation monitoring, personnel safety, and vendor compatibility, which makes it on the most consequential early design choice in an 800VDC facility.

接地方式会连锁影响保护装置数量、故障行为、绝缘监测、人员安全以及供应商兼容性，这使其成为 800VDC 设施中最具影响力的早期设计选择之一。

The Siemens/Nvidia paper “Protections for Data Centers Powered by Direct Current” identifies four options. A ±400V system can use high-resistance grounding (HRG), which tolerates the first ground fault and only requires fast interruption on the second, or solid grounding, which demands immediate clearing of any fault. An 800V monopolar system can float, with insulation monitoring on every branch, or run a solid-grounded return conductor.

西门子与英伟达联合发表的论文《直流供电数据中心的保护方案》（Protections for Data Centers Powered by Direct Current）确定了四种方案。±400V 系统可以采用高电阻接地（HRG），这种方式允许发生第一次接地故障，仅在第二次故障时才需要快速切断；也可以采用直接接地，这要求立即清除任何故障。800V 单极系统可以采用浮地运行，并在每个支路上进行绝缘监测，或者运行一条直接接地的回路导线。

	Highest Fault Voltage	Protection Device Rating	Disconnect Needed on Both + and - ?	Grounnd Fault Mitigation
+/-400V –High Resistance Ground	800VDC	800VDC or 400/800VDC	YES	1st ground fault: Inherently 2nd ground fault: Fast protection devices
+/-400V – Solid Ground	800VDC	800VDC or 400/800VDC	YES	Fast protection devices
800V – Floating/ High Resistance	800VDC	800VDC or 400/800VDC	YES	1st ground fault: Inherently 2nd ground fault: Fast protection devices
800V – Solid Ground	800VDC	800VDC	NO	Fast protection devices

Table 1. Grounding schemes for 800VDC 2-wire system.

Source: Siemens, Nvidia Siemens, Nvidia

The tradeoff is cost. HRG and floating systems need protection devices rated for the full 800VDC on both conductors plus insulation monitoring infrastructure. Solid-grounded-return cuts the protection device count but eliminates galvanic isolation between parallel converters. OCP Diablo 400 permits both ±400V bipolar and 800V monopolar, leaving the call to the operator.

权衡之处在于成本。高阻接地（HRG）和浮地系统需要两个导体上的保护器件均达到800VDC 额定值，并配备绝缘监测基础设施。固定接地回路虽然减少了保护器件的数量，但消除了并行转换器之间的电流隔离。OCP Diablo 400 同时允许 ±400V 双极和800V 单极模式，将选择权留给了运营商。

The reality is that no industry consensus exists. SST and power-electronics vendors are optimizing around different grounding assumptions, which makes the choice a vendor-ecosystem commitment, not just a technical one.

现实情况是，行业内尚未达成共识。固态变压器（SST）和电力电子设备供应商正围绕不同的接地假设进行优化，这使得该选择不仅是一项技术决策，更是一项对供应商生态系统的承诺。

Challenge 2: Cooling and Auxiliary AC Workloads

挑战 2：冷却和辅助交流负载

Cooling is the largest AC load in an 800VDC datacenter, and no vendor sells a DC-native cooling ecosystem. Some vendors like Delta and Danfoss are doing progress. Danfoss's Turbocor compressors, dominant in datacenter chillers, run internally on DC at 700-813V. Danfoss also manufactures the VACON NXP variable-frequency drive, which accepts 640-1200 VDC input directly, placing 800V within its operating range. DCAirco ships 800V DC chillers for e-mobility at 4-8kW, 100-1000x too small for datacenter scale, but proof the refrigeration cycle works at this voltage.

冷却是 800VDC 数据中心中最大的交流负载，目前尚无供应商销售原生支持直流的冷却生态系统。Delta 和 Danfoss 等供应商正在取得进展。Danfoss 的 Turbocor 压缩机在数据中心冷水机组中占据主导地位，其内部运行电压为 700-813V 直流电。Danfoss 还生产 VACON NXP 变频驱动器，可直接接受 640-1200 VDC 输入，使 800V 处于其工作范围内。DCAirco 为电动交通领域提供 4-8kW 的 800V 直流冷水机组，虽然其功率比数据中心规模小 100-1000 倍，但证明了制冷循环在此电压下是可行的。

Beyond cooling, switchgear operating mechanism, lighting, fire suppression pumps, building management sensors and security systems all run on AC. As Nvidia team presented at OCP Global Summit 2025, 800VDC reference architecture will retain an

AC auxiliary bus alongside the 800VDC compute distribution for exactly this reason.

除了冷却系统外，开关柜操作机构、照明、消防泵、楼宇管理传感器和安防系统均在交流电（AC）下运行。正如 Nvidia 团队在 2025 年 OCP 全球峰会上所展示的，正是出于这个原因，800VDC 参考架构将在 800VDC 计算配电网络旁保留一条交流辅助总线。

That said, the supply chain is moving. The Delta CDU noted above is the leading edge, but most auxiliary categories (lighting, fire suppression, security) lack DC variants. With datacenter industrial capex heading toward >$400 billion in 2026 and electrical equipment at 30-35%, the incentive to develop DC-native products is growing.

尽管如此，供应链正在发生变化。上述台达 CDU 处于领先地位，但大多数辅助类别（照明、消防、安防）仍缺乏直流变体。随着 2026 年数据中心工业资本支出趋向于超过 4000 亿美元，且电气设备占比达 30-35%，开发直流原生产品的动力正在日益增强。

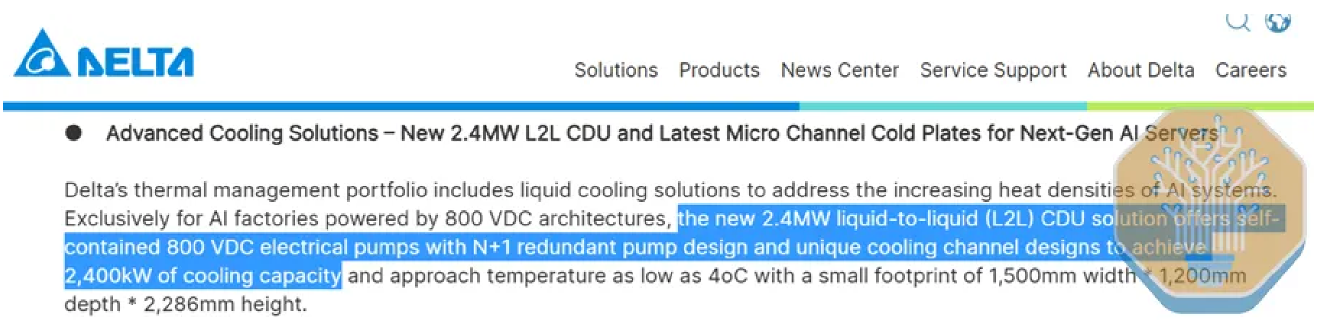

Source: Delta 台达

Challenge 3: Supply Chain Standards

挑战 3：供应链标准

Innovation in DC distribution is ahead of codification, and standards still lag across most 800VDC equipment categories. Busway is a good example of progress. UL 857, the standard governing busway systems, originally capped coverage at 600V and defined values in root-mean-square (RMS). Edition 14, published in 2025, raised the ceiling to 1000VDC, and Edition 15 in development targets 1500VDC. Outside busway, certification paths remain absent, and every installation becomes a custom engineering project where the operator must qualify the product, negotiate conductor ratings, and obtain AHJ approval on a case-by-case basis.

直流配电领域的创新领先于规范制定，且大多数 800VDC 设备类别的标准仍然滞后。母线槽（Busway）是取得进展的一个良好范例。UL 857 是监管母线槽系统的标准，最初将其覆盖范围限制在 600V，并以均方根值（RMS）定义数值。2025 年发布的第 14 版将上限提高到了 1000VDC，而正在制定的第 15 版目标则是 1500VDC。在母线槽之外，认证路径依然缺失，每一次安装都变成了一个定制工程项目，运营商必须逐案对产品进行鉴定、协商导体额定值，并获得管辖当局（AHJ）的批准。

Source: UL Solutions UL Solutions

An OCP white paper targeting 2026 may help, and OCP working groups are coordinating with regulators and certification bodies to land initial standards by year-end 2026, but vendors are already presenting their prototypes. Delta demonstrated 800VDC air-cooled busway at OCP 2025, LS Electric exhibited DC power equipment at DistribuTECH 2026, and in almost all recent conferences the team has been present in, 800VDC-ready prototypes have been without a doubt the protagonists.

一份针对 2026 年的 OCP 白皮书可能会有所帮助，OCP 工作组也正在与监管机构和认证机构协调，力争在 2026 年底前落实初步标准，但供应商们已经开始展示其原型产品。台达（Delta）在 OCP 2025 上展示了 800VDC 风冷母线槽，LS Electric 在 DistribuTECH 2026 上展出了直流配电设备，而在团队近期参加的所有会议中，支持 800VDC 的原型产品毫无疑问都是全场的主角。

Source: LS Electric LS Electric

Challenge 4: Grid Interconnection and Regulatory Pressure

挑战 4：电网互连与监管压力

As we covered in our deep dive on AI training load fluctuations at gigawatt scale, datacenter load-loss events during grid disturbances have become a serious concern for grid operators. 800VDC sharpens the problem by moving grid-facing behavior into software-defined power electronics (SST control algorithms, converter current limits, DC bus capacitance…).

正如我们在深入探讨吉瓦级 AI 训练负载波动时所提到的，电网扰动期间的数据中心负载丢失事件已成为电网运营商严重关注的问题。800VDC 进一步加剧了这一问题，因为它将面向电网的行为转移到了软件定义的电力电子设备中（如 SST 控制算法、变流器电流限制、直流母线电容等）。

Grid operators now have to model and constrain those dynamics, and the regulatory bar is also rising. NERC issued a Level 3 Essential Actions Alert (its highest tier) on

May 2026 covering large computational loads, with a mandatory response deadline of August 3, and has proposed a Computational Load Entity registration for datacenters consuming 1MW+ within a 20MW+ aggregate at 60kV+. ERCOT's NOGRR282 adds voltage and frequency ride-through requirements and mandates both PSS/E and PSCAD electromagnetic transient models for all large loads.

电网运营商现在必须对这些动态特性进行建模和约束，且监管门槛也在不断提高。北美电力可靠性协会（NERC）于 2026 年 5 月发布了针对大型计算负载的 3 级基本行动警报（其最高级别），强制响应截止日期为 8 月 3 日，并提议对在 60kV+ 电压等级下、20MW+ 总容量中消耗 1MW+ 的数据中心进行“计算负载实体”注册。德克萨斯州电力可靠性委员会（ERCOT）的 NOGRR282 增加了电压和频率穿越要求，并强制要求所有大型负载提供 PSS/E 和 PSCAD 电磁暂态模型。

Why 800VDC raises the burden of proof

为什么 800VDC 提高了举证责任

Traditional AC datacenters have a grid-facing vocabulary that planners can model: UPS transfer thresholds, ATS timing, motor-load behavior, generator controls, and composite load models like CMPLDW. None of that captures an 800VDC facility, where the response to a grid voltage dip depends on the SST control algorithm (grid-following vs grid-forming), BESS state of charge, instantaneous GPU load profile, and interactions between multiple parallel SSTs.

传统的交流数据中心拥有一套规划人员可以建模的面向电网的术语体系：UPS 切换阈值、ATS 切换时间、电机负载行为、发电机控制以及像 CMPLDW 这样的复合负载模型。然而，这些都无法准确描述 800VDC 设施，因为在这种设施中，对电网电压跌落的响应取决于 SST 控制算法（跟网型 vs 构网型）、BESS 荷电状态、瞬时 GPU 负载特性，以及多个并联 SST 之间的相互作用。

800VDC also collapses layers of the power stack. In an AC facility, the utility studies the interconnection and aggregate load while the operator engineers UPS, switchgear, and rack distribution independently. In an SST-based 800VDC facility, the same converter controls determine DC bus stability, fault ride-through, current limiting, harmonic injection, and post-fault load recovery. Interconnection therefore becomes an engineering product that requires EPC capabilities bridging power electronics design, grid-level dynamic modeling, and regulatory engagement. This is leading to new entrants like Aran Industries, building AI-native EPCs to deliver PE-stampable

800VDC engineering packages.

800VDC 还简化了电力堆栈的层级。在交流设施中，公用事业部门研究互连和总负荷，而运营商则独立设计 UPS、开关柜和机架配电。在基于 SST 的 800VDC 设施中，相同的转换器控制决定了直流母线的稳定性、故障穿越、限流、谐波注入以及故障后的负荷恢复。因此，互连变成了一种工程产品，需要具备跨越电力电子设计、电网级动态建模和监管参与的 EPC 能力。这正催生出像 Aran Industries 这样的新进入者，他们正在构建 AI 原生 EPC，以交付可获得专业工程师（PE）盖章认证的 800VDC 工程方案包。

Understanding The No-So-Basics: The Physics Behind 800VDC

深入了解非基础知识：800VDC 背后的物理学

Why Going Super Dense Makes Low-Voltage Distribution Break: Heat and Weight

为什么超高密度会导致低压配电失效：散热与重量

At a fixed power level, raising voltage from 54V to 800V cuts current by ~15× and resistive losses by ~220×. That is what makes 800VDC a step-change in copper mass, thermal load, and distribution cost.

在功率水平固定的情况下，将电压从 54V 提高到 800V 可使电流降低约 15 倍，电阻损耗降低约 220 倍。正是这一点，使 800VDC 在铜重、热负荷和配电成本方面实现了阶跃式的变革。

Start with the power equation:

从功率方程开始：

$$P = VxI$$

For a fixed rack power P, raising V reduces I linearly. Lower current means smaller conductors, less copper mass, and easier routing.

对于固定的机架功率 P，提高电压 V 会线性降低电流 I。更低的电流意味着更细的导线、更少的铜质量以及更易于布线。

Ohm’s Law gives the voltage drop across a conductor of resistance R:

欧姆定律给出了电阻为 R 的导体上的电压降：

$$V = IxR$$

That drop is the energy dissipated as heat in the conductor. Substituting into the power equation yields the resistive loss equation:

该压降即为在导体中以热量形式耗散的能量。代入功率方程可得出电阻损耗方程：

$$P_loss = I^2 \times R$$

Current appears squared, so the voltage-to-loss relationship is quadratic, not linear. This is the equation that makes 800VDC inevitable.

电流以平方形式出现，因此电压与损耗的关系是二次方关系，而非线性关系。正是这个方程式使得 800VDC 的趋势变得不可阻挡。

A working example at 600kW rack power (Kyber-class, Vera Rubin Ultra NVL576):

一个 600kW 机架功率（Kyber 级，Vera Rubin Ultra NVL576）的工作示例：

At ~54 VDC (today’s standard):

在约 54 VDC（当今标准）下：

$$I = \frac{P}{V} = \frac{600{,}000\ \text{W}}{54\ \text{V}} \approx 11{,}111\ \text{A}$$

At 800 VDC: 在 800 VDC 下：

$$I = \frac{P}{V} = \frac{600{,}000\ \text{W}}{800\ \text{V}} = 750\ \text{A}$$

That’s a 14.8× reduction in current. Now apply the loss equation. For the same conductor resistance R, the I^2 ratio implies:

电流减少了 14.8 倍。现在应用损耗方程。对于相同的导体电阻 R，I^2 比率意味着：

$$\frac{I_{54V}^2}{I_{800V}^2} = \frac{(11{,}111)^2}{(750)^2} \approx 219\times$$

Resistive heating at 54 V is roughly **219 times higher** than at 800 V for the same conductor. In the more commonly cited comparison using 48 V:

在相同导体下，54 V 时的电阻损耗约为 800 V 时的 219 倍。在更常被引用的 48 V 对比中：

$$I_{48V} = \frac{600{,}000}{48} = 12{,}500 \text{ A} \quad \Rightarrow \quad \frac{(12{,}500)^2}{(750)^2} \approx 278\times$$

In practice, operators do not keep the same conductor and pocket all 219-278× of loss reduction. They shrink the copper, trading loss headroom for reductions in weight, cost, and routing space. Even after right-sizing for 800V, the efficiency gain remains transformative.

在实践中，运营商并不会保留原有的导体并坐享 219-278 倍的损耗降低。相反，他们会缩小铜线尺寸，将损耗余量转化为重量、成本和布线空间的缩减。即便在针对 800V 进行尺寸优化后，其效率提升依然具有变革性。

800VDC vs. ±400VDC: The Topology In Question

800VDC 对比 ±400VDC：备受关注的拓扑结构

'800 VDC' may refer to two distinct electrical configurations, and the distinction matters for deployment strategy, safety engineering, and downstream semiconductor selection. '800 VDC' may refer either to a single-ended 800V bus or a bipolar ±400V bus (800V pole-to-pole):

"800 VDC" 可能指代两种不同的电气配置，这种区别对于部署策略、安全工程以及下游半导体的选择至关重要。"800 VDC" 既可以指单端 800V 总线，也可以指双极性 ±400V 总线（极间电压为 800V）：

Single-ended 800V 单端 800V

In a single-ended 800VDC architecture, the bus is a single 800V rail referenced to return, plus protective earth. At 1MW, the bus carries 1,250A. Lower current means smaller conductors, smaller connectors, and lower I^2R losses throughout the distribution path. The bus structure is also simpler to implement because it does not rely on maintaining symmetry between two rails. Power stages can be designed directly around the full bus voltage using standard high-voltage devices and conventional converter topologies. No midpoint to sense, regulate, or control.

在单端 800VDC 架构中，母线是参考回路的单条 800V 轨道，外加保护地线。在 1MW 功率下，母线电流为 1,250A。较低的电流意味着在整个配电路径中可以使用更小的导线、更小的连接器，并降低 I^2R 损耗。母线结构也更易于实现，因为它不依赖于维持两条轨道之间的对称性。功率级可以直接围绕全母线电压进行设计，使用标准高压器件和常规转换器拓扑。无需对中点进行感测、调节或控制。

Bipolar ±400V 双极性 ±400V

The alternative splits that 800V into two symmetric 400V rails around a grounded midpoint: three power conductors (+400V, midpoint, -400V) plus protective earth. The load still sees 800V across its input, but each rail sits only 400V from ground. The central argument here is not electrical, but economic. 400V power electronics are mature because the EV industry built at scale on 400V platforms. Google stated at OCP EMEA 2025 that selecting 400VDC “allows us to leverage the supply chain established by electric vehicles”. The OCP Diablo 400 specification considers a disaggregated power rack converting 3-phase AC to ±400VDC at 100kW to 1MW per rack. The spec also includes 800VDC monopolar as a design option, leaving the door open.

另一种方案是将 800V 分压为围绕接地中点的两个对称 400V 轨道：三根电源导线（+400V、中点、-400V）外加保护地线。负载在其输入端看到的仍是 800V，但每个轨道距离地电位仅 400V。这里的核心论点并非电气层面，而是经济层面。由于电动汽车（EV）行业在 400V 平台上进行了大规模建设，400V 电力电子设备已趋于成熟。Google 在 OCP EMEA 2025 上表示，选择 400VDC “使我们能够利用电动汽车建立的供应链”。OCP Diablo 400 规范考虑采用解耦式电源机架，将三相交流电转换为 ±400VDC，每个机架功率在 100kW 到 1MW 之间。该规范还将 800VDC 单极方案列为设计选项，为此保留了可能性。

There’s also a tradeoff. That third conductor must be routed, terminated, and protected at every point in the power path. Across thousands of racks, it adds meaningful copper, connector hardware, and installation labor, and it complicates hot-swap connector design where the midpoint must make and break contact in a controlled sequence to avoid transient voltage spikes.

这也涉及一种权衡。第三根导体必须在电源路径的每个点进行布线、端接和保护。在数以千计的机架中，这增加了大量的铜缆、连接器硬件和安装人工，并使热插拔连接器设计变得复杂，因为中点必须按受控顺序进行接触和断开，以避免瞬态电压尖峰。

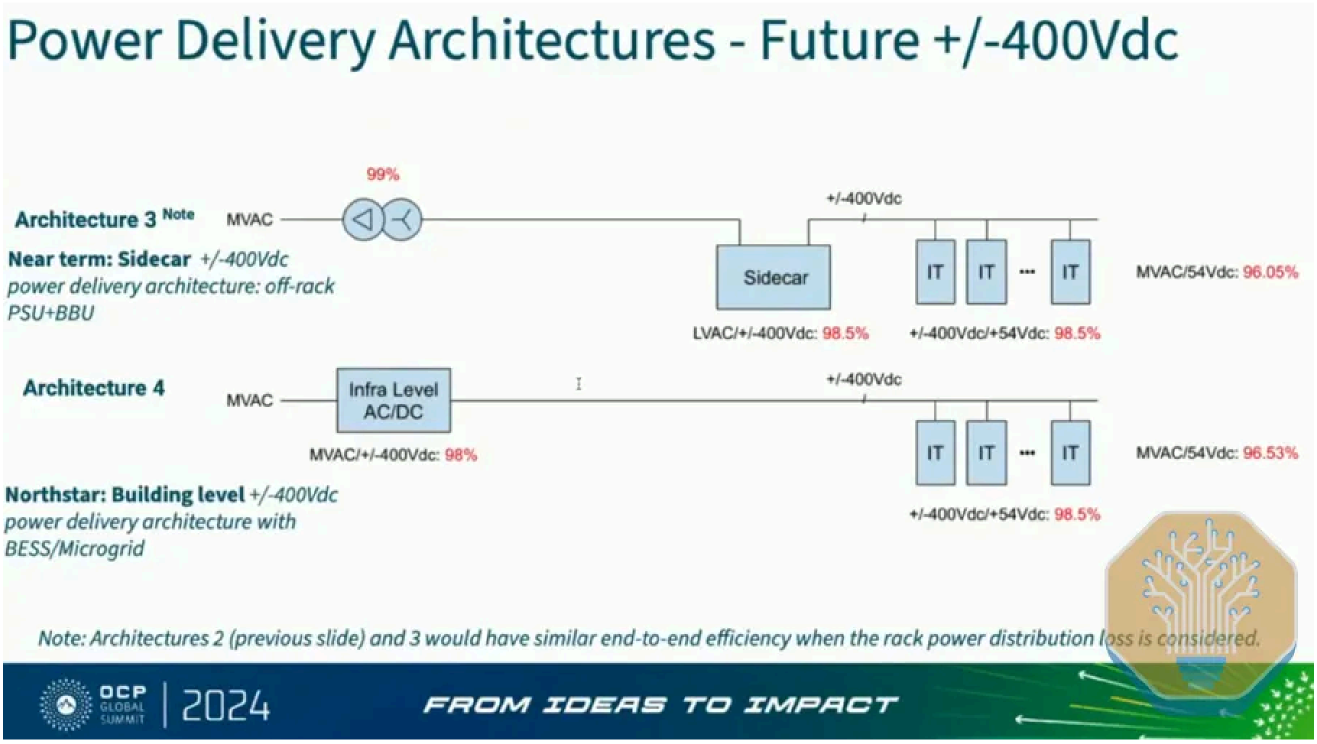

Source: OCP OCPp

Behind paywall we will now discuss the main winners and losers of the 800VDC revolution, and who is better positioned to be benefited from the transition.

在付费墙后，我们将讨论 800VDC 变革的主要赢家和输家，以及谁能更好地从这一转型中获益。

Equipment Supplier Impact

设备供应商影响

White Space VS Grey Space Vendors: Who Wins?

白区与灰区供应商：谁将胜出？

On the shift to HVDC specifically, we like vendors with more white space exposure over gray space. The shift from standard AC/DC to HVDC is mainly a share-shift story, not a meaningful increase in total electrical content per MW, which remains roughly flat at $3.7M–$4.0M/MW across most architectures, aside from the HVDC white space retrofit case at $4.8M/MW. Our Core Research subscribers will receive a more in-depth analysis identifying the specific winners and losers.

在向高压直流（HVDC）转型的过程中，我们更看好在“白区”（IT 设备区）而非“灰区”（动力设施区）拥有更多业务敞口的供应商。从标准交流/直流（AC/DC）向高压直流的转变主要是一个市场份额转移的过程，而非每兆瓦（MW）总电气内容的显著增加。在大多数架构中，每兆瓦的电气成本基本维持在 370 万至 400 万美元之间，唯有高压直流白区改造案例的成本较高，为 480 万美元/兆瓦。我们的核心研究订阅用户将收到更深入的分析报告，以识别具体的受益者和受损者。

White space vendors are better positioned for three reasons. First, the content uplift is large and immediate: white space power electronics vendors shift from selling power shelfs to now entire HVDC power racks, creating a major revenue uplifts. Second, white space vendors are already shipping into 2025–2026 deployments, while gray space vendors are mostly pointing to 2028. Third, white space demand is tied to accelerator shipments and a 3–4 year rack refresh cycle, making it structurally more attractive than gray space spend, which depends on 10–15 year facility cycles and lumpy buildouts.

白区（White space）供应商处于更有利地位的原因有三点。首先，价值含量的提升巨大且立竿见影：白区电力电子供应商从销售电源机架（power shelfs）转型为销售整个高压直流（HVDC）电源机柜，从而带来了大幅的营收增长。其次，白区供应商的产品已经投入 2025-2026 年的部署，而灰区（gray space）供应商大多指向 2028 年。第三，白区需求与加速器出货量以及 3-4 年的机柜更新周期挂钩，这使其在结构上比灰区支出更具吸引力，因为后者依赖于 10-15 年的设施周期和不规则的建设进度。

Gray space incumbents face a less certain outlook. HVDC architectures remove centralized UPS, low-voltage switchgear, and low-voltage transformers, while replacement content like solid-state transformers is still early and lacks a clear winner. White space vendors are also competing for that future content, so gray space incumbents cannot assume they will retain it.

灰色空间领域的现有厂商面临着较不确定的前景。高压直流（HVDC）架构取消了集中式 UPS、低压配电柜和低压变压器，而固态变压器等替代产品仍处于早期阶段，且尚未出现明显的领先者。白色空间供应商也在争夺未来的这些产品份额，因此灰色空间的现有厂商不能理所当然地认为自己能留住这些业务。

That creates risk for names like Legrand and Hammond Power, whose datacenter narratives have benefited from gray space exposure. Vertiv is the standout gray space vendor with proven white space execution at Meta, while Eaton and Schneider have Nvidia reference architecture positions in power racks, though commercial volumes still appear limited.

这对 Legrand 和 Hammond Power 等公司构成了风险，因为这些公司的中心数据中心叙事受益于灰色空间（gray space）的业务敞口。Vertiv 是杰出的灰色空间供应商，在 Meta 拥有成熟的白色空间（white space）执行力；而 Eaton 和 Schneider 虽然在电源机架方面拥有 Nvidia 参考架构地位，但商业销量目前似乎仍然有限。

Delta Electronics (2308 TT): The Structural Winner

台达电子 (2308 TT)：结构性赢家

Delta's core advantage is end-to-end integration: it can deliver a complete 800 V solution across the power shelf, BBU, PCS (including short-duration energy buffering such as supercapacitors), and liquid-cooling systems as one validated package. As rack power scales into the hundreds of kW and toward MW-class deployments, procurement becomes more engineering-led and reliability-driven. A single-vendor, system-level delivery cuts the integration and qualification burden, shortens deployment timelines, and reduces finger-pointing when something fails at ~600 kW per rack. Power shelf ASPs jump from roughly $40k per rack in a standard AC-DC configuration to roughly $400k for an HVDC power rack, a 10x increase driven by scope expansion. That's why Delta's per-rack ASP has meaningful upside in the 800 V transition versus vendors that sell isolated components.

台达的核心优势在于端到端集成：它能够提供涵盖电源机架、BBU、PCS（包括超级电容器等短时能量缓冲装置）以及液冷系统的完整 800 V 解决方案，并作为一个经过验证的整体包交付。随着机架功率扩展至数百 kW 并向兆瓦级部署迈进，采购过程变得更加以工程为导向和以可靠性为驱动。单供应商、系统级的交付降低了集成和认证负担，缩短了部署周期，并减少了在单机架功率约 600 kW 发生故障时责任推诿的情况。电源机架的平均售价（ASP）从标准 AC-DC 配置的每机架约 4 万美元跃升至高压直流（HVDC）电源机架的约 40 万美元，这种 10 倍的增长是由业务范围的扩大所驱动的。这就是为什么在 800 V 转型中，相比销售孤立组件的供应商，台达的单机架 ASP 具有显著的上升空间。

Delta has a high share of AI server rack PSUs and sidecar CDU cooling systems for GB200. Delta's moat is vertical integration across the full power chain: "grid-to-chip." No other player spans this full stack, and Delta is the only player that can credibly supply every major component from the utility interconnect to the VR on the GPU board.

台达在 AI 服务器机架电源（PSU）以及适用于 GB200 的边置式 CDU 冷却系统领域拥有极高的市场份额。台达的护城河在于贯穿整个电源链的垂直整合能力：“从电网到芯片”。目前没有其他厂商能跨越这一完整的技术栈，台达是唯一一家能够可靠地提供从公用事业互连到 GPU 板卡上电压调节器（VR）等所有主要组件的厂商。

Moreover, Delta also supplies CDU systems to MSFT, META, and ORCL through ODM channels (Foxconn, Wiwynn, Wistron). As 800 V HVDC racks shift power hardware out of the IT rack and into in-row power racks, thermal profiles and heat distribution change. Delta can co-optimize power and thermal design in ways pure power or pure cooling vendors cannot.

此外，台达电还通过 ODM 渠道（富士康、纬颖、纬创）向 MSFT、META 和 ORCL 供应 CDU 系统。随着 800 V HVDC 机架将电源硬件从 IT 机架移至列间功率机架，热剖面和热量分布也随之改变。台达电能够以纯电源或纯散热供应商无法实现的方式，对电力和热设计进行协同优化。

We expect most initial power-rack designs to be AC-DC, including early deployments tied to CSPs and Nvidia platforms; AC-DC content is typically higher value than downstream DC-DC content. We expect Delta to be the main supplier for Nvidia, Meta, and Google as early adopters of 800VDC and power racks are set for volume shipments by end of 2026. We see meaningful content-upside for Delta as architectures shift from power shelves to HVDC power racks. Moreover, in a scenario which a dedicated Kyber 800V-50V sidecar is eliminated, Delta could dominate 90% of this market because of its strong existing expertise in in-rack PSUs.

我们预计大多数初始电源机架设计将采用 AC-DC 方案，包括与云服务提供商（CSP）和 Nvidia 平台挂钩的早期部署；AC-DC 内容的价值通常高于下游的 DC-DC 内容。我们预计台达（Delta）将成为 Nvidia、Meta 和 Google 的主要供应商，因为这些 800VDC 和电源机架的早期采用者预计将在 2026 年底前实现批量出货。随着架构从电源层（power shelves）转向高压直流（HVDC）电源机架，我们认为台达将迎来显著的价值增长空间。此外，在取消专用 Kyber 800V-50V 侧柜（sidecar）的情况下，凭借其在机架内电源单元（PSU）领域的深厚专业积累，台达有望占据该市场 90% 的份额。

The bear case for Delta is its limited presence in gray space. UPS and PDU have historically been dominated by Western incumbents such as Vertiv, Schneider Electric, Eaton, and ABB, supported by entrenched relationships and large installed bases. Delta's UPS share in the Americas is minimal. In 800 VDC architectures, UPS and PDU may still sit upstream for protection, switching, and backup power. The bigger incremental gray-space opportunity is likely SST, but that is further out. Delta has cited SST as a future direction and appears to be working toward early readiness.

看空台达（Delta）的理由在于其在灰色空间（gray space）的市场占有率有限。UPS 和 PDU 市场历来由 Vertiv、施耐德电气（Schneider Electric）、伊顿（Eaton）和 ABB 等西方老牌企业主导，这些企业拥有深厚的关系网络和庞大的装机基础。台达在美洲的 UPS 市场份额微乎其微。在 800 VDC 架构中，UPS 和 PDU 仍可能位于上游，负责保护、切换和备用电源。灰色空间更大的增量机会可能是固态变压器（SST），但那还需要更长的时间。台达已将 SST 列为未来发展方向，并似乎正致力于早期准备工作。

Lite-On (2301 TT) 光宝科技 (2301 TT)

Lite-On is the no. 2 player in white space. Despite lower PSU share in NVDA GPU servers (mainly Oracle as a customer), Lite-On has supplied PSU and BBU share in AWS Trainium servers. We expect this to be only an at-most HSD% revenue contributor in the near term, even assuming strong Trainium 3 ramps. We estimate Lite-On's power-supply content spans both PSUs and BBUs. While Lite-On has been vocal about its 800 V portfolio, power shelves, BBUs, PCS, and modular PDU architectures, adoption to date appears more limited. ON Semiconductor (Q4 FY2025) confirmed it has secured rack-level BBU and PSU designs with both Delta and Lite-On, indicating Lite-On retains a seat at the table on next-gen power architectures.

光宝科技（Lite-On）是白牌市场的第二大厂商。尽管在 NVIDIA GPU 服务器中的电源供应器（PSU）份额较低（主要客户为 Oracle），但光宝已在 AWS Trainium 服务器中供应 PSU 和电池备份单元（BBU）。我们预计，即使假设 Trainium 3 强劲放量，短期内这部分业务对营收的贡献最高也仅为高个位数百分比（HSD%）。据我们估计，光宝的电源产品涵盖了 PSU 和 BBU。虽然光宝一直积极宣传其 800V 产品组合、电源机架（Power Shelves）、BBU、电源转换系统（PCS）以及模块化配电单元（PDU）架构，但目前的采用程度似乎较为有限。安森美（ON Semiconductor，2025 财年第四季度）确认已与台达电和光宝共同获得了机架级 BBU 和 PSU 的设计中标，这表明光宝在下一代电源架构中仍占有一席之地。

We read this as a function of transition complexity. As safety standards evolve and racks move to new architectures that require tight power-cooling co-design, customers tend to favor suppliers with proven system-level integration, established hyperscaler validation, and clear accountability for full rack-level delivery. Even with a broad lineup, vendors without a comparable end-to-end deployment track record often see slower uptake.

我们将此解读为转型复杂性的体现。随着安全标准的演进，以及机架向需要紧密电源-冷却协同设计的新架构迁移，客户往往更青睐那些拥有成熟系统级集成经验、已通过超大规模云服务商（Hyperscaler）验证，并能对整个机架级交付承担明确责任的供应商。即使拥有广泛的产品线，如果供应商缺乏同等的端到端部署实绩，其产品的采用速度通常也会较慢。

Lite-On does have meaningful vertical integration across white-space components. Like Delta, it manufactures many power rack components in-house, including PDUs, power controls, chassis, and cabinets. This can support higher margins on integrated systems, and management has pointed to ~30% GM on AI server power versus ~22-24% at the company level. The 800 V power rack also offers substantial ASP uplift versus standalone server power modules.

光宝科技（Lite-On）在数据中心白区组件方面确实拥有显著的垂直整合能力。与台达电（Delta）类似，它自主生产许多电源机架组件，包括电源分配单元（PDU）、电源控制器、机箱和机柜。这有助于提高集成系统的利润率，管理层指出，AI 服务器电源的毛利率约为 30%，而公司整体毛利率约为 22-24%。此外，与独立的服务器电源模块相比，800V 电源机架还提供了大幅的平均售价（ASP）提升。

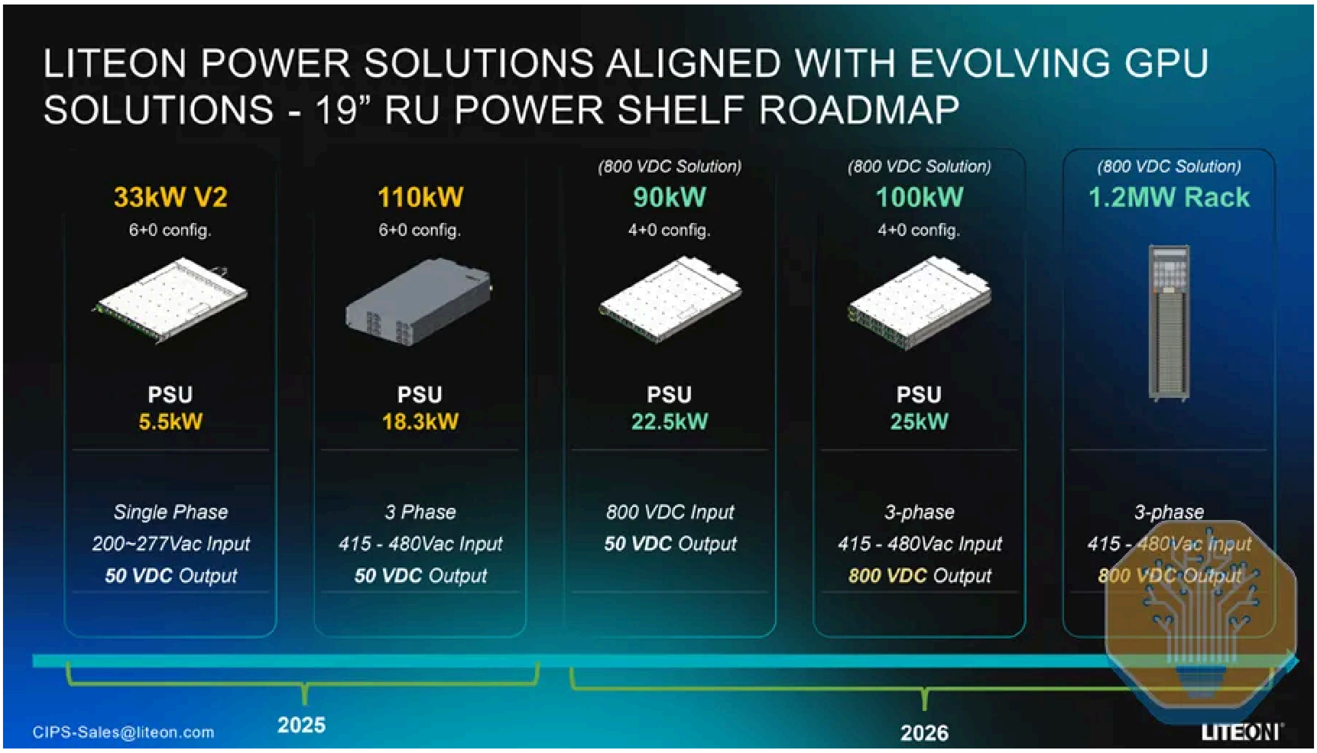

Source: LiteOn 光宝科技

Lite-On is expanding capacity across TW (Kaohsiung), VN (including a $200M capital injection), and the U.S. (TX). However, management is taking a conservative approach to 2026 capacity adds at ~30%, below the ~50% BBU demand growth rate the market expects. That suggests a preference for protecting margins over aggressively chasing incremental share, in contrast to Delta's more aggressive expansion posture.

光宝科技（Lite-On）正在台湾（高雄）、越南（包括 2 亿美元的注资）和美国（德克萨斯州）扩大产能。然而，管理层对 2026 年的产能增幅采取了约 30% 的保守态度，低于市场预期的约 50% 的 BBU（电池备份单元）需求增长率。这表明，与台达电（Delta）更具侵略性的扩张姿态相比，光宝更倾向于保护利润率，而非激进地追求增量份额。

Customer concentration and liquid cooling are the key watch items. Delta has been aggressively taking Lite-On's BBU share at AWS. Lite-On's CDU efforts remain earlier-stage, with products in trial or early production and commercial shipments starting around Q1 2026. Amazon appears to be consolidating CDU spend toward Delta, and the risk is that adjacent power content (PSU, BBU) follows over time, further narrowing Lite-On's share at its most important customer. Delta's stronger design capability and earlier scale in liquid cooling, including meaningful sidecar CDU share for GB200, suggests Lite-On remains behind on this vector and could see a slower ramp in high-margin liquid-cooling content.

客户集中度和液冷技术是关键的观察指标。台达电（Delta）一直在积极抢占光宝科（Lite-On）在 AWS 的 BBU 份额。光宝科在 CDU 方面的投入仍处于早期阶段，产品尚处于试产或初期生产阶段，预计商业出货将于 2026 年第一季度左右开始。亚马逊似乎正在将 CDU 支出向台达电集中，风险在于相关的电源组件（PSU、BBU）随时间推移也可能紧随其后，从而进一步缩小光宝科在其最重要客户中的份额。台达电在液冷领域拥有更强的设计能力和更早的规模化优势，包括在 GB200 中占据显著的 Sidecar CDU 份额，这表明光宝科在这一领域仍处于落后地位，在高利润液冷组件方面的增长速度可能会较慢。

Vertiv (VRT US): The Gray-Space Leader Pushing Into White Space

Vertiv (VRT US)：进军白区市场的灰区领导者

We continue to believe Vertiv should have record order volume (more details in our core research and datacenter models). Vertiv should continue to grow orders meaningfully. We think its customer mix leaves it over-indexed to colo leases and under-represented in hyperscaler self-build. We assume VRT has ~50% exposure to large DCs (hyperscale + colo), ~$3M of VRT content per colo MW, ~$1M per self-build MW, and ~20% share of colo DCs. Vertiv is the top supplier of UPS for most major DC operators, and both AWS and MSFT (along with traditional colo providers) rely heavily on Vertiv infrastructure. That installed base should drive pull-through as operators upgrade power architectures to support higher rack densities.

我们继续认为 Vertiv 的订单量将创下历史新高（更多详情请参阅我们的核心研究和数据中心模型）。Vertiv 的订单量应会继续保持显著增长。我们认为其客户组合使其在托管租赁（colo leases）领域的权重过高，而在超大规模运营商自建（hyperscaler self-build）领域的代表性不足。我们假设 Vertiv 在大型数据中心（超大规模 + 托管）的敞口约为 50%，每兆瓦托管数据中心的 Vertiv 价值量约为 300 万美元，每兆瓦自建数据中心约为 100 万美元，并在托管数据中心市场占有约 20% 的份额。Vertiv 是大多数主要数据中心运营商的首选 UPS 供应商，AWS 和微软（以及传统的托管服务提供商）都高度依赖 Vertiv 的基础设施。随着运营商升级电力架构以支持更高的机架密度，这一庞大的装机基础将推动后续产品的拉动效应。

On 800 V hyperscaler engagement, Vertiv is actively working with META, GOOGL, and MSFT, and vendor selection is still evolving. Vertiv won the META 800 V HVDC power rack program alongside Delta, despite historically low content at META and GOOGL, where BBU-centric architectures have often bypassed central UPS. This win is a meaningful step-up in Vertiv content, from near-zero in whitespace to roughly ~$1M/MW. Vertiv's existing UPS business is not cannibalized in the near term too with move to HVDC. Instead, in the whitespace retrofit HVDC use case outlined above, Vertiv gets additive content: the legacy UPS stays in place (~$1M/MW grey space), and the new power rack (~$1M/MW white space) stacks on top which is a near-term tailwind.

在 800V 超大规模云服务商的合作方面，Vertiv 正积极与 META、GOOGL 和 MSFT 开展合作，供应商选择仍在不断演变中。尽管 Vertiv 在 META 和 GOOGL 的历史参与度较低（这两家公司的架构以 BBU 为中心，通常绕过中央 UPS），但 Vertiv 仍与台达（Delta）一同赢得了 META 的 800V 高压直流（HVDC）电源机架项目。这次中标标志着 Vertiv 价值量的显著提升，从机房白区（whitespace）的近乎为零增长到约 100 万美元/兆瓦。此外，向 HVDC 的转型在短期内并不会蚕食 Vertiv 现有的 UPS 业务。相反，在上述机房白区改造的 HVDC 应用场景中，Vertiv 获得了增量价值：原有的 UPS 保持不变（约 100 万美元/兆瓦的灰区价值），而新的电源机架（约 100 万美元/兆瓦的白区价值）在此基础上叠加，形成了短期的利好因素。

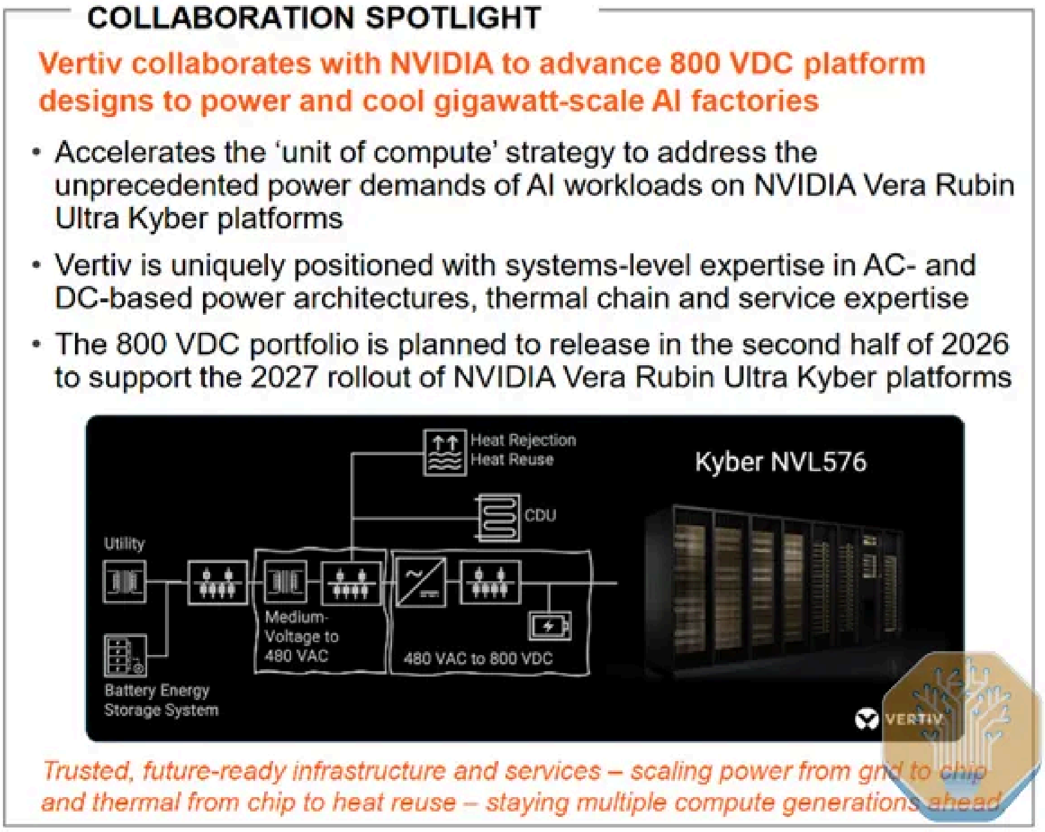

Source: Vertiv 维谛技术

Vertiv also benefits from leading DC thermal management, including liquid cooling. In CDUs, Vertiv is differentiated in that it primarily sells branded CDUs rather than custom designs. Vertiv is mostly working with MSFT and AWS. Near-term contribution remains low, but the longer-term goal is to scale materially. The broader thermal portfolio spans rear-door heat exchangers, direct-to-chip cooling, and the CoolPhase Flex hybrid system. 800 V HVDC pushes more conversion hardware into dedicated power racks which increases the need for coordinated "power + cooling" design at the row level, Vertiv is positioned to sell integrated solutions rather than point products.

Vertiv 还受益于领先的直流热管理技术，包括液冷技术。在 CDU 领域，Vertiv 的优势在于其主要销售品牌 CDU 而非定制设计。Vertiv 主要与 MSFT 和 AWS 合作。虽然短期贡献仍然较低，但长期目标是实现实质性的规模扩张。其更广泛的散热产品组合涵盖了背门热交换器、芯片直接液冷技术以及 CoolPhase Flex 混合系统。800V 高压直流（HVDC）将更多的转换硬件推向专用电源机架，这增加了对行级"电源+散热"协同设计的需求，Vertiv 的定位是销售集成解决方案而非单一产品。

Vertiv's services business (>20% of revenue) is a structural advantage as HVDC architectures scale. 800 V HVDC racks are more complex and carry higher safety and uptime risk, which raises the value of commissioning, maintenance, and lifecycle support. This is an area where Delta, as more of a component/system supplier, tends to capture less of the total wallet over time. It supports Vertiv's ability to defend premium positioning as architectures become more operationally demanding. We think neo-cloud customers typically have servicing needs, while hyperscalers tend to service in-house due to more customized architectures.

Vertiv 的服务业务（占营收 20% 以上）是其在 HVDC 架构规模化过程中的结构性优势。800V HVDC 机架更为复杂，且具有更高的安全和运行时间风险，这提升了调试、维护和生命周期支持的价值。相比之下，台达电（Delta）更多作为组件/系统供应商，随着时间的推移，在整体市场份额中获取的价值往往较低。随着架构在操作要求上变得更加苛刻，这支撑了 Vertiv 捍卫其高端定位的能力。我们认为，新型云客户通常有服务需求，而超大规模云服务商由于架构更加定制化，往往倾向于自行维护。

The main limitation is that Vertiv does not participate in server-side white-space power electronics. It does not make PSUs, BBUs, or DC-DC converter modules for the IT rack, leaving it with effectively zero share of white-space power content in current GB200 racks. If the industry migrates toward power racks that blend gray-space and white-space functions, Vertiv may need to build, partner, or acquire rack-level conversion capabilities to avoid ceding the new profit pool. Vertiv's gray-space stronghold also faces competition from below as new entrants push into power infrastructure and hyperscalers increasingly self-design systems, particularly at META and GOOGL, where custom BBU architecture has historically reduced the role of UPS-centric suppliers.

主要限制在于 Vertiv 并不参与服务器端的白区（white-space）电力电子业务。它不生产用于 IT 机架的电源单元（PSU）、电池备份单元（BBU）或直流-直流（DC-DC）转换模块，这导致其在当前的 GB200 机架中实际上没有白区电力内容的份额。如果行业向融合了灰区（gray-space）和白区功能的电力机架迁移，Vertiv 可能需要通过自建、合作或收购来获取机架级转换能力，以避免失去这一新的利润池。Vertiv 的灰区据点也面临着来自下方的竞争，因为新入局者正涌入电力基础设施领域，且超大规模云服务商（特别是 META 和 GOOGL）越来越多地采用自主设计系统，其定制化的 BBU 架构在历史上已经削弱了以 UPS 为中心的供应商的角色。

Other Western Players 其他西方参与者

A handful of Western "total-solution" vendors have been vocal on 800 VDC, most notably Schneider Electric and Vertiv. Both have published reference architectures and design guides that help the ecosystem standardize layouts and component choices. We still think Delta has the clearest evidence of early, volume exposure, largely because it is more directly embedded in hyperscaler supply chains versus selling packaged end-to-end solutions into colo and neo-cloud. In practice, many reference designs are sold to neo-cloud and colo developers, while hyperscalers often internalize the architecture and multi-source the blocks.

少数几家西方"整体解决方案"供应商一直在积极倡导 800 VDC，其中最引人注目的是 Schneider Electric 和 Vertiv。两家公司都发布了参考架构和设计指南，旨在帮助生态系统实现布局和组件选择的标准化。我们仍然认为 Delta 在早期规模化应用方面拥有最明确的证据，这主要是因为与向托管数据中心（colo）和新云（neo-cloud）销售成套端到端解决方案相比，Delta 更直接地嵌入了超大规模云服务商（hyperscaler）的供应链。在实践中，许多参考设计被出售给新云和托管数据中心开发商，而超大规模云服务商通常会将架构内部化，并对各个模块进行多方采购。

As a result, we're somewhat neutral on the large Electrical Equipment vendors winning content in 800VDC. They sit closer to the assembly level, and near-term 800 VDC demand is still concentrated in a narrower set of early deployments. We don't expect a broad-based step-function uplift for Western integrators in the immediate term.

因此，对于大型电气设备供应商能否在 800VDC 领域赢得市场份额，我们持中立态度。这些厂商更接近组装层面，而短期内 800VDC 的需求仍集中在少数早期部署项目中。我们预计西方集成商在短期内不会出现广泛的阶梯式增长。

Schneider Electric (SU FP)

施耐德电气 (Schneider Electric, SU FP)

Schneider Electric looks structurally behind Delta and Vertiv in the 800V HVDC race. They did showcase their 800VDC sidecar during OCP 2025 capable of up to 1.2MW per rack. Management commentary on shipments has been vague but did note that sidecar will ship "well before" Rubin Ultra's 2027 timeline. We think this is largely dedicated towards the Oberon platform rather than Kyber.

在 800V 高压直流（HVDC）领域的竞争中，施耐德电气（Schneider Electric）在结构上似乎落后于台达（Delta）和维谛（Vertiv）。尽管他们在 2025 年 OCP 大会上展示了单机柜功率高达 1.2MW 的 800VDC 侧置电源柜（sidecar），但管理层对出货情况的表述较为模糊，仅提到该产品将在 Rubin Ultra 2027 年的时间节点“之前很久”出货。我们认为，这主要针对的是 Oberon 平台，而非 Kyber。

Schneider’s guidance at its December 2025 Capital Markets Day for high-single-digit organic revenue CAGR in 2025–30. That outlook assumes they keep outgrowing their underlying market at 6–7% growth in aggregate, with data enters within this at 12–14% annually. Given the breadth of Schneider’s datacenter exposure, that assumption is not unreasonable, even if the 800V messaging remains unclear. Separately, Schneider is the global leader in medium-voltage switchgear and distribution. That position should remain secure through an 800V transition, as the upstream MV infrastructure (11–33 kV) feeding facilities become more complex.

施耐德在 2025 年 12 月的资本市场日上给出的指引显示，2025-30 年有机收入复合年增长率将达到高个位数。这一展望假设其整体业务将继续以 6-7% 的速度超越基础市场增长，其中数据中心业务的年增长率将达到 12-14%。鉴于施耐德在数据中心领域的广泛布局，即便 800V 的相关信息尚不明朗，这一假设也是合理的。此外，施耐德是中压开关柜和配电领域的全球领导者。随着为设施供电的上游中压基础设施（11-33 kV）变得愈发复杂，这一地位在向 800V 过渡的过程中应能保持稳固。

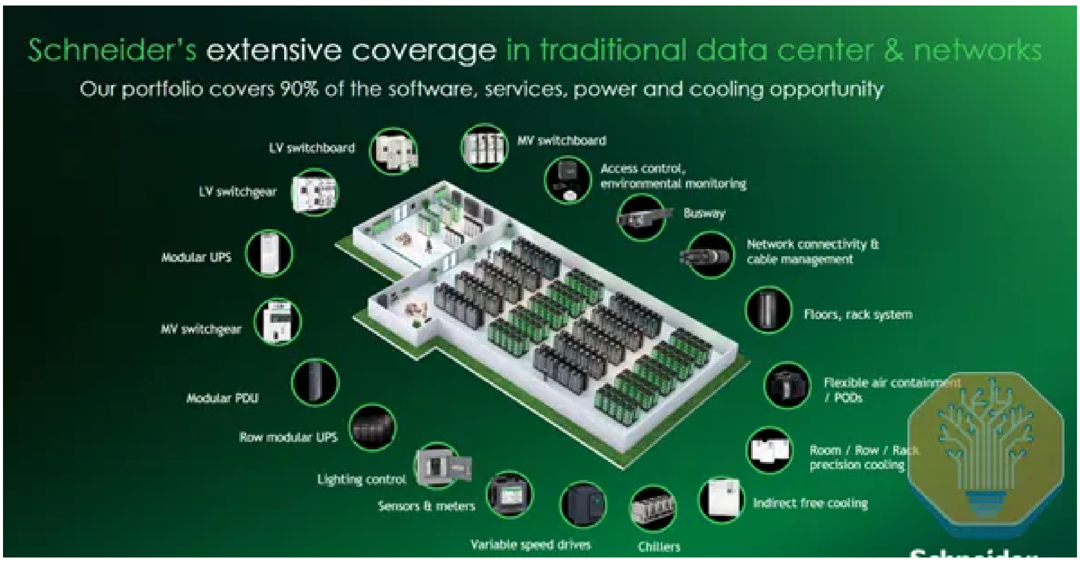

Source: Schneider Electric

施耐德电气

Eaton (ETN US) 伊顿 (Eaton, ETN US)

Eaton is one of the leading global suppliers across much of traditional datacenter grey-space power infrastructure, with a portfolio spanning UPS, switchgear, power distribution, static transfer switches, and prefabricated/modular systems. Eaton's broader datacenter push is increasingly framed around a "grid-to-chip" strategy, and the company has been expanding manufacturing capacity to support that demand. Recent investments include a new Virginia facility for static transfer switches, PDUs, and RPPs, plus broader U.S. transformer capacity expansions in Nacogdoches, Texas and Jonesville, South Carolina.

伊顿是全球领先的传统数据中心"灰区"(grey-space)电力基础设施供应商之一，其产品组合涵盖不间断电源（UPS）、开关设备、配电系统、静态转换开关以及预制化/模块化系统。伊顿在数据中心领域的更广泛布局正日益围绕"从电网到芯片"(grid-to-chip）战略展开，公司一直在扩大产能以支持这一需求。近期的投资包括在弗吉尼亚州新建一座生产静态转换开关、配电单元（PDU）和远程配电柜（RPP）的工厂，此外还扩大了德克萨斯州纳科多奇斯（Nacogdoches）和南卡罗来纳州琼斯维尔（Jonesville）的美国变压器产能。

On 800VDC, Eaton unveiled a reference architecture in October 2025 built in support of Nvidia's 800 VDC architecture for AI factories. The architecture features supercapacitor-based peak buffering using Eaton's XLHV modules (144V, 62.5F, up to 420 kW per module, designed for 20-year lifetime and millions of charge/discharge cycles), busbar integration with the Open Rack V3 (ORV3) standard, and DC connectors optimized for high-current applications.

在 800VDC 领域，伊顿（Eaton）于 2025 年 10 月发布了一套参考架构，旨在支持英伟达（Nvidia）针对 AI 工厂推出的 800 VDC 架构。该架构的特点包括：采用伊顿 XLHV 模块（144V，62.5F，单模块功率高达 420 kW，设计寿命 20 年并支持数百万次充放电循环）的超级电容器峰值缓冲技术、符合开放机架 V3（ORV3）标准的母线集成，以及针对高电流应用优化的直流连接器。

Eaton has little white-space server power business. No PSUs, no BBUs, no DC-DC converters. In current GB200 racks, Eaton captures zero white-space content. The

800VDC transition shifts spending from grey-space (where Eaton dominates) to white-space (where Eaton is absent). Even with the Nvidia reference design position and SST optionality, Eaton's near-term 800VDC monetization path is narrower than Delta's or Vertiv's because the new high-value content sits inside the rack, not outside it.

Eaton 在白区（white-space）服务器电源业务方面几乎没有布局。没有电源单元（PSU），没有电池备份单元（BBU），也没有直流-直流（DC-DC）转换器。在目前的 GB200 机架中，Eaton 捕获的白区价值量为零。800VDC 的转型将支出从灰区（grey-space，Eaton 的主导领域）转移到了白区（Eaton 缺席的领域）。即便拥有 Nvidia 参考设计的地位和固态变压器（SST）的可选性，Eaton 近期的 800VDC 变现路径仍比 Delta 或 Vertiv 更窄，因为新的高价值内容位于机架内部，而非外部。

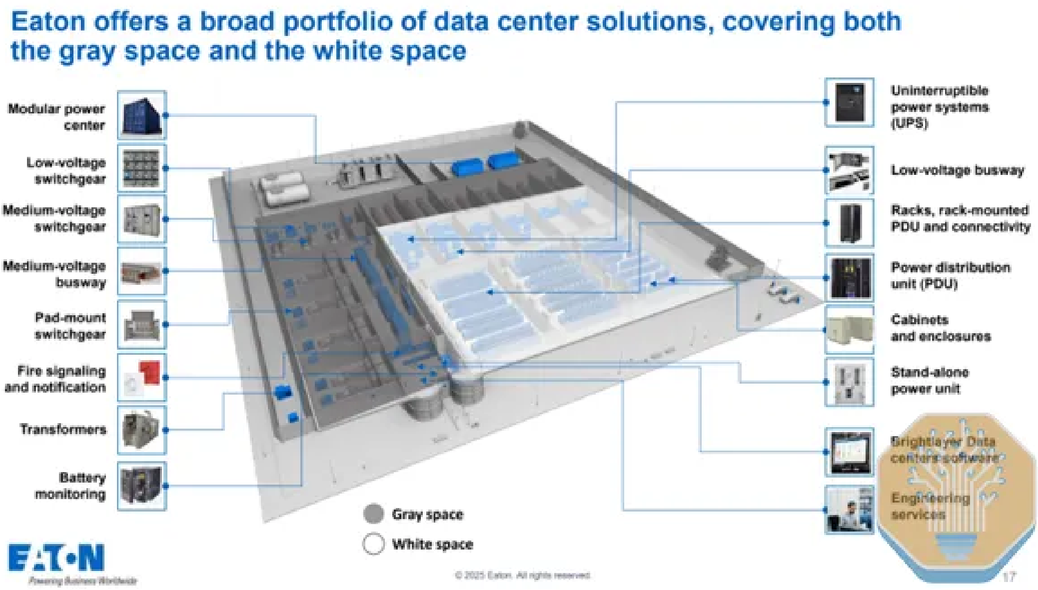

Source: Eaton 伊顿

Eaton's acquisition of Resilient Power Systems ($55M + $95M in earnouts) brings real solid-state transformer IP in-house. Resilient's SSTs combine voltage conversion, power conditioning, and grid support in a single ultra-compact device that connects directly to the medium-voltage distribution grid, potentially eliminating the conventional step-down transformer entirely. If SSTs become the Phase 3 standard for facility-level power delivery, Eaton is positioned ahead of peers with a head start on technology development and a global manufacturing base to scale production (including a $340M investment in a third transformer manufacturing facility in Jonesville, SC, production expected 2027). This is multi-year option value that the market may be underpricing.

Eaton 对 Resilient Power Systems 的收购（5500 万美元 + 9500 万美元的业绩奖励）为其内部带来了真正的固态变压器（SST）知识产权。Resilient 的 SST 将电压转换、功率调节和电网支持整合在一个超紧凑的设备中，可直接连接到中压配电网，从而有可能完全取代传统的降压变压器。如果 SST 成为设施级电力输送的第三阶段标准，Eaton 将凭借在技术开发上的领先优势以及全球制造基地（包括在南卡罗来纳州琼斯维尔投资 3.4 亿美元建设的第三座变压器制造工厂，预计 2027 年投产）而领先于同行。这是一种具有多年期权价值的资产，市场可能对其估值偏低。

ABB (ABB SS)

ABB's datacenter business sits within its Electrification segment, which management identified as the #1 growth driver at their November 2025 CMD. The product portfolio is narrower than Eaton or Vertiv: LV and MV switchgear, breakers, power distribution, MV UPS (their fastest-growing DC product line), gensets, and prefab eHouse solutions. Notably, ABB no longer sells transformers; that business now sits with Hitachi Energy. ABB called out its SAM is ~$2M per MW, roughly 7% of total DC investment, meaningfully below Eaton ($2.9M/MW, $3.4M post-Boyd) and Vertiv ($3-3.5M/MW). ABB's strategy is to be a "best-in-class supplier for discrete packages" rather than a one-stop-shop, which is the right framing for how hyperscalers actually buy, but it also means less content capture per facility.

ABB 的数据中心业务隶属于其电气化业务板块，管理层在 2025 年 11 月的资本市场日（CMD）上将其确定为第一大增长驱动力。其产品组合比 Eaton 或 Vertiv 更窄：包括低压和中压开关柜、断路器、配电系统、中压 UPS（其增长最快的数据中心产品线）、发电机组以及预制化 eHouse 解决方案。值得注意的是，ABB 不再销售变压器；该业务目前归属于日立能源（Hitachi Energy）。ABB 称其可服务市场（SAM）约为每兆瓦 200 万美元，约占数据中心总投资的 7%，明显低于 Eaton（每兆瓦 290 万美元，收购 Boyd 后为 340 万美元）和 Vertiv（每兆瓦 300-350 万美元）。ABB 的战略是成为“离散方案包的最佳供应商”，而非一站式商店，这符合超大规模云服务商（hyperscalers）的实际采购方式，但也意味着单个设施的价值量获取较低。

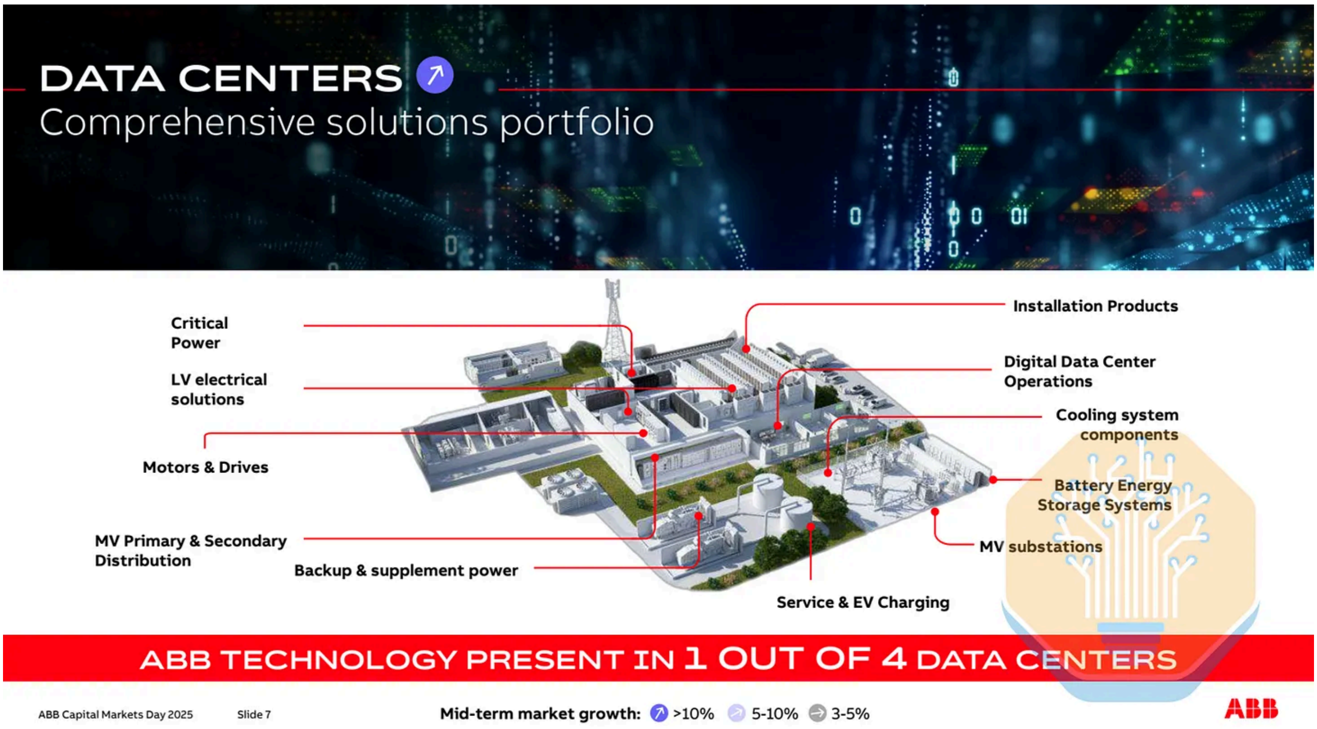

Source: ABB ABB

FY25 was "the best year in ABB's history" with record $4.6B free cash flow, 19% EBITA margin (all-time high), and comparable orders +32% in Q4. Electrification margin targets were raised to 22-26% at the CMD. Datacenter orders are arriving as $100M+ tickets with 12-24 month lead times (Applied Digital's 300 MW Polaris Forge 2 won in Q4 for 2026-2027 delivery). The company has turned away business to maintain delivery credibility. The most compelling near-term story is MV switchgear: lead times of 30-35 weeks (the binding electrical constraint in DC builds), and ABB running three-shift/24-hour operations. Backlog coverage is ~5 months, shorter than Vertiv (10-12) or Eaton (7-9), reflecting a mixed product portfolio with short-cycle LV alongside long-cycle MV.

25 财年是“ABB 历史上最好的一年”，自由现金流达到创纪录的 46 亿美元，息税前利润率（EBITA margin）达到 19%（创历史新高），第四季度可比订单增长 32%。在资本市场日（CMD）上，电气化业务的利润率目标上调至 22-26%。数据中心订单正以超过 1 亿美元的大单形式涌现，交付周期为 12-24 个月（Applied Digital 拥有 300 兆瓦容量的 Polaris Forge 2 项目在第四季度中标，计划于 2026-2027 年交付）。公司为了维持交付信誉，甚至推掉了一些业务。短期内最引人注目的增长点是中压（MV）开关柜：交付周期长达 30-35 周（这是数据中心建设中关键的电气约束条件），ABB 正处于三班倒/24 小时运行状态。积压订单覆盖时间约为 5 个月，短于 Vertiv（10-12 个月）或 Eaton（7-9 个月），这反映了其产品组合中既包含短周期的低压（LV）产品，也包含长周期的中压产品。

On 800VDC, ABB was clear that: “Current strong orders are for existing AC power architecture; new 800-volt DC architecture with Nvidia is a post-2028 opportunity.” They expect DC technology to comprise 40-50% of datacenter capacity by 2030, but nothing in current numbers reflects an 800VDC ramp. ABB announced an Nvidia collaboration on 800V DC architecture in October 2025, but this is positioning, not revenue, and the partnership appears vaguer than Vertiv’s direct Nvidia co-development on power racks.

关于 800VDC，ABB 明确表示：“目前的强劲订单主要来自现有的交流（AC）供电架构；与 Nvidia 合作的新型 800 伏直流（DC）架构是 2028 年以后的机会。”他们预计到 2030 年，直流技术将占数据中心容量的 40-50%，但目前的各项数据并未反映出 800VDC 的增长趋势。ABB 于 2025 年 10 月宣布了与 Nvidia 在 800V 直流架构上的合作，但这属于市场定位而非营收贡献，且该合作伙伴关系似乎比 Vertiv 与 Nvidia 在电源机架上的直接共同开发更为模糊。

Advanced Energy Industries Inc (AEIS US)

AEIS sits in the middle of the 800VDC transition, separate from both the obvious winners and the likely losers. Advanced Energy Industries provides highly engineered power conversion and management systems for AI datacenters, converting utility power into the tightly controlled power required by servers and storage. It entered the datacenter market through its 2020 Artesyn acquisition and now sells AC-DC and DC-DC solutions, PDUs, and power shelves to both hyperscale and enterprise customers.

As of FY25, datacenter accounts for about 30% of revenue.

AEIS 处于 800VDC 转型的中心地带，既不同于那些显而易见的赢家，也不同于可能的输家。Advanced Energy Industries 为 AI 数据中心提供高度工程化的电源转换和管理系统，将公用电网电能转换为服务器和存储设备所需的受控精密电源。该公司通过 2020 年收购 Artesyn 进入数据中心市场，目前向超大规模和企业级客户销售 AC-DC 和 DC-DC 解决方案、PDU 以及电源机架。截至 2025 财年，数据中心业务约占其营收的 30%。

AEIS also supplies [OCP ORv3-compliant power shelves and power management controllers/modules] that coordinate multi-vendor power shelf deployments within a rack. The company is specifically named in the OCP Diablo 400 v0.7.0 specification, where its firmware manages power allocation, load balancing, and fault protection across mixed shelf suppliers. In a traditional 48V AC/DC architecture, that coordination is relatively simple. In an 800VDC architecture, it becomes much more complex. That shift is the key issue: it will determine whether AEIS strengthens its datacenter position or sees its role commoditized.

AEIS 还提供符合 OCP ORv3 标准的电源机架和电源管理控制器/模块，用于协调机架内多供应商电源机架的部署。该公司在 OCP Diablo 400 v0.7.0 规范中被特别提及，其固件负责管理不同机架供应商之间的功率分配、负载均衡和故障保护。在传统的 48V AC/DC 架构中，这种协调相对简单。但在 800VDC 架构中，情况变得复杂得多。这一转变是关键所在：它将决定 AEIS 是巩固其在数据中心的地位，还是眼睁睁看着自己的角色被商品化。

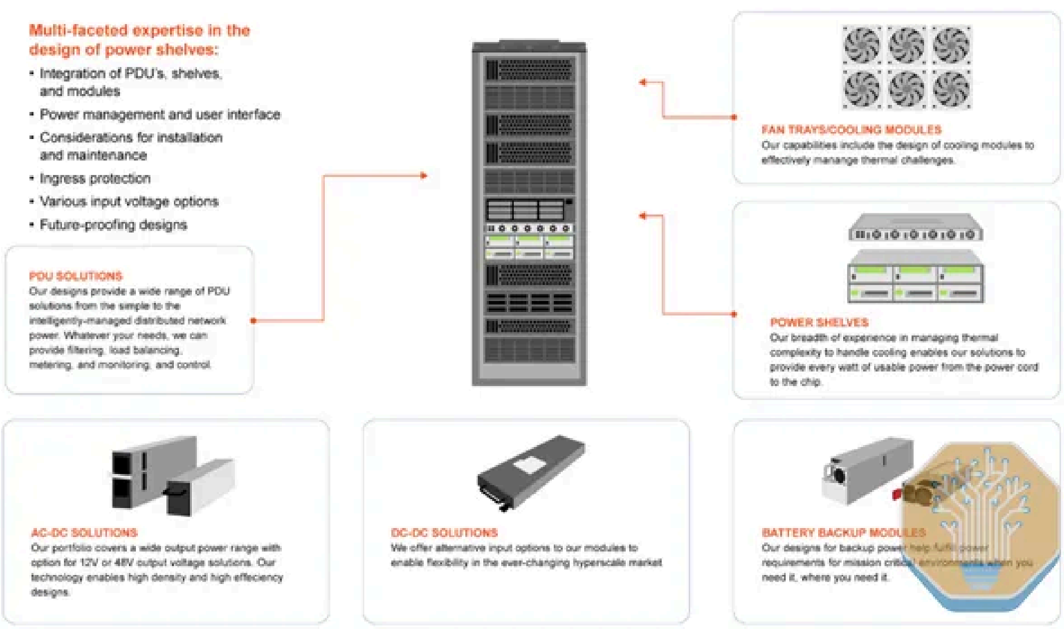

Source: AEIS AEIS

At OCP 2025, AEIS and Delta presented jointly, which AEIS showed the HPR V4 100 kW shelf (18 kW HVDC-to-DC PSUs, >97.5% efficiency), converting HVDC to 50V for compute trays inside the sidecar architecture. AEIS stated at OCP that by Q3 2027, “power shelves will be replaced” as HVDC goes direct to server, giving the current-form shelf a roughly two-year relevance window before AEIS must transition to HVDC-native products. The 2024 acquisition of Airity Technologies (GaN high-voltage power conversion) provides power electronics capability for that transition, positioning AEIS for Phase 3 architectures where 800V flows natively to the server.

在 OCP 2025 大会上，AEIS 与台达（Delta）联合参展，AEIS 展示了 HPR V4 100 kW 机架电源（采用 18 kW 高压直流转直流电源模块，效率 >97.5%），该设备在边柜（sidecar）架构中将高压直流（HVDC）转换为 50V 供计算托盘使用。AEIS 在 OCP 上表示，到 2027 年第三季度，“电源机架将被取代”，因为高压直流将直接输送到服务器。这意味着在 AEIS 必须向高压直流原生产品转型之前，当前形式的电源机架大约还有两年的应用窗口期。2024 年对 Airity Technologies（专注于氮化镓高压功率转换）的收购，为这一转型提供了电力电子技术储备，使 AEIS 能够为 800V 原生输送到服务器的第三阶段架构做好准备。

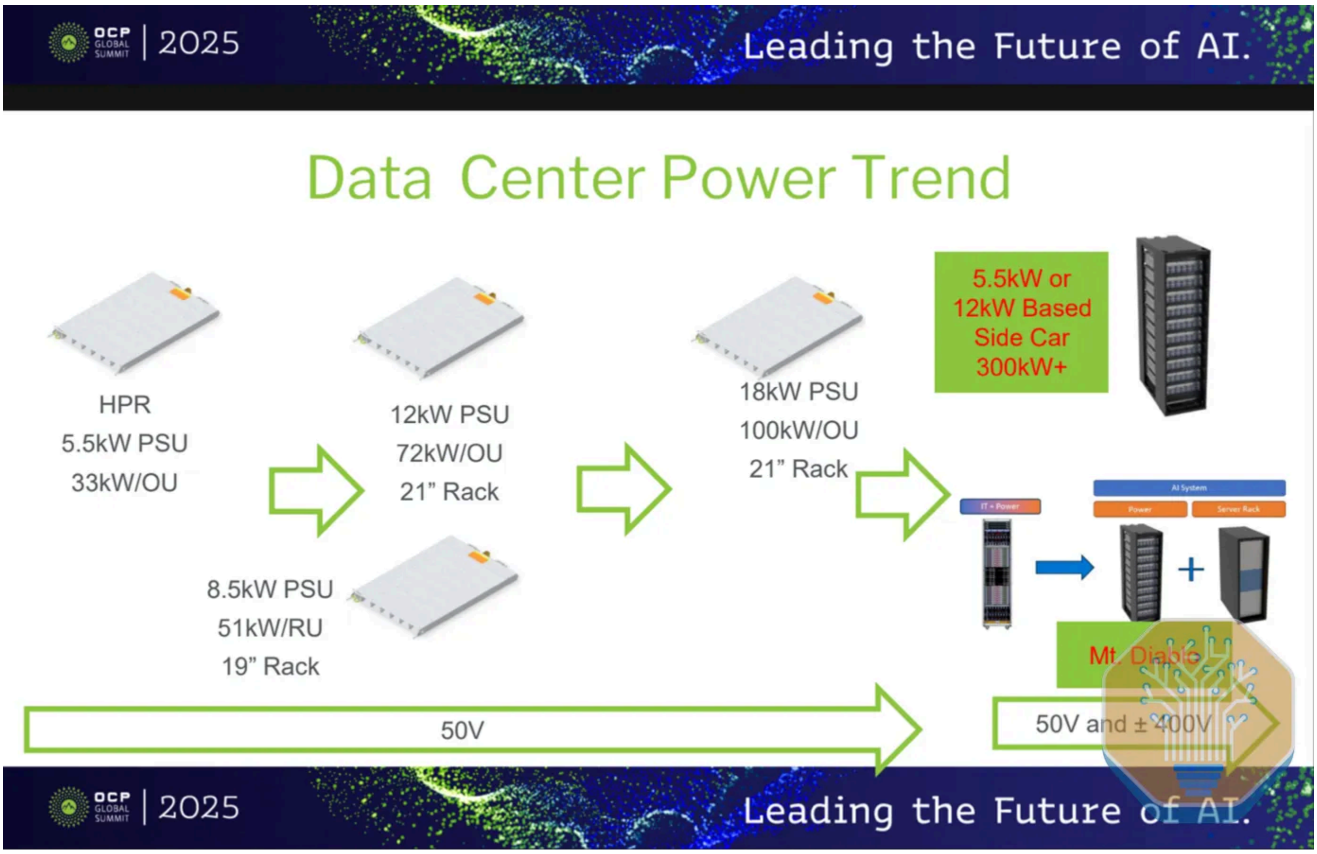

Source: OCP OCPp

The risk for AEIS is not product obsolescence but vertical integration. Delta builds its own PSU shelves, BBUs, capacitor shelves, DC-DC converters, and 800VDC busway, and, with ~75% share in AI server rack PSUs, has the incentive to consolidate the entire sidecar into a single-vendor appliance with proprietary management firmware. Vertiv delivered an integrated HVDC power rack to Meta on a similar basis. In that scenario, the controller function persists but becomes proprietary software optimized for one vendor's hardware, not a standalone product coordinating across multiple suppliers. AEIS also does not appear among Nvidia's 800VDC ecosystem partners. Its defense is the Diablo 400 spec, which standardizes the interface between any vendor's shelf and any controller. As long as hyperscalers enforce multi-vendor procurement (as they have historically) within the power rack, an independent orchestration layer is necessary, and AEIS is the reference implementation. The 800VDC transition is conditionally positive for AEIS as content per rack increases, and the orchestration problem compounds with density, but both outcomes depend on the sidecar remaining a multi-vendor system.

AEIS 面临的风险并非产品过时，而是垂直整合。台达（Delta）自行生产电源机架（PSU shelf）、电池备用单元（BBU）、电容架、直流-直流转换器（DC-DC converter）以及 800VDC 母线槽，且凭借在 AI 服务器机架电源市场约 75% 的份额，台达有动力将整个侧柜（sidecar）整合为带有专有管理固件的单一供应商设备。维谛（Vertiv）也基于类似的模式向 Meta 交付了集成式高压直流（HVDC）电源机架。在这种情况下，控制器功能依然存在，但会变成针对特定供应商硬件优化的专有软件，而非协调多个供应商产品的独立产品。此外，AEIS 也未出现在 NVIDIA 的 800VDC 生态系统合作伙伴名单中。其防御手段是 Diablo 400 规范，该规范标准化了任何供应商的机架与任何控制器之间的接口。只要超大规模云服务商（hyperscalers）像历史上那样在电源机架内强制执行多供应商采购，独立的编排层就是必要的，而 AEIS 正是参考实现。随着单机架价值量的提升，以及编排问题随密度增加而变得更加复杂，800VDC 的转型对 AEIS 而言是有条件的利好，但这两个结果都取决于侧柜是否能维持多供应商系统的格局。

Legrand (LR FP) 罗格朗 (LR FP)

Legrand is a major white-space equipment vendor across busbars, busway, rack PDUs, and IT racks. Datacenters represent about 26% of Legrand's FY25 revenue. Since OCP 2025, management has sent mixed signals on the 800VDC transition, balancing defense of legacy AC margins with development of DC products. Phase 1 architectures introduce the white-space sidecar (HVDC power rack), displacing AC rPDUs and row-level busway. Legrand also recently launched its first DC busbar, disclosed an 800VDC PDU in development, and acknowledged sidecars are where architectures are moving.

罗格朗（Legrand）是母线、母线槽、机架 PDU 和 IT 机架等白区设备的主要供应商。数据中心业务约占罗格朗 2025 财年营收的 26%。自 2025 年 OCP 峰会以来，管理层在 800VDC 转型方面释放了复杂信号，既要捍卫传统交流（AC）业务的利润率，又要开发直流（DC）产品。第一阶段架构引入了白区侧柜（高压直流电源机架），取代了交流机架 PDU 和行级母线槽。罗格朗近期还推出了其首款直流母线，披露了正在研发的 800VDC PDU，并承认侧柜是架构演进的方向。

Management argued on the F2H25 call that only PDUs and UPS, about 20% of DC segment revenue, face 800VDC displacement, and that any losses there will be offset elsewhere in the portfolio. We think that materially understates the risk. By phases 3 and 4 of the transition, we estimate roughly 55% of DC revenue is exposed, including rPDUs and busway, Legrand's highest-margin products. As architectures move toward sidecars and grid-to-chip DC, Legrand's legacy AC distribution risks being designed out and replaced by lower-value DC distribution downstream of a competitor's platform.

管理层在 F2H25 电话会议上辩称，只有配电单元（PDU）和不间断电源（UPS）——约占数据中心（DC）业务板块收入的 20%——面临 800VDC 的替代风险，且该领域的任何损失都将被投资组合中其他业务的增长所抵消。我们认为这严重低估了风险。据我们估算，到转型的第三和第四阶段，大约 55% 的数据中心收入将面临风险，其中包括 Legrand 利润率最高的产品：机架式配电单元（rPDU）和母线槽。随着架构向侧置式（sidecar）和从电网到芯片（grid-to-chip）的直流电转型，Legrand 传统的交流配电方案面临着被设计出局的风险，并可能被竞争对手平台下游价值较低的直流配电产品所取代。

Management frames grid-to-chip as a post-2030 issue, but the direction is clear: once an SST converts AC to DC at the grid edge, power stays DC to the chip, eliminating the inefficient AC-DC-AC conversions in between. The problem for Legrand is that it does not yet appear to have a fully formed 800VDC product portfolio ready for that architecture and seems to trail behind competition. That said, M&A remains a plausible catch-up path given its history of entering adjacent categories through acquisition (For reference, Legrand entered the datacenter space through its acquisitions of Raritan in 2015 and Server Technology in 2017 and has done ~30 datacenter acquisitions over the past 8 years).

管理层将“电网到芯片”视为 2030 年后的议题，但方向已经明确：一旦固态变压器（SST）在电网边缘将交流电（AC）转换为直流电（DC），电能将以直流形式一直传输到芯片，从而消除中间低效的“交流-直流-交流”转换。Legrand 面临的问题是，它目前似乎还没有为这种架构准备好成熟的 800VDC 产品组合，且在竞争中显得落后。尽管如此，考虑到 Legrand 通过收购进入相邻领域的历史，并购仍是一条可行的追赶路径（参考：Legrand 在 2015 年和 2017 年分别通过收购 Raritan 和 Server Technology 进入数据中心领域，并在过去 8 年中完成了约 30 次数据中心相关的收购）。

Source: Legrand 罗格朗

The risks across Legrand's product lines are distinct. Legrand expects to ship DC busbar/busway products by end-2026, ahead of ABB's 2027 timeline, though ABB is developing its DC busway directly with Nvidia while Legrand isn't. The transition from AC to DC busway is primarily a margin compression story; DC busway carries a lower value per MW than its AC equivalent. PDUs face a more direct challenge. Phase 1 sidecars already relegate the rPDU to a downstream accessory, and in Phase 2, Kyber's on-blade power modules eliminate in-rack DC-DC conversion entirely. The sidecar is the deepest gap, capturing the architectural value across Phases 1 and 2 that Legrand cannot participate in. Management has confirmed Legrand has no sidecar product, no development timeline, and no partnerships or acquisitions to close it, saying only that it will look to participate if margins prove attractive. Unlike Delta, Vertiv, and Schneider, Legrand is not part of Nvidia's 800VDC ecosystem.

罗格朗（Legrand）各产品线面临的风险各不相同。罗格朗预计将在 2026 年底前出货直流母线槽产品，领先于 ABB 的 2027 年时间表，尽管 ABB 正直接与 Nvidia 合作开发其直流母线槽，而罗格朗则不然。从交流到直流母线槽的转变主要是一个利润压缩的故事；直流母线槽每兆瓦（MW）的价值低于其对应的交流产品。PDU 面临着更直接的挑战。第一阶段的 Sidecar（侧柜）已经将机架式 PDU（rPDU）降级为下游配件，而在第二阶段，Kyber 的板载电源模块完全消除了机架内的 DC-DC 转换。Sidecar 是最大的缺口，它夺取了罗格朗无法参与的第一阶段和第二阶段的架构价值。管理层已确认罗格朗没有 Sidecar 产品，没有开发时间表，也没有通过合作伙伴关系或收购来弥补这一缺口，仅表示如果利润具有吸引力，将寻求参与。与台达（Delta）、维谛（Vertiv）和施耐德（Schneider）不同，罗格朗并非 Nvidia 800VDC 生态系统的成员。

Source: OCP OCPp

Forgent Power Solutions (FPS US)

Forgent Power Solutions is a pure grey-space electrical distribution supplier: transformers, switchgear, PDUs, and prefabricated enclosures for datacenter and utility grid customers. The company is a Neos Partners roll-up of four legacy businesses (MGM Transformers, PwrQ, States Manufacturing, VanTran) that IPO'd in January 2026. Datacenter accounts for ~42% of revenue, heavily skewed toward colo and neocloud operators.

Forgent Power Solutions 是一家纯粹的灰区（grey-space）配电供应商：为数据中心和公用事业电网客户提供变压器、开关柜、PDU 以及预制机房。该公司是由 Neos Partners 整合四家传统业务公司（MGM Transformers、PwrQ、States Manufacturing、VanTran）而成，并于 2026 年 1 月 IPO。数据中心业务约占其营收的 42%，且客户群严重向托管数据中心（colo）和新型云服务商（neocloud）倾斜。

Forgent exits CY26 with ~$3.5B of spare grey-space manufacturing capacity (~75% headroom versus 15-25% at the majors), translating to 8-20 week lead times on grey-space transformers and switchgear versus 40+ weeks at Eaton, Schneider, and ABB. In a market where equipment availability, not construction speed, is the binding constraint on new datacenter capacity, Forgent has become the fastest path to power for colo and neocloud operators who cannot wait 6-12 months for critical grey-space MV equipment.

Forgent 在 2026 历年结束时拥有约 35 亿美元的闲置灰区制造产能（约 75% 的余量，而主要竞争对手仅为 15-25%），这使得其灰区变压器和开关柜的交货周期仅为 8-20 周，而 Eaton、Schneider 和 ABB 则需 40 周以上。在当前市场环境下，设备可用性而非施工速度才是限制新增数据中心产能的关键瓶颈，对于那些无法为关键灰区中压（MV）设备等待 6-12 个月的托管数据中心和新型云服务商而言，Forgent 已成为获取电力的最快路径。

However, we believe the company is materially behind in 800VDC. We estimate about 35% of Forgent’s datacenter revenue comes from grey-space products that mature 800VDC architectures could eliminate, including UPS eHouses, LV switchgear, ATS, and LV transformers. Forgent does not manufacture white-space UPS systems and only supplies the grey-space enclosures that house them. By contrast, grey-space MV switchgear and substation transformers should remain essential, since MV distribution upstream of the DC conversion point is still required and could become more complex over time. Grey-space LV distribution products such as PDUs, RPPs, and panelboards should also persist with AC-to-DC redesigns; a capability Forgent already supports through DC-rated equipment used in fuel cell deployments.

然而，我们认为该公司在 800VDC 领域实质上处于落后地位。据我们估算，Forgent 数据中心收入中约有 35% 来自灰区（grey-space）产品，而成熟的 800VDC 架构可能会淘汰这些产品，包括 UPS 预制化电力机房（eHouses）、低压配电柜、自动转换开关（ATS）和低压变压器。Forgent 并不生产白区（white-space）UPS 系统，仅提供容纳这些系统的灰区外壳。相比之下，灰区中压配电柜和变电站变压器仍将是必不可少的，因为直流转换点上游的中压配电仍然需要，且随着时间的推移可能会变得更加复杂。灰区低压配电产品（如 PDU、RPP 和配电盘）在经过交流转直流重新设计后也应继续存在；Forgent 已经通过用于燃料电池部署的直流额定设备支持了这一能力。

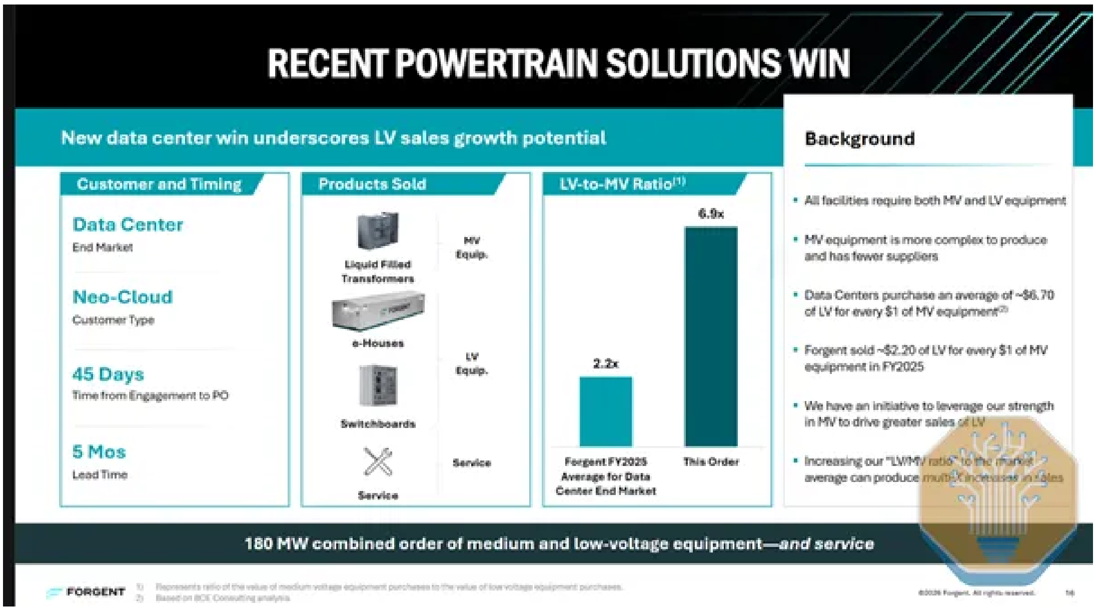

Source: Forgent Forgent

Looking into layered backup power landscape

深入探讨分层备用电源格局

As noted above, the final phase of 800V HVDC transition creates a layered backup power hierarchy with distributed short-duration energy storage near the racks and longer-duration storage upstream. We've gone in depth explaining the technologies in a previous article. In any datacenter architecture (legacy or 800V HVDC), the fundamental need is the same: bridge the gap between a power outage and generator start-up (~10-30 seconds for diesel, longer for gas turbines). What changes is where that bridge energy is stored and how it's delivered.

如上所述，800V 高压直流（HVDC）转型的最终阶段创建了一个分层的备用电源体系，其中包括靠近机架的分布式短时储能以及上游的长时储能。我们在之前的文章中已经深入解释了这些技术。在任何数据中心架构中（无论是传统架构还是 800V HVDC 架构），其根本需求都是相同的：填补停电与发电机启动之间的空白（柴油发电机约为 10-30 秒，燃气轮机则更长）。所改变的是这些桥接能量的存储位置以及交付方式。

Pansonic (6752 JT): The Dominant BBU Supplier Positioning for the 800VDC Transition

松下 (6752 JT)：占据 800VDC 转型先机的电池备用单元 (BBU) 主导供应商

Panasonic Energy (6752.T) claims approximately 80% market share in datacenter Battery Backup Units as of FY25, having shipped over 600 million lithium-ion cells into datacenter applications without a critical safety incident. The company's BBU development has progressed through OCP architectures from 33 kW (HPR V1) to 72 kW (V2), targeting 102 kW next-generation shelves, with underlying cell output rising from 80W to 120W "double type" cells to a target exceeding 200W. Revenue is scaling rapidly: FY25 sales are expected in the "upper ¥200 billion range," with a FY29 target of ¥800 billion at 20%+ ROIC, with over 80% of FY29 sales already secured through customer design wins. BBU demand continues to be revised upward, expanding "ahead of schedule" even relative to December 2025 guidance, prompting Panasonic to re-evaluate its capex allocation. Production expansion is notably capital light: the company is converting underutilized domestic EV battery lines (tripling Japan cell capacity by FY2029) and exploring repurposing its Kansas automotive site, requiring only “double-digit billion yen at most” in incremental investment.

松下能源（6752.T）声称，截至 2025 财年，其在数据中心电池备用单元（BBU）市场占据约 80% 的份额，已向数据中心应用出货超过 6 亿节锂离子电池，且未发生过重大安全事故。该公司的 BBU 开发已随 OCP 架构从 33 kW（HPR V1）演进至 72 kW（V2），并瞄准了 102 kW 的下一代机架，其底层电池输出功率也从 80W 提升至 120W“双型”电池，目标是超过 200W。营收规模正在迅速扩大：2025 财年的销售额预计在“2000 亿日元区间的高端”，2029 财年的目标为 8000 亿日元，ROIC 达到 20% 以上，且 2029 财年超过 80% 的销售额已通过客户定点（design wins）锁定。BBU 需求持续上修，甚至相对于 2025 年 12 月的指引也在“提前”增长，这促使松下重新评估其资本支出分配。产能扩张具有显著的轻资产特征：公司正在转换利用率不足的国内电动汽车（EV）电池生产线（到 2029 财年将日本电池产能翻三倍），并探索重新利用其堪萨斯州汽车电池工厂，仅需“最多两位数亿日元”的新增投资。

Panasonic is developing two next-generation product categories directly tied to the 800VDC transition. First, Capacitor Backup Units (CBUs) use a proprietary supercapacitor developed in-house by combining the industry segment’s chemical/capacitor technology with Energy’s battery assembly expertise. The CBU is designed to be form-factor compatible with existing BBU shelves, allowing operators to mix BBU and CBU modules for flexible customization of backup duration versus fluctuation absorption. Notably, Panasonic’s in-house CBU development positions it as a potential challenger to Musashi Seimitsu. Second, Panasonic is developing high-voltage BBUs specifically for 800V power racks. Panasonic continues to use NCA chemistry over LFP for datacenter, citing energy density requirements incompatible with LFP’s lower volumetric density at the shelf sizes datacenter applications demand.

Panasonic 正在开发两类与 800VDC 转型直接相关的下一代产品。首先，电容备份单元 (CBU) 采用了内部研发的专利超级电容器，该技术结合了工业部门的化学/电容器技术与能源部门的电池组装专业知识。CBU 的设计在尺寸上与现有的 BBU 机架兼容，允许运营商混合使用 BBU 和 CBU 模块，从而在备份时长与波动吸收之间实现灵活的定制化平衡。值得注意的是，Panasonic 的内部 CBU 开发使其成为 Musashi Seimitsu 的潜在挑战者。其次，Panasonic 正在专门为 800V 电源机架开发高压 BBU。在数据中心领域，Panasonic 坚持使用 NCA 化学体系而非 LFP，理由是数据中心应用所要求的机架尺寸与 LFP 较低的体积密度所带来的能量密度需求不相容。

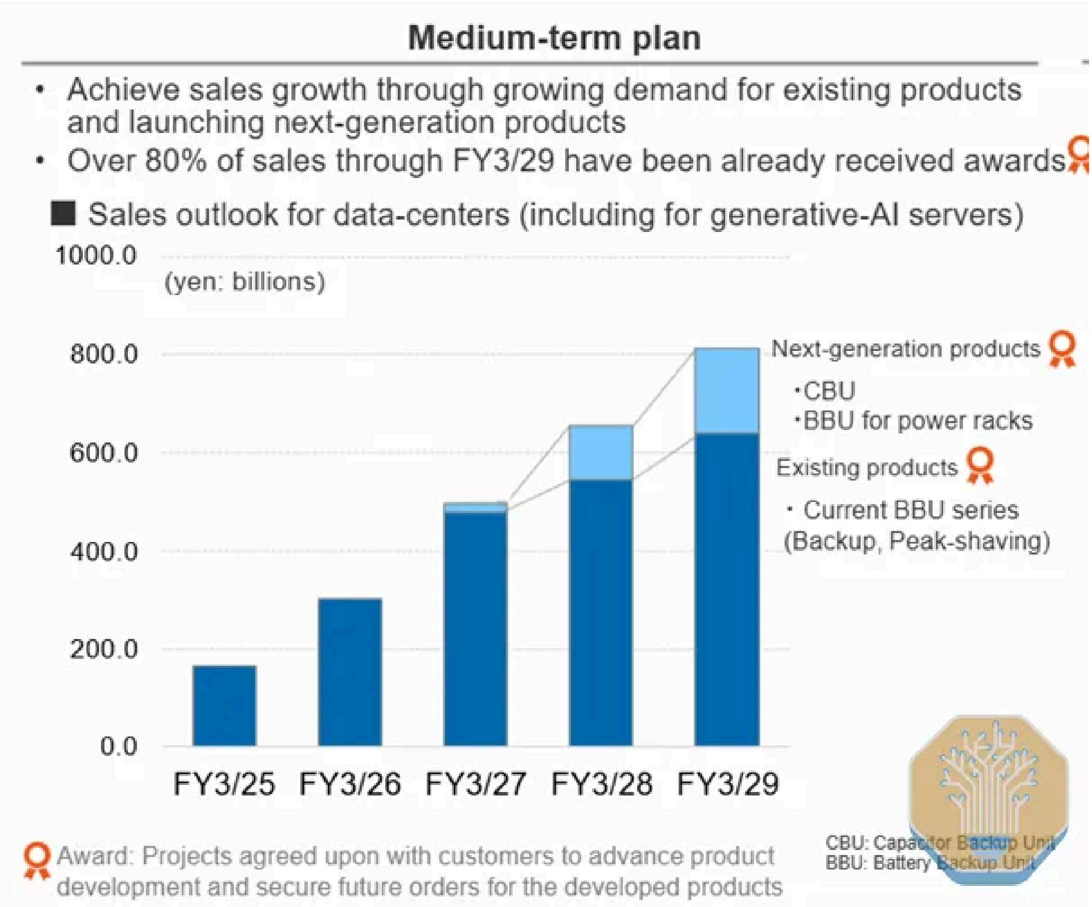

Source: Panasonic 松下

Strengthening supply capabilities	
Japan	**Cell:** • **Expand production capacity to approx. 3 times (vs. FY3/26)** • Increase production lines at existing facilities and modify production lines for automotive applications (start of production scheduled for FY3/27 1Q) **Module:** • Further enhancement of production capacity (with partner companies)
North America	**Cell:** • Utilize part of automotive base in Kansas, US (under review) **Module:** • **Made decision on building a new factory**, in addition to the expansion of existing production lines at the current factory (both in Mexico)

Source: Panasonic 松下

Musashi Seimitsu & Supercapacitors

武藏精密与超级电容器

Musashi Seimitsu (7220 TYO), through its subsidiary Musashi Energy Solutions (MES), holds essentially a monopoly in Supercapacitors. Historically a small-cap Japanese auto parts company, Musashi has been making a high-conviction pivot into AI datacenter energy storage.

武藏精密（Musashi Seimitsu，东京证交所代码：7220）通过其子公司 Musashi Energy Solutions (MES)，在超级电容器领域几乎处于垄断地位。武藏精密原本是一家小型日本汽车零部件公司，目前正坚定地向 AI 数据中心储能领域转型。

Musashi has historically been a Honda-affiliated manufacturer of transmission gears, crankshafts, suspension/steering parts. Musashi's HSC (hybrid supercapacitor) business, adoption will likely scale with 800vdc deployments which won't really see major traction until next year. Musashi has noted that HSC sales are likely ~ ¥10B (LSD% of FY26 earnings).

武藏精密在历史上一直是本田旗下的变速箱齿轮、曲轴、悬挂/转向部件制造商。武藏的 HSC（混合超级电容器）业务的采用规模可能会随着 800VDC 的部署而扩大，而 800VDC 预计要到明年才会真正获得重大进展。武藏指出，HSC 的销售额可能达到约 100 亿日元（约占 2026 财年收益的低个位数百分比）。

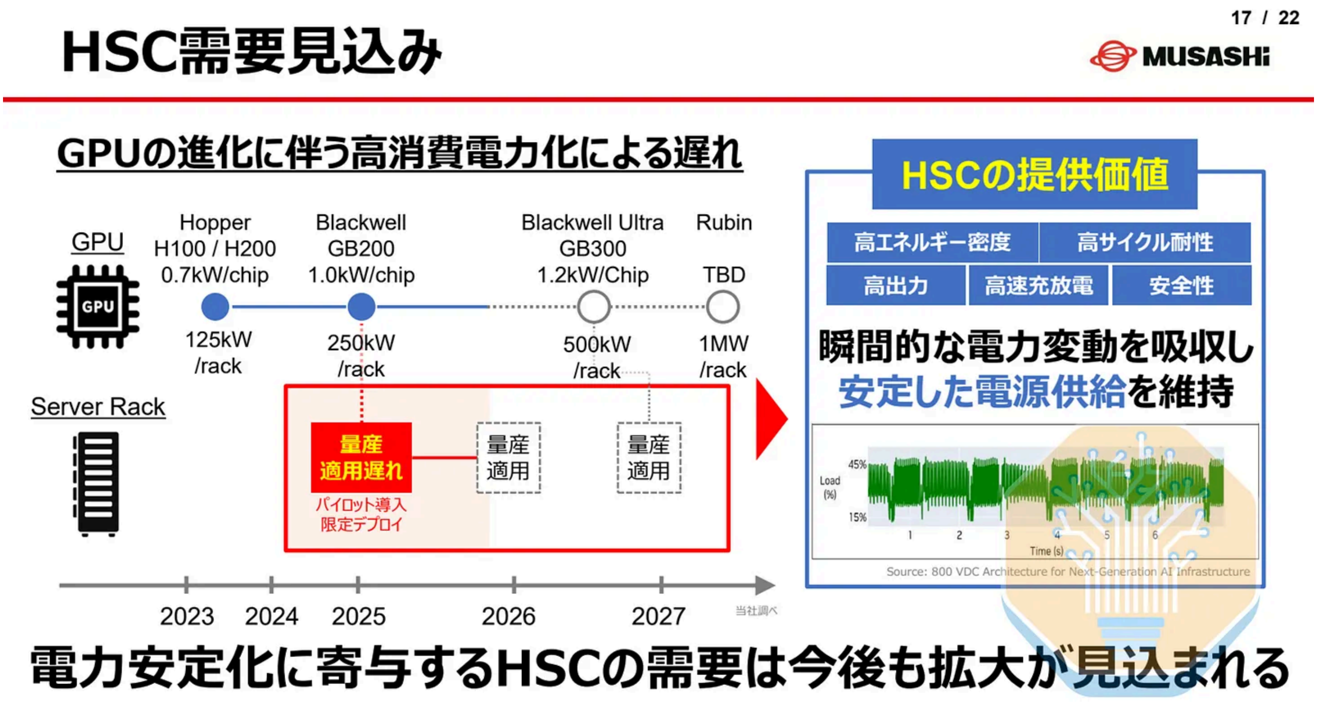

Source: Musashi Musashi

Musashi's HSC is not a standard EDLC (Electric Double Layer Capacitor). It is a hybrid device that combines EDLC positive electrode (activated carbon) with a lithium-ion pre-doped negative electrode. It provides much higher capacitance and energy density than standard EDLC options that other players are building.

武藏精密（Musashi）的锂离子电容器（HSC）并非标准的双电层电容器（EDLC）。它是一种混合型器件，结合了 EDLC 正极（活性炭）与锂离子预掺杂负极。与其它厂商制造的标准 EDLC 方案相比，它能提供更高的电容量和能量密度。

On partnerships, Musashi Seimitsu has already signed contracts with two power supply companies (Flex in the US, Delta in Taiwan). Moreover, at PowerGen 2026, Bloom Energy also highlighted the role of supercapacitors in its “Bloom Energy Stamp” architecture. For data-center deployments, Bloom indicated that ~3 MW of fuel cells are paired with ~2 MW of supercaps, enabling high-frequency load following, suggesting supercap content in grey space as well.

在合作伙伴关系方面，武藏精密已与两家电源公司（美国的 Flex 和台湾的台达电子）签署了合同。此外，在 PowerGen 2026 大会上，Bloom Energy 也强调了超级电容器在其“Bloom Energy Stamp”架构中的作用。针对数据中心部署，Bloom 指出约 3 MW 的燃料电池需与约 2 MW 的超级电容配对，以实现高频负载跟随，这表明超级电容在数据中心灰区（grey space）同样具有应用空间。

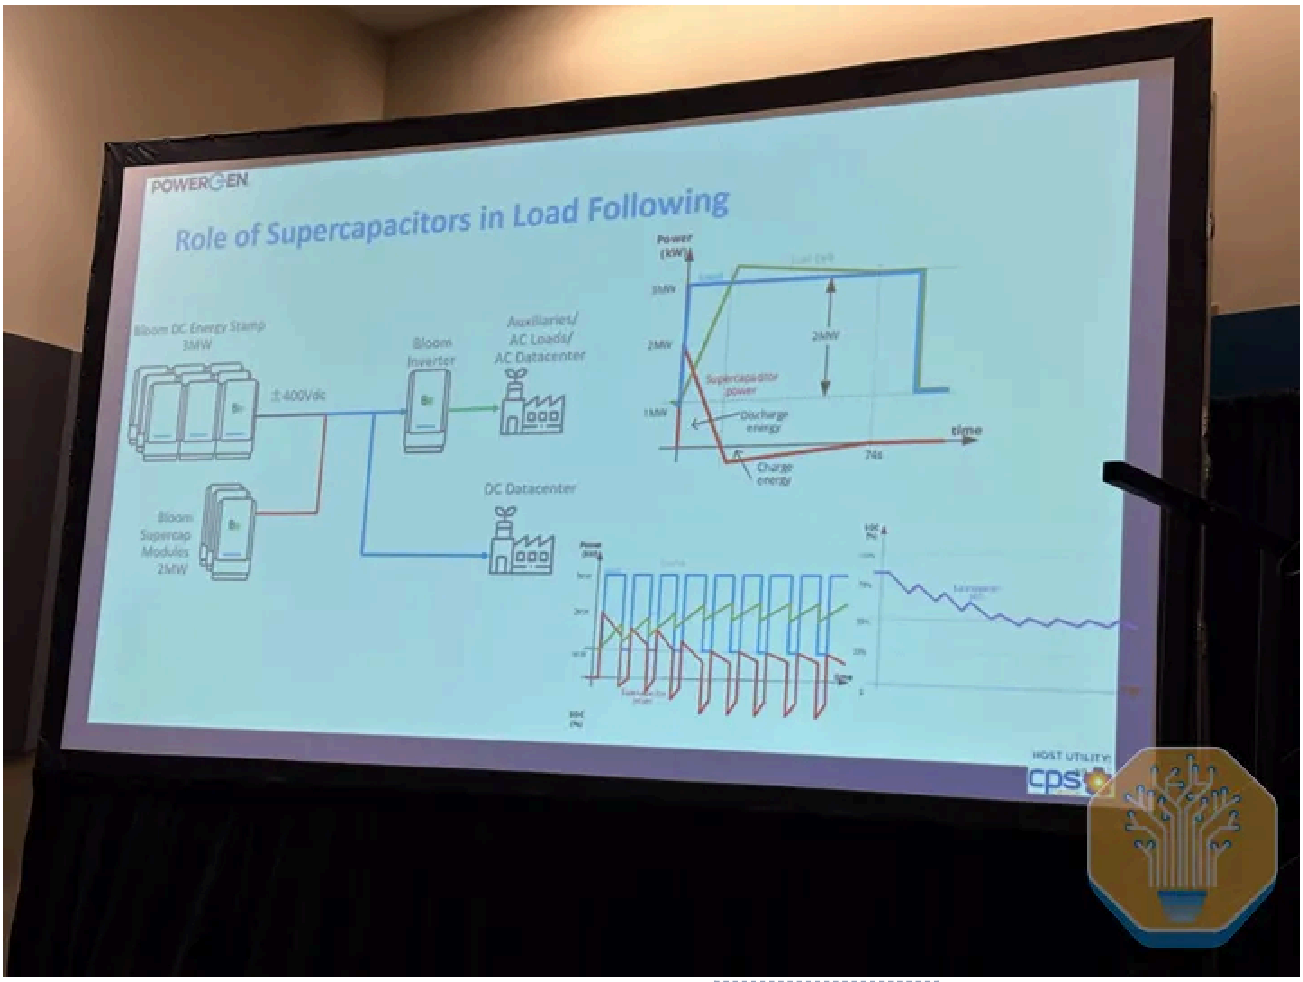

Source: Bloom Energy Bloom Energy

Recommend SemiAnalysis to your readers 向您的读者推荐 SemiAnalysis

Bridging the gap between the world's most important industry, semiconductors, and business.

架起全球最重要的行业——半导体与商业之间的桥梁。

Recommend 推荐

---

21 Likes 21 次赞 · 2 Restacks 2 次转贴

← Previous 上一篇

A guest post by 客座文章作者：

Aran Industries

Building the AI-native EPC for data centers. The world's ability to build has fallen behind its ability to dream. We're closing that gap

打造数据中心专用的 AI 原生 EPC。人类的建设能力已落后于其想象力。我们正在缩小这一差距

Subscribe to Aran 订阅 Aran

A guest post by 特约文章，作者：

Derek Yin

Interested in technology and markets

关注技术与市场

Subscribe to Derek 订阅 Derek

Discussion about this post

Comments Restacks

Write a comment...

© 2026 Dylan Patel · Privacy · Terms · Collection notice
Substack is the home for great culture