Europe Multi-Industry: Expert call takeaways on Nvidia 800VDC implications for electrical equipment



Daniela Costa
+44(20)7774-8354
daniela.costa@gs.com
Goldman Sachs International

Joe Ritchie
+1(212)357-8914
joseph.ritchie@gs.com
Goldman Sachs & Co. LLC

Ines Lefranc
+44(20)7051-8710
ines.lefranc@gs.com
Goldman Sachs International

Today, we hosted a follow-up investor webcast with Mr. Herve Tardy, a 30-year veteran in the datacenter industry with experience at Eaton and Schneider (see our takeaways from our first Webcast with Herve earlier this year here)**.** He gave an updated perspective on the outlook for datacenters with a focus on development of demand and competitive landscape for electrical equipment under the 800VDC architecture proposed by NVIDIA. **Below are our key takeaways from the call:**

- **The adoption of 800VDC power architecture is progressing slower than initially anticipated earlier this year, with the technology still in its infancy.** 800VDC architecture is now expected to account for 21% of new build datacenters in the US in 2030 (down from 25%) and 17% in Europe (down from 20%). This downward revision is primarily due to the stronger-than-expected growth of the datacenter market in AI inference, which does not have the same power density requirements as AI training. Given AI inference requires up to 150kW per rack maximum, traditional datacenter architecture is still viable, making the use of 800VDC for AI inference highly unlikely, with the extent of new architecture for AI inference likely limited to the use of liquid cooling rather than air cooling. Furthermore, full 800VDC architectures are not anticipated for at least 5 years, with most likely near-term development to be 800VDC for powering servers but traditional architecture for other products e.g. networking, storage, lighting. Mr. Tardy also noted that Nvidia Kyber NVL144 rack systems (designed to house Rubin Ultra GPUs) have been pushed to 2028, more than 12 months behind the original 2027 target.
- **Mr Tardy expects non-AI datacenters to remain dominant in 2032, hence there will still be strong demand for traditional datacenter architecture.** He quoted the example of Microsoft's datacenter expansion plans as highlighting a significant long-term growth trajectory for the industry, with the company targeting up to 38GW of power installed by 2032 (c.20% CAGR over the next 6Y). This underpins a highly favorable environment where electrical equipment companies can enjoy robust topline revenue growth and sustain c.20-25%

Goldman Sachs does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. For Reg AC certification and other important disclosures, see the Disclosure Appendix, or go to www.gs.com/research/hedge.html. Analysts employed by non-US affiliates are not registered/qualified as research analysts with FINRA in the U.S.

earnings CAGR. While AI datacenters are projected to grow at a faster rate, traditional power architectures will remain dominant. By 2032, non-AI datacenters are expected to constitute two-thirds of that 38GW installed base (from c.83%, but still implying 17% CAGR) while AI datacenters are expected to constitute only one-third of the installed base (from c.17%, implying 35% CAGR). This highlights that a huge market for traditional power architecture can still be served by players like Legrand which have not announced developments in Solid State Transformers (SSTs). He also expected those traditional areas to remain likely with higher margins than with 800VDC architecture.

- **Electrical equipment manufacturers are taking differing product approaches to the 800VDC transition.** Mr Tardy mentioned Delta Electronics are currently running SST pilots in Asia and targeting a 5MW design which is deemed the sweet spot in terms of size and flexibility for datacenter SSTs. Meanwhile, Eaton is planning to sample its 2MV SST with customers by the end of 2026, leveraging its 2025 acquisition of Resilient Power. By contrast, Schneider Electric, Vertiv and ABB are currently focusing on intermediate products for 800VDC. Schneider is currently pitching the Transformer Rectifier Unit (TRU) and the DC rectifier. Meanwhile, Vertiv and ABB are pitching MV UPS, an intermediate technology that will eliminate certain steps of power conversion. While this marks a step towards 800VDC, it still uses traditional technology. GE Vernova has also announced MV UPS but at 175MW (vs 1-2MW for ABB) and is still vocal about SST development, expecting introduction in 2027. Heron Power, considered the most active start-up, is developing 5MW SST and expects its new production facility to be ready by end of 2027.
- **Solar companies like Enphase and SolarEdge are attempting to enter the SST space but there are considerable barriers to entry.** While these solar companies can leverage their deep understanding of DC distribution, important questions remain regarding their ability to meet the stringent requirements of mission-critical datacenter operations, particularly regarding levels of redundancy, as well as availability of field service provision which they typically have not been as strong with. From a technology standpoint, SolarEdge is developing typical SSTs using silicon carbide, while Enphase is trying to use gallium nitride in parallel in order to reach the MW needed to power a DC.
- **Hyperscalers are sending massive Requests for Proposals (RFPs) for SSTs, likely as price and volume sensitivity exercises.** With current SST pricing remaining high at c.$400-500k per MW due to lack of manufacturing scale, hyperscalers are focusing on running pilots with different vendors to thoroughly test the technologies available. Once preferred product types start to get chosen, the evaluation and ordering cycle is likely to take an additional 12 to 18 months, pushing the earliest window for mass deployment of the SST to 2029.
- **The market is starting to switch its focus to other products that w[illegible] developed for the 800VDC transition, in particular solid state bre[illegible]** circuit breakers for traditional architecture operate within millisecon[illegible] 800VDC architecture necessitates microseconds. This will not only be[illegible] develop but also extremely expensive at c.$600k per MW (5x the pri[illegible] circuit breakers and considerably higher than SSTs at c.$400-500k pe[illegible]



this higher cost of technology, content per MW will increase in the coming years. By 2030, the whole powertrain for AI training is expected to require c.$4mn per MW (up from $1.7-1.8mn per MW in 2025) while AI inference is expected to require $2.6mn per MW (up from $1.4mn per MW in 2025).

- **The ongoing debate surrounding token output optimization is highly critical for AI training, but its relevance is diminishing in the context of AI inference.** Enterprise customers, who run the vast majority of inference workloads rather than hyperscalers, are coming to realise that token maxxing philosophy has been detrimental from a cost-efficiency standpoint. Outside of specialised compute-heavy fields like pharmaceutical drug discovery, typical enterprise applications should not require massive, high-density AI compute clusters.
- **Main catalysts to watch in the next few months include the OCP summit on 12-15th October and the NVIDIA GTC conference on 20-22nd October.** Within the datacenter market, OCP and Nvidia drive the thought leadership while OEMs as a whole are more focused on speed of bringing new technology to the market. There is currently little discussion focus on differentiation; this is likely to transpire in 4-5 years when second-generation SSTs start to be developed.