

AI Data Centers Move Capacitors Back into the Spotlight

Steve Sandler, Picotest

Published, Signal Integrity Journal

August 10, 2025



AI data centers are power hungry, pushing power density to new highs while requiring long-term reliability. The increasing power density, particularly at the board and rack levels, creates challenges for stability and the unpredictable high-current surges resulting from these many-core applications. Interestingly, these difficulties are driving a significant amount of new development in capacitors while also bringing back older technologies to be viewed from a different perspective.

At the rack and tray levels, the OCP Open Rack V3 Power Shelf voltage is 51 V nominal, with a range of 46 V to 52 V, or 54 V nominal, with a range of 52 V to 56 V. The power is connected through four connectors, each rated at 100 A, for a total power of about 20 kW. The current trend is increasing this 20 kW shelf to 30 kW or more, exacerbating these issues.

48 V Power Shelf Capacitance

Capacitance at the 48 V power shelf input is required for three main purposes:

1. Absorb excess energy from transients related to load dumps, inrush, hot swapping, workload demand, etc. The specification requires maintaining ≤ 1 V transient voltage excursions in response to these activities.
2. Provide a low impedance at the 48 V power supply output and 48 V-12 V DCDC converter input for a stable bus. The 48 V input in the converter side is a high-efficiency constant power load, meaning that it presents a negative resistance, which is the heart of an oscillator. We don't want this 48 V bus to be oscillating.
3. Maintain the main Power Shelf supply control loop stability. This is generally controlled by the maximum capacitance specification for the bus. This is currently in the range of about 50,000 uF or so, depending on the version of the specification and the shelf power.

$$ESR \leq \frac{1}{400} = 2.5 \text{ m}\Omega$$

The above must be true to maintain less than 1 V transient from a load dump. This is a conservative approach, as it assumes a 100% load dump instantaneously, with all the energy absorbed by the capacitors; however, it does offer a guide.

Another guide is that the transient voltage could approach a voltage determined by the characteristic impedance of the busbar inductance and the capacitors. The characteristic impedance would be the same (2.5 mΩ).

$$Z_o = \sqrt{\frac{L}{C}} = 2.5 \text{ m}\Omega$$

Three parallel capacitors would achieve the maximum capacitance, present an ESR below 2.5 mΩ, and absorb the energy from a busbar inductance, including clamp-on ferrites, of:

$$L \leq C \cdot Z_o^2 = 225 \text{ nH}$$

Maintain an impedance at the capacitors that is less than the negative real resistance of the Power Shelf's real resistance. For a 20 kW shelf, the high conversion efficiency results in an input current of:

$$I_{in} = \frac{P_{in}}{V_{in}} = \frac{20 \text{ kW}}{V_{in}}$$

This shows that as V_{in} increases, I_{in} decreases, establishing this resistance as negative. Note that the R_{in} is a 2nd order term, so this decreases with increasing V_{in}^2 and increases with P_{in} , meaning that as the shelf power increases, the real negative resistance is lower, imposing a more severe requirement on the capacitor and busbar impedance. As the bus voltage decreases, this can be problematic, particularly because the under voltage lockout will generally allow the power system to operate down to about 35 V.

Hybrid Capacitors

These severe requirements lead the current technology towards hybrid capacitors. The hybrid capacitor was invented by David Evans in 1996, and the patents were assigned to Evans Cap in 2003 (now Quantic Evans Hybrid Capacitor). The technology combined a tantalum anode with a ruthenium oxide cathode, resulting in very high capacitance density, very low ESR, and high maximum voltage ratings. These are all important characteristics for the Power Shelf capacitor, so this technology is being widely rediscovered 25 years later.

The impedance plot of a TDE series 18 mF 75 V Quantic Evans hybrid capacitor is shown in **Figure 1**.




Figure 1. Impedance plot of a TDE 18 mF 75 V Quantic Evans Hybrid Capacitor showing an extremely low ESR.

A comparison of current AI options for hybrid capacitors, shown in **Table 1**, indicates that Quantic Evans hybrid capacitors still hold a significant edge in performance and reliability.

Performance Comparison: Hybrid Capacitors in AI PDNs

Feature	Evans Hybrid (Tantalum –RuO ₂)	Polymer Hybrid (Al–Liquid Electrolyte)	Hybrid Supercapaccitor (EDLC–Li-ion)
ESR	Ultra -low (m0 range)	Low (20–120 m0 typical)	Moderate (depends on chemistry)
Capacitance Density	High	Moderate	Very High
Voltage Range	10–125 V	25–80 V	2,7–3,8 V (typ)
Ripple Current Capability	Excellent	Good	Excellent
Operating Temp Range	–55°C to +125°C	–55°C to +105°C	–30°C to +70°C
Lifetime (MTBF)	10 ⁵ + hours	10 ⁵ –10 ⁵ hours	Millions of cycles
Application Fit	Aerospace, radar, pulsed power	AI PONs, automotive, telecom	AI edge, backup, grid smoothing

Failure Mode Matrix

Stressor	Evans Hybrid (Tantalum –RuO ₂)	Polymer Hybrid Ciaked (Al–Liquid Electrolyte)	Hybrid Supercapacitor (EDLC–Li-ion)
Thermal Fatigue	Minimal (mil-spec rated)	Moderate (ESR drift, dry-out risk)	Low (wide temp tolerance)
Voltage Overstress	Self-healing possible	Risk of electrolyte breakdown	Risk of lithium plating
Humidity	Hermetically sealed	Vulnerable to ingress	Sensitive to moisture
High Ripple Current	Handles well	May degrade ESR over time	Excellent tolerance
Aging / Wearout	Slow ESR drift	Capacitance loss, ESR rise	Cycle life degradation

Table 1. Hybrid Capacitor Comparison

Low Voltage Bulk Capacitors

In the case of a 12 V intermediate bus, at the output of the VRM modules, bulk capacitors are required to provide a stable control loop while absorbing shorter-term dynamic current transients. The decoupling capacitors absorb higher frequency currents while the bulk caps absorb most of the lower frequency current energy. Typical tantalum capacitors offer high density, lower ESR, and a moderate parasitic inductance of about 1.5 nH. This 1.5 nH

then defines the decoupling capacitors, given that decoupling capacitors need to absorb this parasitic inductance related to excess energy.

This is changing due to the demands of AI. As a result, KAVX is developing a new low ESL tantalum capacitor to offer much lower inductance, greatly reducing the demands on the decoupling capacitors. The impedance plots of two of these tantalum capacitors, a 33 uF 35 V sample and a 330 uF 2.5 V sample, are shown in **Figure 2**.

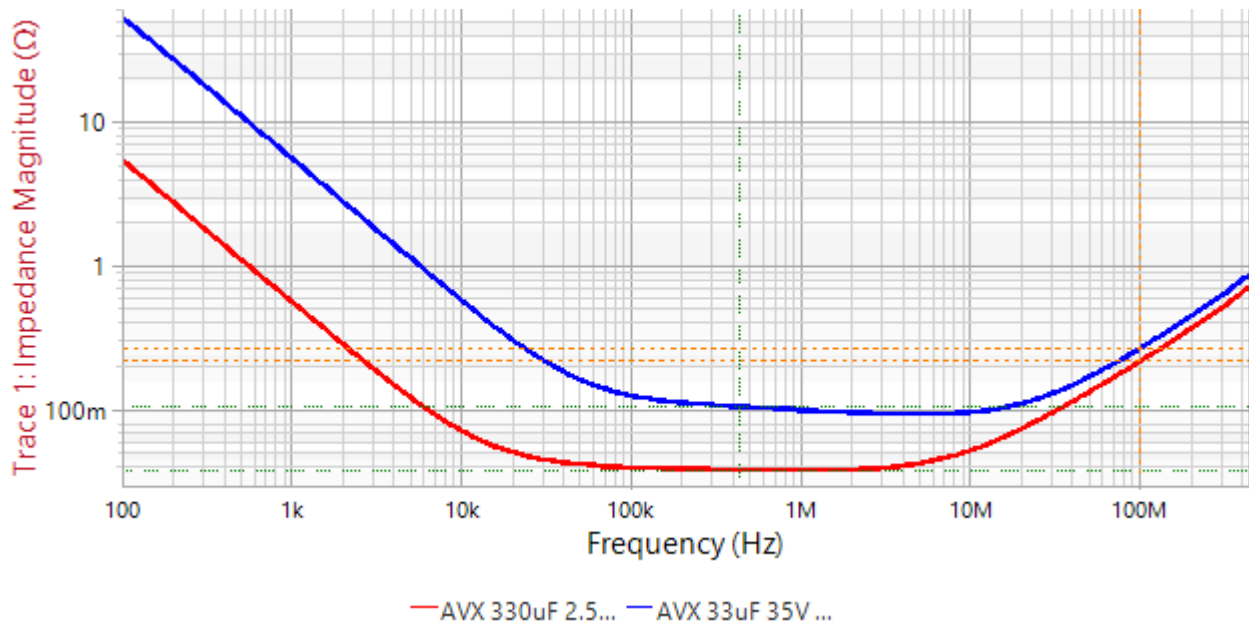


Figure 2. Impedance measurement of 33 uF 35 V and 330 uF 2.5 V TCF 3-terminal tantalum capacitor samples.

The key benefits of the new low ESL concept are the flat, wide bandwidth ESR and the greatly reduced parasitic inductance, similar to an 0805 MLCC capacitor. Since the decoupling capacitance is inversely proportional to this inductance, the decoupling capacitance requirements are decreased by a factor of 4 or 5, the same ratio as the parasitic inductance improvement. Also, the MLCC capacitor has piezo effects, DC and AC bias effects, and thermal expansion issues not seen in tantalum, making the tradeoff attractive.

For comparison purposes, impedance plots of the AVX TPS series 2-terminal tantalum capacitor and the low ESL concept capacitor are shown in **Figure 3**. The two devices have roughly the same capacitance and the same ESR up to about 50 kHz. The 3-terminal TCF capacitor ESR is much flatter than the 2-terminal TPS capacitor. The TCF 3-terminal capacitor is also a much lower inductance at 340 pH vs. 2200 pH for the TPS 2-terminal capacitor.




Figure 3. Impedance comparison of AVX TPS series 2-terminal tantalum capacitor and TCF series 3-terminal capacitor.

Based on Capacitance/Voltage product, in the future, we might expect 16 V components, optimizing the requirements for the 12 V intermediate bus, and 1.5 V components to optimize the core voltage rails. Although a very low ESR can be problematic for PDN, resulting in high Q, we may also see a reasonable reduction in ESR in the future with polymer as well.

There are plans to fabricate some VRM demo boards with low ESL design and some demo boards with AVX TPS series capacitors to examine the difference when everything is equal aside from the bulk capacitor selection.

Low Voltage Decoupling Capacitors

While the MLCC capacitor has been around for decades, pressure is also being exerted on the decoupling capacitor, which is selected based on a balance of capacitance needed to absorb the parasitic inductance of the bulk capacitors and PCB planes and the number of capacitors required to reduce the inductance at the GPU package. This is also becoming challenging as the GPU power increases and the GPU size decreases. This means fewer capacitors can be fit into a smaller space under the GPU package to meet the maximum parasitic inductance demand. On the other side, the package and die caps are required to absorb the parasitic inductance of the board decoupling, making it essential to reduce this decoupling capacitor inductance as much as possible.

To achieve this goal, Quantic Eulex developed the Gap Capacitor, an extremely low inductance, 3-terminal, single-layer, ceramic capacitor. The XG3 3-terminal Gap Capacitor

offers extremely low parasitic inductance, but with a lower capacitance than MLCC capacitors. The ultra-low ESL and physical geometry of the Gap Capacitor are ideal for higher frequency applications, which is what it was designed for. In high bandwidth applications such as transceivers, these capacitors offer benefits in high pulse energy, as well as in microwave frequency applications, such as filters, DC blocks, and other applications up to 25 GHz or higher.

These characteristics also make it attractive for pairing with the AVX 3-terminal tantalum capacitor. Since the lower parasitic inductance of the 3-terminal tantalum capacitor requires less capacitance, the Gap Capacitor might just fit the bill. And with much lower parasitic inductance, fewer capacitors are required to achieve the parasitic inductance allowed at the ever-smaller GPU package. Just how low is the parasitic inductance of the Gap Cap? Typical impedance plots for 0.012 uF and 0.1 uF samples are shown in **Figure 4**.

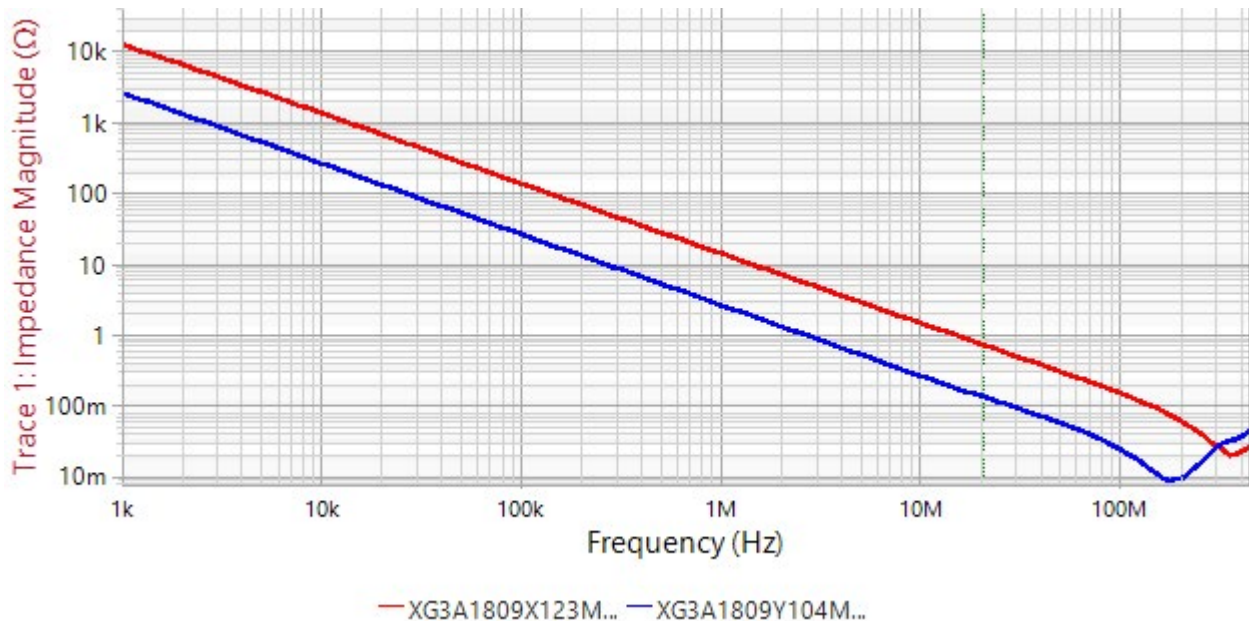


Figure 4. Impedance measurement of 0.012 uF and 0.1 uF Quantic Eulex Gap caps, indicating 6 pH and 16 pH, respectively.

The results of these impedance plots show parasitic inductance magnitudes of 6 pH for the 0.012 uF capacitor and 16 pH for the 0.1 uF capacitor, each more than an order of magnitude lower than a traditional MLCC.

Higher Frequency Capacitors

Technologies are evolving for the package as well. Capacitor companies are now developing and manufacturing capacitors that are capable of being embedded within the substrate core or under the substrate package, as shown in **Figure 5**. They share a common

goal of getting capacitance closer to the die where it is needed, particularly in the case of vertical power delivery, since the PCB under the ASIC isn't available for discrete capacitors.




Figure 5. The E-Cap is intended to be placed on the PCB/package or embedded within it. (Image courtesy of Empower.)

Empower's silicon E-Cap is one example of this high frequency capacitor evolution. The E-Cap is a high-density, ultra-low ESL silicon capacitor. These capacitors are embedded within the core of the package or placed on the landside of the package within the ball-height, offering high capacitance density up to about 2.5 uF/mm² in addition to ultra-low ESL, likely in the single digit pH.

According to an [article](#) published by Charger Lab in June 2025, the Empower EC1004B silicon caps are seen in the teardown of the Xiaomi 15S smartphone, where the capacitors are placed within the solder balls (substrate landside) of the Xring O1 SOC chip, as seen in **Figure 6**. Clearly, they are very thin capacitors to fit in the required height.




Figure 6. The Xiaomi 15S teardown shows the 0.5 mm x 0.64 mm Empower silicon cap embedded between the solder balls of an Xring O1 SOC. (Image courtesy of Empower and ChargerLab.)

An example of a silicon capacitor that has been designed for high-density and ultra-low ESL is the EC1005 shown in **Figure 7**. The EC1005 provides 16.6 uF with sub 1 pH of ESL and 3 mΩ of ESR in a 120-pad configuration to provide the ultimate form of decoupling for a high-performance computer chip. These capacitors are designed to be embedded within a substrate core with the ability to match the thickness of different types of cores.

The Saras Micro Devices STile is another good example of high-density embedded capacitor technology for under-die or embedded decoupling. What sets the STile apart from others is that, according to an [Electronics Engineering Journal](#) article from February 2025, Saras Micro Devices intends to include more passive devices, including the VRM inductors, in embedded printed circuit board layers, as shown in **Figure 7**.




Figure 7. The Empower EC1005 silicon capacitor for substrate core embedding applications for HPC/AI.

One of the issues with the current AI scaling is that the interest in vertical power delivery displaces the decoupling capacitors that are typically placed in this location. Moving these capacitors into the package substrate moves them much closer to the die, where they are the most effective, while clearing the underpackage area for the new technology: low profile VRM modules designed specifically for vertical power delivery.

The Empower E-Cap is another example of high-frequency capacitor evolution. The E-Cap is a high-density, ultra-low ESL silicon capacitor. It can be placed on the substrate or under the package while affording very high capacitance density close to the ASIC. Where there is technology advancement, there are competitors vying for market share, so look for emerging companies offering similar technologies soon.




Figure 8. The Saras Micro Devices STiles are intended to be either board/package mounted or embedded.

What's the Trouble?

While this all sounds like great news, it also presents some new challenges. Specifically, it introduces the challenge of how to accurately measure these low parasitic inductance values. Measuring 10 pH is a significant challenge on its own, but components meant for embedding don't lend themselves to simple measurement. Picotest developed a

component test fixture and a universal calibrator for the measurements of these devices using the OMICRON Lab Bode 500 vector network analyzer. Picotest will make these new devices available soon, allowing anyone to make measurements like these.

Conclusion

The demands of AI data centers are applying pressure to all electronic technologies, from measurements to power converters and capacitors to inductors. The demands are for smaller size, higher frequency, and more efficient solutions. These are forcing (or allowing, depending on perspective) component manufacturers to develop newer technologies. More people are also looking to repurpose old technologies that perhaps didn't fit in the power integrity sector in the past; some might be a great fit for these newer AI data center requirements.

Picotest hopes to show many of these new capacitor technologies in our DesignCon 2026 and APEC 2026 sessions and demonstrations.

<https://www.signalintegrityjournal.com/articles/print/4037-ai-datacenters-move-capacitors-back-into-the-spotlight>