Power Density in AI Data Center PSUs: How Do We Move Above 100 W/In^3

By Charles Bailey, Senior Director of Business Development &
Tao Wei, Sr. Applications Manager, Navitas Semiconductor

Navitas Semiconductor recently developed a 137 W/in^3 power supply for data center servers. This is 40% higher than comparable CRPS power supplies, here Charles and Tao outline the challenges the industry faces in powering the next evolution of generative AI.

Standardization brings many benefits to the data center environment. It enables power supplies (and other components) to be swapped out easily, regardless of vendor, with almost no down time. But it also creates challenges for component manufacturers who need to deliver innovative leaps forward while constrained by specifications that were set several years before.

This can be seen if we take a look at the M-CRPS specification for internal redundant power supplies, which forms part of the Open Compute Project (OCP)'s DC-MHS (Data Center – Modular Hardware System) family of standardized interfaces and form factors, which are used as building blocks for a range of functions from control panel connections to computer host-processor modules and high-speed data interfaces.

The M-CRPS power modules come in a small number of form factors, with notable ones including the CRPS185 (185 x 73.5 x 40 mm) and CRPS265 (265 x 73.5 x 40 mm). Output voltages can be set to 12 V or 54 V and using this latter 54 V standard enhances distribution efficiency in high-power hyperscale / AI workloads.

Recent years have seen rapid growth in both generative AI capabilities and demand. As such, this is placing exceptional demands on the server power networks

The advances underpinning the latest generation of AI processors, such as NVIDIA's Grace Hopper and Blackwell GPUs as well as AMD's Instinct "Antares" MI300X GPU, enable significant energy efficiency improvements compared with their predecessors, with kilowatts per petaFLOPS markedly down, albeit it should be said that despite this, the raw power drawn per processor is simultaneously markedly up.

Featured in Bodo's Power Systems – November 2024

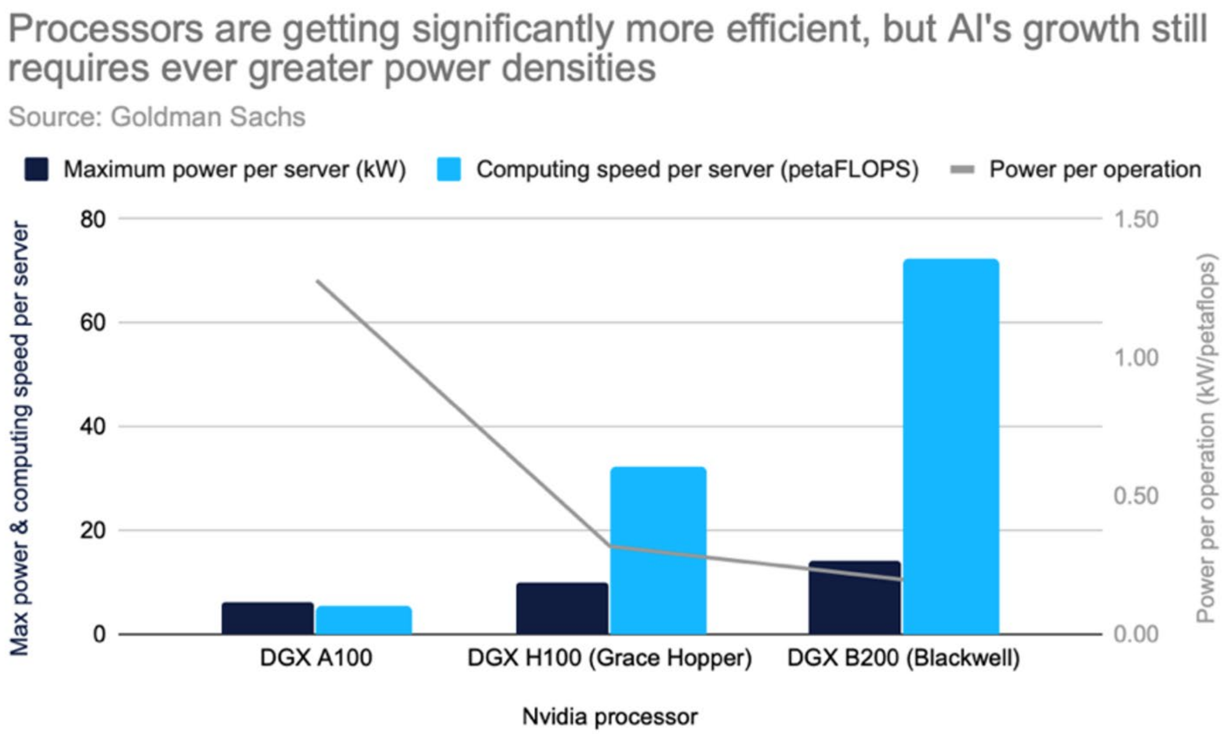

Figure 1: The power efficiencies of NVIDIA's leading AI processor offerings

(Note – specifications for NVIDIA's Rubin processors are not yet available, but are expected to follow the table's trend.)

As per figure 1, these two generational improvements from NVIDIA have delivered a more than 14-fold increase in processing power, and this has been achieved with little more than a doubling of power, which is clearly a great leap forward.

But it's equally clear that to deliver this improved execution performance from a standard CRPS form factor, PSU power density improvements are needed… and needed quickly.

NVIDIA's data center SuperPOD reference designs incorporate six 1U CRPS PSU slots. For the DGXH100 SuperPOD design, the PSUs are configured with a 4+2 redundancy. But, if we look at the DGX B200 (Blackwell) SuperPOD reference documentation, we see redundancy is reduced, with five of the six needing to be energized at any given time: "The system can operate if a single internal power supply unit is de-energized, but will not operate if more than one power supply unit is de-energized, regardless of upstream power redundancies."

With minimal redundancy in place for the current generation, it's a safe bet that the squeeze on power is only set to intensify.

Improving Power Density From 98 to 137 W/In³

With silicon reaching its physical limits, wide bandgap semiconductors – particularly silicon carbide (SiC) and gallium nitride (GaN) – can be applied in PSUs to deliver a higher-density power supply design.

Earlier this year, Navitas Semiconductor developed a reference design for a 54 V CRPS PSU that enabled a 40% increase in delivered power (4.5 kW vs the existing 3.2 kW) from within a standard CRPS185 form factor.

This increases the power density for the reference design from 98 W/in^3 (3.2 kW PSU) to 137 W/in^3.

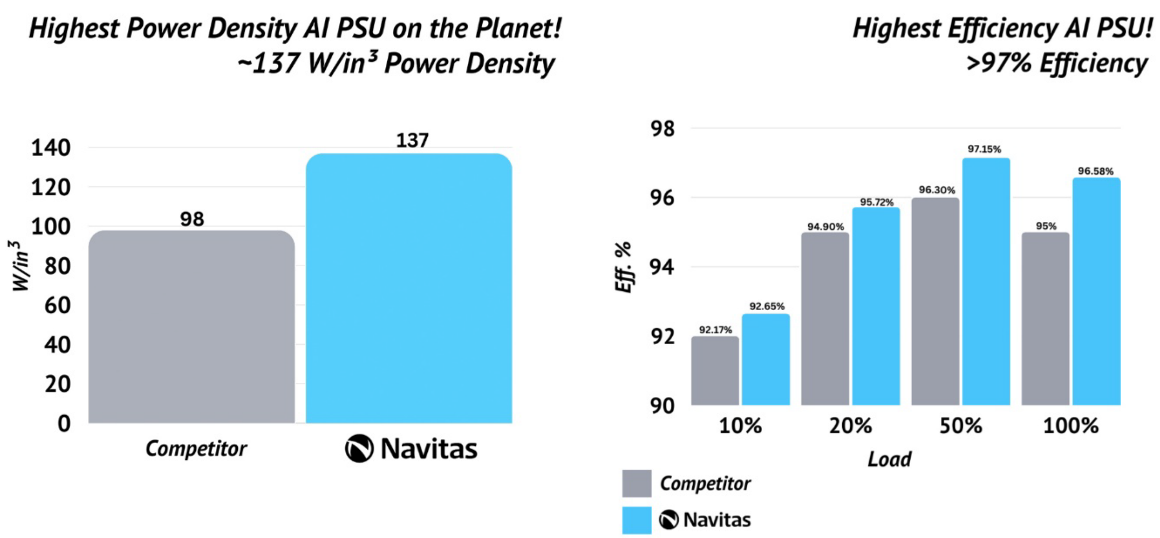

Figure 2: Power densities and efficiencies of the 4.5 kW reference design compared with a commercially available 3.2 kW CRPS185 PSU

Of course, power supplies for data centers must meet efficiency specifications, with 80PLUS Titanium being adopted either voluntarily or (in the case of data centers in the EU) through mandated legislation. In order to demonstrate compliance, PSUs need to meet efficiency targets across the load range – 10%, 20%, 50%, and 100%, with the standard stipulating 96% efficiencies at 50% load. The reference design exceeds requirements across the load range and reaches over 97% at 50% load.

80PLUS Titanium also stipulates PSUs have a power factor of at least 0.95 at lower load levels, which makes active power factor correction (PFC) a necessity.

PFC

As per figure 3, Navitas Semiconductor has adopted a bridgeless interleaved totem-pole PFC for the reference design and this includes a boost stage with steering switches. In comparison with a conventional bridge rectifier, this has the benefit of greatly reducing component losses.

Silicon carbide MOSFETs have been used because of their minimal switching and reverse-recovery losses, which enables the PFC to operate with a loss budget well beyond what would be capable through silicon alone.

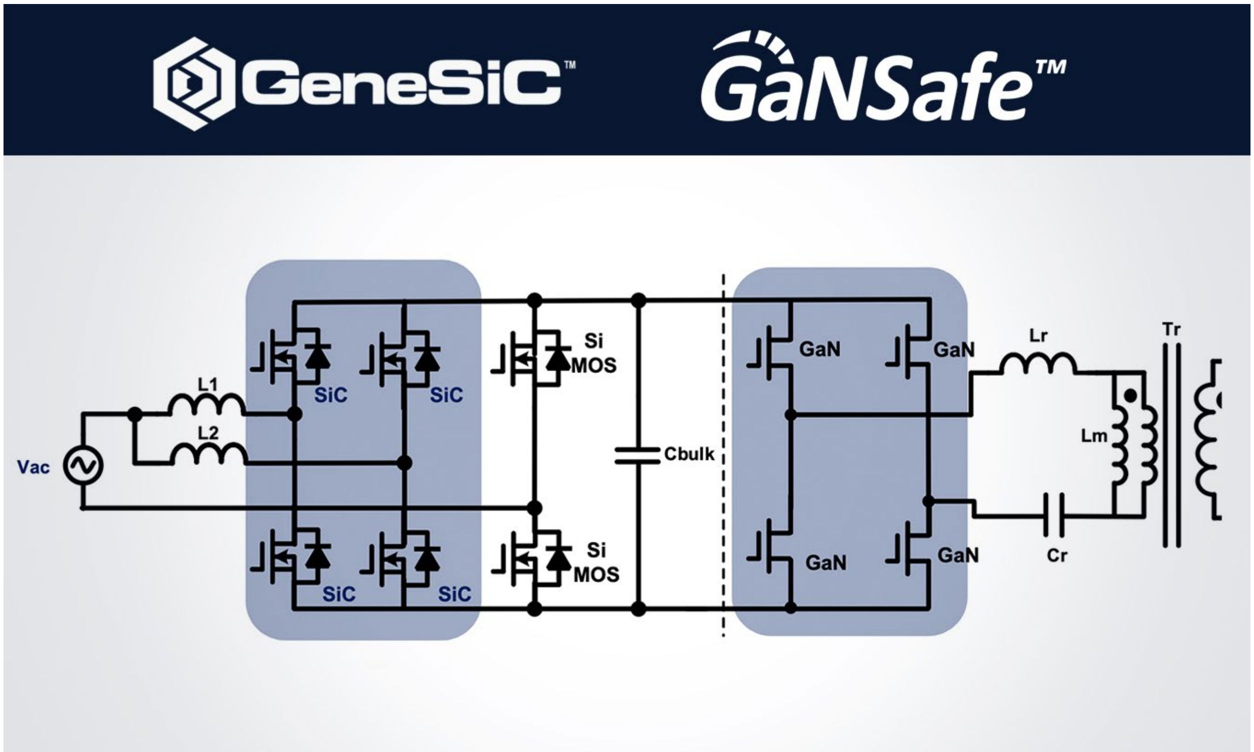

Figure 3: 4.5kW reference design with SiC bridgeless totem-pole PFC and GaN full-bridge LLC

LLC

This PFC stage in turn powers an LLC resonant converter with a full-bridge square wave generator in order to excite the resonant tank circuit, with the stable 54 V output delivered on the transformer secondary side through a CR filter and GaN rectifiers.

By using zero-voltage switching (ZVS) of the full-bridge transistors at the resonance frequency of the tank circuit, we are able to further improve efficiencies. However, the resonant components, as well as the associated circuitry and output filter must be able handle a greater current within the same overall form factor with output current for a 54 V PSU delivering the full-load power of 4.5 kW being 83 A.

The full bridge for the 4.5 kW reference design is built using 650 V GaNSafe ICs, with the reference design's power density requiring a selected switching frequency of 300 kHz. This is roughly double that of the most powerful silicon-based CRPS units (c.150 kHz), and while silicon's switching frequency cannot significantly go beyond 150 kHz, properties such as the output capacitance and gate charge of the GaN power transistor enable it to operate efficiently well beyond 300 kHz.

GaN's Gate Fragility

At this point, GaN's fragile gate structure should be noted. To mitigate this and protect against negative voltage spikes and ringing, the gate-drive circuit's design is critical.

To protect the gates, GaNSafe ICs integrate an optimized driver that enables a carefully controlled inductance and resistance between the output and the gate. This can also be achieved through through discrete components but doing so would introduce extra design challenges as well as require an increased PCB area (and therefore reduced power density).

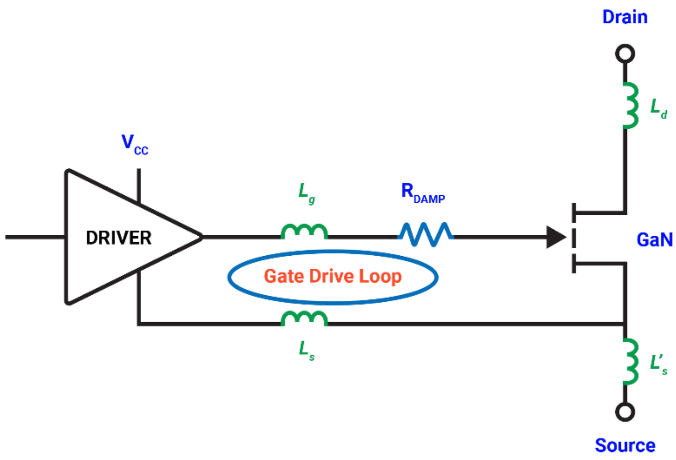

Figure 4: The integrated driver permits controlled gate-loop inductance

Conclusion

With the rise in demand for AI has come a huge rise in data center power consumption. Constrained by fixed form factors, PSU manufacturers must significantly improve the power density of its supplies if the industry is to be able to meet these evolving needs.

Through this reference design, we have shown this is possible through a combination of SiC totem-pole PFC and GaN high-frequency LLC that both drive efficiency to the 80PLUS Titanium stipulated level and reach a power density well beyond the capabilities of ordinary silicon devices.

As the next generation of GPUs for AI data centers enter operation, the need for ever more powerful PSUs will continue and Navitas Semiconductor has therefore set out a roadmap to reach 8.5 kW per PSU before the end of 2024, and 10 kW thereafter.