ISSUE: August 2026

Planar Design Tool Speeds Simulation Of 350-W Transformer Design For 12-mm Enclosure

by Marcos Luna, Frenetic, Madrid, Spain

This article walks through the design of a 350-W planar transformer for a 48-V isolated dc-dc converter. The converter switches at 400 to 500 kHz with GaN devices, and the whole module has to stay under 12 mm. That height limit is what makes the design interesting. So is the way we meet this requirement.

Here we demonstrate how to use Frenetic's planar tool to get from requirements to a part we can build much faster than with conventional approaches. The workflow is short: pick a core from the library and customize its dimensions if you need to, set the gap, choose the material, build the winding stack-up layer by layer, load the current waveform, and run. Five to ten minutes later you have the full loss breakdown, the leakage inductance, the parasitic capacitance and the temperature. We use COMSOL underneath.

After some discussion of the transformer design challenge, we review the benefits of a planar design versus a wound component. Low profile is not the only reason for choosing planar. We then identify two possible topologies for consideration, working through the details of core and winding design, and determining losses and temperature rises for the transformer design for each topology using Frenetic's planar tool. Finally, we perform a thermal analysis using Frenetic's new thermal model to evaluate the performance of the two transformer designs to see which one wins.

The Design Challenge

Where The Converter Sits

Racks in AI data centers now draw tens of kilowatts, and 48-V distribution is how that power reaches the board without losing it in copper. Most of that conversion is done with low-profile bricks: standard footprints that save board area and can be paralleled to scale the power. This article is about the transformer inside one of them, a 350-W isolated stage taking 48 V in and giving 12 V out.

Why 400 To 500 kHz

GaN will switch a lot faster than this. Output capacitance is lower and the transitions are quicker, so switching loss stays manageable well past a megahertz, where silicon would already be in trouble. We stopped at 400 to 500 kHz because the ceiling here is the magnetics, not the devices.

Higher frequency makes the magnetics smaller. The volt-seconds the core has to support scale inversely with frequency, so we can use a smaller core and fewer turns.

The cost is in the copper. Skin and proximity effects exist at any frequency, but they grow with it, and at these frequencies they become the dominant loss in most designs. Most of the design time goes into the winding, not the core.

Why Planar

Planar is not automatically better than a wound transformer. The usual argument is that planar has lower losses, but that is not what decided it here. We chose planar for two reasons: repeatability and heat.

Repeatability. In a wound transformer the geometry depends on the operator, and small changes in how the wire sits move the leakage inductance, the dc resistance, the ac resistance and the parasitic capacitance. With a PCB winding all of those are set by the layout, so we can put a tolerance on them. That matters because the converter design assumes specific values for leakage and winding resistance; if they drift, ZVS and efficiency drift with them. On losses alone there is no clear winner, since with Litz you reduce the strand diameter and with a PCB you change the copper weight.

Heat is the reason we went planar. A PCB winding dissipates much better than a wound one, because the copper is spread over a large flat area and it is in contact with the core over most of its surface. This does not

© 2026 How2Power. All rights reserved.

mean fewer losses (a planar part is usually smaller, so the loss density is higher). What we get is a lower temperature for the same loss, and that is what lets us shrink the part.

The last reason is mechanical. Our maximum height is 12 mm and we cannot exceed it—a wound transformer has a hard time getting there. Round wire dissipates worse than a flat layer bonded to a board, so it has to run at a lower current density, which means more copper cross-section and a taller bobbin. The height limit and the winding technology are not independent choices.

The Constraints

The power converter and application dictate the contraints given in Table 1.

Table 1. Electrical and mechanical design specs for transformer.

Parameter	Value
Output power	350 W
Input voltage	40 to 60 V (48 V nominal)
Output voltage	12 V
Switching frequency	400 to 500 kHz
Turns ratio	Set by topology
Winding technology	Multilayer PCB
Maximum module height	12 mm
Ambient	25°C, natural convection

The brick mounts as a horizontal card, so the module height is set by the board-to-board spacing. Standard brick envelopes land in the same place: an open-frame eighth brick is specified at around 12.2-mm installed height, a quarter brick with baseplate at 12.7 mm. We designed to 12 mm to keep some margin.

There is room left between the core and the chassis, and in a real product we would use it. An aluminium plate on top of the core spreads the heat, and a thermal gap filler between that plate and the chassis carries it out. We would expect that to be worth 20°C to 30°C, although we have not simulated it. Every temperature in this article is natural convection with no plate and no conduction path to the chassis, so they are worst case numbers rather than the numbers a cooled assembly would give.

Two Topologies

For a 48-V intermediate bus converter the fixed-ratio LLC is the common choice, and we started there. The problem is the gap. An LLC needs the magnetizing inductance set to a particular value, and a value low enough to build the resonant tank around means gapping the core. In a planar transformer that gap ends up a fraction of a millimeter from the copper, and the fringing field around it goes straight through the layers next to it.

Another option for this converter is the PSFB (phase shifted full bridge). Like the LLC, this is an isolated topology that can be implemented with a similar parts count to the LLC. It also offers reasonably good efficiency.

This does not mean an LLC cannot be built in planar (it can, and in some applications it is still the better option). For this design it was not, and in the discussion that follows we'll compare the numbers for an LLC versus a PSFB.

From Requirements To An Initial Design

The tool is organized as five tabs, and the design is built by going through them in order: Core, PCB, Windings, Operating conditions and Simulation. The rest of this section follows that order. If you have never opened this tool, there are short video walkthroughs that cover the same ground.[1,2,3]

Core

Core selection is the first decision, and it is mechanical before it is magnetic. The core has to fit under 12 mm together with the board and the clearance, leave enough window for the copper, and keep the flux density down at 400 to 500 kHz.

© 2026 How2Power. All rights reserved.

Everything in this part is done on the Core tab: pick the family and the part from the library, choose the material, and set the gap. Individual dimensions can be overridden when the real part is machined rather than a catalog item, which is what we did for the EQ 30.

Both candidates are EQ cores, but they are closed in different ways, which also gave us a chance to compare the two approaches. Closing with a flat I plate (PLT) rather than a second full EQ half saves height, and that is what we did on the EQ 25. On the EQ 30 we used two machined halves instead.

Table 2. Comparing two cores for the LLC.

	EQ 25/5.6/18 + I 25/2.3/18	EQ 30/8/20, cut halves
Footprint	25 × 18 mm	30 × 20 mm
Assembled height	7.9 mm	10.7 mm
Ae	89.7 mm²	108 mm²
Le	26.4 mm	31.5 mm
Ve	2370 mm³	3400 mm³

The EQ 25 is the standard part closed with an I 25/2.3/18 plate. The EQ 30 is two machined halves with no plate, cut so the pair adds up to the same total height as an EQ 30 with a plate would. An EQ 30 with a plate might improve things, but that would need to be studied.

The EQ 25 Did Not Work For The LLC

We did not start with two cores. Both the LLC and PSFB designs began on the EQ 25, because it is the smaller part and it fits the height with room to spare.

However, the LLC did not survive there. The first stack-up we tried was 4:1 in sixteen layers with 3-oz on the primary and a 0.15-mm gap, and at 40 V in it came out at 138°C on the winding and 119°C on the core, with 2.74 W in the copper and 1.22 W in the core. Nothing we changed in the winding brought that down far enough (Fig. 1).

The core was too small for the flux the LLC puts through it at low line, and the core loss says so: 1.22 W here against 0.41 W once we moved to the EQ 30, on the same material.

So the LLC moved to the EQ 30 and the PSFB stayed on the EQ 25. This is the first place where the tool changed a decision instead of confirming one, and it took about six minutes to find out.

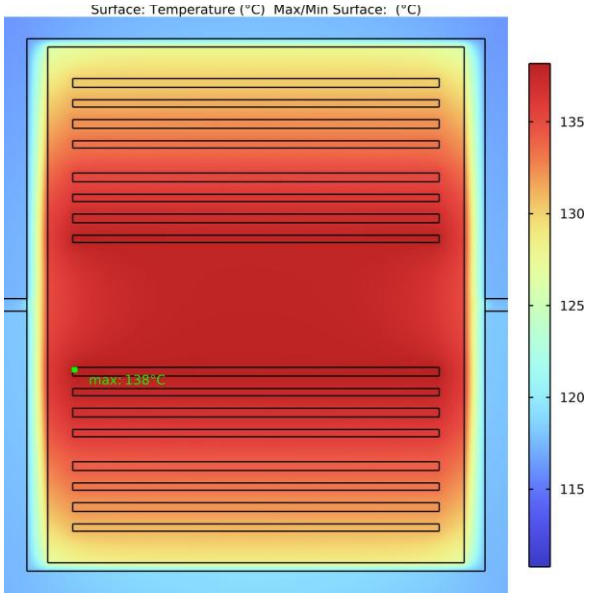

Fig. 1. The first LLC attempt, on the EQ 25 at 40 V in. The winding reaches 138°C and the whole window is hot, because the core is too small for the flux at low line. No change to the winding brought this down far enough.

© 2026 How2Power. All rights reserved.

Both fit under 12 mm, so height on its own does not choose between them. Temperature does.

The PSFB transformer runs ungapped, which is one of the advantages of the topology: the magnetizing inductance only has to be high enough, not a specific value, so we can leave the core closed and keep fringing flux away from the copper.

The LLC needs a gap, and where we put it matters as much as how big it is.

The obvious move is to gap the center leg only and push the winding sideways, away from it. We tried that first, and it works less well than it looks. The fringing field reaches around the copper and the layer ends nearest the center leg still take the worst of it, close to 95 A/mm² against about 15 A/mm² along the rest of the same layers.

Seeing that sent us the other way: split the gap, and separate the winding vertically instead of horizontally. Instead of 0.35 mm in the center leg we used 0.2 mm in all three legs, which gives the same magnetizing inductance, and we split the stack into two groups of eight layers with 1.5 mm of space in the middle, aligned with the gap.

Fringing spreads out roughly in proportion to the length of the gap it comes from, so three small gaps put less of it into the copper than one large one, and the vertical separation puts the copper where the field is weakest. The peak comes down to about 25 A/mm². Machined halves make the three-leg gap easy to do (Fig. 2).

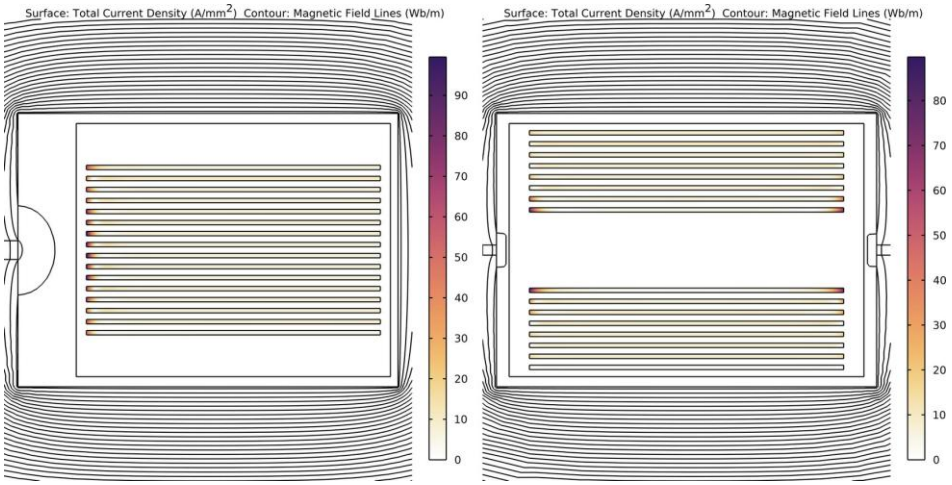

Fig. 2. Two ways of dealing with the gap, both LLC at 48 V in. On the left a single gap in the center leg with the winding pushed horizontally away from it: the ends of the layers nearest the center leg still reach about 95 A/mm². On the right the gap split 0.2 mm across all three legs, with the stack separated vertically into two groups: the peak drops to about 25 A/mm². Moving copper sideways is not enough on its own.

Material

Both designs use 3F36. It is aimed at 400 to 500 kHz, the losses are where we need them in that range, and it is standard enough to be easy to source. A TDK equivalent grade would work as well.

Turns Ratio

The turns ratio is set by the topology and it is not the same for both. The PSFB is 3:1. A 4:1 ratio does not work here, because at the bottom of the input range 40 V in leaves 10 V on the secondary before the phase shift and there is no way back up to 12 V from there. With 3:1 there is margin at 40 V and the phase shift covers the rest of the range. The LLC is 4:1.

Winding Arrangement

On the Windings tab you add one winding per port, then set N. turns, N. turns per layer and N. parallels for each. The number of layers, the trace width and the current density are recalculated as you type, and the tool flags the current density when it gets high (Fig. 3). Copper weight goes in as oz or mm, layer-to-layer spacing can be applied uniformly or layer by layer, and Customize Layer Order is where the interleaving is set.

© 2026 How2Power. All rights reserved.

The PSFB winding is twelve copper layers on a single board, 2.5 oz each, with 0.12 mm between layers:

- Primary: Three turns, one turn per layer, three layers per parallel, two parallels. Six layers, 10 A total.
- Secondary: One turn, six parallels. Six layers, 30 A total.
- Fully interleaved, so every primary layer has a secondary layer next to it.
- The stack occupies about 2.6 mm of the window.

At 400 kHz the skin depth in copper is about 115 µm at the 70°C the simulation assumes, so 2.5 oz (87 µm) sits just under one skin depth. There is no point going thicker.

The current waveform, the number of harmonics, the ambient temperature and the cooling condition go in on the Operating conditions tab. Then Run Complete Simulation on the Simulation tab returns everything at once.

Full interleaving is what gives 3.45 nH of leakage inductance referred to the primary, and 194 pF of open-circuit primary capacitance is what we pay for it. With 41.8 µH of primary inductance that puts the self-resonance at 1.77 MHz, about four times the switching frequency, which is far enough away here.

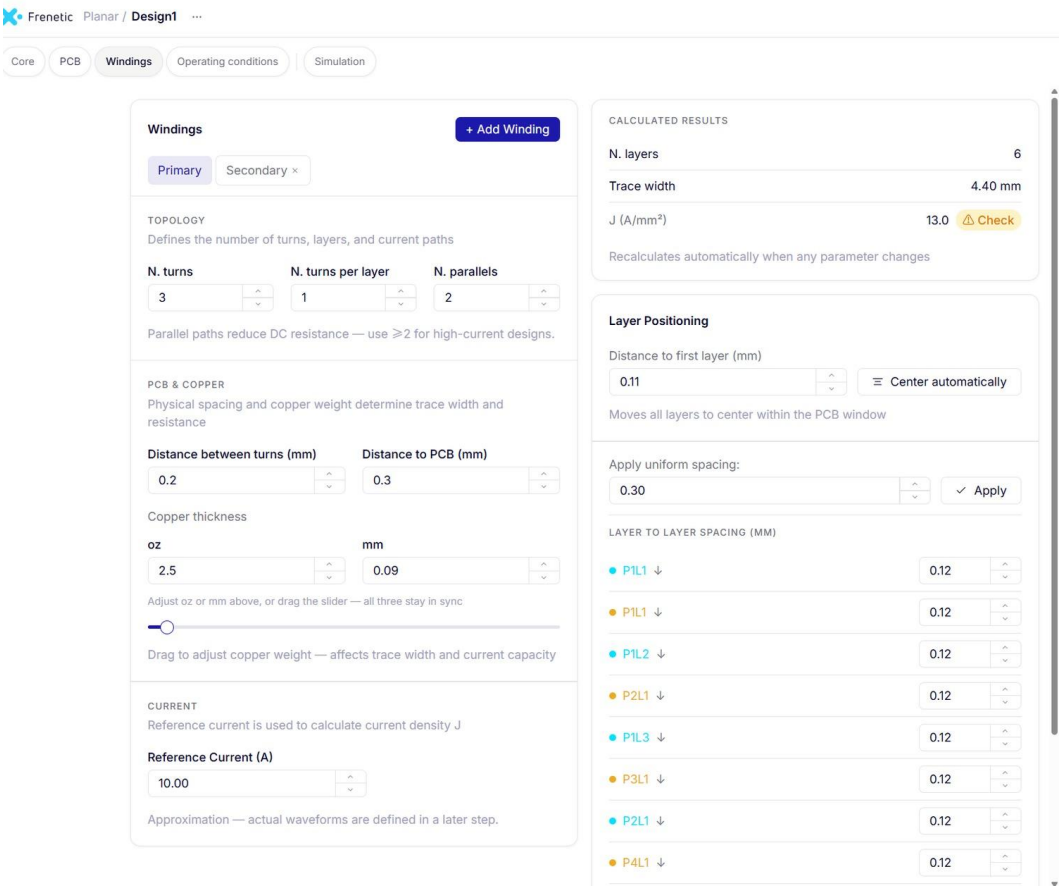

Fig. 3. The Windings tab. Turns, turns per layer and parallel paths at the top left, copper weight in oz or mm below it, layer-to-layer spacing on the right, and the resulting layer count, trace width and current density recalculated as you type.

Comparing The Two Designs

The Thermal Model

Until now the tool gave us losses, leakage and capacitance, and we estimated temperature separately from that. The new thermal model closes the loop: it takes the loss distribution layer by layer, the geometry of the stack, the core, the board and the boundary condition, and solves for temperature.

Three things about it matter for a design like this one. It reports the winding hotspot and the core hotspot separately, and they are neither the same number nor in the same place; in the PSFB at 60 V in the winding

© 2026 How2Power. All rights reserved.

reaches 85.5°C and the core 73.6°C, so a single figure would hide 12°C. It treats the board as a body that conducts, which matters most in planar, because the dielectric and the copper planes spread heat sideways and into the core. And it works from the harmonic content of the real current waveform rather than from an RMS value, with ten harmonics in every run here.

Setting up a thermal run costs almost nothing on top of the loss calculation. The inputs are the same ones the losses already need—the geometry, the stack-up and the waveform, plus the ambient temperature and whether there is forced air. Everything else comes from the material definitions. One press of Run Complete Simulation returns core loss, winding loss per winding, the leakage matrix, the inductances, the parasitic capacitance and the temperature, and each result has a toggle between the numbers and the 2D map (Fig. 4).

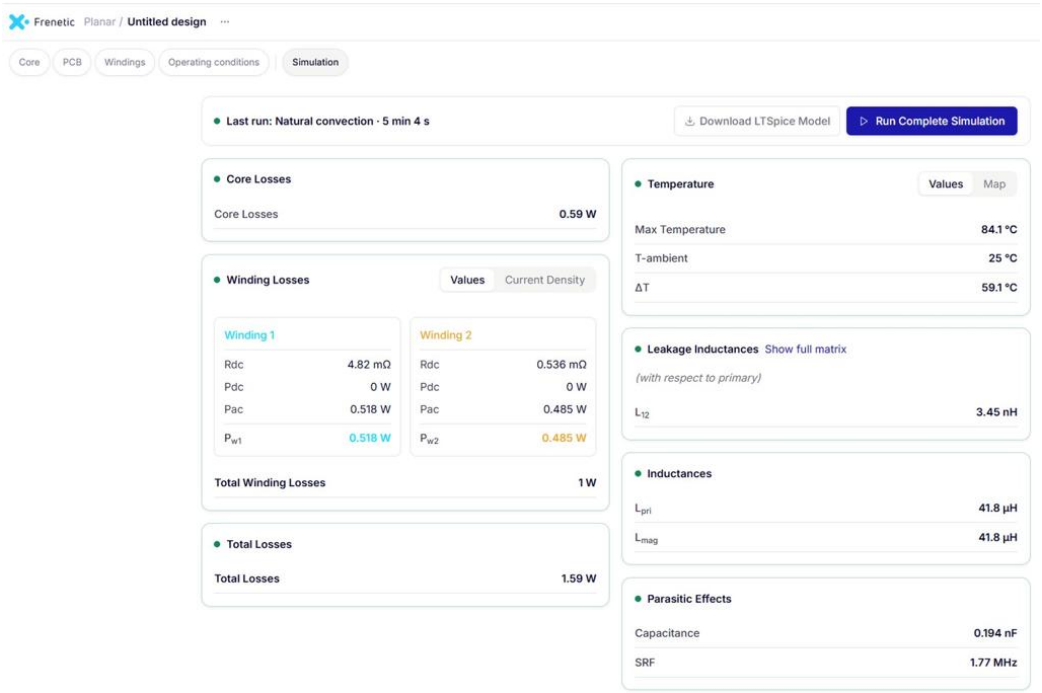

Fig. 4. The Simulation tab after a run. Core loss and winding loss per winding on the left, temperature, leakage inductances, inductances and parasitic capacitance on the right. Losses and temperature each toggle between the numbers and the 2D map.

Across The Input Range

Simulation of losses and temperature were run for the two transformer designs across the input voltage range. Conditions include three input points, natural convection, 25°C ambient, no conduction path to the chassis. The PSFB is the EQ 25 ungapped, while the LLC is the EQ 30 with 0.2 mm in each of the three legs.

Table 3. Loss and hot spot calculations for the PSFB and LLC transformer designs.

Design	Vin	f (kHz)	P wind.	P core	P total	Winding hotspot	Core hotspot
PSFB	40 V	400	0.95 W	0.48 W	1.44 W	78.4 °C	67.6 °C
PSFB	48 V	400	0.99 W	0.53 W	1.51 W	81.7 °C	70.3 °C
PSFB	60 V	400	1.03 W	0.59 W	1.62 W	85.5 °C	73.6 °C
LLC	40 V	285	2.43 W	0.41 W	2.84 W	108.3 °C	89.9 °C
LLC	48 V	369	0.98 W	0.49 W	1.47 W	71.3 °C	64.6 °C
LLC	60 V	532	0.76 W	0.43 W	1.18 W	62.6 °C	57.7 °C

The PSFB Is Flat, The LLC Is Not

The PSFB winding moves 7°C across the whole input range. It switches at a fixed 400 kHz and the only thing that changes with input voltage is the duty cycle, so nothing dramatic happens anywhere in the range.

© 2026 How2Power. All rights reserved.

The LLC moves 46°C. All of it is copper: winding loss goes up by a factor of three from high line to low line while core loss stays between 0.41 W and 0.49 W. The core is not the problem.

Two things move at low line. The RMS current in the windings goes up, and the fringing field from the gap is there on top of it. We have not separated the two contributions, but the current density maps show where the loss ends up, and the gap is visible even at the nominal point (Fig. 5).

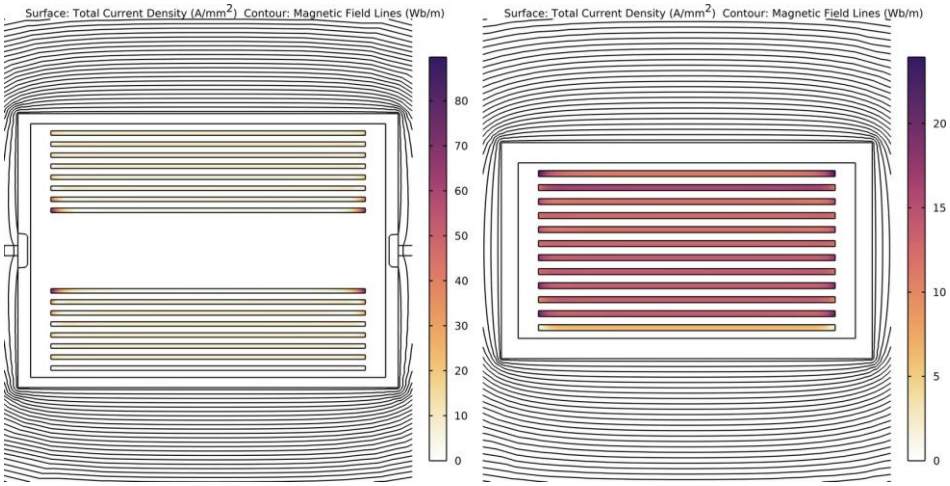

Fig. 5. Current density at 48 V in, the same operating point for both transformers. The LLC on the left is the split-gap design, the PSFB is on the right. In the LLC the two layers facing the split in the stack carry several times the current density of the rest of the winding, which is the fringing field from the gap. The PSFB is even across all twelve layers and only rises at the layer ends, which is a termination effect.

One detail worth pointing out: core loss across the LLC range is not monotonic, and the maximum sits at the nominal point rather than at either extreme. Lower frequency raises the flux, but the material loses less per unit of flux, and the two curves cross in the middle of the range.

Why One Operating Point Is Not Enough

Look only at 48 V in and the LLC comes out ahead, 71.3°C against 81.7°C on slightly lower total loss. That is where the converter spends most of its life, so it is a tempting place to stop.

It is also the wrong place to decide. At 48 V the LLC sits at resonance, which is its best case. Move to the bottom of the input range and it goes to 108.3 °C, while the PSFB has barely moved.

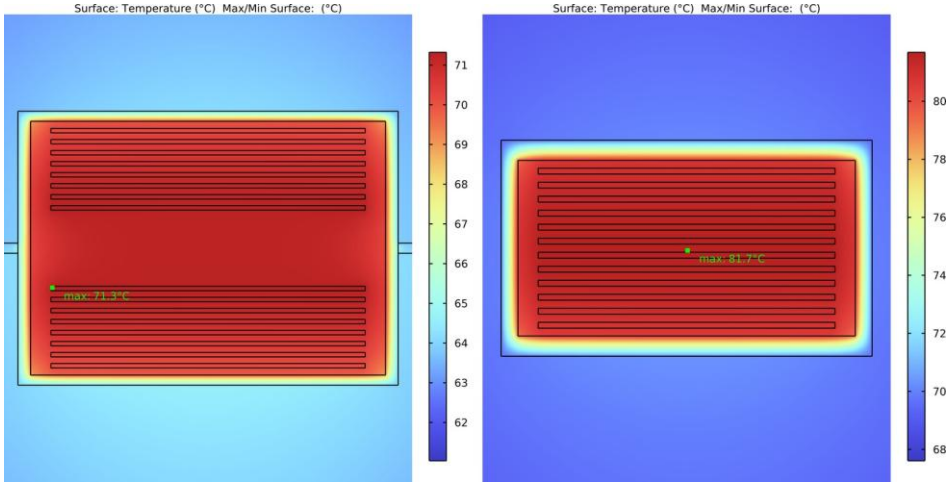

Fig. 6. Temperature at 48 V in, the nominal point. LLC is on the left at 71.3°C, PSFB is on the right at 81.7 °C.

© 2026 How2Power. All rights reserved.

Final Choice And Lessons

Table 4 lists the specifications for the final design of the PSFB on the EQ 25. Fig. 7 shows the worst-case temperature simulatiuons that determined the choice of the PSFB design.

Table 4. PSFB specifications.

Parameter	Value
Core	EQ 25/5.6/18 with I 25/2.3/18, 3F36, no gap
Assembled core height	7.9 mm
Footprint	25 × 18 mm
Turns ratio	3:1
Winding	12 layers, 2.5 oz, fully interleaved, 0.12-mm spacing
Primary inductance	41.8 µH
Leakage to primary	3.45 nH
Open-circuit primary capacitance	194 pF
Self-resonant frequency	1.77 MHz
Rdc primary / secondary	4.8 mΩ / 0.53 mΩ
Worst case total loss	1.62 W at 60 V in
Worst case winding hotspot	85.5°C at 60 V in, 25°C ambient
Loss as fraction of output	0.46%

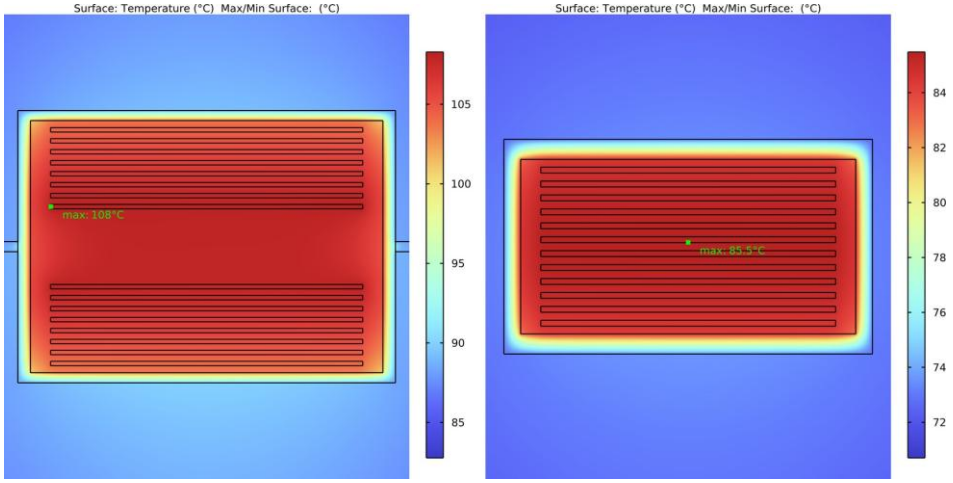

Fig. 7. Each transformer at its own worst point. LLC at 40 V in on the left, 108.3°C. PSFB at 60 V in on the right, 85.5 °C. This is the comparison that decided the design.

The core assembly is 7.9 mm, so 4.1 mm is left under the 12-mm limit. Part of that goes to the board the transformer mounts on. The rest is useful: an aluminium plate on top of the core to spread the heat, and a thermal gap filler closing the distance to the chassis. That is where the 20°C to 30°C we expect to gain would come from, and it is the obvious next thing to simulate. The 85.5°C above is with none of it.

What Is Left To Do

The design is settled, the layout is not, and a few things about the layout are decided by the winding arrangement rather than after it.

The layer distribution drives the layout. Twelve layers fully interleaved is the right answer electrically, but it constrains how the turns route and where the vias go. The stack can also be split across several boards in parallel, which makes each board simpler at the cost of a step in assembly.

© 2026 How2Power. All rights reserved.

The copper weight is not a problem, since 2.5 oz is standard and any board house will quote it. Insulation is not a problem either: we kept minimum spacings and at these voltages that is enough, because 48 V in and 12 V out is a low-voltage application and the prepreg between layers does the job on its own, with no barrier or margin tape.

Terminations depend on the current. On the secondary, at 30 A, blade terminals inserted into the board make sense. On the primary, at 10 A, ordinary pins are enough. The board can be glued to the core.

Lessons

Simulate the range, not the operating point. At 48 V nominal the LLC is 10°C cooler than the PSFB and would have been the obvious choice. At 40 V it is at 108°C. One operating point tells you very little about a converter whose frequency moves.

Total loss does not tell you the temperature. The winding hotspot and the core hotspot differ by 12°C here, and they are not in the same place. Two designs with the same total loss do not run at the same temperature.

A gap is expensive in a planar transformer. The fringing field around it goes through the copper sitting next to it. If the topology forces a gap, split it across the legs and keep copper away from it. If the topology does not force a gap, that is worth something on its own.

There is no point making the copper thicker than about one skin depth. At 400 kHz that is roughly 115 µm at operating temperature, so 2.5 oz is already at the limit. Go past it and the current stops using the middle of the conductor, so the copper carries less than its cross section suggests while you pay for all of it, in money and in height.

Start from the mechanical envelope. The height limit picked the core family before any magnetic argument did.

The whole comparison cost six simulations. Six runs at about six minutes each, so an afternoon of work. The alternative is finding out in hardware: order the boards, wait for them, wind and assemble the part, put it on the bench, and discover at low line that it runs more than 20°C hotter than you expected. Then order the other one and do all of it again. That is weeks of schedule and two hardware iterations paid for, against an afternoon of simulation.

References

1. "Designing a USB Type-C Flyback Planar with Frenetic," Frenetic video, YouTube, May 14, 2026.
2. "Frenetic Planar Tool: From Conventional to Planar Transformer" by Marcos Luna, Frenetic video, YouTube, May 8, 2026.
3. "Two LLC 300W designs, 1 standard VS 1 planar—which one wins?" by Sotiris Zorbas and Julia Mimigiani, Frenetic video, YouTube, March 12, 2026.

About The Author

Marcos Luna Fernández is power electronics engineering manager at Frenetic. He has more than seven years of experience in power electronics, including four years at Frenetic focused on the design, simulation, and validation of high-frequency magnetic components. Marcos works with companies across industries such as automotive, aerospace and defense, industrial manufacturing, semiconductors, and medical devices. His background also includes technology consulting, research and development, project management, and engineering team leadership.

For more on magnetics design, see these How2Power Design Guide search results.

© 2026 How2Power. All rights reserved.