

My name is Jeff Jull, and I have worked on CPU power delivery for almost 25 years. For a decade, I was responsible for enabling the Intel IMVP power delivery ecosystem, which powered over 300 million new PCs each year. For the last 6 years, I have fine-tuned that expertise within MPS to create leading-edge power management products in the PC market.

The PC power delivery segment is distinctive in the overall power delivery market, in that each supplier's solution is a unique and complex chipset with its own custom features and pinout (or footprint). The chipset controller communicates with the CPU, then executes complex control of the voltage, which is supplied by smart power stages (SPS) or driver + MOSFET power stages (DrMOS), to the CPU (see Figure 1). In addition, the chipset provides an assortment of telemetry information.

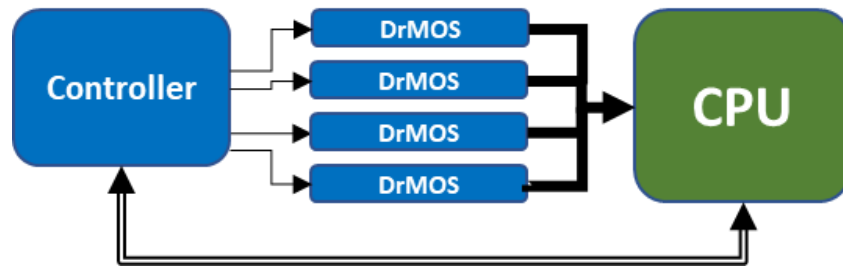


Figure 1: Chipset Controller Communication

The chipsets are also unique in that each customer platform requires labor-intensive tuning to maximize performance and efficiency — both of which affect the battery life, which makes these factors critical in the laptop PC market. The diversity of customer needs and solutions has created opportunities for each power delivery vendor to thrive, since each solution is fit for an application that meets the customer's specifications.

More importantly, each novel solution funds the development of the next generation's solutions within the PC power management IC market. A fundamental rule in the PC industry is, "You pay for performance." With this in mind, a few vendors have focused on the high-performance and higher-priced commercial PC market, which focuses on B2B solutions. Other vendors create more cost-effective solutions used in the consumer market.

The worldwide silicon shortage seen in recent years has impacted the PC power delivery market like any other market, but there is one significant difference. The continuity of supply (COS) issues in the CPU power delivery market have prompted PC manufacturers to create multi-source solutions in a market where sole-source solutions were the norm. The seemingly obvious resolution to the COS issue is the implementation of a common footprint (CFP) strategy.

As part of the CFP strategy, each CPU power delivery solution — consisting of 1 controller and 3 to 9 DrMOS components — would use the same pinout for each type of device. In theory, this would allow for parts used in the manufacturing process to be easily replaced, as any customer could mix and match available parts from vendors. The COS would be maintained based on the assumption that every CPU power delivery vendor would produce a certain number of parts for the market.

The CFP model is supported by its success in the server CPU power delivery market. Although server applications differ in size and capabilities compared to PCs, there are still common features (e.g. a controller and several DrMOS components) that make it simple to assume that models in the server market are applicable to the PC market. However, there are critical differences between these markets that mean the CFP model could negatively impact the PC market. Specifically, CFP strategies can hurt reliability, performance, battery life, and the very thing CFP wishes to improve most: supply continuity.

CFP in the Server CPU Power Delivery Ecosystem

The server CPU power delivery ecosystem has several elements that make CFP an effective solution for COS:

- Long development schedules
- Long product lifecycles
- High component costs (driven by current capability, physical size, and reliability requirements), which prevent race-to-the-bottom price wars

These elements are described in greater detail below.

Long Development Cycles

The product development cycle of a server platform can take years. This relatively long process is required due to the rigorous validation required to achieve the high reliability targets demanded in a quality server. Tests are run on the CPU power delivery system with every possible variation for input power, output voltage, and load transition. Guaranteeing reliability is a thorough process that cannot be shortened.

Because this validation rigor is performed across the entire platform, there is ample time to validate more than one CPU power delivery solution, which makes a CFP model possible. Without the common footprint, multiple cost-prohibitive boards would need to be designed and produced to validate other solutions. Instead, because the server power management market values high reliability and long development cycles, the CFP strategy has the potential to keep the market healthy while driving innovations for the best possible products (see Figure 2).




Figure 2: Server CPU Ecosystem

Long Product Lifecycle

The product features already discussed (e.g. high reliability and high cost), plus a relatively small total available market (TAM) (compared to the PC market), means that original design manufacturers (ODMs) can and must sell each server product for a very long time to make each product design profitable. The typical server manufacturing lifecycle is about 5 years, whereas a PC design is manufactured for only 18 months. A server's long product lifecycle demands COS for its CPU power delivery, so CFP is a promising strategy to provide assurance.

In a typical server CPU power delivery CFP implementation, an ODM first chooses the suppliers it plans to use. Then the ODM provides assurances regarding the sales volumes of the product lifespan, as well as the expected total sales. With this information, the supplier can build a product to meet the customer's demand, while simultaneously building additional inventory for the first few years. This inventory provides a buffer for the customer in case one of the other suppliers has an issue with their CFP parts. If this inventory is not consumed over the initial years, the supplier stops production and uses its inventory to fulfill demand for the later years of the product lifecycle. Across a 5-year product lifecycle, CFP is a necessary and beneficial strategy.

High Component Costs

Servers are an expensive, and therefore lucrative, market. Everything about a server platform is expensive, from the final product down to the CPU, motherboard, cooling solutions, and power delivery. These high costs are driven by high performance and 99.99% reliability requirements. This requirements also increase the physical size of the critical power delivery components, which further increases costs.

For server implementations, the impact of all of these cost drivers is that only the best and most reliable suppliers and components are used, which greatly limits the competition. Supplier reliability is determined by experience and reputation in an ecosystem where mishaps cannot be kept secret.

Over time, this market has determined a hierarchy of supplier quality and reliability. Certain suppliers are known to always provide a steady supply of parts, while others have a reputation for pulling out of markets suddenly if profit margins drop. Still, other suppliers have a reputation for cheap parts of questionable reliability, which disqualifies them from the general server market. In this limited field of reputable competitors, a CFP strategy has little impact on the high prices of power delivery components.

CFP in the PC CPU Power Delivery Ecosystem

The PC power management market will first be examined for the same three elements that make CFP so successful in the server market: development lifecycles, product lifecycles, and component costs (see Figure 3).



Figure 3: PC CPU Market

Short Development Cycles

The PC market comes out with a new platform generation every year or so. The schedule is blazing fast, and it forces OEMs and their suppliers to generate new products at a seemingly unsustainable pace. With virtually every generation of PC, there is a new power requirement for the CPU, forcing the development of new controllers and DrMOS components.

When this short turnaround is combined with solutions' increasing current capability and the constant push to reduce the size of the motherboard and power delivery, it is fair to say that CPU power delivery suppliers and their customers are constantly busy validating a single solution. A CFP strategy means that there are additional solutions to validate, along with the mixed/matched validation if components from different suppliers are used in one solution. The only way to effectively execute a CFP strategy would be to cut corners on all of the validation efforts, thus leaving all of the solutions at risk due to unknown violations in production. The short development cycles of the PC ecosystem leave no time for a CFP model.

Short Product Lifecycle

As already mentioned, the PC market churns out new products every 12 to 18 months. In most cases, this is also the production lifecycle for the OEM PC factory, as they often change for the next generation of products. The goal of a CFP strategy is to maintain COS during those 18 months of production.

In the beginning of the server market CFP cycle, the supplier can build inventory in the earlier years and sell it off in the later years, but this is not possible in the PC market. The production volumes being demanded ramp up and down very quickly, so there is no time to build inventory as it is all being consumed. In addition, there is no guarantee that inventory will be sold off. This leaves the supplier with parts that they must sell; otherwise, they risk losing money. The result is that there is no ready supply of parts available from any of the other suppliers. Additionally, processing new orders typically takes 6 months, which is roughly one-third of the entire product lifecycle. The short PC ecosystem lifecycle does not provide enough time for CFP to accomplish the stated goal of COS.

Mixed Price PC Power Delivery Component Costs

Where the server market has a high-priced (\$\$\$\$) market for power delivery components, the PC market has a mix of both low (\$) and higher-priced (\$\$) parts. Just like high- and low-priced PCs, you always pay for performance, meaning there are plenty of options within different prices ranges. The two pricing brackets in the PC market are known as “commercial” (e.g. speedy and reliable PCs for businesses, or the PC you receive from the IT department) and “consumer” (e.g. what you buy on Black Friday at a large store).

The commercial segment uses higher-priced PC CPU power deliver components provided by the silicon industry and technology leaders in the market. These devices are made by the tier 1 companies that add novel features to their products — features that customers never knew they “just had to have.”

These suppliers invest in high-quality field engineers who spend an inordinate amount of time with every PC design-win. Engineers in this segment program and tune the power delivery components to achieve the best possible performance for every supported PC. This robust and growing PC ecosystem provides better user experience and advances the next generation of technology with the following symbiotic relationship:

- The PC manufacturer’s commitment to use a customized, unique solution for power delivery
- The supplier’s commitment to provide the best technology and world-class support

For consumer PC power delivery, suppliers develop parts that only meet the CPU requirements, and special tuning is not required because the part simply has to work. Factors such as performance and battery life are not considered at this level. To make the power delivery solution as cheap as possible, some PC manufacturers will remove capacitors one at a time until the PC stops running, then put one back. These CPU power components are inexpensive, but they only meet the bare minimum.

In the commercial PC power management market, the CFP strategy would open the door for parts that are simply “good enough” to enter the commercial market and would eliminate the need for commitments between the PC manufacturer and supplier. The end of this symbiotic relationship changes the ecosystem to a race-to-the-bottom price, which stagnates the advancement of PC technologies regarding power delivery performance and battery life.

Proponents of the CFP strategy in the PC market are trying to allow for new technology, which can result in higher prices due to technological demands. However, the recent silicon constraint has thrust unlimited authority on the PC manufacturer’s procurement teams to get components by any means necessary to ensure that PCs stay in production. These teams almost always select components with the lowest price possible. When silicon becomes more available in the second half of 2022 and entering 2023, the “pay anything to get parts” mentality will quickly revert back to “get the cheapest parts.” For PCs, the CFP strategy feeds the race-to-the-bottom model.

The Final Conundrum

There is a very recent example of a CFP shortage in the notebook PC market, where each platform requires a 3.3V DC/DC converter. There are at least five companies that supply the more than 200 million required pieces each year. They all have the same form, fit, and function properties, so a notebook design can use the part from any of these suppliers. Recently, the low-cost (and consequently) market share leader abruptly stopped producing their 3.3V converter — presumably in favor of high-margin products.

This major disruption in the COS should not be an issue under the CFP model, since there are at least four other sources for the part. However, because of the industry constraint on silicon, none of the other suppliers carried the inventory needed to fill the gap created by the lead supplier leaving the market. New inventory to meet the demand would take 6 months to produce. Additionally, this 3.3V converter has been in use for almost 5 years, and it will continue to be used for at least the next 2 years (if not longer). With this in mind, the suppliers could have built up a significant inventory if they had been guaranteed total sales at the beginning of the lifecycle. By any measure, a viable CFP strategy should have prevented such a large gap in the ecosystem.

The missing piece that resulted in the 3.3V converter shortage is the commitment-based relationship between the manufacturer and the supplier. Because there were so many suppliers, the PC manufacturers assumed that there would always be a supply from somewhere. On the other hand, each supplier assumed that other suppliers would keep the ecosystem stable, so they cut back on manufacturing the low-margin 3.3V converter in favor of more profitable opportunities.

The final conundrum with implementing a CFP strategy in the PC power management market is that the number of suppliers will decrease over time, which damages COS rather than ensuring its integrity. When the inexpensive, consumer CPU power components are tested in the commercial market, the empowered procurement management teams will compare the prices between suppliers and determine that they are content ordering parts that are “good enough.” Then the ASPs will fall and force the technology leaders (roughly 35% of the suppliers) out of the market.

Conclusion

The silicon constraint of the last 2 years has hit the industry and the world very hard. All components in the PC world, as well as electronics everywhere, have been affected. It is understandable that the highest levels of management are driving for COS. These stakeholders want all of their components to be CFP compliant, as they erroneously believe it should increase the number of suppliers and (presumably) the number of available parts.

This analysis of the unique PC CPU power chipset ecosystem determined that the two subsystems of the PC power management market (commercial and consumer) do not fit in the CFP model. In fact, applying a CFP strategy to this market will more than likely create a market that settles for the bare minimum, while risking power loss and shorter battery lives, which will negatively impact user experience.

Rather than relying on the old paradigm of CFP, the PC power delivery market needs to evolve and adopt a new model for COS. One such approach is for the existing suppliers to license their IP, which would enable new vendors to manufacture that suppliers’ parts. This would provide a new supply source of parts identical to the unique design adopted by the tier 1 suppliers, and would not require any additional validation. The PC market is driven by revolutionary changes, and it is time the supply chain changes with it.