AMD and Sony’s PS6 chipset aims to rethink the current graphics pipeline

It feels like it was just yesterday that Sony hardware architect Mark Cerny was first teasing Sony’s “PS4 successor” and its “enhanced ray-tracing capabilities” powered by new AMD chips. Now that we’re nearly five full years into the PS5 era, it’s time for Sony and AMD to start teasing the new chips that will power what Cerny calls “a future console in a few years’ time.”

In a quick nine-minute video posted Thursday, Cerny sat down with Jack Huynh, the senior VP and general manager of AMD’s Computing and Graphics Group, to talk about “Project Amethyst,” a co-engineering effort between both companies that was also teased back in July. And while that Project Amethyst hardware currently only exists in the form of a simulation, Cerny said that the “results are quite promising” for a project that’s still in the “early days.”

Mo’ ML, fewer problems?

Project Amethyst is focused on going beyond traditional rasterization techniques that don’t scale well when you try to “brute force that with raw power alone,” Huynh said in the video. Instead, the new architecture is focused on more efficient running of the kinds of machine-learning-based neural networks behind AMD’s FSR upscaling technology and Sony’s similar PSSR system.

> *From the same source. Two branches. One vision.*

> *My good friend and fellow gamer @cerny and I recently reflected on our shared journey — symbolized by these two pieces of amethyst, split from the same stone.*
>
> *Project Amethyst is a co-engineering effort between @PlayStation and… pic.twitter.com/De9HWV3Ub2*
>
> *— Jack Huynh (@JackMHuynh) July 1, 2025*

While that kind of upscaling currently helps let GPUs pump out 4K graphics in real time, Cerny said that the "nature of the GPU fights us here," requiring calculations to be broken up into subproblems to be handled in a somewhat inefficient parallel process by the GPU's individual compute units.

To get around this issue, Project Amethyst uses "neural arrays" that let compute units share data and process problems like a "single focused AI engine," Cerny said. While the entire GPU won't be connected in this manner, connecting small sets of compute units like this allows for more scalable shader engines that can "process a large chunk of the screen in one go," Cerny said. That means Project Amethyst will let "more and more of what you see on screen… be touched or enhanced by ML," Huynh added.

https://arstechnica.com/gaming/2025/10/amd-and-sony-tease-new-chip-architecture-ahead-of-playstation-6/