



The Quest for Power

Mar 1, 2003 12:00 PM, by Ned Mann

We all want more power. Admit it. When it comes to the gear we all lust after — whether it's a Massenburg EQ or the hottest new plug-in — it never seems to be enough. The same can be said for raw DSP power: As soon as we finish installing the latest, greatest CPU, software developers figure out something to overwhelm it.

The concept that a single computer will eventually run everything in your entire studio is slowly evolving into a growing awareness that the tasks will need to be divided up between several CPUs dedicated to specific tasks. Before you reach for the Advil, consider that most studios today are already dependent on multiple CPUs. From console automation to hardware synths to reverbs, studios are actually running dozens of CPUs — we just don't think of it that way.

In my previous article (October 2002), I detailed some ins and outs of setting up a high-end DAW and the awesome power that can be harnessed on a single CPU. Correctly configured, the single CPU can handle most users' needs. However, I also touched on some limitations (conflicts between different programs and the eventual taxing of even the fastest CPU) that users run into when trying to expand systems to illogical extremes.

For users who want unlimited power, native sound sources and instant recall, the multiple-CPU scenario is the current state of the art. Multi-CPU setups take many forms, but, in general, break down into two basic groups: users who want to isolate sampler responsibilities and those who want to integrate additional sequencer environments.

The first scenario (an additional CPU to handle virtual synths), though uncommon even a few years ago, is now becoming a *de facto* standard. GigaStudio has ushered in a world where 1-gig piano samples and libraries that stretch into the teraflops are now commonplace. Having vast libraries and hundreds of voices available in real time all but requires a dedicated CPU, as most serious composers have found out. While GigaStudio was the first to be able to stream samples directly from the hard drive, it is now being joined in the act by Steinberg's Halion, Emagic's ESX 24 and Native Instruments' Kontakt. These soft samplers allow users to access the sought-after Giga libraries directly from within the sequencing environment, which begs the question: Why do I need to dedicate an entire computer to run GigaStudio? If you are running a small number of tracks and want to play a few Giga sounds, then the answer is: You don't. However, consider for a moment what many large productions actually entail. Can we expect one CPU to play back 100-plus audio tracks — each with EQ, compression and effects — and stream 160 voices in real time? I think not.

Many companies are now filling this niche market with custom-built PCs preconfigured to run GigaStudio. Although slightly more expensive than doing it yourself, it takes all of the guesswork out of setting up and configuring a system and provides a true source of tech support. In fact, many film composers who were initially resistant to adding a GigaStudio PC now find that they can't live without the concept of dedicated computer/dedicated task; some even broke down and installed three or four of them! With the advent of VST Shells (such as Steinberg's new V-Stack), PCs can access 16 VST and DirectX plug-ins directly without a host-sequencer application. This will allow CPUs set up to run GigaStudio to double as dedicated soft-synth players, as well. Amazing!

On a simpler level, many users who are upgrading their existing G3 or Pentium III find that the computer that they are replacing can serve many uses. The most obvious would be setting it up as a dedicated Internet portal,