

## HPC User Forum Update

# Interviews with Steering Committee Members: Jeff Broughton, NERSC

Thomas Gerard and Steve Conway  
September 2020

### IN THIS UPDATE

---

After the global pandemic forced Hyperion Research to cancel its 2020 HPC User Forums, we decided to reach out to the HPC community by publishing a series of interviews with select HPC thought leaders. Our hope is that these leaders' perspectives on HPC's past, present and future will be interesting and beneficial to others. To conduct the interviews, Hyperion Research engaged insideHPC Media. We welcome comments and questions addressed to Steve Conway, [sconway@hyperionres.com](mailto:sconway@hyperionres.com) or Earl Joseph, [ejoseph@hyperionres.com](mailto:ejoseph@hyperionres.com).

This interview is with Jeff Broughton, Deputy for Operations at the National Energy Research Scientific Computing Center (NERSC). NERSC is the production scientific computing center for the U.S. Department of Energy's Office of Science. At NERSC, Broughton has responsibility for acquiring, installing and operating all computational, networking, and storage equipment for NERSC and the Joint Genome Institute. During his time at NERSC, his duties have included acting as the project director for the Edison system, supporting the acquisition and installation of the Hopper and Cori systems, and serving as the point person for the construction of Shyh Wang Hall, NERSC's home on the Lawrence Berkeley National Lab's main campus. Jeff has 30 years prior experience in HPC and related fields, including positions at Lawrence Livermore National Laboratory, Amdahl, Sun Microsystems Laboratories, and the startups, Key Research and PathScale. He has tackled projects in multiple disciplines as both an engineer and manager, including networking, computer-aided design, processor design, compilers, and operating systems. Jeff also holds multiple patents. His inventions include optimistic concurrency protocols, distributed cache coherence protocols, domain partitioning mechanisms, and software methods for cycle-based logic simulation.

The HPC User Forum was established in 1999 to promote the health of the global HPC industry and address issues of common concern to users. More than 75 HPC User Forum meetings have been held in the Americas, Europe and the Asia-Pacific region since the organization's founding in 2000.

## JEFF BROUGHTON INTERVIEWED BY DAN OLDS, insideHPC

---

**Olds:** Hello. On behalf of Hyperion Research and InsideHPC, I'm Dan Olds and we're going to be interviewing Jeffrey Broughton today. Let's talk about the course of your career and how you got started in HPC.

**Broughton:** I started back in 1979. I went to work at Lawrence Livermore National Lab on a project called the S1 Project, which was an attempt to develop and design, literally from scratch, a high-performance computing system that at that time would have been in the tens-of-megaflops range.

**Olds:** That's a pretty big number back then.

**Broughton:** Yes, it was a pretty big number back then. The project was kind of everything from soup to nuts. We actually developed the CAD technology to design the machine. We developed a number of architectural innovations, many of which are now seen in traditional processors, although in much more elaborate forms, and we actually physically built the machine as well. My responsibility, originally, was working on the software for it but ultimately, I moved into a management role.

**Olds:** That's a big project for a young man.

**Broughton:** Yes, it was.

**Olds:** So, what happened next?

**Broughton:** So, after about a decade there I went to work in industry, joining a startup called Key Computer, which was also targeting and doing a high-performance processor. I was one of many of the folks from the earlier project who decided it was about time to go into Silicon Valley and make our fortunes. Key lasted as an independent entity about a year-and-a-half and we were acquired by Amdahl Corporation. This was a time when risk was big. We went on to do some internal developments at Amdahl, which ultimately, truthfully, never saw the light of day, as happens with many acquisitions. I stayed at Amdahl for many years, both in a technical role and I actually spent some time on the dark side in marketing as well.

**Olds:** Oh no! Don't say that. That tainted you. At least for a while.

**Broughton:** I'm sorry. I saw the light. I came back to the straight and narrow path, and after that I went to work for Sun Microsystems for about three years. We were developing a massively-parallel simulator whose purpose was specifically to do logic simulation for future generations of Sun Systems, and it was a thing to behold because it was a massively-parallel simulator. We were actually able to get about 80,000-way parallel simulation of the logic which, of course, itself is massively parallel. It was truly a thing to behold. This was a really interesting machine because individual processors had no more than 4,000 instructions, there were just a lot of them. It had a network between it, but it was entirely statically scheduled by the compiler. That was a really interesting project, it had some benefits internally, but I left and went off at that point and became a founder of PathScale.

**Olds:** Oh, yes. I've heard of PathScale.

**Broughton:** PathScale was kind of interesting. We were looking at building technology for high-performance computers. We had somewhat more grandiose plans to cover the entire spectrum but we started by doing compilers and the niche we were looking at was to do compilers for AMD. We brought

in the team and actually adapted the SGI compilers, which at that point were open source and applied that to AMD systems and actually were able to achieve quite good performance on that and also on Intel machines. But the major part of it was that we were also developing a high-performance computing interconnect and we ultimately ended up targeting InfiniBand, so we became the alternative to Mellanox.

**Olds:** That's what I remember.

**Broughton:** It had some secret sauce in it that was designed for extraordinarily low latency. I was there three-and-a-half generations of interconnect until we were finally acquired by QLogic, and I stayed with QLogic for a number of years before deciding to work at NERSC. Actually, it was rather refreshing. Instead of being the vendor getting beaten up by the customers, I could be the customer beating up on the vendors.

**Olds:** Much more fun, isn't it?

**Broughton:** Much more fun in many ways. So, I've been at NERSC now for about 11 years. I've taken on a number of roles. I was originally head of the systems department and completed the installation of the Hopper machine, which was the first petaflop machine at NERSC. Then I ran the procurement of the Edison machine, which was literally serial no. 1 of the Cray XD series. And then later on I took over responsibility for building our new data center on the Lawrence Berkeley National Lab campus and brought up our new data center there and have been working to expand that ever since.

**Olds:** Fantastic track record. So, what are some of the biggest changes you've seen in HPC over the years?

**Broughton:** You know, the interesting thing about HPC is that in many ways it hasn't changed a whole lot. If you look at the people who have been involved with it, people become enamored of high-performance computing, and although some of the people have moved from vendor to customer, from one site to the next, many of the leading people in this have really been in this more-or-less continually, in many cases for decades in one form or another.

**Olds:** It's a close-knit community.

**Broughton:** It's a close-knit community with really much of the same goals associated with it. So, in that sense I don't think it has really changed a whole lot. Certainly, technology has evolved, right? We went from the highly specialized machines into leveraging the commodity infrastructure. In fact, I would say that's probably the single biggest change I've seen, the introduction and leveraging of commodity technology with a little bit of specialized magic here and there to make it actually work.

**Olds:** Do you kind of see us heading back towards the specialist route with so many processors under development now that are specialized for single-purpose?

**Broughton:** I think the commodity-based systems have a lot of life left in them, but I think there is tremendous opportunity for some of the specialized processors that we are seeing, especially in the machine learning area, that will come to the fore.

**Olds:** Both training and inference.

**Broughton:** Yes. I think NERSC's general view is that what we are going to start to see is large supercomputers that have a variety of processing elements in them that are tuned to different things:

GPUs especially for large-scale simulations, commodity processors with lots of I/O that'll be applied to data analytic technology, and then, as I said, the specialized things, either FPGAs, custom processors, or application-specific processors. For some of the more complex, especially the large, experimental projects, having that mix of capabilities within a single tightly-knit computing complex will be very beneficial.

**Olds:** I totally agree and that's kind of the way I see it going as well. So, looking down the road do you have any concerns about where HPC is going? Anything that needs to be avoided, for example?

**Broughton:** Interesting.

**Olds:** Any trends that you think might be wrong directions that people are tempted towards, for example.

**Broughton:** I cannot be a pessimist, so it's hard to say that there is anything that I'm really particularly afraid of. I do think, maybe, that we have come to the end of the usefulness of the Top500 list. The exascale project has accomplished an enormous amount to push us forward, but it's also an enormous investment on the part of the Department of Energy. While I think there are places for singular, very large machines, I think, as I talked about earlier, having things which have more diverse capabilities is very important. To go to the flip side: what really encourages me?

**Olds:** Yes, that's what I want to hear next.

**Broughton:** I think one of the fallout of the exascale process is you really can see the notion of having petaflop machines that sit in the size of one or two conventional racks that'll cost less than a million bucks and that looks a lot like a VAX in some way, right? The ability to deliver that to organizations which may not have been able to own their own hundred-million-dollar supercomputer in the past may end up being very beneficial for science as well. Of course, we get into the areas of, can you actually run a machine like that efficiently without having some skill in how to optimize for energy efficiency? and so forth. That's a different question.

**Olds:** And it's interesting in that I guess some vendors are calling it the, "democratization," of HPC, but the prices have come down so fast for huge amounts of compute that I think there is going to be some sort of a revolution there.

**Broughton:** Well, certainly GPUs have done an enormous amount for that as well. Even small multiway systems with a couple of GPUs. That would've been unbelievable when I started.

**Olds:** Absolutely. It's such a change, right? And on your point about the Top500 list. I would like to see it split into about five different categories like CFD, Monte Carlo, different kinds of simulations and different kinds of compute, memory intensive, I/O intensive, CPU intensive, that sort of thing and let's do a list that way. Does that make sense?

**Broughton:** We have started to evolve the Green500 list as an alternative, but that's still basically an HPL application. I think probably having three or four metrics that target some of the areas, that may well be of particular interest.

**Olds:** We might see an AI list soon.

**Broughton:** Yes.

**Olds:** And there is an HPCG list out there and that's a nice benchmark because it scales and it takes about half-an-hour to run, not depending on the system.

**Broughton:** I don't know much about that myself so I won't comment.

**Olds:** Well, it would run on your laptop. I think they've got a laptop version of it and if they don't Jack Dongarra will spin one up. Well this has been great. Thank you so much for your time, really appreciate talking to you and thank you all out there for watching.

**Broughton:** Okay, thank you very much.

## About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

## Headquarters

365 Summit Avenue

St. Paul, MN 55102

USA

612.812.5798

[www.HyperionResearch.com](http://www.HyperionResearch.com) and [www.hpcuserforum.com](http://www.hpcuserforum.com)

---

## Copyright Notice

Copyright 2020 Hyperion Research LLC. Reproduction is forbidden unless authorized. All rights reserved. Visit [www.HyperionResearch.com](http://www.HyperionResearch.com) or [www.hpcuserforum.com](http://www.hpcuserforum.com) to learn more. Please contact 612.812.5798 and/or email [info@hyperionres.com](mailto:info@hyperionres.com) for information on reprints, additional copies, web rights, or quoting permission.