Recent Interview with PathScale CEO Scott Metcalf

Scott Metcalf
President & CEO

Scott Metcalf is a proven senior executive with a 30-year history in the electronics and computer systems industry. Scott has served in a variety of operational and business management roles, including being the President/CEO of three successful start-up companies. Scott spent 5 years each at Sun Microsystems and WaveTek in senior operational roles before becoming President and CEO of Dynabook Computers, which he sold to Unisys in 1991. Scott joined start-up HaL Computers, which delivered the first 64-bit SPARC microprocessors, compilers and 64 bit Solaris operating system as part of its line of supercomputing workstations and servers. Scott served as President and CEO of HaL before selling the company to Fujitsu in 1993.

Q: What is the primary mission of PathScale?

A: PathScale develops software and hardware solutions that enable users of Linux clusters to achieve higher application performance and better cluster efficiency. We provide core technologies and products that allow clusters to scale to unprecedented levels. Our goal is to help scientific and engineering users solve existing problems quicker and to solve new classes of problems that could not be cost-effectively solved before.

HPC users today are looking for better returns on their investments. They are trying to speed up “time to problem resolution” using the most cost-effective computing solutions available. The vast majority of HPC applications today still reside on large-scale symmetrical multi-processing servers. Not only are these servers expensive and difficult to maintain, but more importantly they don’t scale into the hundreds or thousands of processors. This means individual jobs are typically limited to running on dozens of CPUs. HPC users often have problems that can take an unlimited number of CPU cycles and are limited, literally, only by their budget.

While in theory there is no limit to the scaling of a cluster other than budget, in reality many problems are limited in other ways; by the precision of the processor, by their addressing schemes, by the ability to create and tune the communications between the distributed portions of the application solution, and by the latency of the interconnections between the processors in the cluster. Commodity priced CPU’s, like the highly successful AMD Opteron, solve part of the first problem with their 64-bit architectures. PathScale products will address all of the rest.

Q: Why aren't all HPC applications running on clusters today?

A: There are still a number of inhibitors to ubiquitous cluster adoption. Although low-cost 64-bit cluster nodes now exist, highly optimized 64-bit Linux-based compilers have been lagging. This has held back 64-bit cluster application development. Another significant inhibitor has been low-cost, high-bandwidth, low-latency cluster interconnect. A cost-effective cluster interconnect that scales to thousands of nodes is not available. The other remaining challenge is the cluster programming environment. Writing software programs for clusters that can truly scale to hundreds or thousands of nodes is extremely difficult. The standard programming environment for high-end clusters is called MPI, or message passing interface. Writing efficient MPI programs is time consuming and requires expert skills not commonly available. This is why many people refer to MPI as the "assembly language of clusters". The big challenges come down to: 1) how do I get my applications to cost-effectively scale, and 2) how do I easily develop, deploy and tune MPI-based applications?

PathScale will be solving a number of these problems for HPC users with a combination of software and hardware technologies. Again, our goal is to help scientific and engineering users solve existing problems quicker and to solve new classes of problems that could not be cost-effectively solved before.

Q: So, is it fair to say that PathScale's goal is to provide high performance, affordable scalability for people using Linux clusters?

A: Yes, that's an accurate statement. PathScale enables scientists and engineers who need high performance computing to achieve superior results without having to buy expensive SMPs or supercomputers. For people who are truly doing performance-sensitive applications, large-scale SMP solutions are inadequate. They would rather operate in a cluster environment, and that's why clusters are so popular today. Everyone wants to go to clusters because the cost of computation is about one fifth to one-tenth the cost of a typical large-scale SMP system. This is the inherent economic reason that so many people want to move to clusters. After migrating to clusters, users often unfortunately discover that they have scalability and programming challenges. Low capital expense should not translate into high operating costs, but that is too often the case today. PathScale will help solve this problem.

Q: And was that the vision of PathScale's founders?

A: The vision of our company founders was to reduce the cost of ownership of high performance computing systems and accelerate the rise of cluster computing across all high performance applications. This is not limited to the scientific and engineering HPC segments we're pursuing today, but extends to commercial computing segments such as financial modeling, online transaction processing systems, decision support systems and business intelligence applications.

Our vision is that all high-end computing in the future will be done on clustered architectures due to the inherent cost, scalability and reliability advantages of clusters. PathScale's long-term goal is to be an essential technology supplier of hardware and software technologies that enable the cluster paradigm to dominate high-end computing.

Q: PathScale's first product announcement was a 64-bit Linux compiler suite for AMD Opteron-based systems. What is the status of that product now?

A: We are now making production shipments of the PathScale EKO Compiler Suite to customers around the world. We experienced an immense amount of interest from beta customers for the compilers -- more than 400 requests for beta software. And although we did not have the resources to support all of these requests, we did receive very valuable feedback from the approximately 50 beta sites that tested our software during the beta period. Our goal for the PathScale EKO Compiler Suite was to be the highest performance compiler for the AMD64 architecture. Feedback from our beta sites and tests on industry-standard benchmarks has proven this to be the case. The PathScale compilers have proven to be between 10% and 40% higher performance than the competitive compilers and while the performance advantages will vary, many customers are seeing even better results on real applications.

Q: Why did you decide to "open source" your compilers?

We are strong proponents of the "open source" approach and releasing our EKO Compilers under an open source subscription license is a demonstration of that commitment. We hope that universities and others will enhance and extend the code base. Anything that makes it easier to move applications into cluster environments is good for all of us. At the same time, we can offer customers the same type of quality support for Linux compilers that they have had for Linux itself and that concept has been very well received by our customers.

Q: What types of environments and users did you choose for the beta tests?

A: PathScale offers C, C++ and FORTRAN compilers, so we looked for people using those different languages. We also wanted a broad mix of different applications, such as CFD codes, life sciences codes, molecular modeling codes and weather codes. Our performance engineering team is working on a variety of applications across a variety of industries to give us a good sense of how we perform in each area. At PathScale, we treat lack of performance leadership as a "bug" and we are committed to fixing any performance bug.

Q: Why do you believe there has been so much interest in PathScale's compilers?

A: Demand for AMD64-based servers is very high. HPC end-users are flocking

to Opteron and Athlon64-based systems. They can deliver outstanding price performance on both 32-bit and 64-bit applications. The AMD Opteron, with its on-chip memory controller and on-chip HyperTransport offers an architecture that is ideally suited to HPC applications.

However, to this point there have not been mature compilers available that offer Opteron users the ability to create applications that fully exploit the capabilities of that outstanding architecture. PathScale EKO compilers offer these users the alternative that they have been seeking; a mature compiler suite that is easy to use, yields correct results, is available to users in binary or source form and fully exploits the AMD64 architecture. This allows them to reduce the time to develop or port applications and offers significantly higher performance on the resultant code. The fact that we have a made a version of our compilers available as open source with full gnu front-end compatibility has also made the PathScale compilers very attractive to the gcc developers.

Q: That was a very nice endorsement of AMD. Does PathScale have a formal relationship with AMD?

A: We have a multi-faceted relationship with AMD in both technology and marketing. AMD is important to PathScale because we currently do all of our core development around their processor. We utilize engineering support from them to tune the performance of our compilers and we use their help in benchmarking and quality testing.

Q: And in marketing?

A: Yes, we are also very strong marketing partners. AMD attends all of the major HPC trade shows and they have consistently invited PathScale to participate with them in their booth. Their Web site has heavy traffic and they have PathScale profiled there. AMD is engaged in dialogues with large customers, OEMs and all of the system vendors. These relationships are very important to PathScale. From both a technology and a marketing standpoint, AMD is a very important partner for PathScale.

Q: What value does your relationship bring to AMD?

A: AMD has great processors but, unlike Intel, they don’t have an in-house compiler development team. They have to rely on third parties, like PathScale, to complete their software development ecosystem. PathScale is filling an important role by providing the highest performance, easiest-to-use 64-bit compilers that are completely optimized for the AMD64 architecture. PathScale makes AMD processors more competitive by making software run significantly faster than it would if only previously available compilers were used. This makes AMD processors even more attractive. Some of the other unannounced technologies we have in development will showcase the benefits of the AMD64 architecture and help further grow the Opteron market.

Q: How do the AMD chips compare to Intel chips, such as the Itanium?

A: We think the AMD Opteron compares favorably with any competitive commodity processor. It has the added advantage of efficiently running existing x86 software. Just about anybody who wants to use Itanium has to re-write their applications to take advantage of it. The market has spoken and Opteron is winning. We believe that the announcement by Intel of the 64-bit Xeon processor is recognition of these facts. Bringing out a 64-bit Intel product with an architecture that is compatible with AMD64 will allow Intel to again compete more effectively.

Q: Will PathScale support the new Intel 64-bit Xeon?

A: Yes. Since Intel is now endorsing AMD's 64-bit strategy, PathScale plans to fully support the Xeon 64 when it is available. It appears that Intel's 64-bit extensions are binary compatible with the AMD64 architecture. If this is true, the PathScale compilers will work just fine on the new Intel chips without modification. We haven't tested Intel samples yet, but we expect to receive them shortly and expect to have the best compilers available in market for both Opteron and Xeon-64. Intel always builds a great product, however AMD has implemented the Opteron chip with a direct on-chip memory controller and high-performance I/O right on the chip. This remains unique to the Opteron and offers significant advantages for HPC users. PathScale will be leveraging these advantages in some of our future products.

Q: PathScale has claimed that it has the world's fastest 64-bit compilers for Linux-based Opteron systems. What benchmarks have you used to validate this?

A: We have run a large number of both industry standard benchmarks and actual customer benchmarks, both internally and at customer sites. The results have shown that the PathScale performance advantage for 64-bit applications typically ranges from 20 to 40 percent faster than competitive alternatives. Some of the industry standard benchmarks that we have run include the SPEC suite, Livermore Loops, EuroBen, Polyhedron and NAS parallel benchmarks. Some of our beta sites have also run real application codes such as Amber, Gaussian and Charmm. The PathScale compilers have consistently proven to be the performance leader. As always, we say "your mileage may vary," because every application is unique. We now have a 30-day trial version of our compiler suite available through our Web site for those users who would like to "test drive" our compilers.

Q: How are your compilers sold? Do you sell them direct?

A: Our distribution model for the PathScale Compiler Suite is primarily an

indirect one. Our preferred method is for customers to purchase our compilers through our PathScale FastPath Authorized Reseller partners. We currently have signed over 15 system OEM partners to help us bring our compilers to market, including IBM, Linux Networx, Appro, Microway, Angstrom Microsystems, Rackable Systems, PSSC labs and seven international partners. We will be adding additional resellers shortly to expand our international coverage. Customers who are building homegrown clusters or prefer not to buy from a reseller may purchase directly from the PathScale Web site.

Q: How many resellers are you planning to have in total?

A: We are building a worldwide network of resellers who are all highly competent Linux-oriented system integrators. They all understand how to implement complex Linux clusters. When our program is complete, we should have at least one reseller serving every major metropolitan market through the world -- North America, South America, Europe and the Pacific Rim.

Q: Can you give us a brief history of PathScale?

A: PathScale was founded in 2001 as Key Research and its original mission was to develop clustered Linux server solutions based on a low-cost 64-bit design. We evolved to our current business model in early 2003 when we saw how successful the AMD Opteron could be. We refocused our efforts at that time on technologies that enhance the Opteron ecosystem. For example, high-performance 64-bit compilers were a critical need missing in the AMD solution set, as well as other cluster technologies that we now have in development but have not yet announced.

Our people all have very strong systems backgrounds and we understand what it takes to build the world's best clustered Linux servers. We have a highly-talented technical team that understands how to create the tools and enhancements needed to enable Linux clusters to scale to new levels of performance and efficiency.

One might say the seeds of the company were sown over 20 years ago at the Lawrence Livermore National Lab. Four of the company's seven founders all worked together building the S1 supercomputer back in the early 1980s. Our chairman and CTO, Dr. Tom McWilliams, had the initial idea for the company and incorporated in July of 2001. He added three of his LLNL colleagues, Jeff Rubin, Jeff Broughton and Dr. Fred Chow to the company shortly thereafter. Tom McWilliams had previously been a company founder at Valid Logic and Key Computer and has worked at SGI, Sun and Amdahl. Dr. Chow was formerly chief scientist for compilers at SGI and MIPs and is recognized as one of the world's leading authorities on compiler technology.