



Hewlett Packard
Enterprise

Evolving the Supercomputer

HPE Apollo 70 is a new model for innovation in the HPC market





Table of contents

Not just a statement—a conversation

Collaborative from the start

Arm-based HPC

Solid partner ecosystem

A new thought leader in Supercomputing





Not just a statement—a conversation

We hear so much these days about industry disruption we assume it's happening everywhere. While it might be true for consumer technology and even enterprise IT, it's a concept that hasn't really extended to the rarified air of the Supercomputing world.

Until now. Instead of taking a “build it and they will come” attitude, thought leaders at HPE decided to build an entirely new **supercomputing platform** by doing something much more interactive, more democratic.

Known until recently by its internal code name, “Comanche,” the team at HPE wanted to build the next generation of Supercomputer by finding out what its customers wanted.

It was the difference between making a statement and having a conversation.

As a result of this back-and-forth, 9-month dialog between HPE, its customers, and partner ecosystem, the team designed a brand new Supercomputing platform from the ground up.

This dense, scalable platform—and HPE's first Arm®-based HPC system—brings more choice and flexibility to HPC customers. It provides easy access to HPC technology with support for standard HPE provisioning, cluster management and performance software.

The **Apollo 70**, using Cavium's 64-bit Armv8-A ThunderX2™ Server Processor, is purpose-built for memory intensive **HPC workloads** and delivers up to 33% more memory bandwidth than today's industry standard servers.

The Apollo 70 also provides access to HPE's partnership ecosystem delivering key HPC components including Red Hat® Enterprise Linux®, SUSE® Linux Enterprise Server for Arm, and Mellanox® high-speed InfiniBand & Ethernet **fabric solutions**.





Collaborative from the start

The idea of giving customers early access to hardware isn't a new one. HPE and other industry leaders regularly participate in beta, POC, and discovery lab programs that grant customers early access to hardware. But this was different.

"All the pre-launch work we've done on Apollo 70 is totally unique in the HPC space," explains Kelly Pracht, HPE Senior Engineering Manager, Advanced Technology Engineering, who led the engineering for HPE on the Comanche collaboration. "One unique aspect is that it's a customer-funded program—the customers actually participated in a shared funding model with us—which represents not just a monetary commitment, but a shared vision."

It's an idea that inspired HPE's target customers to get involved and accelerate the development of the new ecosystem.

"To be brutally honest, my first reaction was: 'It's about time,'" recalls Gary Grider, HPC Division Leader at Los Alamos National Lab. "We've been involved in trying to bring an Arm-based environment for the **HPC** community in the Department of Energy for quite some time, because the number of ecosystems available to us is very limited. So the idea of collaborating on building a totally new, open-source HPC system is attractive to us."

The depth and timing of the involvement is unprecedented, and a key to the value of the collaboration. "At the lab we've always got new hardware and software coming in, but this was different," says Ross Miller, Systems Integration Programmer at Oak Ridge National Laboratory. "There was this great involvement and back-and-forth between the labs. It's not just getting hardware a month or two early—HPE is really listening much more carefully to what we have to say."





According to Rick Stevens, Associate Laboratory Director for the Computing, Environment and Life Sciences Directorate at Argonne National Laboratory, the opportunity to create something totally new was an important incentive.

“Inducing competition is a critical part of our mission and our ability to meet our users’ needs,” Stevens said at the Apollo 70 launch event at the SC17 Conference. “We have to build the pipeline for future systems, too. Industry partnerships are critical to our ability to do our job—which is to provide extreme-scale computing capabilities for solving some of the biggest challenges facing the world today.”





Arm-based HPC

So why build an entirely new HPC environment on Arm processors? According to the HPE team, the whole idea behind Apollo 70 was to help the entire HPC community to gain knowledge, create opportunities for competitive advantage, and enable new paradigms for computing.

The move to Arm is a big step for HPE and possibly a tectonic shift for the larger industry. “Having Arm come in to HPC is good for competition, it’s good for the Department of Energy, and it’s good for the U.S.,” says Grider at Los Alamos.

From Grider’s perspective, the move presents an opportunity for the HPC community to approach its problems differently. “A decade ago, when we saw Moore’s Law begin to run out of gas, we needed to figure out how to continue our advancement in order to do the science we wanted to do.”

Instead of investing billions to affect a new chip design, HPC leaders turned to the gaming industry, leveraging GPUs for the massive number of cycles they needed. But it still wasn’t an ideal solution. “It got us somewhere, but it wasn’t necessarily where we wanted to be. It meant we needed to turn our problems into graphics problems,” Grider recalls.

What the Comanche collaboration presented was something totally different. “This is the opposite approach. This is the opportunity to have a say at the hardware level, and actually design the hardware to be tuned to what our problems are, instead of tuning our problems to someone else’s hardware,” explains Grider. “It’s really exciting for us to step into that world and realize this is the beginning of a new architecture that is designed specifically for HPC. This is something that hasn’t happened since 1975.”





Solid partner ecosystem

In the early phases of the Comanche project, HPE team leaders realized they would need both customers and partners if the collaboration was going to be a success. Because the Apollo 70 was created from scratch for the HPC industry, it needed partners to build out the software stack.

“A year ago, if you asked why there wasn’t an Arm-based Supercomputer, the answer would be that there’s no ecosystem for it,” recalls Mike Vildibill, Vice President of Advanced Technology at HPE. “So we set out to change that.”

After setting goals and establishing a transparent legal and collaborative framework, HPE identified its dream team: Cavium for the 64-bit Armv8-A ThunderX2 Server Processor; Red Hat for Enterprise Linux, SUSE for Linux Enterprise Server for Arm, and Mellanox for its high-speed InfiniBand & Ethernet fabric solutions.

“So often, these products come to market, and they’re just the result of a single company giving us what they think we need,” Los Alamos’ Grider explains. “That’s not the case with Apollo 70. It’s a collaboration between all these big players in HPC—and I’m really glad it happened that way.”

For Ross at Oak Ridge, the partner collaboration experience has been eye opening. “I’m certainly not used to having industry leaders ask me what I want and being interested in what my experience has been,” Ross explains. “It’s been very positive so far. It’s obvious that someone did a lot of work on this platform before I became aware of it—the Linux kernel worked, the compilers were already there for a large part of it—someone put in some work to make the software this good.”



A new thought leader in Supercomputing

With all the groundwork, research, input, and fluidity that's been invested in the Apollo 70, it's clear HPE didn't just go into this project to build and market a new product. "We set out to change expectations about what could be done in Supercomputing," HPE's Pracht relates. "And we believe this crowd-funded, collaborative approach can be applied to other programs within HPE."

Taking those tangible steps to be inclusive at the outset, to demonstrate a commitment to the HPC community, and accelerate the development of a new platform has taken HPE on a journey that's been both educational and transformative. The effort hasn't gone unnoticed.

"I think it's just wonderful. This idea that we could actually have an effect on the development of new HPC-specific hardware is fascinating. It hasn't happened in a long, long time, and the promise of it is incredibly exciting," Grider from Los Alamos says. "We've spent the last 15 years leveraging other people's hardware, and now for the first time, we may not have to do that. HPE is listening to us."

Learn more at
[**hpe.com/supercomputing**](https://hpe.com/supercomputing)

 Make the right purchase decision. Click here to chat with our presales specialists.



Sign up for updates



© Copyright 2018 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Arm is a registered trademark of Arm Limited. All other third-party trademark(s) is/are property of their respective owner(s).

a00040908ENW, February 2018