

# The Journey to Frontier: The Story of How the Exascale Era Began

*How the Frontier team broke the exascale barrier to launch a new supercomputing era for scientific discovery*

**By Matt Lakin**

Oak Ridge National Laboratory

OAK RIDGE, Tenn. — Just before dawn, Scott Atchley woke up for the third time, took another sip of coffee and sat down at his computer to watch the next failure.

It was the morning of May 27, 2022. Atchley and fellow scientists had spent months tuning and tweaking [Frontier](#), the \$600 million supercomputer installed at the Department of Energy's Oak Ridge National Laboratory. The machine, built in spite of the COVID-19 pandemic and an international supply-chain crisis, had already set a new worldwide record for computing speed, but that wasn't enough.

The team didn't want to just set a record. They wanted to break the exascale barrier — the scientific milestone that would herald the next generation of supercomputing, with machines that could run a quintillion calculations per second.

Each test run brought the goal a little closer, only for Frontier to stall and crash. Atchley watched all night, sneaking naps between disappointments.

Now he was ready to log the latest failure and puzzle again over what might still stand in the way. Atchley sat down, glanced at the screen and snapped to attention.

Frontier had reached the end of its run. As he watched, the supercomputer shot past its day-old record and broke through the [exascale barrier](#), a feat deemed far-fetched and impractical by some of the world's leading scientists as recently as five years before.

“It felt like a miracle,” said Atchley, an ORNL distinguished scientist and chief technical officer for the Frontier project. “As a scientist, I'm not supposed to believe in miracles. But I was there a year ago when the vendor showed me a list of critical parts we couldn't build the machine without and said there was no way to get them.

“I was there six months ago when we finally had a supercomputer on the floor but couldn't keep it running. Just a week earlier, we didn't believe this was going to happen yet. Now we've done what we set out to do, and we're the first. Others will get here, but we're the first, and the whole world will benefit from this machine and what it can do.”

Frontier's speed clocked in at an average of 1.1 exaflops — that's 1.1 quintillion, or a billion billion, calculations per second. Each floating point operation, or flop, represents a mathematical calculation such as addition or multiplication.

“That exascale number has always been one of those magic thresholds for us,” Atchley said. “At an exaflop, the equivalent would be if everybody on earth — man, woman and child — worked on the same problem at the same time with a calculator and could do at least one addition or multiplication per second. We have about 8 billion people on earth. They would need four years to do what Frontier can do every second.”

Researchers hope to harness that number-crunching power to drive the next generation of scientific innovations: everything from [cures for cancer](#) and [prediction of natural disasters](#) to [cleaner-burning gasoline engines](#), [nuclear reactors that fit on a tabletop](#), and inquiries into the [origin of the universe](#).

“It’s a moonshot,” said Justin Whitt, who directed the [Oak Ridge Leadership Computing Facility](#) as Frontier came online. “The science that will be done on this machine will be truly game-changing for countless researchers and will make long-term impacts on society that we can’t even begin to foresee.”

The odds were against the effort from the start. Frontier’s a 296-ton machine of roughly 60 million parts, each essential to operation, that initial estimates projected would require more power to run than a mid-sized U.S. city.

That machine had to be not only built but assembled, booted up, tuned, primed and finally proven to perform.

“Given all the obstacles and events that stood in our way, you can look at the probability and show mathematically this success story should not have happened and a machine as complicated as Frontier should not work,” said Thomas Zacharia, now retired, who led development of ORNL’s supercomputing capability and served as the lab’s director during the launch of Frontier and its predecessor, Summit.

“Frontier, when you come down to it, is a story of human beings rising to the challenge and doing the impossible because it is worth doing. We believe it can be done — not out of blind belief, but belief based on experience, on a shared vision, on a mutual determination that science and the pursuit of knowledge should never stop. Frontier serves as a reminder of what we at this laboratory can achieve when we all come together for a great purpose.”

## **The magic number**

Exascale computing’s promise rests on the ability to synthesize massive amounts of data into detailed simulations so complex that previous generations of computers couldn’t handle the calculations. The faster the computer, the more possibilities and probabilities can be plugged into the simulation to be tested against what’s already known — how a nuclear reactor might respond to a power failure, how cancer cells might respond to new treatments, how a [3D-printed design](#) might hold up under strain.

The process helps researchers target their experiments and fine-tune designs while saving the time and expense of real-world testing. Scientists from around the world compete for time on

Frontier through DOE's [Innovative and Novel Computational Impact on Theory and Experiment](#), or INCITE, program.

“The bigger the idea, the bigger the simulation,” Atchley said. “A machine of Frontier’s power can let you tackle exponentially larger problems because it can do the math faster than any other machine.”

Simulating the universe? Exascale allows for not just more planets and stars but whole galaxies. Simulating quantum mechanics? A machine like Frontier allows for more particles. Simulating climate or weather? Frontier allows global modeling at a size, scale and level of accuracy over a longer time frame than ever before.

“It’s like the difference between ‘Donkey Kong’ and ‘Grand Theft Auto,’” Atchley said. “Because Frontier is so much faster, we can perform simulations in minutes, hours or days that would take years or even decades to complete on other machines — which means they wouldn’t be done.”

Scientists initially questioned whether exascale computing could be done at all. The discussion jumped from theoretical to practical in 2008 after the [Roadrunner](#) supercomputer at [Los Alamos National Laboratory](#) achieved petascale speed with a run clocked at 1 petaflop, or 1 quadrillion calculations per second. The OLCF’s [Jaguar](#) logged nearly double that speed a year later.

Could high-speed computing make the leap to the next order of magnitude? Not likely, experts warned.

“Just one major challenge threatened to be a showstopper,” said Al Geist, an ORNL corporate fellow and chief technology officer for the Frontier project. “We identified four, and all of them would have to be overcome — power consumption, reliability, data movement and parallelism.”

The biggest roadblock, power consumption, towered over all others. Jaguar and every major supercomputer before it relied on [central processing units](#), or CPUs, for computational power. A CPU essentially acts as a computer’s brain by performing calculations, retrieving information and executing instructions.

Jaguar in 2009 generated an average of 328 megaflops for every watt of power consumed, for a total energy footprint of up to 7 megawatts, or 7 million watts. Based on that ratio of power to performance, a supercomputer using state-of-the-art technology for the time would have required not just millions but [billions of watts](#) to crank out a single exaflop.

“The electric bill alone for that kind of consumption would be about \$600 million per year,” Geist said. “And even projecting technological advancements that might save power, we were still looking at unsustainable costs.”

The next obstacle: reliability. Would this megawatt-munching machine keep running long enough to crunch a single set of numbers? Projections said no.

“We’d seen from earlier supercomputers about how long they would last before crashing, and we figured what if it’s a thousand times bigger?” Geist said. “How long could it stay up? Maybe just for a few minutes. That’s not even long enough to get to the end of a science problem. Nobody’s going to build a computer that burns that much power and crashes every few minutes.”

The challenges got no easier. If the OLCF could afford the multimillion-dollar energy bill to operate an exascale computer — and if engineers could keep it running long enough to solve a science problem — could its billions of circuits move enough data from memory to processing to storage quickly enough to keep up with the speed of calculations? And could the computer’s army of processors work in parallel to break a massive problem into small enough chunks to solve each equation and reassemble the results into a verifiable answer?

The outlook wasn’t good.

“The largest number of pieces we had ever broken a problem up into at that point was about 100,000,” Geist said. “That’s parallelism. To get to an exaflop, our calculations showed we would have to find a way to achieve billion-way parallelism — break the problem into a billion pieces, solve each piece at the same time and then put all the pieces back together in the right order. We didn’t even know how to get to million-way parallelism yet.”

The eventual solution to all four challenges came about almost by accident. ORNL counted on a chip to be built by [Intel](#) to power Jaguar’s successor, [Titan](#). Intel canceled the design as inefficient and never built a single chip.

That shift left OLCF engineers scrambling to find a new solution to drive an even faster supercomputer. Scientists settled on a gamble: incorporating graphics processing units, or GPUs — traditionally used to render images and video for computer games — as accelerators to free up CPUs and power the next generation of scientific supercomputing.

“It was a big bet at the time,” Geist said. “But lo and behold, that decision took us down the path of accelerator-based computing nodes that ultimately led to Frontier.”

GPUs excel at performing the kind of repetitive, parallel calculations necessary for high-resolution graphics. That processing power also helps to deliver high-resolution simulations and to power artificial-intelligence algorithms such as machine learning. The shift to GPUs, together with advances in microchip development that resulted in more efficient CPUs and increased reliability in memory and hard-drive components, helped deliver two exponentially more powerful petascale computers — Titan at 27 petaflops and [Summit](#) at 200 petaflops — and finally Frontier.

Today’s exascale supercomputer not only keeps running long enough to do the job but at an average of only around 30 megawatts. That’s a little more than three times the power consumption of Jaguar in 2009 for more than 500 times the speed. For every watt Frontier burns, the system wrings out more than 62 gigaflops of performance, or 62 billion-plus calculations per second.

“The combination of GPUs and CPUs turned out to hit a sweet spot that’s not only energy-efficient but tremendously effective,” Geist said. “It wasn’t the only key to reaching exascale. As we developed these successive generations of supercomputers, we hit on new approaches and tweaks to the architecture, such as adding more GPUs and memory. So it was a series of incremental steps that got us here, rather than the giant leap we thought it might take.”

## Navigating a pandemic

Overcoming the technical obstacles didn’t put Frontier on the data center floor. DOE ultimately launched the Exascale Computing Initiative in 2016 in response to orders from the Obama White House.

That effort, focused on building an exascale supercomputer, also created the [Exascale Computing Project](#), charged with building a comprehensive system of target applications, a software stack and accelerated hardware technology primed to take full advantage of exascale speeds on the day of launch. The total project assembled thousands of researchers from across academia, federal agencies, private industry and DOE’s national laboratory system to translate the vision into reality.

In May 2019, DOE [awarded computing manufacturer Cray](#) the contract to build Frontier and chip-maker [AMD](#) the contract to build the CPUs and GPUs that would power Frontier. [Hewlett Packard Enterprise](#), or HPE, acquired Cray later that year.

By early 2020, the project looked on track. Then news broke of a new, potentially deadly virus raging in China and slowly spreading across Europe.

“When we [first talked about the coronavirus pandemic](#) and what might happen, it was the last week of February 2020, and the virus was still seen as just a problem overseas,” said Jason Hill, the risk manager for the Frontier project at the time. “Some of us wondered then, ‘Why are we dedicating so much time to this?’ Then came the order in March to work from home. That’s when it became apparent to everyone this pandemic was going to be something that could really endanger the project.”

Initial plans called for Frontier to be delivered and installed at ORNL by July 2021. The worldwide shutdown in response to the virus threatened to delay the project by years.

“We knew if the lab shut down completely, our construction crews couldn’t get the work done to get the data center ready for Frontier,” Hill said. “We knew if all our computing staff had to work from home, we introduced potential reliability and stability issues with our electronic infrastructure. And all our vendors along the supply chain were doing the same thing.”

The OLCF team had barely a week and a half to prepare. The slim lead time proved enough to make the difference. Small steps — such as shoring up network security, creating an electronic signature process, increasing advance orders and drawing up safety protocols for crews onsite — allowed work to continue at ORNL to prepare for Frontier’s arrival.

In Chippewa Falls, Wisconsin, the factory workers building Frontier took similar steps.

“We had to change the whole way we worked,” said Gerald Kleyn, HPE’s vice president for high-performance computing and artificial intelligence systems.

### **Missing pieces**

By spring 2021, the workarounds looked like a success. Scott Atchley, who led the hardware evaluation team for the Frontier project, visited Chippewa Falls that May for an update on the supercomputer’s delivery, still set for summer.

“They showed us the first three racks of Frontier, we talked about the production schedule, and everything sounded fine,” Atchley said. “Then the plant manager and head of purchasing took me off for a side meeting and showed me a list.”

The list went on for pages — 150 parts, each vital for Frontier to run. None could be found.

Atchley felt his hopes sinking.

“I said, ‘You can’t get these today. Can you get them next week or next month?’

“They said, ‘We don’t know when we can get them — or if we can get them at all.’ ”

The shutdown caused by the pandemic [disrupted international supply chains](#) across nearly every industry. Some of the parts needed for Frontier doubled, tripled or more in price. Some no longer existed on the open market.

“We all remember the shortages started with [toilet paper](#), right?” said Kleyn, the HPE vice president. “Then everyone started running out of computer parts and chips. We did a lot of work to redesign some of our systems to try to get around these gaps, but some of these parts we couldn’t do without.”

The [Defense Production Act](#) and advance orders allowed the Frontier team to fast-track most orders for CPUs, GPUs and other high-value parts. But the top-shelf components turned out not to be the problem.

Supplies of the everyday parts typically bought by the box, bucket or pallet — bolts, gaskets, voltage regulators, oscillators, resistors — dried up. The lack of nickel-and-dime essentials threatened to sink the multimillion-dollar project.

HPE’s procurement office detailed a team of 10 people to search full-time for each necessary component. The team called up dealers, wholesalers and competitors across the world to plead, coax and bargain for every piece, down to the last bolt.

“They scoured the planet for these parts,” Kleyn said. “The pandemic was still going on and not only could we not always get these parts, we couldn’t always get them to the right places on time. We couldn’t get trucks, we couldn’t get drivers, shipments would be delayed by someone getting the coronavirus and be rerouted. It was like a jungle we had to hack our way through every day.”

### **Some assembly required**

Frontier’s July delivery date turned to August, and August to September. A few onlookers let out a quiet cheer the morning of Sept. 24, 2021, when the first of Frontier’s 74 cabinets arrived at the OLCF data center.

“That was the moment we all felt it was finally happening, that after all the obstacles this was real,” said Gina Tourassi, then director of the National Center for Computational Sciences, which oversees the OLCF. “I was there to see it when the first 8,000-pound beast came off the truck, because I wanted to watch the magic happen.”

The last of Frontier’s 74 cabinets finally arrived at ORNL on Oct. 18, 2021, a day after the crew in Chippewa Falls screwed on the final part — the last oscillator on the factory floor.

Paul Abston, data center manager for the OLCF, oversaw the cabinets’ arrival and installation. Each cabinet, 4 feet by 7 feet, weighed in at 4 tons, so heavy that crews had to slide them across the floor on metal sheets to keep from digging troughs in the tiles.

“Trucks were backing up to the dock at all hours,” Abston said. “Each of those trucks can carry only three or four cabinets at a time because of the weight, and they have to arrive in the right order. Most days we got done in 12 hours, sometimes 16.”

As the cabinets arrived, final assembly began. Frontier’s not so much a single supercomputer as a constellation, with each of its more than 9,400 nodes a self-contained 150-teraflop supercomputer of one CPU and four GPUs, capable of up to 150 trillion calculations per second. More than 90 miles of cables connect the nodes in a high-speed network, Slingshot.

“Picture the most elaborate, complex thing you’ve ever taken apart,” said Joe Glenski, an HPE senior distinguished technologist who helped oversee the Frontier installation. “Now try to put all the pieces back together and make it work. Now multiply that by 60 million.”

### **Trial and error**

Each node had been assembled for brief testing on the factory floor before shipment. But the entire system came together for the first time only after arriving in Oak Ridge.

Now the real testing began.

“Everything in this machine is brand new, and some of these parts are going to fail immediately,” Atchley said. “It’s just like when you plug in a light bulb and it pops. We have 60 million of these parts to test — CPUs, GPUs, network cards, cables — and one or two out of every batch are inevitably going to be bad. We’re plugging everything in to see what breaks and why. Eventually we replace the bad parts, tighten the loose connections and it all stabilizes, but one failure at any time will take down the whole machine. This goes on as long as necessary.”

The Frontier team hoped to deliver the machine by November 2021, but November turned to December, and December to January.

In May a new TOP500 list would be released, ranking the fastest supercomputers in the world — an artificial deadline. But announcing a top ranking with the world’s first exascale system would grab headlines and deliver an unparalleled opportunity to highlight Frontier’s potential for breakthrough science.

As 2022 wore on, those hopes began to fade.

“By February, we started to have our doubts,” said Whitt, the OLCF director. “These millions of parts worked fine individually, but scaling them up to work in parallel was pushing the limits of the technology. I had to tell our lab director, ‘We have three months left, and we’re not going to make it on our current trajectory.’”

The lab director wasn’t having it.

“Imagine Tom Brady being down a touchdown in the fourth quarter of the Super Bowl and saying, ‘Maybe next game,’” Zacharia recalled. “We have a track record at ORNL of committing to impossible things and delivering them. This was going to be the first time we’ve fallen short?”

Whitt and the team got the message: Exascale was the only acceptable result. Get there no matter what.

A meeting in Houston — initially scheduled to discuss plans for Frontier’s successor — turned into an emergency summit between the Frontier team and top executives from HPE and AMD. The installation team at ORNL became a war room of around 50 engineers and other experts working in 16-hour shifts seven days a week to resolve every bug.

“That was a real sea change,” said Glenski, the HPE scientist. “Here’s this group of people from three different organizations who’ve never met each other in person before, all working together and brainstorming every idea we can think of to pinpoint these stubborn problems deep in the machine. I don’t remember a single time during those three months I heard anyone say, ‘No, I’m not going to help.’ We all stepped up. And the clock was ticking the whole time.”

### **Racing the clock**

The team had until May 13 to submit a speed for Frontier to the TOP500 committee. As the deadline approached, tensions rose.



Most of the early problems with the system had been resolved. But calculation speeds still clocked in well short of an exaflop, at barely half Frontier's promised speed.

"We started to second-guess ourselves at that point," said Whitt, the OLCF director. "Maybe there really were fundamental limits to the technology that we just didn't anticipate. It was like racing a car that should be able to go 100 mph, but the fastest you could get was 60. We're not used to coming in second, but we had a week to go and we couldn't find what was holding us back."

Scientists measure supercomputing speed by how quickly each machine can solve the [High-Performance Linpack](#), or HPL, benchmark test, a complicated series of lengthy algebraic equations that stresses all elements of the system.

"It's like qualifying for the Olympics," Atchley said. "The HPL score doesn't necessarily measure what kind of science can be done on the system. But it's a fantastic benchmark for exercising the machine, and it puts so much stress on the network that it can help you spot any faulty parts. We would run it on every node and see one GPU here or there might be slower, so we would replace it. The test works like a relay race, where the team's only as fast as the slowest runner. If any node or server drags, it bogs down the whole machine, and every little failure kills you."

No matter how the team tweaked the code, no matter how many processors they replaced, the system hit a wall. The fastest speed clocked by the deadline, 614 petaflops, topped any other supercomputer in the world. But it wasn't exascale.

"We could have stopped there," Zacharia said. "But that wouldn't be ORNL. The day we decide 'good' is enough — that's the day we no longer aspire to keep moving forward. I told the Frontier team the deadline didn't matter. 'Don't stop. Keep pushing. Even if we miss the top spot, we're going to deliver the first exascale machine.'"

### **The fatal flaw**

As the team pored over data readouts, engineers noticed a bizarre pattern. A typical supercomputer cycle kicks off with a sharp spike in power consumption that gradually tapers down as the run finishes. Frontier's power profile looked like a sawtooth: a straight spike, then a sudden drop, repeated every four minutes.

The team studied the individual processors, then the nodes and servers. Nothing. The sawtooth appeared only at full scale when the nodes ran together.

The team tweaked various algorithms, changed the way nodes and servers communicated. The sawtooth pattern faded, but speed took a nosedive.

“We’d never seen a pattern like that before,” Atchley said. “This was one of the most frustrating and depressing points of my professional life. We knew the machine should be doing better than this.”

Phil Roth, an algorithms and performance analyst, took a closer look at the performance data from one of the sawtooth runs and wrote an algorithm that plotted the details. Each of Frontier’s processors was numbered and could communicate with any other. Each message should travel anywhere in the system at the same speed.

But the higher a sender’s number, the data showed, the longer a recipient took to process the message. The lower a sender’s number, the faster the message went through.

“That’s the aha moment,” Roth said. “Instead of the message being processed immediately, the recipient would start at zero to look up the sender’s address and search one-by-one, whether the message came from No. 10 or No. 10,000. It’s like caller ID, but instead of the name and number popping up automatically you’re opening the phone book, starting at ‘A’ and searching page-by-page, line-by-line to look for Zebediah’s Auto Shop. That’s why the higher numbers saw the longest delays.”

The team discovered a software library’s debug option, turned on by default, caused the search hangup. They switched it off. On the next run, just after 3:30 a.m. on Thursday, May 26, Frontier clocked in at 939.8 petaflops — within sight of the exascale goal.

But the number still dangled just out of reach. Frontier gained speed but lacked stability. The system continued to crash, one run after another.

“We knew we were close,” Atchley said. “The power readout would shoot up, Frontier would start to run and the sawtooth was gone. It’s running really well and we’re all starting to believe, then the job crashes. We start up another and watch it go. This one dies after a few minutes. Over and over and over.”

### **Breaking the barrier**

By the evening of May 26, the system reached the hour mark, then two hours — still crashing every time. Atchley set an alarm for every two hours and settled in for the longest night of his career.

The hours dragged by — 11 p.m., 1 a.m., 3 a.m.

“I wanted to believe we would get there,” Atchley said. “I wanted to believe every time. If we can get an HPL score tonight, there’s still a chance we can make the TOP500 and make history. It’s like an Olympic team. We feel like we represent the U.S. in this competition, and we’re never going to have a better chance than now.”

Just after 5 a.m., he woke again and sat down to watch the next run.

The system was winding down. He saw the power readout — the spike, the plateau, the gradual slump as the run neared its end, this time closer than before.

Atchley leaned forward. His eyes widened. The team’s chat channels fell silent.

“Now I’ve gone from groggy to wide awake,” he said. “My heart’s racing. Everything looked so good, we’re afraid to say anything for fear we’ll jinx it.

“At the very end of the HPL run, there’s a last blast of effort before the job completes. We’re all waiting for it to crash. All of a sudden everything drops off, and then we realize we made it. We got to the end, and we broke the exascale barrier for the first time.”

Members of the team, scattered across offices and chatrooms, let out a chorus of cheers. Phones erupted with texts of celebration.

Some of the team still tear up at the memory.

“I woke my kids,” said Kleyn, the HPE vice president. “I’d fallen asleep on the couch and when I saw the message, I let out a scream so loud they came running in. It was a really great day, one of the best of my career, and I still get chills down my back just thinking about it.”

Tourassi, the NCCS director, rolled over in bed and handed the phone to her husband.

“I told him, ‘Check this and tell me I’m not dreaming,’” she said. “He told me. It was no dream. The rest of the day felt like walking on a cloud.”

## **Telling the world**

The TOP500 committee [announced Frontier’s record-breaking results](#) May 30, 2022, at the International Supercomputing Conference in Hamburg, Germany, to worldwide acclaim. Zacharia personally accepted the certificate for first-place status.

“It was a proud moment, but the accomplishments to come on Frontier will make us all prouder,” he said. “There’s so much more work to be done. Frontier will continue to power the imagination that drives our conquest of the next frontier in computing.”

Frontier opened for full user operations in April 2023.

“We had a lot of help from early users on the system who helped test it and find bugs,” said Veronica Melesse Vergara, who oversaw the user acceptance process. “Frontier was definitely the most challenging system of all the leadership computing systems to bring online, but every extra moment it takes to vet a system like this pays off later.”

Plans have begun for Frontier’s successor, which remains in development. The next speed barrier to break would be zettascale — 1 sextillion calculations per second, or a trillion billion — and that moment could be a long time coming.

“Right now no one sees a way to get there,” said Geist, Frontier’s chief technology officer. “We didn’t see a way to exascale initially. But we’d have to overcome the same problems presented by exascale, plus boost the speed by a factor of a thousand. I think we’re more likely to see a blending of computing technologies such as artificial intelligence or [quantum computing](#) to find new ways of solving problems that make more efficient use of exascale speeds.”

Frontier [still holds the title](#) of world’s fastest supercomputer after new TOP500 lists came out in November 2022, June 2023, and November 2024, and OLCF engineers expect further tuning to [coax even faster speeds](#) from its processors. But the team that broke the exascale barrier knows that distinction can’t last forever.

“It’s an ephemeral honor,” Atchley said. “We expect others to follow. We’d be disappointed if they didn’t. But the number is just a measurement. I’m more excited now to see what kind of science Frontier can do and to see what we can do next because of Frontier.”

*UT-Battelle manages ORNL for the Department of Energy’s Office of Science, the single largest supporter of basic research in the physical sciences in the United States. The Office of Science is working to address some of the most pressing challenges of our time. For more information, please visit [energy.gov/science](https://energy.gov/science). The Oak Ridge Leadership Computing Facility is a DOE Office of Science user facility located at ORNL.*

Read more:

[Super speeds for super AI: Frontier sets new pace for artificial intelligence](#)  
[Super construction at super scale: How ORNL built new home for Frontier.](#)

Watch the [documentary](#).