

Case Study

High Performance Computing (HPC)
Intel® Xeon® Scalable Processor



NASA Ames Builds its Next-Generation Flagship Supercomputer Aitken on Modular Architecture

Aitken's first of 12 modules is built on Intel Xeon Gold processors (3.69 petaFLOPS peak), using local climate-based air cooling, resulting in 1.03 PUE

Aitken Highlights:

- Built on modular HPC architecture
- Intel Xeon Gold 6248 processor (20 cores)
- 1,152 nodes (first module buildout, only 1/3 populated)
- Built by Hewlett Packard Enterprise (HPE) based upon the HPE SGI 8600 system
- 1.03 PUE leveraging local climate for cooling



Executive Summary

Having outgrown its current High Performance Computing (HPC) data center, NASA Ames' [NASA Advanced Supercomputing \(NAS\) Division](#) began the buildout of a large petascale system called [Aitken](#), its next flagship supercomputer, using a containerized architecture with 2nd Generation Intel Xeon Scalable processors. The new Intel processors will allow scientists to accomplish research not possible on existing machines, such as running models longer and getting faster turnaround on detailed simulation results. The modular computing design process involved prototyping a new approach with container-based computing systems using a prototype cluster called [Electra](#). The new design takes advantage of the performance and scalability of the 2nd Generation Intel Xeon Gold 6248 processors and the local temperate climate to cool the systems. The new data center space now houses containers with Electra and the early stages of Aitken, which will be expanded over the next several years.

Challenge

Big science requires big computing. With missions that span aeronautics, astronomic, oceanographic, and human space exploration, NASA Ames does a lot of very big computing. Its [Pleiades](#) supercomputer installed in 2008 was one of the fastest in the world at the time. Having undergone many expansions since then and now at 7.24 petaFLOPS (theoretical peak), it continues to be the NAS facility's most capable cluster.

NASA supports many other High Performance Computing (HPC) resources besides Pleiades. With the need for ever-deeper understanding of the world and universe and more advanced missions, such as Artemis—the mission to the moon—and the



With the need for ever-deeper understanding of the world and universe and more advanced missions, NASA needed additional computing capability and opted for an HPE SGI 8600 system powered by Intel Xeon Gold processors.
(Photo courtesy of NASA)

Space Launch System (SLS) to take humans there, NASA needed additional computing capability.

The NAS current facility is a six-megawatt building, and it is full. To add CPU cycles, the team led by Bill Thigpen, assistant division chief of HPC operations in the NAS Division, determined they would have to consider where to remove some cycles to make room. That means they cannot take full advantage of all the procurements they have made so far. After studying several options, they concluded they would need a new facility.

Solution

About one-third of the power the current NAS facility consumes goes to cooling—chillers, air handlers, etc.—and the rest to computing. That is a Power Usage Effectiveness (PUE) of about 1.33, which is acceptable with today's technologies according to NAS managers. But, given NASA AMES' location on the West Coast of the U.S. with ideal climate conditions, they believed they could improve power usage by utilizing their environmental advantages.

Container Architecture and Air Cooling Cuts Costs

So, NASA looked at both a new site and a modular architecture, using specialized computing containers to house their HPC systems. The new facility would be located on a one-acre site with 30 megawatts of power available. They learned that a modular architecture using containers could possibly give them finer control over the computing environment using the surrounding temperate climate, allowing them to eliminate the use of chillers—but they also needed to determine if they could build a tightly coupled cluster across containers. There were their two important questions: could they cool using the surrounding environment and eliminate chillers, and could a modular system be built as a tightly coupled cluster for running the largest-capability computing jobs? A prototype would help answer those questions.

The NASA team approached the vendor community to design a solution that included everything—from the containers to the power and computing system. HPE won the project with Schneider Electric providing the power connectivity.

The prototype system, named Electra after one of the stars in the star cluster Pleiades, was built in two computing modules (each container is a module). Each container could support

four individual computing clusters called E-cells. The first module, built in 2016, allowed them to study whether they could cool without chillers. It housed 1,152 nodes of dual-socket Intel Xeon E5-2680 v4 processors and consumed 500 kilowatts of computing power with a PUE of 1.025. The second module, built in 2017 using 1,152 dual-socket nodes of Intel Xeon Gold 6148 processors, delivered much more dense computing of about 1.2 megawatts with a PUE of about 1.04. With the second module, they evaluated if they could tightly couple multiple modules into a single high-end capability computer. The Electra prototype showed the team that a modular architecture met both objectives, while saving approximately 91 percent of power and 96 percent of water compared to other facilities.

Today Electra spans two containers and comprises a total of 78,336 cores, 368 TB of memory, and 29 petabytes of RAID storage. The system provides 4.79 petaFLOPS (theoretical peak) of tightly coupled computing across the two containers. Many scientific and engineering projects have run on Electra.

Aitken Arises

With the modular approach proven successful, NASA moved to the next phase of the NAS facility expansion project—building the agency's next flagship supercomputer called Aitken, named after American astronomer Robert Grant Aitken. The new system would be built in stages over time, utilizing the latest advanced technologies available to enable very large capability computing.

Each module can house twelve E-cells containing 288 computing nodes each and providing its own cooling. The first four E-cells of Aitken were deployed in 2019. It uses 20-core Intel Xeon Gold 6248 processors with 40 cores per node, for a total of 46,080 cores and providing 3.69 petaFLOPS (theoretical peak) computing capacity.


With 30 megawatts available at the new facility, they have room for expansion. Aitken today provides just under four petaFLOPS of theoretical compute in a single compute cell. Each module can support four of these, and Aitken is planned for twelve modules.

That means Aitken today is only 1/36th its planned size and already providing nearly four petaFLOPS. Aitken will be expanded over the next several years, just as Pleiades has continually evolved since 2008. If the cluster were fully populated today with the same capability it currently has, each module would offer 11.07 petaFLOPS of computing with the entire cluster delivering over 133 petaFLOPS. That can support a lot of big science.

Result

Aitken was built for a wide range of modeling and simulation work. With 2nd Gen Intel Xeon Gold 6248 processors, researchers can accomplish capability computing not possible on older technologies. They can improve the accuracy of their models for greater insight into the problems they are studying. And with these latest-generation processors, scientists can run jobs longer for more detailed results.

Aitken is still in its infancy. While NASA does not target a specific computer for a certain science, there's been a lot of science run on Aitken and Electra to date. The people working on Artemis and the new SLS have run jobs on these clusters.



Preliminary results from a high-resolution GEOS/ECCO simulation showing evaporation (red colors) and precipitation (blue colors) over the ocean. (Photo courtesy of NASA)

Case Study | NASA Ames Builds its Next-Generation Flagship Supercomputer Aitken on Modular Architecture

And a lot of aeronautics research, ranging from noise reduction to new types of aircraft and advanced air mobility has been run on Electra.

Work on Electra Advances Aeronautics

NASA research scientist Neil Chaderjian has used Electra to better understand why rotorcraft, such as helicopters, are restrained in forward flight because they reach a dynamic stall speed caused by blade vortex interaction (BVI). His work with advanced computational fluid dynamics (CFD) [simulations on rotorcraft](#) using Electra has revealed some first-time observations and discoveries about BVI. His work is shared with the rotorcraft design communities in order to help them design safer, faster, and more energy-efficient rotorcraft.

Electra has also supported [aeronautical research for next-generation aircraft designs](#), such as the transonic truss-braced wing (TTBW), NASA's X-57 electrically powered aircraft, and the X-59 QueSST designed with quiet supersonic technology.

NASA scientists are also using Electra to develop next-generation Earth system models. For example, they combined two flagship NASA models, the Goddard Earth Observing System (GEOS) and the Estimating the Circulation and Climate of the Ocean (ECCO), to simulate Earth's weather and climate with unprecedented detail for some early results. Electra and Aitken will allow the GEOS and ECCO teams to substantially advance NASA capabilities for seamless weather and climate simulation, estimation, and prediction.

Today, Electra is still bigger than Aitken, but as the new flagship system expands, it will become the most capable computing resource NASA has. With the modular architecture, NASA will spend fewer dollars on cooling, so they can deploy more computing capability, giving scientists computational resources to perform some novel and detailed simulations that are not possible today.

Solution Summary

NASA Ames had outgrown its computing facility. It is expanding its HPC resources using a modular compute architecture built by HPE that leverages the temperate climate of the San Francisco Bay area to build a more efficient computing center with PUE at 1.03 to 1.04. The first system built in the new facility was a prototype system named Electra, housed in two containers and providing 4.79 petaFLOPS (theoretical peak). With the new approach proven,

NASA began building Aitken, the next-generation flagship supercomputer to be expanded over several years. Aitken's first module is 1/3rd populated with Intel Xeon Gold 6248 processors. It delivers 3.69 petaFLOPS (theoretical peak) performance. When finished, Aitken will occupy 12 compute modules and provide a level of high-performance computing not possible with previous technologies.

Where to Get More Information

Learn more about the [NASA Advanced Supercomputing \(NAS\) Division](#).

Find out more about [Electra](#) and [Aitken](#).

Learn more about [Intel Xeon Scalable processors](#).

Explore the [HPE SGI 8600 system](#).

Aitken Supercomputer Ingredients

- Intel Xeon Gold 6248 processor (20 cores)
- 1,152 nodes (first module buildout, only 1/3 populated)
- Built by Hewlett Packard Enterprise (HPE)
- Leverages local climate for cooling enabling 1.03 PUE



Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Results have been estimated or simulated.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.