DELL EMC

intel

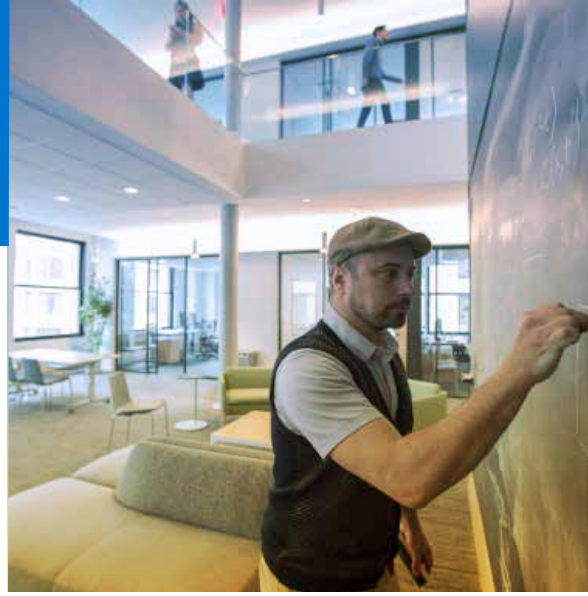

Pure science

Flatiron Institute delivers the computational power to fuel leading-edge scientific research.

FLATIRON INSTITUTE

Scientific computing | United States

Business needs

The Flatiron Institute needs a scalable high-performance computing environment to accelerate research in astrophysics, biology, mathematics and quantum physics.

Solutions at a glance

- Dell PowerEdge™ servers with Intel® Xeon® Scalable Processors
- Dell EMC DSS storage
- Intel® Omni-Path Architecture for the interconnect
- Bright Cluster Manager

Business results

- Driving advances in computational sciences
- Providing a permanent home for professional scientists

The Flatiron Institute's HPC cluster in New York City has more than

10,000 cores

Collectively, the institute's distributed infrastructure gives researchers access to

30,000 cores

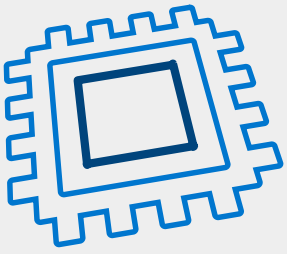

A research-driven mission

The Flatiron Institute is a large, interdisciplinary scientific hub within the Simons Foundation. The institute is dedicated to the advancement of scientific research through computational methods, including data analysis, modeling and simulation. The institute addresses the need for new computational approaches to analyze the vast experimental and observational datasets being generated and collected in biology, chemistry and physics.

Flatiron, which began operations in 2016, provides a permanent home for professional scientists who are driven by the problem of developing, deploying and supporting state-of-the-art computational methods for modeling and analysis. The institute supports research centers devoted to computational astrophysics, computational biology, computational mathematics and computational quantum physics, each of which is led by a distinguished scientist.

SIMONS FOUNDATION

Advancing the frontiers of math and science

The Simons Foundation, the parent company of the Flatiron Institute, works to advance the frontiers of research in mathematics and the basic sciences. The private foundation supports basic — or discovery-driven — scientific research, undertaken in pursuit of understanding the phenomena of our world. The foundation was co-founded by Jim and Marilyn Simons in 1994 and is based in New York City.

Scientists at Flatiron have the opportunity to dedicate themselves to long-term research that is not driven by standard funding cycles. The institute currently serves the needs of 160 scientists, and has plans to grow to 250 scientists over time.

A growing IT environment

The Scientific Computing Core (SCC) is the technical backbone of the Flatiron Institute. The SCC develops, deploys and maintains computational infrastructure — from supercomputers to desktop PCs — dedicated solely to the use of Flatiron researchers. The SCC also creates and disseminates software tools for the wider computational science community, with the goal of ensuring that technology is not the limiting factor of scientific progress.

The SCC's main computing resource is a high-performance computing cluster housed in the institute's data center in New York City's Flatiron District. This system is based on Dell EMC PowerEdge™ servers that collectively have more than 10,000 Intel® Xeon® Scalable Processor cores, all connected via Dell H-Series networking based on Intel® Omni-Path Architecture (Intel® OPA) and managed with Bright Cluster Manager.

On the storage side, this steadily growing HPC system uses Dell DSS 7500 servers to hold thousands of disks and nearly 20 petabytes of storage, which is accessed via a Ceph file system. The mix of files in the system's archive ranges from millions of tiny files to thousands of 100-GB files.

In addition to the resources in the New York City data center, Flatiron gives researchers access to two off-site supercomputers. One of these resources is located at Brookhaven National Laboratory on Long Island, New York, and the other is at the San Diego Supercomputer Center at the University of California, San Diego. Collectively, the institute's distributed infrastructure gives researchers access to 30,000 processing cores.

DELL EMC

The HPC facilities offered through the Scientific Computing Core are intentionally diverse, because that's what it takes to meet the diverse requirements of the Flatiron Institute's scientists, according to Ian Fisk, co-director of the SCC. For example, the SCC maintains a small cluster with GPU hardware for researchers who want to use GPUs to accelerate machine learning or other data-intensive applications.

"One of the things that is special about our nodes is that they have a tremendous amount of memory," Fisk says. "That's because we have to run applications that span several scientific domains. One of the ways that we meet this need is by making the nodes as fast as we can afford in terms of memory. Then we don't have to worry about partitioning the cluster for various science problems. Every node is intended to be able to handle any research problem."

This high memory capacity is one of the many essential features of Dell EMC PowerEdge platforms, and a capability found in the Intel® Xeon® Scalable Processor family. The Dell EMC PowerEdge family has platforms that support 768 GB of memory to 6 TB of memory.

High-powered research

The computationally driven research conducted at the Flatiron Institute spans the gamut of astrophysics, biology, mathematics and quantum physics. Researchers use the Scientific Computing Core systems to analyze neuronal activity, predict the dynamic behavior of materials and molecules, analyze astronomical events and piece together the mysteries of the origins of life — to highlight just a few examples of the diverse research taking place every day at the Flatiron Institute.

"We have a variety of projects in genomics, neuroscience, systems biology, which have a lot of correlations between microbiome projects and conditions," Fisk says. "We have tasks in complex quantum systems, looking at molecular-level interactions at the quantum level. We have a large astronomy group that deals with simulations of galaxy formation, black-hole simulations and exoplanet searches. And then we have our new applied math group working on accelerating algorithms, machine learning and computer science."

A common thread through this diverse range of research is the need for HPC systems that can quickly process huge amounts of data to accelerate time to insight. That's what the Flatiron Institute is all about.

"One of the things that is special about our nodes is that they have a tremendous amount of memory."

Ian Fisk
Co-director of the SCC

Researchers use the Scientific Computing Core systems to analyze neuronal activity, predict the dynamic behavior of materials and molecules, analyze astronomical events and piece together the mysteries of the origins of life.

DELL EMC

Working with Dell EMC

The Flatiron Institute has worked closely with Dell EMC from its startup days through to the present.

"Dell EMC has been a fantastic partner from the beginning in a variety of areas," Fisk says. "The first one was that they had a product line that spanned all of our needs, from when we were an incredibly small organization all the way up to where we are now."

In addition to its ability to provide technology ranging from the desktop to the data center, Dell EMC offered valuable personnel and expertise for a young organization on a fast growth path. For example, technical experts from Dell EMC helped the institute map out its requirements for servers, processors, I/O, networking, power density, cooling and other components of a robust HPC cluster. "Dell EMC has helped us grow from a very small group up to where we are now by providing access to technical experts," Fisk says. "They helped us determine what would be the appropriate cluster for our physical space and what infrastructure we needed to meet the scientific challenges that we were facing."

That kind of assistance has been one of the keys to the success of the Flatiron Institute and its Scientific Computing Core. Fisk notes that without that expert support, the institute would have had to go looking for a lot of expertise, from thermal specialists to networking professionals.

"Dell EMC has been a very solid partner, basically since we have been an organization," Fisk says. "They have facilitated our growth. They have supported our vision. And they have provided us with not just services and technology but also architectures and expertise. That has allowed us to support the scientific mission of the institute. Without them, we would not be the organization that we are right now."

Intel Inside®. Powerful Solution Outside.

FLATIRON INSTITUTE

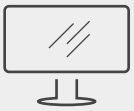

Learn more about Dell EMC advanced computing

Unlock the value of data with artificial intelligence

Share this story

DELL EMC

intel