Unlocking cosmological secrets

Drawing on the power of high performance computing, Durham University and DiRAC scientists are expanding our understanding of the universe and its origins.

Scientific Research | United Kingdom

Business needs

Durham University and the DiRAC HPC facility need leading-edge high performance computing systems to power data- and compute-intensive cosmological research.

Solutions at a glance

- Dell EMC PowerEdge C6525 servers
- AMD® EPYC™ 7H12 processors
- CoolIT® Systems Direct Liquid Cooling technology
- Mellanox® HDR 200 interconnect

Business results

- Providing DiRAC researchers with world-class capabilities
- Accelerating computational cosmology research
- Enabling a much greater level of discovery
- Keeping Durham University at the forefront cosmology research

The COSMA8 cluster has as many as

76,000 cores

The AMD EPYC processor used in the COSMA8 cluster has

64 cores per CPU

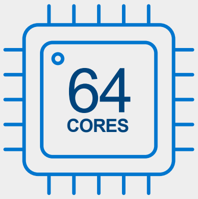

Researching very big questions

In scientific circles, more computing power can lead to bigger discoveries in less time. This is the case at Durham University, whose researchers are unlocking insights into our universe with powerful high performance computing clusters from Dell Technologies.

What is the universe? What is it made of? What is dark matter? What is dark energy? These are the types of questions explored by the distinguished researchers working with Durham University's Institute for Computational Cosmology (ICC). This level of research requires massive amounts of computational power, and Durham University gets exactly that in its legendary COSMA supercomputer.

The name COSMA, which is short for Cosmology Machine, reflects the core mission of this evolving system — researching very big cosmological questions. Cosmology is just one of the research areas of DiRAC, the UK national HPC facility of which COSMA is a part.

A purpose-built machine

Durham University and DiRAC rolled out the latest iteration of its COSMA supercomputer, known as COSMA8, in September 2020. The cluster is based on Dell EMC PowerEdge C6525s. With four servers in 2U, each chassis provides up to 512 processing cores, with up to 3200MT/s memory speed to reduce latency and PCIe Gen4 to transfer data faster.

The PowerEdge C6525 server configuration used in COSMA8 includes Direct Liquid Cooling (DLC) technology from CoolIT Systems. This groundbreaking technology leverages the exceptional thermal conductivity of liquid to provide dense, concentrated cooling to targeted areas. By using DLC and warm water, the dependence on fans and expensive air handling systems is drastically reduced. DLC technology enables higher rack density, overall reduced power usage and significantly higher performance potential.

The first installation of COSMA8 has 32 compute nodes, used for testing and benchmarking, and for getting codes up to scratch. This is now being expanded to 360 nodes, with plans to increase up to 600 nodes, over the next year.

"This will primarily be used for testing, for getting code up to scratch," says Dr. Alastair Basden, technical manager for the COSMA high performance computing system at Durham University. "But then, coming very soon, we hope to place an order for another 600 nodes."

Deploying the full 600 nodes would provide a whopping 76,800 cores. "It's a big uplift," Dr. Basden notes. "COSMA7 currently is about 12,000 cores. So COSMA8 will have six times as many cores. It will have more than twice as much DRAM as previous systems, which means that we'll be able to more than double the size of the simulations. And having six times as many cores means we'll be able to get through these datasets much more quickly."

Dr. Basden is specifying 1TB of memory per server, giving a sizeable 8GB memory per core. That memory will help accelerate the data- and compute-intensive work that comes with cosmological studies. The speed from having eight-channel 3,200MHz memory is valuable too, Dr. Basden notes.

"Because we have a large amount of data, it's important to be able to pull it through the processors quickly," he says.

COSMA8's requirements are specific to the kind of workloads supported by the DiRAC Memory Intensive service. "It's what we call a capability system," Dr. Basden explains. "This is a system that's designed to offer a capability that wouldn't otherwise be available unless you go to a much larger setup. It's a capability system primarily because of the large amounts of RAM. If you're doing large-scale cosmological simulations of the universe, you need a lot of RAM. They can have runtimes of months, and after they've produced their data, including snapshots of the universe at different time-steps and red-shifts, years are spent in processing and analysis."

An epic system

In the course of developing the specifications for COSMA8, Dr. Basden was particularly interested in AMD EPYC processors, due to their higher core density and single NUMA memory domain.

"We got access to the Dell Technologies HPC & AI Innovation Lab in Austin so that we could do benchmarking on these processors," Dr. Basden says. A large fraction of COSMA's simulation work is performed with software called SWIFT (SPH With Inter-dependent Fine-grained Tasking). Dr. Basden's team used sample cosmology datasets with SWIFT to test the processor performance.

A national HPC facility

The memory-intensive service Durham University is part of the U.K.'s DiRAC, a distributed national HPC facility that encompasses five academic institutions: Cambridge, Durham, Edinburgh, Leicester and University College London. DiRAC, which stands for "Distributed Research utilising Advanced Computing," is devoted to theoretical modeling and HPC-based research in particle physics, astrophysics, cosmology and nuclear physics. DiRAC arms researchers with a variety of compute resources, matching machine architecture to the algorithm design and requirements of the research problems to be solved.

"It ran as we would hope," he says, noting that tests were performed against other processors. "The performance of a single core was better than COSMA7, and when you've got more cluster cores, it's a no-brainer. The extra core count means the AMD EPYC processors could be used to crunch scientific problems significantly faster," Dr. Basden says.

For COSMA8, Dr. Basden opted for dual 280-watt AMD EPYC 7H12 processors per node with a 2.6GHz base clock frequency and 64 cores, installed in a Dell EMC PowerEdge C-series server.

"We wanted a large number of cores per node, because we could cut down on the amount of internode communication," Dr. Basden explains. "But because there are parts of the code that don't parallelize 100 percent, we also wanted high clock rates, so that the lower-threaded parts of the code stood up well, which meant the 7H12 processor would be the best option."

Unlocking the secrets of the universe

With its rich mix of technologies, COSMA8 delivers the robust high performance compute and storage resources that scientists require when they are trying to unlock the secrets of a very big and complex universe. This work involves enormous amounts of data that must be processed at high speeds. A single simulation can produce hundreds of terabytes of data.

"About 75 percent of the universe is made up of dark matter that we don't understand," Dr. Basden explains. "By running simulations, we are able to find out more about it. And, of course, when we do that, we begin to understand more about what the universe is made from."

Ultimately, COSMA8 will enable a much greater level of discovery, in comparison to simulations run on earlier-generations of the COSMA supercomputer. The increased core density, larger amounts of faster DRAM, and faster PCIe connectivity together add up to a potent boost in performance.

"It will mean we can run massively more detailed simulations, which we can compare much better with observations from telescopes," Dr. Basden explains. "This will help us to understand the meaning of the universe, dark matter, dark energy and how the universe was formed. It's really going to help us drill down to a fundamental understanding of the world that we live in."

Learn more about Dell EMC advanced computing

Unlock the value of data with artificial intelligence

Share this story