

Sharing the wealth of HPC-driven apps with flexible as-a-service models

[HPC wire hpcwire.com/2021/10/04/sharing-the-wealth-of-hpc-driven-apps-with-flexible-as-a-service-models/](https://www.hpcwire.com/2021/10/04/sharing-the-wealth-of-hpc-driven-apps-with-flexible-as-a-service-models/)

October 4, 2021



October 4, 2021

Sponsored Content by Dell Technologies | Intel

Growing numbers of IT shops are delivering high performance computing systems and the applications they support as on-demand services accessed via cloud connections.

In years past, high performance computing shops functioned as resources dedicated to the needs of scientists and engineers whose work required the computational power of a supercomputer. But not so today. In modern HPC shops, IT administrators function as service providers who cater to the needs of many rank-and-file HPC users who need more processing power than they can get on a desktop or laptop system.

Take the case of the University of Michigan, where the HPC specialists run an academic supercomputing center like an enterprise. They make the resources of their Great Lakes Supercomputer available to approximately 2,500 users who run hundreds of different

applications. Even better, the reach of these HPC resources, based on systems from Dell Technologies, extends out into the community. For example, system users include MCity, a public-private initiative that brings together industry, government and academia to advance transportation safety, sustainability and accessibility.

The IT leaders at the University of Florida act in a similar manner with their UF Innovate hub, the University's technology business incubator. Among other support functions, UF Innovate gives start-up companies access to a Dell Technologies supercomputer for high performance computation, data visualization and analysis. This HPC system, known as HiPerGator, accelerates diverse research workloads with the power of more than 46,000 CPU cores.

Bringing HPC to the masses — via cloud connections

So how do today's IT shops extend the goodness of HPC to thousands of users? Increasingly, the answer is a multi-cloud environment that gives users web-based access to a wide variety of HPC systems, both in on-premises private clouds and in public clouds.

The University of Michigan, for example, provides its academic community with easy access to on-premises HPC clusters via its [Open OnDemand](#) program and to public cloud computing platforms via its [ITS Cloud Services](#) program.

The University of Florida does much the same via its version of the [Open OnDemand](#) service, which connects the university community to the HiPerGator cluster to accelerate compute- and data-intensive scientific workloads. To work on HiPerGator, users connect to the system from a local computer via an SSH terminal session or through web application interfaces provided by the UF Research Computing team.

Meanwhile, out on the West Coast, the San Diego Supercomputer Center (SDSC) at the University of California, San Diego, offers its cloud-based [OnDemand system](#) to users who require immediate access to a supercomputer for event-driven science. Urgent applications that might make use of the OnDemand system range from making movies of earthquakes to providing near real-time warnings based on predictions about the path of tornados, hurricanes and toxic plumes.

Automating processes and educating users

This brings us to another question. With such a vast variety of jobs coming in from all over the place, how do HPC shops direct those jobs to the right infrastructure? In a word: automation.

SDSC, for example, automates the process with software that automatically determines where a job will run, matching the right job with the right IT resources. To enable this process with its Expanse supercomputer from Dell Technologies, SDSC is pioneering [composable HPC systems](#) that dynamically allocate resources tailored to individual workloads.

In addition to automating access to the right supercomputing resources, HPC shops are focusing heavily on education and training to help their users gain the greatest value from HPC clusters. This includes detailed online instructions, expert guidance and tips from people who have been there before. They all have teams to help people optimize their code, optimize the infrastructure and optimize the results.

This is the case at the University of Michigan, where the [Advanced Research Computing – Technology Services](#) team provides workshops on such topics as GPU programming, writing machine learning code, choosing machine learning tools, and building and training deep learning models.

Everything is an HPC workload

Research computing is all about collaborating to make the next big scientific discovery or technological innovation. And to make this research happen, we need high performance computing systems. Whether it's unlocking the secrets of a deadly virus like SARS-CoV-2 or simulating the consequences of an earthquake, HPC is now an essential tool for scientific research.

HPC is also an essential tool for new applications for data analytics, artificial intelligence, machine learning and deep learning. For these and other applications, HPC technologies serve as the engine under the hood, helping us turn raw data into valuable insights. To that end, HPC systems provide big memory to enable applications like image recognition, visualizations and molecular dynamics simulations. They offer GPUs for training deep learning workloads, and CPUs for machine learning and inferencing jobs that perform well in HPC systems.

And increasingly, it's all available under as-a-service approaches that incorporate things like data pipeline tools for ingesting and processing data from a variety of sources, systems tailored to calculations that require large amounts of physical memory, and storage clusters that hold petabytes of data for solving scientific problems. You name it, and you can probably get it as a service, via a cloud interface.

Gaining hands-on experience

To reduce the risk associated with new technology investments and improve speed of implementation, Dell Technologies invites customers to experience HPC-driven solutions firsthand in a global network of dedicated facilities. These [Customer Solution Centers](#) are trusted environments where world-class IT experts collaborate to share best practices, facilitate discussions of effective business strategies, and use briefings, workshops and proofs-of-concept to accelerate IT initiatives.

Other Dell Technologies resources available to support modern computing initiatives include the [HPC & AI Innovation Lab](#), [HPC & AI Centers of Excellence](#) and the [Dell Technologies HPC Community](#).

For a closer look the HPC environments of the institutions discussed here, check out these assets:

- University of Michigan case study: “[The University of Michigan Gets Serious About Supercomputing ROI](#).”
- University of Florida video: “[Florida HiPerGator gives researchers and start-ups a supercomputing boost](#).”
- San Diego Supercomputer Center case study: “[A Supercomputing Condo](#).”