

February 12, 2020 | By Jan Zverina


# Second GPU Cloudburst Experiment Yields New Findings

## Results help pave the way for large-scale cloud computing

In late 2019, researchers at the San Diego Supercomputer Center (SDSC) and the Wisconsin IceCube Particle Astrophysics Center (WIPAC) caught the attention of the high-performance computing community and top commercial cloud providers by successfully completing a bold experiment that marshalled all globally available-for-sale GPUs (graphics processing units) for a brief run which proved it is possible to elastically burst to very large scales of GPUs using the cloud, even in this pre-exascale era of computing.

That experiment also proved that it is possible to conduct a massive amount of data processing within a very short period – an advantage for research projects that must meet a tight deadline. It also showed that such bursting of massive amounts of data – in this case data generated by the [IceCube Neutrino Observatory](#), an array of 5,160 optical sensors deep within a cubic kilometer of ice at the South Pole – is suitable for solving a wide range of challenges across not only astronomy but many other science domains.

Completed just prior to the opening of the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC19) in Denver, CO last November, the first experiment was coordinated by Frank Würthwein, SDSC lead for high-throughput computing; Igor Sfiligoi, SDSC's lead scientific software developer for high-throughput computing; Benedikt Riedel, global computing coordinator for the IceCube Neutrino Observatory and computing manager at [WIPAC](#); and David Schultz, a production software manager with IceCube.



*Using only the most cost-effective instances in spot/preemptible mode, SDSC and WIPAC researchers were able to provision and sustain 15,000 GPUs from cloud regions all over the world, for an equivalent of 170 PFLOP32s, significantly more than the largest National Science Foundation-funded supercomputers. Image courtesy of Igor Sfiligoi, San Diego Supercomputer Center, UC San Diego.*

Fast forward to early February 4, 2020, when the same research team conducted a second experiment with a fraction of the remaining funding left over from a modest [National Science Foundation EAGER grant](#). As before, the researchers used cloud resources from [Amazon Web Services \(AWS\)](#), [Microsoft Azure](#), and [Google Cloud Platform \(GCP\)](#), but limited themselves to only the most effective cloud instance types in either spot or preemptible mode. Moreover, instead of using a dedicated cloud only setup, this time they added cloud resources to the existing ‘on-prem’ resources available through the [Open Science Grid \(OSG\)](#), the [Extreme Science and Engineering Discovery Environment \(XSEDE\)](#), and the [Pacific Research Platform \(PRP\)](#). [HTCondor](#) was again used as the workload management system.

“We drew several key conclusions from this second demonstration,” said SDSC’s Sfiligoi. “We showed that the cloudburst run can actually be sustained during an entire workday instead of just one or two hours, and have moreover measured the cost of using only the two most cost-effective cloud instances for each cloud provider.”

The team managed to reach and sustain a plateau of about 15,000 GPUs, or 170 PFLOP32s (i.e. fp32 PFLOPS) using the peak fp32 FLOPS provided by NVIDIA specs. The cloud instances were provisioned from all major geographical areas, and the total integrated compute time was just over one fp32 exaFLOP hour. The total cost of the cloud run was roughly \$60,000.

In the second experiment, which was about eight hours long versus less than two hours in the first experiment, the IceCube Neutrino Observatory processed some 151,000 jobs, up from about 101,000 in the first burst.

“This means that the second IceCube cloud run produced 50% more science, even though the peak was significantly lower,” explained Sfiligoi, who also noted that the latter experiment added OSG, XSEDE, and PRP’s [Kubernetes resources](#), effectively making it a hybrid-cloud setup unlike the first time, when it was purely cloud-based.

The experiment also showed that the most cost-effective cloud instances are those providing the NVIDIA Tesla T4 GPUs. They are about three times more cost-effective for the IceCube project compared to the next best option, the instances providing the NVIDIA Tesla V100 GPUs.

Moreover, unlike the earlier experiment, the team fetched data directly from IceCube’s home location at the University of Wisconsin–Madison using straight HyperText Transfer Protocol (HTTP), the standard application-level protocol used for exchanging files on the World Wide Web.

“This second cloud burst was performed on a random day, during regular business hours, adding cloud resources to the standard IceCube on-prem production infrastructure,” said SDSC’s Würthwein. “It’s the kind of cloud burst that we could do routinely, a few times a week if there were principal investigators willing to pay the bill to do so.”

Both experiments were funded by NSF grants [OAC-1941481](#), [MPS-1148698](#), [OAC-1841530](#), [OPP-1600823](#), [OAC-190444](#), and [OAC-1826967](#).

---

#### MEDIA CONTACT

**Jan Zverina**, 858-534-5111, [jzverina@sdsc.edu](mailto:jzverina@sdsc.edu)

**Madeleine O’Keefe**, WIPAC Communications, 608-515-3831 [mokeefe@icecube.wisc.edu](mailto:mokeefe@icecube.wisc.edu)

UC San Diego’s [Studio Ten 300](#) offers radio and television connections for media interviews with our faculty, which can be coordinated via [studio@ucsd.edu](mailto:studio@ucsd.edu). To connect with a UC San Diego faculty expert on relevant issues and trending news stories, visit <https://ucsdnews.ucsd.edu/media-resources/faculty-experts>.