

DEPARTMENT: LEADERSHIP COMPUTING

The COVID-19 High-Performance Computing Consortium

Jim Brase, *Lawrence Livermore National Laboratory, Livermore, CA, 94550, USA*

Nancy Campbell, *IBM Research, Yorktown Heights, NY, 10598, USA*

Barbara Helland and Thuc Hoang, *Department of Energy, Washington, DC, 20585, USA*

Manish Parashar, *University of Utah, Salt Lake City, UT, 84112, USA*

Michael Rosenfield  and James Sexton, *IBM Research, Yorktown Heights, NY, 10598, USA*

John Towns, *University of Illinois, Urbana, IL, 61801, USA*

In March of 2020, recognizing the potential of High Performance Computing (HPC) to accelerate understanding and the pace of scientific discovery in the fight to stop COVID-19, the HPC community assembled the largest collection of worldwide HPC resources to enable COVID-19 researchers worldwide to advance their critical efforts. Amazingly, the COVID-19 HPC Consortium was formed within one week through the joint effort of the Office of Science and Technology Policy (OSTP), the U.S. Department of Energy (DOE), the National Science Foundation (NSF), and IBM to create a unique public-private partnership between government, industry, and academic leaders. This article is the Consortium's story—how the Consortium was created, its founding members, what it provides, how it works, and its accomplishments. We will reflect on the lessons learned from the creation and operation of the Consortium and describe how the features of the Consortium could be sustained as a National Strategic Computing Reserve to ensure the nation is prepared for future crises.

In March of 2020, recognizing the potential of High-Performance Computing (HPC) to accelerate understanding and the pace of scientific discovery in the fight to stop COVID-19, the HPC community assembled the largest collection of worldwide HPC resources to enable COVID-19 researchers worldwide to advance their critical efforts. Amazingly, the COVID-19 HPC Consortium was formed within one week through the joint effort of the Office of Science and Technology Policy (OSTP), the U.S. Department of Energy (DOE), the National Science Foundation (NSF), and IBM. The Consortium created a unique public-private partnership between government, industry, and academic leaders to provide access to advanced HPC and cloud computing systems and data

resources, along with critical associated technical expertise and support, at no cost to researchers in the fight against COVID-19. The Consortium created a single point of access for COVID researchers. This article is the Consortium's story—how the Consortium was created, its founding members, what it provides, how it works, and its accomplishments. We will reflect on the lessons learned from the creation and operation of the Consortium and describe how the features of the Consortium could be sustained as a *National Strategic Computing Reserve* (NSCR) to ensure the nation is prepared for future crises.

CREATION OF THE CONSORTIUM

As the pandemic began to significantly accelerate in the United States, on March 11 and 12, 2020, IBM and the HPC community started to explore ways to organize efforts to help in the fight against COVID-19. IBM had years of experience with HPC, knew its capabilities to help solve hard problems, and had the vision of organizing the HPC

Members and Affiliates

Academia

- Massachusetts Institute of Technology
- MGHPC
- Rensselaer Polytechnic Institute
- University of Illinois
- University of Texas at Austin
- University of California - San Diego
- Carnegie Mellon University
- University of Pittsburgh
- Indiana University
- University of Wisconsin-Madison
- Ohio Supercomputing Center
- UK Digital Research Infrastructure
- CSCS – Swiss National Supercomputing Centre
- SNIC PDC – Swedish National Infrastructure for Computing, Center for High Performance Computing

Federal Agencies

- NASA
- National Science Foundation
 - XSEDE
 - Pittsburgh Supercomputing Center
 - Texas Advanced Computing Center (TACC)
 - San Diego Supercomputer Center (SDSC)
 - National Center for Supercomputing Applications (NCSA)
 - Indiana University Pervasive Technology Institute (IUPTI)
 - Open Science Grid (OSG)
 - National Center for Atmospheric Research (NCAR)

Industry

- IBM
- Amazon Web Services
- AMD
- D.E.Shaw Research
- Dell
- Google Cloud
- Hewlett Packard Enterprise
- Intel
- Microsoft
- NVIDIA

Department of Energy National Laboratories

- Argonne National Laboratory
- Idaho National Laboratory
- Lawrence Berkeley National Laboratory
- Oak Ridge National Laboratory
- Lawrence Livermore National Laboratory
- Los Alamos National Laboratory
- Sandia National Laboratories

International

- Korea Institute of Science and Technology Information (KISTI)
- Ministry of Education, Culture, Sports, Science and Technology (MEXT) - Japan
 - RIKEN Center for Computational Science (R-CCS)

Affiliates

- Atrio
- Data Expedition
- Flatiron
- Fluid Numerics
- Immortal Hyperscale InterPlanetary Fabrics
- MathWorks
- Raptor Computer Systems
- SAS
- The HDF Group

FIGURE 1. Consortium members and affiliates as of July 7, 2021.

community to leverage its substantial computing capabilities and resources to accelerate progress and understanding in the fight against COVID-19 by connecting COVID-19 researchers with organizations that had significant HPC resources. At this point in the pandemic, the efforts in the DOE, NSF, and other organizations within the U.S. Government, as well as around the world, were independent and *ad hoc* in nature.

It was clear very early on that a broader and more coordinated effort was needed to leverage existing efforts and relationships to create a unique HPC collaboration.

Early in the week of March 15, 2020, leadership at the DOE Labs and at key academic institutions were supportive of the vision: very quickly create a public-private consortium between government, industry, and academic leaders to aggregate compute time and resources on their supercomputers and to make them freely available to aid in the battle against the virus. On March 17, the White House OSTP began to actively support the creation of the Consortium, along with DOE and NSF leadership. The NSF recommended leveraging their Extreme Science and Engineering Discovery Environment (XSEDE) Project¹ and its XSEDE Resource Allocations System (XRAS) that handles nearly 2000 allocation requests annually² to serve as the access point for the proposals. Recognizing that time was critical, a team, now comprising IBM, DOE, OSTP, and NSF, had been formed with the goal of creating the Consortium in less than a week! Remarkably,

the Consortium met that goal without formal legal agreements. Essentially, all potential members agreed to a simple statement of intent that they would provide their computing facilities' capabilities and expertise at no cost to COVID-19 researchers, that all parties in this effort would be participating at risk and without liability to each other, and without any intent to influence or otherwise restrict one another.

From the beginning, it was recognized that communication and expedient creation of a community around the Consortium would be key. Work began on the Consortium website^a the following day. The Consortium Executive Committee was formed to lay the groundwork for the operations of the Consortium. By Sunday, March 22, the XSEDE Team instantiated a complete proposal submission and review process that was hosted under the XSEDE website^b and provided direct access to the XRAS submission system, which was ready to accept proposal submissions the very next day.

Luckily, the Consortium assembled swiftly because OSTP announced that the President would introduce the concept of the Consortium at a news conference on March 22. Numerous news articles came out after the announcement that evening. The Consortium

^a<https://covid19-hpc-consortium.org>

^b<https://www.xsede.org/covid19-hpc-consortium>

Consortium organizational structure

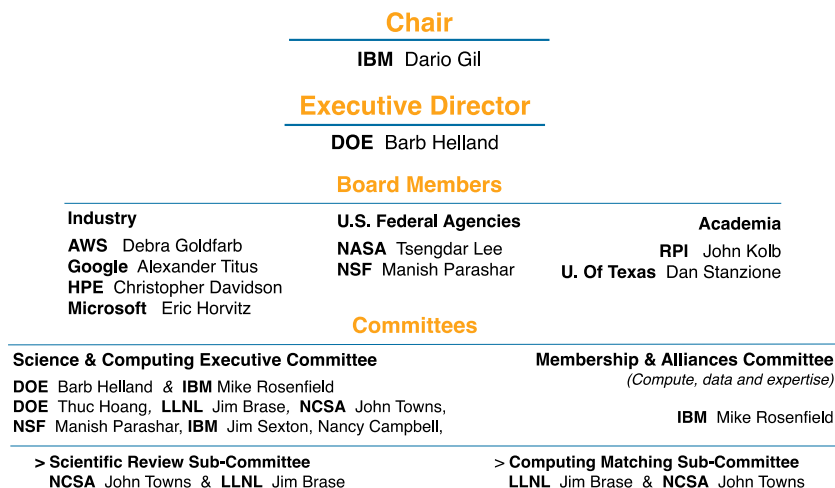


FIGURE 2. Consortium organizational structure as of July 7, 2021.

became a reality when the website^c went live the next day, followed by additional press releases and news articles. The researchers were ready—the first proposal was submitted on March 24, and the first project was started on March 26, demonstrating our ability to connect researchers with resources in a matter of days—an exceptionally short time for such processes typically. Subsequently, 50 proposals were submitted by April 15 and 100 by May 9.

A more detailed description of the Consortium's creation can be found in the IEEE Computer Society Digital Library at <https://doi.ieeecomputersociety.org/10.1109/MCSE.2022.3145608>. An extended version of this article can be found on the Consortium website.^a

CONSORTIUM MEMBERS AND CAPABILITIES

The Consortium initially provided access to over 300 petaflops of supercomputing capacity provided by the founding members: IBM; Amazon Web Services; Google Cloud; Microsoft; MIT; RPI; DOE's Argonne, Lawrence Livermore, Los Alamos, Oak Ridge, and Sandia National Laboratories; NSF and its supported advanced computing resources, advanced cyberinfrastructure, services, and expertise; and NASA.

Within several months, the Consortium grew to 43 members (see Figure 1) from the United States, and around the world (the complete list can be found at <https://covid19-hpc-consortium.org/>) representing ac-

cess to over 600 petaflops of supercomputing systems, over 165,000 compute nodes, more than 6.8 million compute processor cores, and over 50,000 GPUs, representing access to systems worth billions of dollars. In addition, the Consortium collaborated with two other worldwide initiatives: The EU PRACE COVID-19 Initiative and a COVID-19 initiative at the National Computational Infrastructure Australia and Pawsey Supercomputing Centre.^d The Consortium also added nine affiliates (also listed and described at websites^{a,c}) who provided expertise and supporting services to enable researchers to start up quickly and run more efficiently.

GOVERNANCE AND OPERATIONS

Even though there were no formal agreements between the Consortium members, an agile governance model was developed as shown in Figure 2. An Executive Board, comprised of a subset of the founding members, oversees all aspects of the Consortium and is the final decision-making authority. Initially, the Executive Board met weekly and now meets monthly. The Board reviews progress, reviews recommendations for new members and affiliates, and provides guidance on future directions and activities of the Consortium to the Executive Committee. The Science and Computing Executive Committee, which reports to the Executive Board, (see also Figure 2) is responsible for day-to-day operations of the Consortium, overseeing the review and computer matching process, tracking project progress, maintaining/updating the website, highlighting the Consortium results (for

^c<https://covid19-hpc-consortium.org/news>

^d<https://covid19-hpc-consortium.org/collaborations>

example, with blogs and webinars), and determining/proposing next steps for Consortium activities.

The Scientific Review and the Computing Matching Sub-Committees play a crucial role in the success of the Consortium. The Scientific Review team—comprised of subject matter experts from members of the research community and coming from many organizations^e—reviews proposals for merit based on the review criteria and guidance^b provided to proposers, and recommends appropriate proposals to the Computing Matching Sub-Committee. The Computing Matching Sub-Committee team, comprised of representatives of Consortium members providing resources, matches the computer needs from recommended proposals with either the proposer's requested site or other appropriate resources. Once matched, the researcher needs to go through the standard onboarding/approval process at the host site to gain access to the system. Initially, we expected that the onboarding/approval process would be time consuming (since this was the only time where actual agreements had to be signed), but those executing the onboarding processes with the various member compute providers worked diligently to prioritize these requests, and thus, it typically takes only a day or two. As a result, once approved, projects are up and running very rapidly.

The Membership Committee reviews requests for organizations and individuals to become members or affiliates to provide additional resources to the Consortium. These requests are in turn sent to OSTP for vetting, with the Executive Committee making final recommendations to the Executive Board for approval.

PROJECT HIGHLIGHTS

The goal of the Consortium is to provide state-of-the-art HPC resources to scientists all over the world to accelerate and enable R&D that can contribute to pandemic response. Over 115 projects have been supported, covering a broad spectrum of technical areas ranging from understanding the SARS-CoV-2 virus and its human interaction to optimizing medical supply chains and resource allocations, and have been organized into a taxonomy of areas consisting of basic science, therapeutic development, and patients.

Consortium projects have produced a broad range of scientific advances. The projects have collectively produced a growing number of publications, datasets, and other products (more than 70 as of the end of calendar year 2021), including two journal covers.^f A more

detailed description of the Consortium's Project Highlights and Operational Results can be found at <https://covid19-hpc-consortium/projects> and <https://covid19-hpc-consortium.org/blog>, respectively.

While Consortium projects have contributed significantly to scientific understanding of the virus and its potential therapeutics, direct and near-term impact on the course of the pandemic has been mixed. There are cases of significant impact, but, overall, the patient-related applications that have the most direct path to near-term impact have been less successful. It may be possible to attribute this to the lower level of experience in HPC that is typical of these groups, but patient data availability and use restrictions and the lack of connection to front-line medical and response efforts are also important factors. These are issues that will need to be addressed in planning for future pandemics or other crisis response programs.

LESSONS LEARNED FROM THE COVID-19 HPC CONSORTIUM

The COVID-19 pandemic has shown that the existence of an advanced computing infrastructure is not sufficient on its own to effectively support the national and international response to a crisis. There must also be mechanisms in place to rapidly make this infrastructure broadly accessible, which includes not only the computing systems themselves, but also the human expertise, software, and relevant data to rapidly enable a comprehensive and effective response.

The following are the key lessons learned.

- ▶ The ability to leverage existing processes and tools (e.g., XSEDE) was critical and should be considered for future responses.
- ▶ Engagement with the stakeholder community is an area that should be improved based on the COVID-19 experience. For example, early collaboration with the NIH, FEMA, CDC, and medical provider community could have significantly increased impact in the patient care and epidemiology areas. Having pre-negotiated agreements with these and similar stakeholders will be important going forward.
- ▶ Substantial time and effort are required to make resources and services available to researchers so that they can do their work. A standing capability to support the proposal submission and review process, as well as coordinating with service providers to provide the necessary access to resources and services, would have been helpful.
- ▶ It would have been beneficial to have had use authorizations in place for the supercomputers

^e<https://covid19-hpc-consortium.org/who-we-are>

^f<http://doi.org/10.1021/acs.jcim.0c00929> and <http://doi.org/10.1039/d0cp03145c>

and resources provided by U.S. Government organizations.

- › While the proposal review and award process ran sufficiently well, there was no integration of the resources being provided and the associated institutions into an accounting and account management system. Though XSEDE also operates such a system, there was no time to integrate the resources into that system. This would have greatly facilitated the matching and onboarding processes. It also would have provided usage data and insight into resource utilization.
- › Given the absence of formal operating and partnership agreements in the Consortium and the mix of public and private computing resources, the work supported was limited to open, publishable activities. This inability to support proprietary work likely reduced the effectiveness and impact of the Consortium, particularly in support for private-sector work on therapeutics and patient care. A lightweight framework for supporting proprietary work and associated intellectual property requirements would increase the utility of responses for similar future crises.

NEXT STEP: THE NSCR

Increasingly, the nation's advanced computing infrastructure—and access to this infrastructure, along with critical scientific and technical support in times of crisis—is important to the nation's safety and security.^{g,h} Computing is playing an important role in addressing the COVID-19 pandemic and has, similarly, assisted in national emergencies of the recent past, from hurricanes, earthquakes, and oil spills, to pandemics, wildfires, and even rapid turnaround modeling when space missions have been in jeopardy. To improve the effectiveness and timeliness of these responses, we should draw on the experience and the lessons learned from the Consortium in developing an organized and sustainable approach for applying the nation's computing capability to future national needs.

We agree with the rationale behind the creation of an NSCR as outlined in the recently published OSTP Blueprint to protect our national safety and security by establishing a new public-private partnership, the NSCR: a coalition of experts and resource providers (compute, software, data, and technical expertise)

^gThe U.S. needs a National Strategic Computing Reserve, *Scientific American*, June 2, 2021. [Online]. Available: <https://www.scientificamerican.com/article/the-u-s-needs-a-national-strategic-computing-reserve/>

^h<https://covid19-hpc-consortium.org/blog/national-strategic-computing-reserve>

spanning government, academia, nonprofits/foundations, and industry supported by appropriate coordination structures and mechanisms that can be mobilized quickly and efficiently to provide critical computing capabilities and services in times of urgent needs.

Figure 3 shows a transition from a pre-COVID *ad hoc* response to crises to the Consortium and then to an NSCR.ⁱ

Principal Functions of the NSCR

In much the same way as the Merchant Marine^j maintains a set of “ready reserve” resources that can be put to use in wartime, the NSCR would maintain reserve computing capabilities for urgent national needs. Like the Merchant Marine, this effort would involve building and maintaining sufficient infrastructure and human capabilities, while also ensuring that these capabilities are organized, trained, and ready in the event of activation. The principal functions of the NSCR are proposed to be as follows:

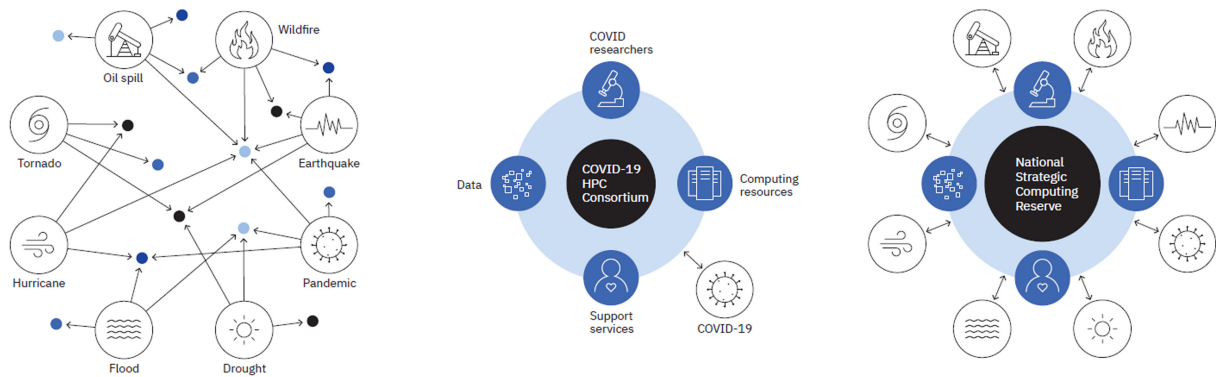
- › recruit and sustain a group of advanced computing and data resource and service provider members in government, industry, and academia;
- › develop relevant agreements with members, including provisions for augmented capacity or cost reimbursement for deployable resources, for the urgent deployment of computing and supporting resources and services, and for provision of incentives for nonemergency participation;
- › develop a set of agreements to enable the Reserve to collaborate with domain agencies and industries in preparation for and execution of Reserve deployments;
- › execute a series of preparedness exercises on some frequency basis to test and maintain the Reserve;
- › establish processes and procedures for activating and operating the national computing reserve in times of crisis;
- › during a crisis,
 - › execute procedures to review and prioritize projects and to allocate computing resources to approved projects;
 - › track project progress and disseminate products and outputs to ensure effective use and impact;
 - › participate in the broader national response as an active partner.

ⁱ<https://covid19-hpc-consortium.org/blog/national-strategic-computing-reserve>

^jUnited States Merchant Marine – 46 U.S.C. §§ 861-889 Merchant Marine Act.


National Strategic Computing Reserve

Accelerating our Nation's Ability to Respond



Pre-COVID (legacy)

During a crisis incident, providers are found ad-hoc:



The Consortium (current)

The Covid-19 High Performance Computing Consortium provides an organized resource for:

- rapid access to expertise and experience
- research projects, computing, and services
- the most powerful, advanced computing

This environment:

- provides integrated, shared resources
- allows open data and development
- accelerates the pace of discovery
- breaks barriers between government, industry, and academia
- enables and informs a national response

NSCR Proposal (future)

The NSCR expands and builds upon the Consortium model for future crises:

- investigates and implements approaches
- leverages existing relationships
- updates frequently
- gathers, curates, maintains critical databases
- builds on successful experience
- executes preparedness exercises
- recruits and sustains members
- develops agreements with members
- executes procedures and tracks progress
- sustained, coordinated investment

FIGURE 3. Potential path from a pre-COVID to the NSCR.

CONCLUSION

The COVID-19 HPC Consortium has been in operation for almost two years^k and has enabled over 115 research projects investigating multiple aspects of COVID-19 and the SARS-CoV-2 coronavirus. To maximize impact going forward, the Consortium has transitioned to a focus on the following:

- 1) proposals in specific targeted areas;
- 2) gathering and socializing results from current projects;
- 3) driving the establishment of an NSCR.

New project focus areas target having an impact in a six-month time period and the Consortium is particularly, though not exclusively, interested in projects focused on understanding and modeling patient response to the virus using large clinical datasets;

learning and validating vaccine response models from multiple clinical trials; evaluating combination therapies using repurposed molecules; mutation understanding and mitigation methods; and epidemiological models driven by large multimodal datasets.

We have drawn on our experience and lessons learned through the COVID-19 HPC Consortium, and on our observation of how the scientific community, federal agencies, and healthcare professionals came together in short order to allow computing to play an important role in addressing the COVID-19 pandemic. We have also proposed a possible path forward, the NSCR, for being better prepared to respond to future national emergencies that require urgent computing, ranging from hurricanes and earthquakes to pandemics and wildfires. Increasingly, the nation's computing infrastructure—and access to this infrastructure along with critical scientific and technical support in times of crisis—is important to the nation's safety and security, and its response to natural disasters, public health emergencies, and other crises.

^k<https://covid19-hpc-consortium.org/blog/a-year-on-hpc-consortium-national-strategic-computing-reserve>

ACKNOWLEDGMENTS

The authors would like to thank the past and present members of the Consortium Executive Board for their guidance and leadership. In addition, the authors would like to thank Jake Taylor and Michael Kratsios formerly from OSTP, Dario Gil from IBM, and Paul Dabbar formerly from DOE, for their key roles in helping make the creation and operation of the Consortium possible. The authors also would like to thank Corey Stambaugh from OSTP for his leadership role on the Consortium membership committee. Furthermore, the authors would also like to thank all the members and affiliate organizations from academia, government, and industry who contributed countless hours of their time along with their compute resources. In addition, the service provided by researchers across many institutions as scientific reviewers are critical is selecting appropriate projects and their time and efforts are greatly appreciated, and, of course, they also want to thank the many researchers who did such outstanding work, leveraging the Consortium, in the fight against COVID-19.

REFERENCES

1. J. Towns *et al.*, "XSEDE: Accelerating scientific discovery," *Comput. Sci. Eng.*, vol. 16, no. 5, pp. 62–74, Sep./Oct. 2014, doi: [10.1109/MCSE.2014.80](https://doi.org/10.1109/MCSE.2014.80).
2. R. Light, A. Banerjee, and E. Soriano, "Managing allocations on your research resources? XCRAS is here to help!" in *Proc. Pract. Experience Adv. Res. Comput.*, 2018, pp. 1–4, doi: [10.1145/3219104.3229238](https://doi.org/10.1145/3219104.3229238).

JIM BRASE is currently a Deputy Associate Director for Computing with Lawrence Livermore National Laboratory (LLNL), Livermore, CA, USA. He leads LLNL research in the application of high-performance computing, large-scale data science, and simulation to a broad range of national security and science missions. Jim is a co-lead of the ATOM Consortium for computational acceleration of drug discovery, and on the leadership team of the COVID-19 HPC Consortium. He is currently leading efforts on large-scale computing for life science, biosecurity, and nuclear security applications. In his previous position as an LLNL's deputy program director for intelligence, he led efforts in intelligence and cybersecurity R&D. His research interests focus on the intersection of machine learning, simulation, and HPC. Contact him at brase1@llnl.gov.

NANCY CAMPBELL is responsible for the coordinated execution of the IBM Research Director's government engagement agenda and resulting strategic partnerships within and across industry, academia, and government, including the COVID-19

HPC Consortium and International Science Reserve. Prior to this role, she was the program director for IBM's COVID-19 Technology Task Force, responsible for developing and delivering technology-based solutions to address the consequences of COVID-19 for IBM's employees, clients, and society-at-large. Previously, she led large multidisciplinary teams in closing IBM's two largest software divestitures for an aggregate value in excess of \$2.3 billion, and numerous strategic intellectual property partnerships for an aggregate value in excess of \$3 billion. Prior to joining IBM, she was a CEO for one of Selby Venture Partners portfolio companies and facilitated the successful sale of that company to its largest channel partner. She attended the University of Southern California, Los Angeles, CA, USA, and serves as an IBM's executive sponsor for the USC Master of Business for Veterans program. Contact her at nncampbe@us.ibm.com.

BARBARA HELLAND is currently an Associate Director of the Office of Science's Advanced Scientific Computing Research (ASCR) program. In addition to her associate director duties, she is leading the development of the Department's Exascale Computing Initiative to deliver a capable exascale system by 2021. She was also an executive director of the COVID-19 High-Performance Computing Consortium since its inception in March, 2020. She previously was an ASCR's facilities division director. She was also responsible for the opening ASCR's facilities to national researchers, including those in industry, through the expansion of the Department's Innovative and Novel Computational Impact on Theory and Experiment program. Prior to DOE, she developed and managed computational science educational programs at Krell Institute, Ames, IA, USA. She also spent 25 years at Ames Laboratory working closely with nuclear physicists and physical chemists to develop real-time operating systems and software tools to automate experimental data collection and analysis, and in the deployment and management of lab-wide computational resources. Helland received the B.S. degree in computer science and the M.Ed. degree in organizational learning and human resource development from Iowa State University, Ames. In recognition for her work on the Exascale Computing Initiative and with the COVID-19 HPC Consortium, she was named to the 2021 Agile 50 list of the world's 50 most influential people navigating disruption. Contact her at barbara.helland@science.doe.gov.

THUC HOANG is currently the Director of the Office of Advanced Simulation and Computing (ASC), and Institutional Research and Development Programs in the Office of Defense Programs, within the DOE National Nuclear Security Administration (NNSA), Washington, DC, USA. The ASC

program develops and deploys high-performance simulation capabilities and computational resources to support the NNSA annual stockpile assessment and certification process, and other nuclear security missions. She manages ASC's research, development, acquisition and operation of HPC systems, in addition to the NNSA Exascale Computing Initiative and future computing technology portfolio. She was on proposal review panels and advisory committees for the NSF, Department of Defense, and DOE Office of Science, as well as for some other international HPC programs. Hoang received the B.S. degree from Virginia Tech, Blacksburg, VA, USA, and the M.S. degree from Johns Hopkins University, Baltimore, MD, USA, both in electrical engineering. Contact her at thuc.hoang@nnsa.doe.gov.

MANISH PARASHAR is the Director of the Scientific Computing and Imaging (SCI) Institute, the Chair in Computational Science and Engineering, and a Professor with the School of Computing, University of Utah, Salt Lake City, UT, USA. He is currently on an IPA appointment at the National Science Foundation where he is serving as the Office Director of the NSF Office of Advanced Cyberinfrastructure. He is the Founding Chair of the IEEE Technical Consortium on High Performance Computing, the Editor-in-Chief of *IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS*, and serves on the editorial boards and organizing committees of several journals and international conferences and workshops. He is a Fellow of AAAS, ACM, and IEEE. For more information, please visit <http://manishparashar.org>. Contact him at mparasha@nsf.gov.

MICHAEL ROSENFELD is currently a Vice President of strategic partnerships with the IBM Research Division, Yorktown Heights, NY, USA. Previously, he was a vice president of Data Centric Solutions, Indianapolis, IN, USA. His research interests include the development and operation of new collaborations, such as the COVID-19 HPC Consortium and the Hartree National Centre for Digital Innovation as well as future computing architectures and enabling accelerated discovery. Prior work in Data Centric Solutions included current and future system, and processor architecture and design including CORAL and exascale systems, system software, workflow performance analysis, the convergence of Big Data,

AI, analytics, modeling, and simulation, and the use of these advanced systems to solve real-world problems as part of the collaboration with the Science and Technology Facility Council's Hartree Centre in the U.K. He has held several other executive-level positions in IBM Research including Director Smarter Energy, Director of VLSI Systems, and Director of the IBM Austin Research Lab. He started his career at IBM working on electron-beam lithography modeling and proximity correction techniques. Rosenfeld received the B.S. degree in physics from the University of Vermont, Burlington, VT, USA, and the M.S. and Ph.D. degrees from the University of California, Berkeley, CA, USA. Contact him at mgrosen@us.ibm.com.

JAMES SEXTON is currently an IBM Fellow with IBM T. J. Watson Research Center, New York, NY, USA. Prior to joining IBM, he held appointments as a lecturer, and then as a professor with Trinity College Dublin, Dublin, Ireland, and as a postdoctoral fellow with IBM T. J. Watson Research Center, at the Institute for Advanced Study at Princeton and at Fermi National Accelerator Laboratory. His research interests include span high-performance computing, computational science, and applied mathematics and analytics. Sexton received the Ph.D. degree in theoretical physics from Columbia University, New York, NY, USA. Contact him at sextonjc@us.ibm.com.

JOHN TOWNS is currently an Executive Associate Director for engagement with the National Center for Supercomputing Applications, and a deputy CIO for Research IT in the Office of the CIO, Champaign, IL, USA. He is also a PI and Project Director for the NSF-funded XSEDE project (the Extreme Science and Engineering Discovery Environment). He holds two appointments with the University of Illinois at Urbana-Champaign, Champaign. He provides leadership and direction in the development, deployment, and operation of advanced computing resources and services in support of a broad range of research activities. In addition, he is the founding chair of the Steering Committee of PEARC (Practice and Experience in Advanced Research Computing, New York, NY, USA). Towns received the B.S. degree from the University of Missouri-Rolla, Rolla, MO, USA, and the M.S. degree and astronomy from the University of Illinois, Ames, IL, both in physics. Contact him at jtowns@ncsa.illinois.edu.