The COVID-19 High Performance Computing Consortium

James Brase, Nancy Campbell, Barbara Helland, Thuc Hoang, Manish Parashar, Michael Rosenfield, and John Towns

March, 2022

Introduction

In March of 2020, recognizing the potential of High Performance Computing (HPC) to accelerate progress in the fight to stop COVID-19, the HPC community asked two simple questions: How can we leverage the nation's and the world's most powerful HPC capabilities and resources to accelerate understanding and the pace of scientific discovery? And, how can we provide those resources to help COVID-19 researchers worldwide to advance their critical efforts?
The answer to both questions: the COVID-19 HPC Consortium. Within one week, the Office of Science and Technology Policy (OSTP), the US Department of Energy (DOE), the National Science Foundation (NSF), and IBM created a unique public-private partnership between government, industry, and academic leaders to provide, at no cost to researchers in the fight against COVID-19, a single point of access to advanced HPC and Cloud computing systems and data resources along with critical associated technical expertise and support. This paper is the Consortium's story—how the Consortium was created, who were the founding members, what it provides, how it works, and what it has accomplished. Finally, we will reflect on the lessons learned from the creation and operation of the Consortium for over eighteen months since late March 2020 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.

The Creation of the Consortium

As the pandemic began to significantly accelerate in the United States, on March 11 and 12 of 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 and knowing its capabilities to help solve hard problems. Dario Gil, Senior Vice President and Director of IBM Research had the vision of organizing the HPC community to leverage the substantial computing capabilities and resources to accelerate progress and understanding in the fight against COVID-19 and connecting COVID-19 researchers with those 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. For example,

- NSF issued a Dear Colleague Letter on March 05, 2020 (https://nsf.gov/pubs/2020/nsf20055/nsf20055.jsp) making its advanced cyberinfrastructure including NSF-funded HPC resources such as the Frontera supercomputer – the fastest supercomputer deployed on an academic campus – available to advance COVID-19 research.

- Rensselaer Polytechnic Institute provided access, for COVID-19 research, to their Artificial Intelligence Multiprocessing Optimized System, or AiMOS, supercomputer. (https://cci.rpi.edu/aimos)
- Microsoft was standing up several efforts in parallel at the time: AI for Health, Studies in Pandemic Preparedness, and a grant to the University of Washington School of Medicine.
- Across Google, several efforts were underway, including offering $20M in cloud credits through the Harvard Global Health Institute, standing up the COVID-19 public dataset program, collaborating with Harvard Medical School for large-scale drug screening, and developing new COVID-19 forecasting models.
- Amazon had launched The AWS Diagnostic Development Initiative (DDI)
- DOE was in the process of establishing the National Virtual Biotechnology Laboratory (https://science.osti.gov/nvbl) to provide access to DOE user facilities, including its HPCs, to address key challenges in responding to the COVID-19 threat. DOE's HPC facilities were engaged early, particularly focusing on structural biology studies characterizing the virus drug screening, therapeutic antibody design, and epidemiology.
- NASA computing centers were doing COVID-19 environmental impact studies (reduced air pollution due to reduced airline flights) as part of their Earth Science efforts

However, it was clear very early on that a broader and more coordinated effort was needed to leverage existing efforts and relations to create a unique HPC collaboration.

Over the weekend of March 14-15, 2020, Michael Rosenfield, Vice President of Strategic Partnerships at IBM Research, shared IBM's vision with collaborators in the DOE National Labs and discussed DOE's plan to help in the fight against COVID-19. At the time, the DOE operated the two most powerful, and four of the top ten most powerful, supercomputers in the world (see https://top500.org/lists/top500/2019/11/ **)**. Given that the DOE Labs have a long history of applying HPC to important scientific problems, representatives from Oak Ridge, Lawrence Livermore, and Argonne National Labs, were very supportive of the idea of creating what would become the Consortium and helped further develop the concept and vision. Discussions were also held with other early supporters from the Massachusetts Institute of Technology (MIT) and Rensselaer Polytechnic Institute (RPI) who also expressed great interest and provided guidance on the overall concept.

By March 16, 2020, the team had developed the vision for the Consortium: to 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. Given the interest, the need, and the urgency, on March 17, 2020, Gil from IBM approached Michael Kratsios, then the Chief Technology Officer of the United States in the White House Office of Science and Technology Policy (OSTP). Kratsios recognized the significance of the proposed public-private partnership and tasked Jake Taylor, Assistant Director for Quantum Information Science at OSTP, to facilitate the creation of the Consortium along with Rosenfield from IBM. In addition, Kratsios shared the Consortium's vision with Paul Dabbar, then

Undersecretary of Science for DOE. Dabbar tasked DOE Office of Science and Barb Helland, Associate Director, Advanced Scientific Computing Research and Thuc Hoang, Director of the Office of Advanced Simulation and Computing in DOE's National Nuclear Security Administration (NNSA) to drive the creation of the Consortium within DOE. Taylor contacted Manish Parashar, Office Director of NSF's Office of Advanced Cyberinfrastructure, who recommended leveraging NSF's eXtreme Science and Engineering Discovery Environment (XSEDE) project **(**https://doi.org/10.1109/MCSE.2014.80**)** and its XSEDE Resource Allocations System (XRAS) that handles nearly 2,000 allocation requests annually **(**https://doi.org/10.1145/3219104.3229238**)** 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!

From the beginning, it was recognized that communication and expedient creation of a community around the Consortium would be key. Work began at IBM on the Consortium website, https://covid19-hpc-consortium.org the following day. On March 19, 2020, Taylor and Parashar contacted John Towns, Executive Director, Science & Technology, National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign, and Principal Investigator of the XSEDE project, to arrange for support from XSEDE and use of the XRAS platform, and to invite his participation in what would eventually be the Consortium Executive Committee as it was being formed.

During the rest of the week, the team shared their vision with other federal agencies, universities and HPC cloud providers. By the end of the week, Rosenfield at IBM; Taylor at OSTP; Helland at DOE; Hoang at DOE/NNSA; Parashar at NSF; Towns at Illinois/NCSA, Jim Brase (who was also involved in the initial discussions earlier that week), Deputy Associate Director, Computing, Lawrence Livermore National Laboratory; Jim Sexton, Director of Data Centric Systems at IBM Research; Ajay Royyuru, Vice President of Healthcare and Life Sciences at IBM Research formed the initial Consortium Executive Committee and worked on laying the groundwork for the operations of the Consortium. In addition, NASA, MIT, RPI, Amazon Web Services, Hewlett Packard Enterprise (HPE), Google, and Microsoft joined the Consortium as founding members along with OSTP, DOE, NSF, and IBM. Remarkably, the Consortium was formed within a week 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.

By Sunday, March 22nd, Towns and the XSEDE team instantiated a complete proposal submission and review process, supported by the XRAS service **(**https://doi.org/10.1145/3219104.3229238**),** which was ready to accept proposal submissions. Towns and Brase, with help from Sexton and Royyuru, set up the XSEDE website which would form the foundation of the Consortium. In addition, to facilitate the proposal process the team developed guidance to proposal submitters that detailed the expectations of the proposers and of the proposals they submitted. This guidance delineated the review criteria, the format of the proposals and provided information of the various resources to which they could request access. This content was posted to a companion website customized for the Consortium and hosted by XSEDE (https://www.xsede.org/covid19-hpc-consortium), providing direct access to the XRAS submission system.

According to the guidance on the website, the proposals were to be evaluated on the following criteria:

- potential benefits for COVID-19 response
- feasibility of the technical approach
- need for high-performance computing
- high-performance computing knowledge and experience of the proposing team
- estimated computing resource requirements

This and other guidance have been updated as appropriate on the XSEDE-hosted web page.

The main Consortium website design was finalized and seamlessly linked to the XSEDE website and the XRAS service. Luckily, the Consortium was ready because OSTP announced that the President would introduce the concept of the Consortium at a news conference on March 22nd. Numerous news articles came out after the announcement that evening (see https://covid19-hpc-consortium.org/news). The Consortium became a reality when the website, went live the next day on the 23rd of March, followed by additional press releases and news articles (https://covid19-hpc-consortium.org/news**).** The researchers were ready – the first proposal was submitted on March 24th and the first project was started on March 26th, demonstrating our ability to connect researchers with resources in a matter of days – an exceptionally short time for such processes typically. Subsequently, fifty proposals were submitted by April 15th and 100 by May 9th.

Consortium Members and Capabilities

The Consortium initially provided access to over 300 petaflops of supercomputing capacity provided by the initial 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 (Figure 1) from the U.S. and around the world (the complete list can be found at https://covid19-hpc-consortium.org/) representing access 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 (NCI) Australia and Pawsey Supercomputing Centre (https://covid19-hpc-consortium.org/collaborations). The Consortium also added nine affiliates (also listed and described at https://covid19-hpc-consortium.org and https://www.xsede.org/covid19-hpc-consortium) who provided expertise and supporting services to enable researchers to start up quickly and run more efficiently.

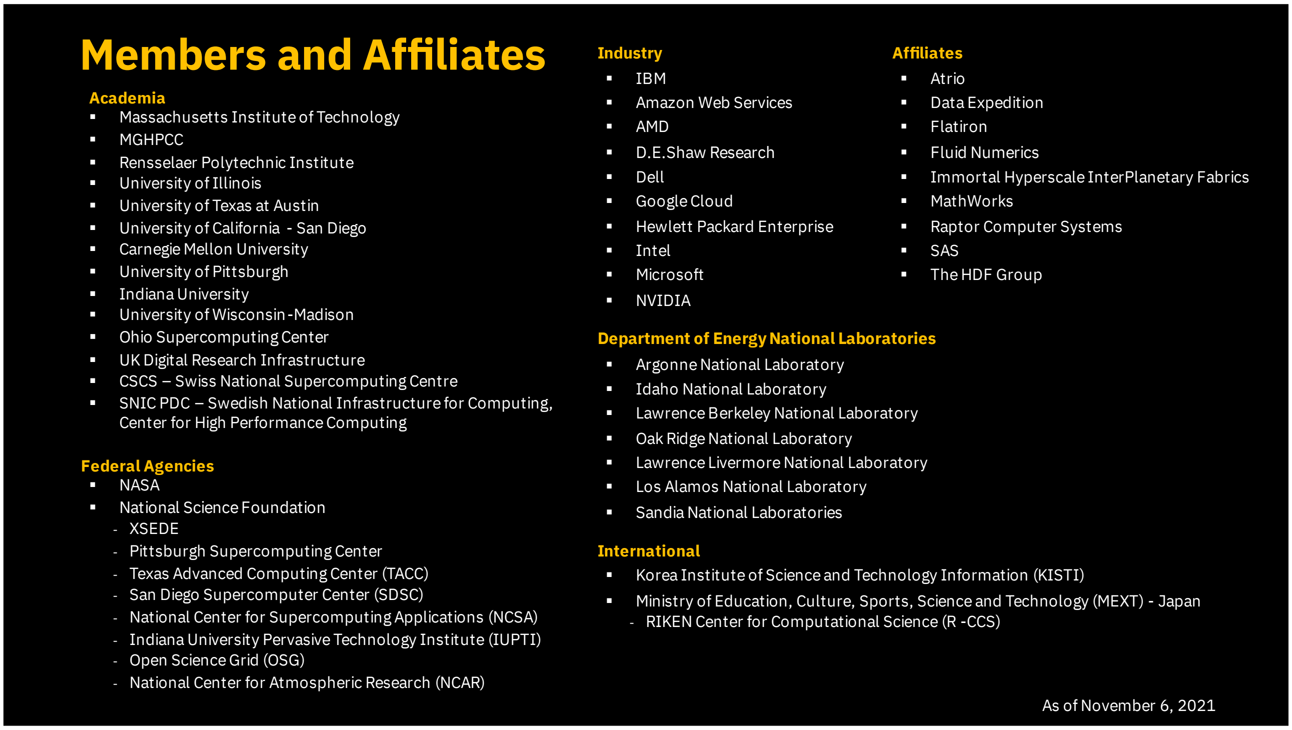

Figure 1: Consortium Members and Affiliates as of November 6, 2021.

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

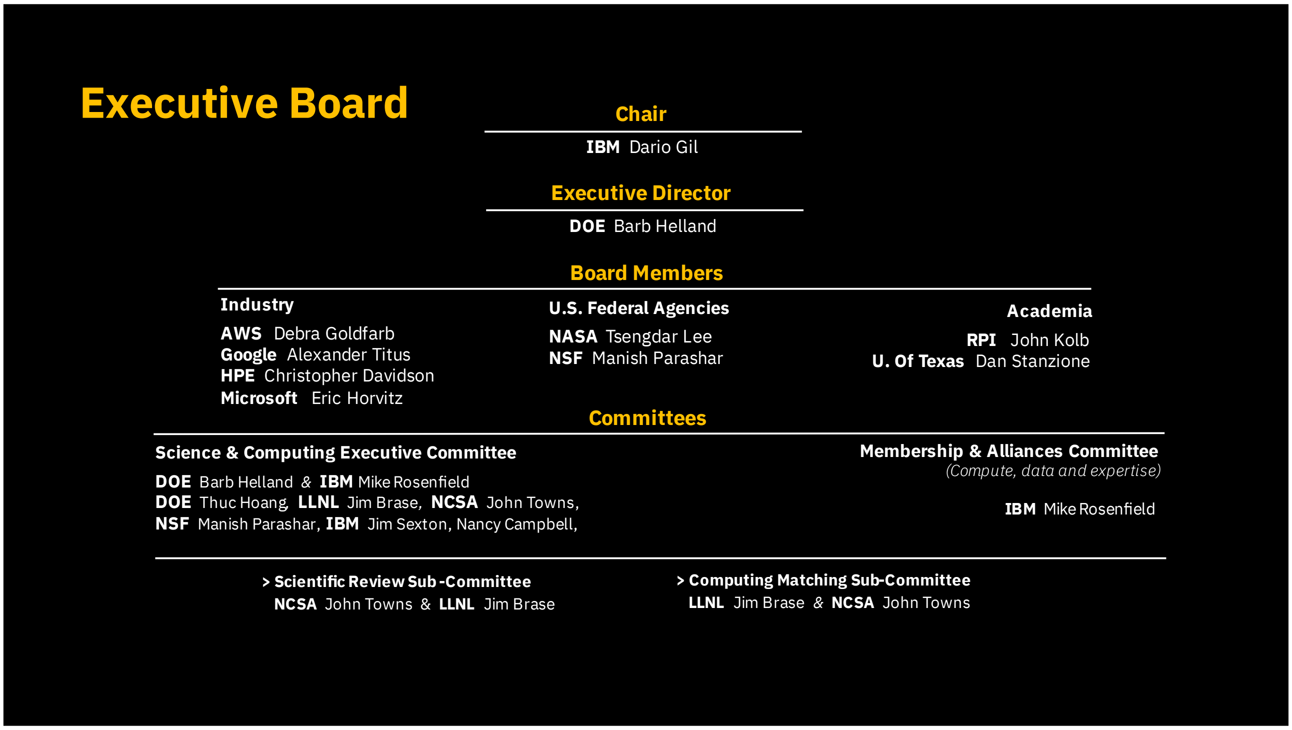

Figure 2. The Consortium Organizational Structure as of November 6, 2021

The Scientific Review Sub-Committee and the Computing Matching Sub-Committee, which report to the Executive Committee, are critical to 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 (see acknowledgements at https://covid19-hpc-consortium.org/who-we-are)—reviews proposals for merit based on the review criteria and guidance provided to proposers noted above, and recommends appropriate proposals to the Computing Matching Sub-Committee. Researchers can request specific supercomputers or cloud providers on which to run their applications and analyses (for example, if they already have experience with a particular system or cloud provider, they could specifically request that platform). The Computing Matching Sub-Committee team, comprised of representatives of Consortium members providing resources, matches the computer needs from recommended proposals with either the requested site or other appropriate resources. Once matched, the researcher needs to go through the standard on-boarding/approval process at the host site to gain access to the system—working with the matching sub-committee member(s) affiliated with the system(s) to which they have been provided access. Initially, we expected that the on-boarding/approval process would be time consuming (since this was the only time where actual agreements had to be signed), but those executing the on-boarding process with the various member compute providers worked diligently to prioritize these requests and thus this is typically only taking a day or two. As a result, once approved, projects are up and running very rapidly.

The Membership Committee, which also reports to the Executive Committee, reviews requests for organizations and individuals to become members or affiliates. Requests are sent to OSTP for vetting. Requestors are asked to describe what they intended to offer at no cost, to provide a web link to what was being offered, and to describe the relevant subject matter experts who would facilitate the researchers' use of the offered resource. The Executive Committee then makes recommendations to the Executive Board to approve or reject the requests. Once approved, the Executive Committee works with the Member or Affiliate member to finalize the description of the offering and capabilities subsequently posted on the Consortium and XSEDE websites

Project Highlights

The goal of the Consortium is to provide state-of-the-art high-performance computing 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 as shown in Figure 3:

BASIC SCIENCE		THERAPEUTIC DEVELOPMENT		PATIENTS	
Total	**37**	**Total**	**54**	**Total**	**26**
Viral structure and function	17	Target discovery	6	Patient trajectory and outcomes	6
Viral-human interaction	11	Small molecule design	24	Medical technologies	1
Viral evolution	4	Antibody, vaccine, protein design	13	Supply chain and resource allocation	2
Environmental effects	3	Drug repurposing	10	Epidemiology	8
Science tools	2	Development technologies	1	Environmental effects	4
				Social interaction analytics	3
				Detection and diagnostics	2

Figure 3. Distribution of Consortium projects across the taxonomy of pandemic-related technical topics as of 11/19/21.

The basic science and therapeutic development communities have significant capability and experience in using HPC and responded rapidly to utilize Consortium resources. The patient care and analysis community clearly had less intersection with HPC early in the Consortium operations, but this area has grown as more groups have developed computing-intensive approaches, large data sets, and relevant projects.

Consortium projects have produced a broad range of scientific advances. The projects have collectively produced a growing number of publications, datasets, and other products (over 70 of these by the end of calendar year 2021) including the covers of two journals: The Journal of Chemical Information and Modeling (http://doi.org/10.1021/acs.jcim.0c00929) and Physical Chemistry Chemical Physics (http://doi.org/10.1039/d0cp03145c).

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 .

A few notable highlights include the following:

- Harel Weinstein, Weill Cornell Medical College. Provider: RPI. Simulations of molecular mechanisms of SARS-CoV-2 interactions with membranes are opening new targets for development of therapeutics.
- Gouwei Wei, Michigan State. Provider: IBM. Analysis and predictive modeling of potential SARS-CoV-2 mutations and their impact on diagnostic testing, vaccines, and therapeutics. 10 accepted publications, 2 in review. Prediction of new variants and their properties is emerging as an important priority.
- Giulia Palmero, UC Riverside. Provider: Microsoft. New versions of the CRISPR gene-editing system are being harnessed as a fast, yet reliable diagnostic tool against SARS-CoV-2 infections. Improved diagnostics continues to be a critical concern. This research aims at characterizing how the CRISPR-associated proteins recognize viral genetic material. Selected as an American Chemical Society "Editor's Choice" paper.
- Aaron Morris, PostEra. Running on AWS. Postera and the COVID-19 Moonshot have provided free chemical synthesis to projects all over the world. Over 300 assay results for crowdsourced have been completed with 20 showing potency against the SARS-CoV-2 main protease.
- Jeremy Smith, University of Tennessee, Knoxville. Running on Summit and Google. Experimental testing of predicted compounds has so far led to the discovery of several new inhibitors of viral proteins and the identification of three already approved drug compounds that inhibit the infectivity of live coronavirus. Two of the identified compounds are currently in clinical trials.
- Jennifer Diaz, Mt. Sinai. Running on Summit. Combination therapy design is an important approach to rapid response but data and computationally challenging. Using a new transcription-based drug-pair synergy approach, the project team completed 35B predictions to identify ~10 drug pairs predicted to target the COVID-19 protein interactions. In-vitro validation underway.
- John Davis, UCSD. Running at SDSC. Published a multi-county COVID-19 transmission model on California Department of Public Health's CalCAT website (https://calcat.covid19.ca.gov/cacovidmodels/). The model forecasts hospital and ICU beds as well as mortality stratified by age group for each county.
- Amanda Randles, Duke. Provider: Microsoft. Created a parallel computing model that simulates air flow through a ventilator to enable splitting the ventilator across patients whose lungs have different levels of resistance. The HPC Consortium compute allocation on Azure enabled the Duke team to submit for FDA authorization for emergency use.
- Eve Wurtele, Iowa State. Provider: PSC. Completed analysis of differential gene expression among COVID-19 patient populations. Population dependence has important implications for treatment planning.

- Som Dutta, Utah State. Provider: TACC. New high-fidelity turbulence-resolved simulations of aerosol transport in indoor environment are being used to study residence time and deposition pattern of virus-laden aerosols in classrooms, an important application for planning re-opening activities.

While Consortium projects have contributed significantly to scientific understanding of the virus and potential therapeutics, direct and near-term impact on the course of the pandemic has been mixed. There are cases of significant impact (as noted in the above examples) 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 high-performance computing 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 pandemic or other crisis response programs.

Operational Results and 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. As described, the Consortium has been successful in bringing together the U.S. Federal government, industry, and academic leaders to provide access to the world's most powerful high-performance computing resources in support of COVID-19 research. Within its first week of existence the Consortium instantiated an operational framework for providing advanced computing and data resources and services for rapid crisis response. The consortium effectively:

- Worked together to overcome a myriad of institutional and organizational boundaries within government, industry, and academia to create a common portal to access advanced computing resources and helped coalesce *ad hoc* and defocused efforts in smaller "consortia" around the country and the world;
- Ramped up quickly, and in an agile way, to meet urgent and presumably short-term needs that could not be easily addressed within ongoing industry, agency, and academic programmatic opportunities, without formal partnership agreements;
- Developed review, matching and on-boarding processes that have reviewed over 200 scientific proposals with more than 100 projects running or completed with results beginning to be publicized;
- Set up the framework for a worldwide community via the Consortium website, webinars, international alliances, blogs, and press releases;
- Established new groups of computing users able to utilize a variety of advanced computing and data resources effectively for crisis response; and
- Accelerated programs in both basic understanding of the virus and its host interactions and early-stage drug development, mainly within university groups and, to a more limited extent, in small companies.

The Consortium was able to take advantage of the processes and tools (XSEDE) that were in place at the NSF for managing and reviewing proposals. Key lessons learned:

- The ability to leverage existing processes 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 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 not 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.

The Next Step: The National Strategic Computing Reserve

Computing has become a central element of many of our most important capabilities that are essential to properly respond to national crises, ensuring public health and safety, and protecting critical resources and infrastructure. Advanced computing resources and services are strategic national assets and can serve as important tools for crisis response. To effectively mobilize these assets in a crisis, an organized program of planning and preparation is urgently needed.

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. 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 a National Strategic Computing Reserve (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, technical expertise) spanning government, academia, non-profits/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.

The diagram below, Figure 4, shows a transition from a pre-COVID ad-hoc response to crises to the Consortium and then to a National Strategic Computing Reserve.

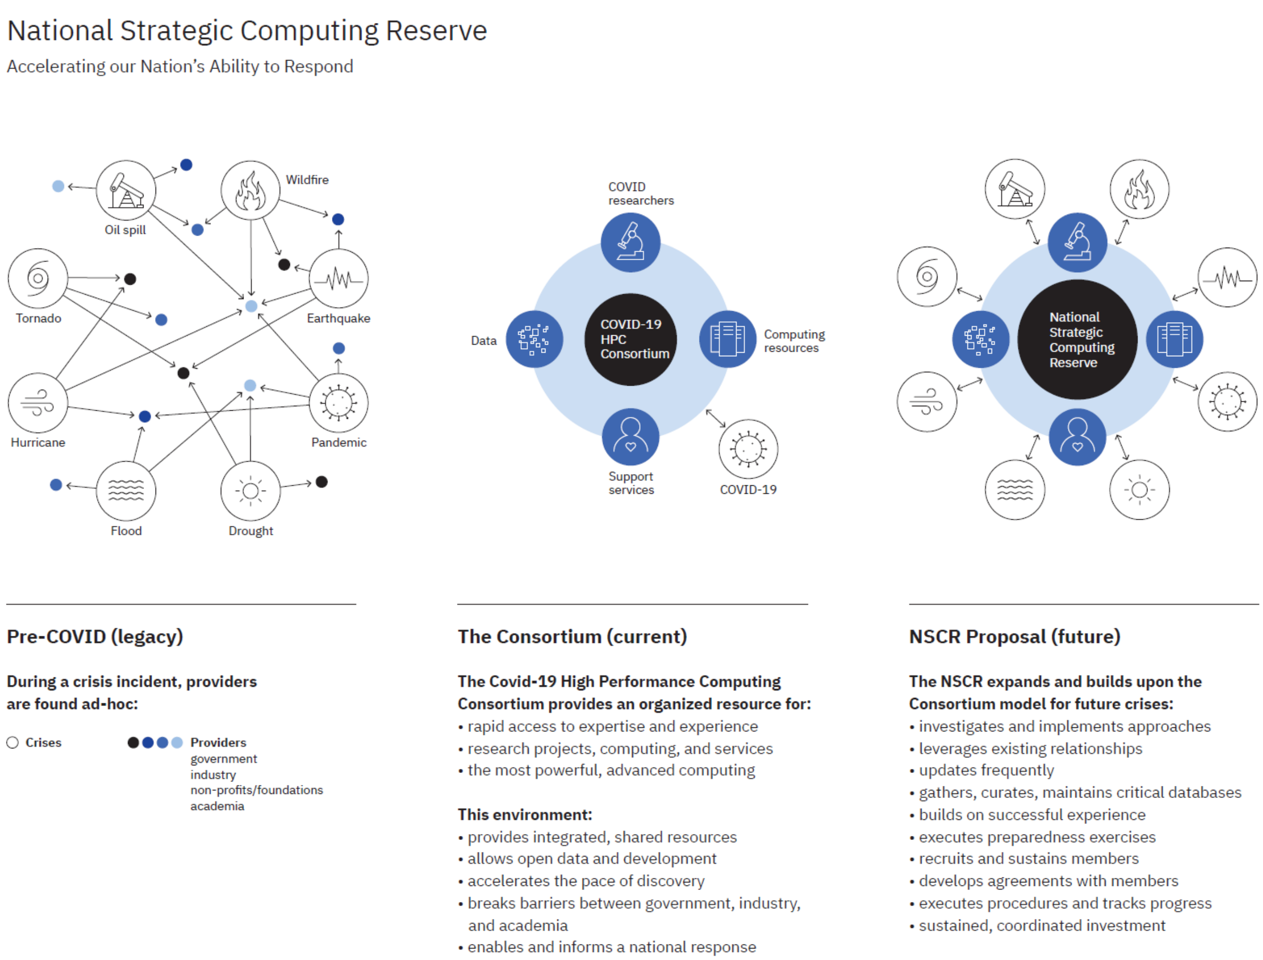

Figure 4: A potential path from Pre-Covid to the NSCR

A National Strategic Computing Reserve should be analogous to the Civil Reserve Air Fleet[1] and the United States Merchant Marine[2] (among others) that can be called upon in times of crisis and should aim to make advanced computing and data resources and services available to address urgent needs. The NSCR would be composed of a set of domain experts working with computing resource providers who are committed to make resources available to respond to crises.

Please also see: https://covid19-hpc-consortium.org/blog/national-strategic-computing-reserve

Principal Functions of the NSCR

In much the same way as the Merchant Marine maintains a set of "ready reserve" resources that can be put to use in wartime, the NSCR would maintain reserve computing capabilities that could be made available in times of urgent national need. 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:

- 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 non-emergency participation;
- Development of a set of agreements to enable the Reserve to collaborate with domain agencies and industries in preparation for and execution of Reserve deployments;
- Execution of 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; and
  - Participate in the broader national response as an active partner.

Conclusion

The COVID-19 High Performance Computing Consortium has been in operation for over a year (https://covid19-hpc-consortium.org/blog/a-year-on-hpc-consortium-national-strategic-computing-reserve) and has enabled over 100 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 (1) proposals in specific targeted areas (2) gathering and socializing

[1] Civil Reserve Air Fleet – 10 U.S.C §§ 9511-9515 Civil Reserve Air Fleet
[2] United States Merchant Marine – 46 U.S.C. §§ 861-889 Merchant Marine Act

results from current projects, and (3) driving the establishment of a National Strategic Computing Reserve. New project focus areas target having an impact in a 6-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 multi-modal datasets.

In March of 2020, we asked ourselves two very simple questions: How can we leverage the nation's and the world's most powerful HPC capabilities and resources to accelerate understanding and the pace of scientific discovery? And, how can we provide those resources to help COVID-19 researchers worldwide to advance their critical efforts? Our vision was to provide COVID-19 researchers worldwide with access to the world's most powerful high performance computing resources to significantly advance the pace of scientific discovery in the fight to stop the virus. In the time span of approximately one week, we created a unique public-private partnership between government, industry, and academic leaders to aggregate free compute time and resources on their HPC systems.

As demonstrated by the response to COVID-19, advanced computing infrastructure is a strategic national asset and can serve as an important tool for crisis response. To effectively mobilize this asset in a crisis, a well-planned and organized program is needed.

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 National Strategic Computing Reserve, 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.

Acknowledgements

The authors would like to thank the past and present members of the Consortium Executive Board for their guidance and leadership. In addition, we 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. We also would like to thank Corey Stambaugh from OSTP for his leadership role on the Consortium membership committee. Further, we would 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 is critical is selecting appropriate projects and their time and efforts are greatly appreciated. And, of course, we want to thank the many researchers who did such outstanding work, leveraging the Consortium, in the fight against COVID-19.

The authors would also like to thank the IEEE Computing in Science and Engineering (https://www.computer.org/csdl/magazine/cs) journal for publishing a version of this paper in their Jan-Feb 2022, Volume 24, Issue 1 edition:
https://www.computer.org/csdl/magazine/cs/2022/01/09734778/1BLn0zdzhIY

Author Biographies

Jim Brase is the Deputy Associate Director for Computing at Lawrence Livermore National Laboratory (LLNL). 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 co-lead of the ATOM Consortium for computational acceleration of drug discovery, and on the leadership team of the COVID-19 HPC Consortium. Jim's research interests focus on the intersection of machine learning, simulation, and high-performance computing. He is currently leading efforts on large-scale computing for life science, biosecurity, and nuclear security applications. In his previous position as LLNL's Deputy Program Director for Intelligence, Jim led efforts in intelligence and cybersecurity R&D.

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 High Performance Computing Consortium and International Science Reserve. Prior to this role, Nancy served as 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, Nancy 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, Nancy served as CEO for one of Selby Venture Partners portfolio companies and facilitated the successful sale of that company to its largest channel partner. Nancy attended the University of Southern California and serves as IBM's Executive Sponsor for the USC Master of Business for Veterans program.

Barbara Helland is currently the 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 has also served as the Executive Director of the COVID-19 High Performance Computing Consortium since its inception in March, 2020. Ms. Helland previously served as 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 (INCITE) program.   Prior to DOE, Ms. Helland developed and managed computational science educational programs at Krell

Institute. 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. She received B.S. in Computer Science and a M. Ed in Organizational Learning and Human Resource Development from Iowa State University. In recognition for her work on the Exascale Computing Initiative and with the COVID-19 High Performance Computing Consortium, Ms. Helland was named to the 2021 Agile 50 list of the world's 50 most influential people navigating disruption.

Thuc Hoang is 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 Department of Energy National Nuclear Security Administration (NNSA). 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. Ms. Hoang manages ASC's research, development, acquisition and operation of high-performance computing (HPC) systems, in addition to the NNSA Exascale Computing Initiative and future computing technology portfolio. She has served on proposal review panels and advisory committees for the National Science Foundation, Department of Defense and DOE Office of Science, as well as for some other international HPC programs. Ms. Hoang holds a Bachelor of Science in Electrical Engineering from Virginia Tech and a Master of Science in Electrical Engineering from Johns Hopkins University.

Michael Rosenfield is Vice President of Strategic Partnerships at the IBM Research Division in Yorktown Heights, NY. Previously, he was Vice President of Data Centric Solutions. His current focus areas include the development and operation of new collaborations such as the COVID-19 High Performance Computing 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 UK. Mike 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. He has a BS in Physics from the University of Vermont and a MS and Ph.D from the University of California, Berkeley.

James Sexton is an IBM Fellow at IBM T. J. Watson Research Center in New York. Dr. Sexton received his Ph.D. in Theoretical Physics from Columbia University, NY. His areas of research span High Performance Computing, Computational Science, Applied Mathematics and Analytics. Prior to joining IBM, Dr. Sexton held appointments as Lecturer then Professor at Trinity

College Dublin, and as postdoctoral fellow at IBM T. J. Watson Research Center, at the Institute for Advanced Study at Princeton and at Fermi National Accelerator Laboratory.

John Towns holds two appointments at the University of Illinois at Urbana-Champaign. He is the Executive Associate Director for Engagement at NCSA (the National Center for Supercomputing Applications, http://www.ncsa.illinois.edu), and Deputy CIO for Research IT (http://researchit.illinois.edu) in the Office of the CIO at Illinois. He is also PI and Project Director for the NSF-funded XSEDE project (the Extreme Science and Engineering Discovery Environment, http://XSEDE.org). Towns 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, http://www.pearc.org). He earned M.S. degrees in Physics and Astronomy from the University of Illinois and a B.S. in Physics from the University of Missouri-Rolla.