


Viewpoint

Transforming Science through Cyberinfrastructure

NSF's vision for the U.S. cyberinfrastructure ecosystem for science and engineering in the 21st century.

ADVANCED CYBERINFRASTRUCTURE (CI) is critical to science and engineering (S&E) research. For example, over the past two years, CI resources (including those provided by the COVID-19 HPC Consortium^a) enabled research that dramatically accelerated efforts to understand, respond to, and mitigate near- and longer-term impacts of the novel coronavirus disease 2019 (COVID-19) pandemic.^b Computer-based epidemiology models informed public policy in the U.S., and in countries throughout the world, and newly studied transmission models for the virus have been used to forecast resource availability and mortality stratified by age group at the county level.^c Artificial intelligence and machine learning approaches accelerated drug screening to find candidate medicines from trillions of possible chemical compounds,^d and differential gene expressions among COVID-19 patient populations have been analyzed with important implications for treatment planning.^e Structural modeling of the virus has led to new insights, speeding the development of



vaccines and antigens. One such effort earned the ACM's Gordon Bell Prize Recognizing Outstanding Achievement in High-Performance Computing.^f

CI encompasses more than the computing resources themselves. Rather—and as the response to the pandemic illustrates—CI constitutes an expansive ecosystem, comprising these resources as well as data, software, networking and security, coordination and user support, and connections to instrumentation and large-scale infrastructure. Realizing such a CI ecosystem requires blending fundamental and translational research in computer and computational science, research infrastructure,

and private-sector innovations to ensure continuous refresh of the ecosystem to align with evolving use cases and needs.

Within the U.S., the conceptualization, design, and implementation of such an advanced CI ecosystem for S&E research and education is led by the National Science Foundation's (NSF) Office of Advanced Cyberinfrastructure (OAC). Over the past two decades, OAC (and its predecessors) developed a balanced portfolio of complementary CI investments,^g and funded and

^a See The COVID-19 High Performance Computing Consortium; <https://bit.ly/39AL8WO>

^b See Why are supercomputers so important for COVID-19 research?; <https://bit.ly/3b94R0e>; and Harnessing Computing Power to Fight COVID-19; <https://bit.ly/3y0Yjtl>

^c See <https://bit.ly/3mUoknY>

^d See <https://bit.ly/3O54dPF>

^e See <https://bit.ly/39zCWGd>

^f See <https://bit.ly/3b9BnPL>

^g OAC's portfolio is balanced across the different elements of the CI ecosystem as well as the S&E communities served. See 2018 Office of Advanced Cyberinfrastructure (OAC) Committee of Visitors (COV) Report; <https://bit.ly/3QrRTL3>

coordinated exploration, development and provisioning of advanced CI resources, facilities, and services (see the figure here). Collectively, these investments have laid the groundwork for today's S&E advances.

In this Viewpoint, we look at the current and emerging landscape and provide a vision for an integrative, holistic advanced CI ecosystem, together with a balance of foundational and translational research and innovation, to drive the nation's S&E enterprise.

A Landscape of Disruptive Changes

Dramatic changes in the availability of data and computation, in the nature, scale, and urgency of applications, and in technology landscapes have profound implications for strategic priorities and investments for CI.

Disruptive Application Pulls. A series of disruptive applications are prompting the need for innovations throughout the CI ecosystem:

- *High-fidelity modeling* of phenomena across multiple scales and physics is resulting in high-resolution, dynamic, coupled simulation workflows capable of running at extreme scales.

- *Increasing availability and scales of experimental and observational data* are creating unprecedented volumes of data, expanded requirements for data management and integration, and concerns and trade-offs regarding data sharing and data privacy.

- *Online (urgent) data processing* requirements are resulting from the proliferation of data sources, new opportunities for online and near-real-time monitoring, data processing (including at the edge), and actuation, and new classes of data-driven applications.

- *A growing "long tail" of applications* (that is, applications requiring relatively moderate to small scale national resources) is increasing in complexity and scale and are quickly dominating overall computational workloads.

- *Novel data-centric applications* using deep and machine learning are complementing more traditional workloads. Increasing use of novel software and tools not part of the traditional software stack, such as digital notebooks, machine-learning libraries and containers, are becoming more common.

- *The central role of CI* as a key enabler of robust S&E research is positioning

transparency, traceability, reproducibility, and security as important CI concerns with closely-related policy concerns around open science, which seek to expand access to research results and the supporting data and software.^h

Disruptive Technology Pushes.

Along with these application "pulls" are a set of emerging and disruptive technologies pushing innovations in CI:

- *Unprecedented technological advances* coupled with the exploration of technologies and paradigms beyond Moore's Law are resulting in increasing processing speeds, new classes of processors with increased parallelism and purposeful hardware accelerators, novel storage technologies and deeper storage hierarchies, faster communication fabrics, software-defined (programmable) architectures and systems, and extreme-scale systems.

- *High-bandwidth/low-latency networks* deployed at campus, regional, national and international scales, together with the growing availability of next-generation wireless networks and systems, are providing global connectivity and access.

- *Distributed and federated institutional CI* available across campuses and through large projects and facilities has grown in capacity and capability to be-

come a significant provider of resources.

- *Increasing edge and in-network capabilities*, for example, those provided by the growing spectrum of "smart" edge devices as well as advanced network functions are enabling sophisticated, real-time data processing closer to the data sources and/or users.

- *Growing availability of commercial cloud services and hybrid models* are playing an important and complementary role, lowering barriers to access, supporting new access modes (for example, on-demand, elastic), and providing unique capabilities (for example, purposeful accelerators).

- *Growing concerns about energy consumption and failures* are making energy efficiency and fault tolerance first-class CI design concerns.

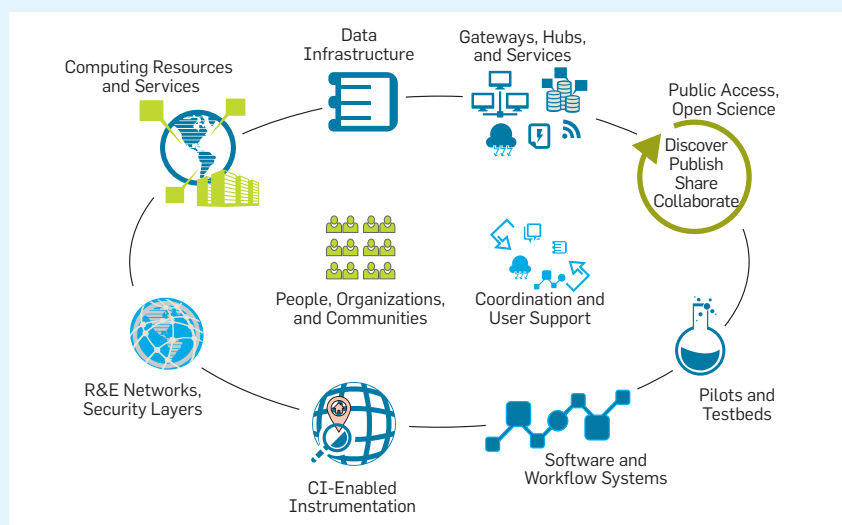
A Vision for the NSF-Funded CI Ecosystem

In this landscape of disruptive changes, NSF envisions an agile, integrated, robust, trustworthy and sustainable CI ecosystem that drives new thinking and transformative discoveries in all areas of S&E research and education.ⁱ This CI ecosystem builds on fundamental research advances in academia and industry, responds to evolving S&E

^h See Public Access to Results of NSF-funded Research: <https://bit.ly/3zLrZw6>; and NSF 18-060, Dear Colleague Letter: Advancing Long-term Reuse of Scientific Data; <https://bit.ly/3QnMZ1L>

ⁱ See Transforming Science Through Cyberinfrastructure: NSF's Blueprint for a National Cyberinfrastructure Ecosystem for Science and Engineering in the 21st Century; <https://bit.ly/39Gjrf0>

The current NSF-funded CI ecosystem. NSF embraces an expansive view of CI motivated by research priorities and scientific process.



needs, integrates different CI dimensions, and engages partnerships with other U.S. agencies, industry, and international funders.

Key Principles of the Vision

Our strategy going forward centers on six interrelated principles:

► **View CI holistically:** Computational and data-enabled S&E research requires multiple elements of CI, including compute, data, software, networking, and security resources, tools, services, and expertise, to come together to enable advanced application workflows.

► **Support translational research:** Translational research, in the context of evolving application requirements and technology landscapes, can advance research results to practice along a continuum spanning: catalyzing essential core CI innovations; fostering the development of community tools and frameworks; and enabling the deployment of sustainable production-quality CI services.

► **Balance innovations with stability:** Availability of robust, stable and dependable CI services must be balanced against the evolution of the CI ecosystem to incorporate new requirements and ongoing innovations.

► **Co-design/coupling cycles of innovation:** S&E applications and CI are continually innovating and evolving. Consequently, application-CI co-design and a tight coupling of the cycles of discovery and innovation are essential, enabling the necessary bidirectional information/knowledge flow: CI innovations enable new scientific formulations, which in turn drive further CI innovations.

► **Focus on usability:** New levels of usability can ease the pathways for discovering, accessing, understanding and utilizing powerful CI capabilities and services, democratizing availability of and access to the CI and enhancing scientists' productivity and science impact.

► **Diverse, skilled workforce:** A diverse, skilled CI workforce, supported by investments in learning and workforce development is a critical need.

Toward an Integrated CI Ecosystem: NSF Blueprints for the Road Ahead. NSF has already begun to evolve its programs to align with this vision. Programs provisioning advanced computational resources balance innovation with production operation. The OAC

core research program^j focuses on translational CI research; an integrated data and software program^k supports the development and deployment of tightly coupled software and data services; the campus CI program^l encourages cloud integration; and OAC leads a foundation-wide program^m on learning and workforce development. NSF works closely across S&E communities to couple the cycles of discovery and innovation to address new challenges and opportunities, ensure usability, and enhance scientists' productivity and the overall impact on S&E outcomes.

Additionally, NSF has published a series of blueprints,ⁿ informed by the community through workshops, meetings, requests for information, and surveys. These blueprints focus on different aspects of the CI ecosystem; each one describes a specific vision for that aspect of the CI ecosystem, along with detailed plans for achieving that vision. Further, collectively, the blueprints offer a path for realizing NSF's vision of an integrated CI ecosystem. Specifically:

► The **Computational Ecosystem** blueprint outlines plans for a establishing an integrated, balanced, and scalable computational ecosystem;

► The **Coordination Services** blueprint presents a vision for a fabric of CI coordination services aimed at providing agile and scalable structures and services;

► The **Data and Software Services** blueprint describes a national data and software CI ecosystem to enable and accelerate S&E research;

► The **Learning and Workforce Development** blueprint presents a vision for fostering, broadening, and nurturing a diverse, recognized, and skilled CI workforce that can accelerate and amplify the transformative impact of CI across all S&E research and education; and

► The **International Research and Education Network Connections (IRNC)** blueprint describes strategic international network connections providing core capabilities for international collaborations, driven by current

and future S&E needs.

These blueprints are released as drafts, enabling the community to provide feedback before NSF translates them into programs and solicitations.

From Vision to Action

Research CI is critical to the Nation's and the Foundation's strategic priorities, which in turn help define the nature and structure of the CI ecosystem. As demonstrated in the past year, CI is also essential to U.S. preparedness and responsiveness to crises and to the nation's resilience to future natural disasters. In the near- and mid-term, NSF, through OAC, is focused on:

► Addressing needs holistically across multiple research disciplines, domains, and priority areas such the NSF major facilities to encourage more and more diverse researchers to use CI resources and services;

► Fostering CI innovation through investment in foundational and translational CI research and education activities;

► Achieving robustness, accessibility and responsiveness of the CI ecosystem that builds on previous and ongoing investments and maximally leverages new technologies;

► Continuing investments in education and training, developing a broad and diverse CI community, promoting coordination and exchange between the CI and research communities; and

► Facilitating dissemination of best practices for design, development, and operation of CI resources and capabilities.

Your engagement, as members of the S&E community, is critical as NSF moves toward achieving our goal of realizing an integrated CI ecosystem that transforms all S&E research and education. **C**

Manish Parashar (mparasha@nsf.gov) is Director of the Office of Advanced Cyberinfrastructure at the U.S. National Science Foundation, Alexandria, VA, USA.

Amy Friedlander (afriedlander9@gmail.com) was Deputy Director (retired) of the Office of Advanced Cyberinfrastructure at the U.S. National Science Foundation, Alexandria, VA, USA.

Erwin Gianchandani (egiancha@nsf.gov) is Assistant Director for Technology, Innovation and Partnerships, and was previously Deputy Assistant Director for Computer and Information Science and Engineering at the U.S. National Science Foundation, Alexandria, VA, USA.

Margaret Martonosi (mmartono@nsf.gov) is Assistant Director for Computer and Information Science and Engineering at the U.S. National Science Foundation, Alexandria, VA, USA.

Copyright held by authors.

^j See <https://bit.ly/3mXTURD>

^k See <https://bit.ly/3mZMzB0>

^l See <https://bit.ly/3tJBFDm>

^m See <https://bit.ly/2SKhCjC>

ⁿ See <https://bit.ly/39Gjrf0>