Pioneering the Future Advanced Computing Ecosystem: A Strategic Plan

A Report by the
SUBCOMMITTEE ON FUTURE ADVANCED COMPUTING ECOSYSTEM

COMMITTEE ON TECHNOLOGY

of the
NATIONAL SCIENCE & TECHNOLOGY COUNCIL

NOVEMBER 2020

About the National Science and Technology Council

The National Science and Technology Council (NSTC) is the principal means by which the Executive Branch coordinates science and technology policy across the diverse entities that make up the Federal research and development enterprise. A primary objective of the NSTC is to ensure science and technology policy decisions and programs are consistent with the President's stated goals. The NSTC prepares research and development strategies that are coordinated across Federal agencies aimed at accomplishing multiple national goals. The work of the NSTC is organized under committees that oversee subcommittees and working groups focused on different aspects of science and technology. More information is available at http://www.whitehouse.gov/ostp/nstc.

About the Office of Science and Technology Policy

The Office of Science and Technology Policy (OSTP) was established by the National Science and Technology Policy, Organization, and Priorities Act of 1976 to provide the President and others within the Executive Office of the President with advice on the scientific, engineering, and technological aspects of the economy, national security, homeland security, health, foreign relations, the environment, and the technological recovery and use of resources, among other topics. OSTP leads interagency science and technology policy coordination efforts, assists the Office of Management and Budget with an annual review and analysis of Federal research and development in budgets, and serves as a source of scientific and technological analysis and judgment for the President with respect to major policies, plans, and programs of the Federal Government. More information is available at http://www.whitehouse.gov/ostp.

About the Subcommittee on Future Advanced Computing Ecosystem

The NSTC Subcommittee on Future Advanced Computing Ecosystem (FACE) coordinates the Federal agencies to pioneer, sustain, and enhance the advanced computing ecosystem necessary for U.S. scientific, technological, and economic leadership. It is guided by the objectives, priorities, and recommendations outlined in the 2019 report, *National Strategic Computing Initiative: Pioneering the Future of Computing*, in alignment with the Networking and Information Technology Research and Development (NITRD) Subcommittee and other subcommittees as appropriate.

About this Document

This document outlines a Federal strategic plan for a whole-of-nation approach to pioneering the future national advanced computing ecosystem and establishes the operational and coordination structure to support the implementation of its objectives.

Copyright Information

NATIONAL SCIENCE & TECHNOLOGY COUNCIL

Chair

Kelvin K. Droegemeier, Director, OSTP

Staff

Tracie Lattimore, Executive Director

Grace Diana, Deputy Executive Director

COMMITTEE ON TECHNOLOGY

Co-Chairs

Walter Copan, Under Secretary of Commerce for Standards and Technology and Director, NIST

Paul Dabbar, Under Secretary for Science, DOE

Michael Kratsios, Deputy Assistant to the President, OSTP, and Chief Technology Officer of the United States

SUBCOMMITTEE ON FUTURE ADVANCED COMPUTING ECOSYSTEM

Co-Chairs

Manish Parashar, Assistant Director for Strategic Computing, OSTP

Steve Binkley, Principal Deputy Director, DOE/SC

Amy Friedlander, Acting Office Director, NSF/CISE

Erwin Gianchandani, Deputy Assistant Director, NSF/CISE

Zachary Goldstein, CIO and Director of HPCC, NOAA

Barbara Helland, Assistant Director, DOE/SC

Frank Indiviglio, Deputy Director of HPCC, NOAA

Margaret Martonosi, Assistant Director, NSF/CISE

Irene Parker, Assistant CIO of Satellite and Information Service, NOAA

Executive Secretary

Ji Hyun Lee, NCO/NITRD

Members

Tali Bar-Shalom, OMB

Stephen Dennis, DHS

Susan Gregurick, HHS/NIH

William Harrod, ODNI/IARPA

Peter Highnam, DoD/DARPA

Thuc Hoang, DOE/NNSA

Jalal Mapar, DHS

Jill Marlowe, NASA

Dave Mountain, DoD/NSA

Raju Namburu, DoD

Yi Pei, OMB

Kamie Roberts, NCO/NITRD

Charles Romine, DOC/NIST

Keel Ross, DOI/USGS

Tara Schwetz, HHS/NIH

Jesse Taylor, DOJ/FBI

Table of Contents

Abbreviations

AAII	American Artificial Intelligence Initiative
AI	artificial intelligence
AI/ML	artificial intelligence/machine learning
CISE	Computer and Information Science and Engineering
CMOS	complementary metal–oxide semiconductor
DARPA	Defense Advanced Research Projects Agency
DHS	Department of Homeland Security
DOC	Department of Commerce
DoD	Department of Defense
DOE	Department of Energy
DOI	Department of the Interior
DOJ	Department of Justice
EO	Executive Order
FACE	Future Advanced Computing Ecosystem (NSTC Subcommittee)
FAIR	findable, accessible, interoperable and reusable (data)
FBI	Federal Bureau of Investigation
FTAC	Fast Track Action Committee
GPU	graphics processing unit
HHS	Department of Health and Human Services
HPC	high performance computing
HPCC	High Performance Computing and Communications Program (NOAA)
IARPA	Intelligence Advanced Research Projects Activity
IotF	Industries of the Future
ML	machine learning
M&S	modeling and simulation
NASA	National Aeronautics and Space Administration
NCO	National Coordination Office
NIH	National Institutes of Health
NIST	National Institute of Standards and Technology
NITRD	Networking and Information Technology Research and Development Program
NNSA	National Nuclear Security Administration (DOE)
NOAA	National Oceanic and Atmospheric Administration
NQCO	National Quantum Coordination Office
NQI	National Quantum Initiative
NSA	National Security Agency
NSCI	National Strategic Computing Initiative
NSF	National Science Foundation
NSTC	National Science and Technology Council
ODNI	Office of the Director of National Intelligence
OMB	Office of Management and Budget
OSTP	Office of Science and Technology Policy
PB	petabytes
PPP	public–private partnerships
QIS	quantum information science
R&D	research and development
S&E	science and engineering
SC	Office of Science (DOE)
STEM	science, technology, engineering, and mathematics
USGS	United States Geological Survey

Executive Summary

This strategic plan envisions a future advanced computing ecosystem that provides the foundation for continuing American leadership in science and engineering, economic competitiveness, and national security. It develops a whole-of-nation approach based on input from government, academia, nonprofits, and industry sectors, and builds on the objectives and recommendations of the 2019 *National Strategic Computing Initiative Update: Pioneering the Future of Computing*.[1] This strategic plan also identifies agency roles and responsibilities and describes essential operational and coordination structures necessary to support and implement its objectives.

The plan outlines the following strategic objectives:

- Utilize the future advanced computing ecosystem as a strategic resource spanning government, academia, nonprofits, and industry.
- Establish an innovative, trusted, verified, usable, and sustainable software and data ecosystem.
- Support foundational, applied, and translational research and development to drive the future of advanced computing and its applications.
- Expand the diverse, capable, and flexible workforce that is critically needed to build and sustain the advanced computing ecosystem.

This strategic plan also complements the objectives and activities of other initiatives and national priorities, including the National Quantum Initiative (NQI),[2] the American Artificial Intelligence Initiative (AAII),[3] and the Industries of the Future (IotF),[4] and envisions synergistic relationships with these initiatives and priorities.

[1] https://www.whitehouse.gov/wp-content/uploads/2019/11/National-Strategic-Computing-Initiative-Update-2019.pdf

[2] https://www.congress.gov/bill/115th-congress/house-bill/6227

[3] https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/

[4] https://www.whitehouse.gov/briefings-statements/america-will-dominate-industries-future/

1 Introduction

The national advanced computing ecosystem remains essential to the Nation's leadership in science and engineering (S&E), global economic competitiveness, and national security. Emerging and future applications aligned with key national priorities and having great societal importance will depend on the capabilities and services provided by America's future advanced computing ecosystem (e.g., see Figure 1).

> ***Figure 1.*** Examples of Federal Agency Use Cases for a Future Advanced Computing Ecosystem
>
> ***DOE and NIH:*** These agencies have a joint effort to develop new therapies and improve the ability to detect cancer at an early stage. Because more than 30 percent of all cancers are driven by a change in one biomolecular protein family (RAS), therapeutic solutions will need to employ adaptive machine learning (ML) that can help drive multiscale physical simulations to give a more realistic view into RAS cancer biology. This requires a computing ecosystem that brings together huge datasets, massive computational and simulation power, novel artificial intelligence (AI)-based data analytics, and the relevant expertise. This advanced computing ecosystem could effectively enable cancer biologists to gain insights into the biology of RAS and the effects of potential drug therapies.
>
> ***NASA***: The goals of NASA's Artemis program are to land American astronauts, including the first woman and the next man, on the Moon by 2024; establish a sustainable human presence in space; and send American astronauts to Mars. This requires solutions for complex coupled physics challenges (such as lightweight galactic cosmic radiation protection; entry descent and landing systems for a spacecraft through the thin Martian atmosphere; and habitats pre-staged for Astronaut arrival), and necessitates a future advanced computing ecosystem to enable the next generation of human space exploration.

The National Strategic Computing Initiative (NSCI) was launched by Executive Order (EO) 13702[5] in July 2015 to advance U.S. leadership in high performance computing (HPC). The NSCI was designed as a whole-of-nation effort to create a cohesive, multiagency strategic vision and Federal investment strategy, executed in collaboration with academia and industry, to maximize the benefits of HPC for the United States.

The computing ecosystem is rapidly evolving across multiple dimensions due to the introduction of new and potentially disruptive technologies and paradigms, increased hardware heterogeneity, data volumes, software complexity, and novel approaches such as those based on AI and ML. These changes coincide with the unique requirements of new classes of compute-intensive, data-driven application workflows, such as those mentioned above, and the growing demands for improved robustness, agility, usability, and productivity. In order to address these emerging challenges and needs and to maintain U.S. leadership in advanced computing, it is essential to ensure continual research, development, and deployment of innovative paradigms and technologies with increased capabilities and capacities. This will catalyze emerging and future applications of real-world significance.

The Office of Science and Technology Policy (OSTP), through the National Science and Technology Council (NSTC) Committee on Technology and Subcommittee on Networking and Information Technology Research and Development (NITRD), created a Fast Track Action Committee (FTAC) on Strategic Computing in June 2019 to tackle the myriad challenges of this emerging computing landscape and to update the 2016 *National Strategic Computing Initiative Strategic Plan*.[6]

[5] https://www.govinfo.gov/content/pkg/CFR-2016-title3-vol1/pdf/CFR-2016-title3-vol1-eo13702.pdf
[6] https://www.whitehouse.gov/sites/whitehouse.gov/files/images/NSCI%20Strategic%20Plan.pdf

In their report, *National Strategic Computing Initiative Update: Pioneering the Future of Computing*,[7] the FTAC explored current and emerging challenges and opportunities. They revisited the objectives of the NSCI with an overarching goal of realizing an advanced computing ecosystem that combines heterogeneous computing systems (from extreme-scale to edge-centered systems and beyond) with the networking, software, data, and expertise required to support national security and defense as well as U.S. scientific and economic leadership. Inputs to the FTAC included public responses to a Request for Information on Update to Strategic Computing Objectives,[8] and stakeholder observations from a Future Computing Community of Interest meeting.[9] The resulting refocused objectives and recommendations identified in the 2019 FTAC report are:

- Pioneer new frontiers of digital and nondigital computation to address the scientific and technological challenges and opportunities of the 21st century.
- Develop, broaden, and advance the Nation's computational infrastructure and ecosystem.
- Forge and expand partnerships for the future of computing to ensure American leadership in science, technology, and innovation.

The NSTC Subcommittee on Future Advanced Computing Ecosystem (FACE) was established following the recommendations of the FTAC to coordinate the Federal agencies in a strategic approach to realizing the future advanced computing ecosystem, in alignment with the NSTC NITRD Subcommittee and other subcommittees as appropriate. This Subcommittee provides a venue for unclassified and classified discussions of both R&D and strategic considerations, facilitates cross-agency opportunities and collaborations, ensures synergies with related national initiatives and activities, and maintains awareness of the global computing ecosystem.

The strategic plan presented in this document was developed by the FACE Subcommittee to respond to the 2019 FTAC objectives and build toward a future advanced computing ecosystem that provides the foundation for American leadership in S&E, economic competitiveness, and national security in the coming decades. It defines agency roles and responsibilities and develops a whole-of-nation approach to advancing this vision. The plan also describes essential operational and coordination structures necessary to support and implement its objectives.

This plan is informed by the August 2020 OSTP Convening: Pioneering the Future Advanced Computing Ecosystem[10] that brought together key stakeholders from government, academia, nonprofits, and industry. The workshop highlighted the critical importance of an integrated advanced computing ecosystem, the central role of software and data, the essential investments in foundational and translational research, the crosscutting challenges associated with fostering and sustaining a skilled workforce, the risks associated with ensuring robust and reliable supply chains, and the need for novel partnerships among the ecosystem's various stakeholders. The figures in this plan illustrate Federal agency programs that contribute to and rely on a sustainable, scalable, and secure future advanced computing ecosystem.

---

7 https://www.whitehouse.gov/wp-content/uploads/2019/11/National-Strategic-Computing-Initiative-Update-2019.pdf

8 https://www.Federalregister.gov/documents/2019/06/18/2019-12866/request-for-information-on-update-to-strategic-computing-objectives

9 https://www.nitrd.gov/nitrdgroups/index.php?title=FC-COI-2019

10 https://www.nitrd.gov/nitrdgroups/index.php?title=Pioneering-the-Future-Advanced-Computing-Ecosystem

2 Application Drivers

The landscape of applications essential to American leadership, economic competitiveness, and national security is rapidly evolving. This evolution is driven by new S&E frontiers; national priorities such as IotF; and the growing scale and availability of computing, data, and novel technologies such as artificial intelligence/machine learning (AI/ML) and edge processing. These critical applications present new requirements for, and will be enabled by, the future advanced computing ecosystem. Summarized in the following subsections are the key classes of emerging and future applications that will have broad impact on our Nation, science, and society.

2.1 Extreme Scale Modeling & Simulation

Extreme scale modeling and simulation (M&S) applications can model phenomena more holistically across multiple scales and physics and require high-resolution, dynamic, coupled simulation workflows. Furthermore, these applications are increasingly integrating data analytics as well as new AI/ML approaches to improve modeling, often as part of *in situ* workflows. (Please see Figure 2.)

Figure 2. Example Application Drivers: Extreme Scale M&S

Coupled simulations based on multiscale, multiphysics application codes: Examples include modeling and simulation of hypersonic and space vehicles (DoD, NASA); physics-based simulations for stockpile stewardship (DOE/NNSA); exploration of the harsh environments of space (NASA); modeling of multiple earth system components to increase fidelity and skill of predictions (NOAA); and extremely short design analysis cycles for highly complex and ultrahigh-reliability systems for aviation and human space exploration (NASA).

Simulation workflows with integrated data analytics combining simulations with (*in situ*) data analytics: Examples include the modeling of complex physical systems (DoD, DOE); complex biological/biomedical processes and systems with coupled large-scale, heterogeneous, and streaming data analytics (DARPA); comprehensive situational awareness, anytime, anywhere, and across the battlefield (DoD); coupled simulations and data analytics for weather and climate prediction and coastal predictions (NOAA); multi-messenger astrophysics involving advanced computing and rapid data exchange among a multidisciplinary array of observational instruments (NSF); and large-scale data analytical capabilities, coupled with advanced computational simulations, for discoveries of important biomedical interactions, beyond cell-membrane-initiated signaling cascades at unprecedented scale and fidelity (NIH and DOE).

The complexity and scale of the resulting M&S application workflows presents new challenges and requirements for the future advanced computing ecosystem. The integration of coupled models, data, and analytics requires composable execution environments that integrate appropriate resources and capabilities. Furthermore, these applications are significantly impacted by changes in the technology landscape and warrant long-term foundational, applied, and translational research.

The above challenges are further complicated by the growing complexity of algorithms and programming models as well as the need to modernize current code bases[11] to adapt to new methods and modern development paradigms. The usability and manageability of complex extreme scale systems, as well as the need to ensure overall productivity, are added challenges.

2.2 Data-Intensive Application Workflows

Growth in the scale and availability of data in all domains, and the desire to transform the data into knowledge, insights, and actions, is resulting in data-intensive application workflows across all of S&E.

[11] A codebase (or code base) is the complete body of source code for a given software program or application.

These workflows require extreme-scale and data-intensive computing as well as novel approaches to data management and sharing. These requirements are further amplified by the desire to effectively leverage innovations in technologies, data science, and analytics techniques such as AI/ML.

Concurrently, data production from next-generation instrumentation, such as “long tail” data (large in number, heterogeneity, and application) is creating a need to process streaming data as well as to create federated distributed networks of data storage, preservation, and analysis systems. Furthermore, standard interfaces for data discovery and access, interdisciplinary data integration, scalable data analytics using AI/ML, data management and curation, and data dissemination are essential for ensuring the data is FAIR (findable, accessible, interoperable, and reusable).[12] Managing and processing protected/sensitive data presents additional challenges.

Together, the requirements of data-intensive application workflows will significantly influence the structure and operation of the future advanced computing ecosystem, necessitating new foundational, applied, and translational research; robust, secure and sustainable software stacks; as well as scalable and composable execution environments. (Please see Figure 3.)

Figure 3. Example Application Drivers: Data-Intensive Computing

Mining massive volumes of data. Examples include mining of critical observational data to better understand planet Earth, the solar system, and the universe (NASA); time-critical analysis of geospatial data to monitor, analyze, and predict current and evolving dynamics of complex human and natural Earth-system interactions (USGS); and identifying the coronavirus sequence from over 10 years of data (32 petabytes [PB] total) collected by the global research community (NIH).

Predictive modeling. Data-intensive applications that leverage data from numerous sources, at multiple levels of classification, and training large AI/ML systems using large-scale accelerator platforms are critical. Examples include decision aids (DARPA), advanced data-driven biomedical technologies (NIH), and data-intensive electromagnetic cyber analytics, military systems performance, and digital twins (DOD).

Federated networks of data storage/preservation/analysis. Examples include the Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability Initiative partnerships with commercial cloud providers to make available over 80 PB of federated datasets and advanced computational infrastructure, tools, and services (NIH).

2.3 End-to-End and Real-Time Application Workflows

End-to-end workflows and real-time applications enable data-driven and timely decision making to support accurate and rapid responses. Such workflows are essential to a broad range of application areas including disaster response, complex missions, and transformational science. They are particularly challenging when decisions must be made in minutes, if not seconds, and involve heterogeneous and unreliable system components that span the globe.

Similarly, transformational science accelerates innovations across multiple domains by exploiting end-to-end workflows and integrating online data sources with modeling and simulation. For example, with the emergence of the Internet of Things, highly instrumented physical systems are now available to provide massive amounts of real-time data.

This class of application workflows present stringent demands on the structure, composition, and operation of the future advanced computing ecosystem. For example, urgent, real-time application

[12] Wilkinson, M., Dumontier, M., Aalbersberg, I., et al. The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data* 3, 160018 (2016). https://doi.org/10.1038/sdata.2016.18

workflows might include intelligent adaptive networks, intermediate reconfigurable computing architectures, innovative algorithmic methods, and a sustainable software and data ecosystem. These are needed to improve the data-to-information balance, but also to verify the information that impacts real-time diagnostic, predictive, and prescriptive applications. Similarly, transformational science workflows must effectively leverage capabilities and resources that are distributed across a high-performance programmable network. (Please see Figure 4.)

Figure 4. Example Application Drivers: End-to-End Workflows

Disaster situations. Examples include requiring real-time diagnostic, predictive, and prescriptive computations to enable coordinated situation awareness and focused response via secure analysis of multiparty data from a large number of edge sources challenged by changing network and computational latencies (DHS, DoD).

Transformational science. Examples involve using highly instrumented physical systems to provide massive amounts of real-time data, which can be used in conjunction with simulated data to steer experiments, manufacturing processes, engineering development testing (e.g., wind tunnel testing), changes in biosphere integrity, biogeochemical flows, land system changes, biotechnology, and therapeutic developments. (DOE/SC, NASA, NSF, USGS).

Complex missions. Examples include all-domain command and control with comprehensive cyber situational awareness, anywhere, and at any time (DoD); large-scale, faster-than-real-time modeling and simulation capability and associated advanced analytics that support rapid system-of-systems analysis (DARPA); live, constructive, and virtual distributed experimentation environments that support large-scale advanced warfighting architecture development (DARPA); exascale-to-edge integration to support autonomous teams comprising both human decision makers and unmanned platforms (DoD); complex, real-time analytics to support situational awareness and operations in uncertain environments and ensure reliability, efficiency and recoverability (DoD); and diagnosing biothreats, including harsh environmental effects, and assessing therapeutic effects (DoD).

3 Pioneering the Future Advanced Computing Ecosystem

3.1 Strategic Objective 1: Advanced computing ecosystem as a strategic national asset

Today's advanced computing systems and services are the result of long-term, strategic, and aligned investments that represent a strategic national asset. This asset empowers American S&E leadership, fuels innovation in nearly all sectors of the U.S. economy, and enhances national security and preparedness. (Please see Figure 5.)

The emerging applications highlighted in Section 2, including those underlying IotF, are giving rise to profoundly new and diverse requirements for the future advanced computing ecosystem. For example, there is an unprecedented need to stitch together, in real-time, streaming data from distributed and heterogeneous sources for AI-based application workflows. At the same time, this ability is essential for continued S&E discovery, critical societal decision making, and more generally, advancement of American ingenuity and competitiveness on the global stage.

The future advanced computing ecosystem envisioned in this strategic plan will be comprised of a federated set of resources and services that are heterogeneous in architecture (e.g., classical and quantum); resource type (e.g., HPC, cloud, accelerators, network-centric, embedded and real-time, and edge); and usage mode. This will allow for seamless alignment between the varied resource offerings and individual use cases (e.g., data and computational experiments). Key attributes of this ecosystem

Figure 5. Examples of Current Activities Supporting Strategic Objective 1

- **DoD, DOE/NNSA, DOE/SC, NIH, NSF**: Creating an integrated fabric of advanced computing systems and services spanning HPC, cloud, hybrid, and other emerging models.
- **DARPA, DHS, DoD, FBI, NASA, NIH, NOAA, NSF**: Exploring novel computing paradigms to enhance data and computation.
- **DoD, DOE/SC, NASA, NIST, NSF, USGS**: Leveraging real-time streaming observational and experimental data sources such as instruments and sensor networks to pursue advanced computing, including data analytics, AI, and M&S; and accelerating scientific breakthroughs and predictive decision making.
- **DARPA, DHS, FBI, NASA, USGS**: Catalyzing compute environments that foster partnerships among government, academia, nonprofits, industry, and international counterparts to align new and emerging platforms with mission-specific problems.
- **NASA, NOAA, USGS**: Working with vendors to specify and pursue requirements for Earth system prediction models while understanding emerging technologies that will improve capabilities of the end-to-end prediction system (observations to post-processing).
- **DARPA, DHS, DoD, NIST, NSF**: Pursuing security, trust, verifiability, and usability in advanced computing systems and services.
- **DARPA, DOE/NNSA, DOE/SC, NIST, NSF**: Securing the microelectronics and photonics supply chains—including the development of next-generation microelectronics materials and devices—underlying the advanced computing ecosystem from chip design through system integration.

include reconfigurability, programmability, reliability, and energy-efficiency; it will address security and privacy at its core.

The future advanced computing ecosystem will collectively represent a strategic national asset spanning government, academia, nonprofits, and industry. Elements of this ecosystem have served as a bedrock of U.S. competitiveness for decades, but in the future, this ecosystem will be the foundation upon which American leadership will be built at the frontiers of S&E, including the IotF. It will power economic competitiveness and national security preparedness well into the 21st Century and serve as a proving ground for new and emerging technology capabilities that include neuromorphic, bioinspired, quantum, analog, hybrid, and probabilistic computing. Through shared platforms and testbeds, agencies will be able to leverage synergies and efficiencies to evaluate novel technology concepts and facilitate the hardening and eventual translation of these concepts to practice.

3.1.1 Objectives

- **Federate a spectrum of capabilities and capacities** (including data, software, networking and security) that can be used collectively as a strategic national asset.
- **Address the needs of emerging application workflows** (e.g., IotF) that have diverse advanced computing requirements as well as natural affinities to specific innovative technologies, system architectures, and usage modes.
- **Promote/support the availability, integrity, and security of critical advanced computing components** in the international software and hardware supply chains.
- **Accelerate access to innovative computing paradigms, technologies, and capabilities** (e.g., post-exascale, neuromorphic, bioinspired, quantum, analog, hybrid, and probabilistic computing), while integrating and sustaining existing advanced computing systems critical to agency missions.
- **Leverage crosscutting synergies and efficiencies** across government, academia, nonprofits, and industry, and with like-minded international counterparts.

3.1.2 Execution Plan

- **Federate a spectrum of capabilities and capacities** through enhanced coordination mechanisms (e.g., consortia models and shared user agreements) to enable greater access to a federated set of advanced computing systems and services that allow for seamless alignments between resource types and application use cases.
- **Address the needs of emerging application workflows**:
  - Support future, data-driven, AI-enhanced workflows (e.g., IotF and urgent science) that necessitate a composition of different usage modes and resource types, including HPC, cloud, accelerator, graphics processing unit (GPU), network-centric, and edge.
  - Develop models for sustaining and sharing the software and services essential to national priority applications. This includes agreement on standard sets of platform-agnostic, interoperable software and workflows to support emerging application activities.
  - Establish a common, secure, multiparty computing protocol for information and data sharing.
- **Promote/support the availability, integrity, and security of critical advanced computing components**:
  - Support approaches that ensure a secure, verifiable, provable, reliable, and adaptable advanced computing ecosystem. This will require integrating hardware underpinnings that provide roots of trust, full open-source stacks, and secure configurations of massive heterogeneous systems.
  - Introduce technologies to ensure supply chain security into chip designs, wafer foundries, packaging facilities, and integration facilities for the building blocks of advanced computing. Technologies include chip design verification and circuit defect and anomaly detection that are in-line to production workflows, as well as decision-driven approaches to secure supply chain integrity.
  - Reestablish domestic materials, equipment, and expertise supply chains, and invest in domestic foundries.
- **Accelerate access to innovative computing paradigms, technologies, and capabilities**:
  - Use consortia models to help develop platforms and testbeds for distinct novel technologies, in at-scale configurations, and across the computing stack. Testbeds would be made available to government, academia, nonprofits, and industry researchers to ensure the entire U.S. research community can develop relevant approaches based on at-scale experiences and expertise.
  - Support reliability and portability across diverse platforms, and encourage research on end-to-end software, algorithms, and workflow approaches with distinct compiler targets to demonstrate these capabilities.
- **Leverage crosscutting synergies and efficiencies**:
  - Establish partnerships across government, academia, nonprofits, and industry, and with like-minded international counterparts.
  - Support strategic approaches for emergency-response scenarios, including coordination services and mechanisms, as exemplified by the COVID-19 HPC Consortium (see Figure 9 in Section 4.2).
  - Maintain and expand public-private partnerships (PPPs) to support access to resources, novel experimental workflows, and technology innovation and translation.

3.2 Strategic Objective 2: Robust, sustainable software and data ecosystem

A trusted, verified, usable, and sustainable software and data ecosystem is essential for translating innovations in science and technology into national S&E and economic leadership. Software supporting the future advanced computing ecosystem must balance the following attributes while ensuring robustness and correctness: efficiency of development, debugging, verification and validation; usability, reproducibility, manageability, extensibility, and sustainability; security, privacy, and trust; and performance and scalability. Software must be capable of operating in multiple modes and with high degrees of parallelism, with efficient management of memory and input/output, while also supporting workflow composability and execution. Emerging computing technologies present new opportunities (e.g., real-time response, embedded processing of data from sensors or actuators, and steering of models and experiments), but they also require new algorithms, models of computation, data, programming environments, and software stacks. Significant disruptions from emerging technologies and paradigms are evident at all levels of the software stack; the growing complexity of algorithms and programming models (designed to overcome these disruptions and integrate emerging technologies and paradigms into existing applications) further exacerbate software challenges. (Please see Figure 6.)

Figure 6. Examples of Current Activities Supporting Strategic Objective 2

- **DARPA, DoD, DOE/NNSA, DOE/SC, NASA, NIST, NOAA, NSF**: Modernizing legacy codes in the exascale era and beyond, including mapping algorithms to new hardware architectures and investing in software modernization for coupled community modeling efforts.
- **DARPA, DoD, DOE/NNSA, DOE/SC, NIST:** Conducting foundational research in and developing software verification and validation approaches for nondeterministic systems and uncertainty quantification.
- **DoD, DOE/SC, NASA, NSF**: Developing sustainable software/data services at the element, framework, and institute levels built on sound software engineering and architecting; designing and prototyping scalable end-to-end data cyberinfrastructure; and operating shared-use production software and data services.
- **NASA**: Developing and deploying an open-source, commercial-cloud-based ingest, archive, and distribution system using a community ecosystem development model; and consolidating and streamlining data stewardship processes across multiple Earth Science disciplines using cross-functional and distributed teams.
- **DHS, FBI:** Creating development environments that maximize usability and utility of available resources and employee productivity.
- **NIST:** Developing and deploying characterization and measurement science for technologies beyond CMOS (complementary metal–oxide semiconductor) and alternate paradigms, including quantum and neuromorphic.
- **DoD, DOE/SC, NASA, NIH, NOAA, NSF:** Curating data to ensure that it is findable, accessible, interoperable, and reusable and available to the agency, government, and academia.
- **NASA, NIH, USGS:** Deploying trusted digital repositories, science data catalogs/metadata, and data lifecycle management; and generating new biomedically relevant datasets amenable to ML analysis.
- **DARPA, DoD, DOE/SC**: Creating multidisciplinary codesign teams focused on a more holistic design approach from device to the applications, and supporting rapid software development and fielding, as well as providing support throughout the lifecycle of complex hardware and software systems.

Software development, including adapting large legacy applications to new systems, continues to be a challenge. Existing application codebases represent many years of development effort and have increasingly large sustainment costs, including personnel and maintenance. Investments in modernizing legacy codebases will require significant development resources and will incur additional risks in both time and cost. These investments are crucial, allowing the community to take advantage of existing and emerging processing capabilities. This will require new approaches for balancing stability and maintenance (e.g., libraries and other tools used by applications, with innovation and software evolution).

Open-source software provides an important and growing contribution to existing and future software environments. Continued fostering of the open-source community will be needed to ensure not only a growing body of software, libraries, programming environments and tools for research, but also a sustainable, robust, and trusted environment for operational practice. Furthermore, there are new opportunities for "smarter" software that leverage advances in AI/ML to address software and system complexities.

Large data volumes produced by next-generation instrumentation, as well as exponentially growing long-tail data, are creating a demand for a distributed but federated network of data preservation and processing systems. Solutions to many global challenges require interdisciplinary data integration based on agreed-upon standard interfaces for data discovery, access, compatibility, and reusability. International organizations are advancing these capabilities, and U.S. and European policies exist to make publicly funded research data openly available for technologies such as AI/ML. These technologies need access to large bodies of carefully curated, interoperable, and validated training data.

3.2.1 Objectives

- **Ensure a robust and sustainable software ecosystem** that will translate technology innovations into national S&E leadership.
- **Support needs for novel software development** (e.g., algorithms, programming systems, and runtimes).
- **Ensure a robust data ecosystem** that includes collaborative data management platforms for real-time processing, curation, analysis, and sharing of data across hardware platforms and geographic locations; increased availability of data across government, academia, nonprofits, industry, and the public; and an accelerated pace of discovery.
- **Develop, deploy, operate, and promote trusted services and capabilities** that ensure secure and effective management and high utilization of resources.
- **Explore inovative models for PPP** aimed at developing models for software and data innovation and sustainability.

3.2.2 Execution Plan

- **Ensure a robust and sustainable software ecosystem** by supporting novel software development as well as by modernizing legacy application codes, libraries, and software tools:
  - Support multidisciplinary codesign teams focused on a more holistic approach for novel software development (e.g., algorithms, programming systems, runtimes, and workflows) driven by emerging technologies such as AI/ML, new computing paradigms, and "smart" applications.
  - Foster the use of agile software development, continuous design, and integration processes for both classified and unclassified networks that support rapid software development and fielding. This includes providing support throughout the lifecycle of complex hardware and software systems.
  - Deploy exascale-capable software and exascale-ready applications to usher in the exascale ecosystem era; plan for the post-exascale era that includes "smarter" software and software processes.
  - Support open software libraries and software sharing.
  - Create and deploy a collaboration environment to share best practices, analysis, implementation issues, and results.

- **Support needs for novel software development** that are driven by emerging technologies and computing paradigms and that address their unique requirements, including access to novel hardware platforms and testbeds.
- **Ensure a robust data ecosystem**:
  - Promote and practice robust, proactive information-security procedures to ensure appropriate stewardship of protected/sensitive data while at the same time enabling scientific advances.
  - Pursue ubiquitous implementation of FAIR principles for open data.
  - Support distributed data storage and secure multiparty sharing and processing.
  - Deploy trusted digital repositories, science data catalogs/metadata, and data lifecycle management to ensure that data adheres to FAIR principles.
  - Support high-performance data analytics/data science software in areas of graphs and advanced statistical analyses, quantum information science (QIS), and ML including deep learning, especially training of models with large numbers of parameters against very large data sets.
  - Develop and deploy collaboration environments that facilitate secure sharing of data, analysis techniques, and results across organizations.
- **Develop, deploy, operate, and promote trusted services and capabilities**:
  - Continually refine investments in emerging technologies to inform the development, testing, and hardening of the components of the software ecosystem.
  - Increase use of emerging cybersecurity technologies (e.g., use formally verified microkernels to reduce the attack surface).
  - Continue fostering a growing open-source software community with additional focus on translating software and software processes to trusted operational practice.
- **Explore innovative models for PPP** such as consortia or centers of excellence that are focused on innovation and sustainability of software and data solutions and provide the foundational components of an open science ecosystem to increase collaboration and efficiency and ensure that scientific results are reproducible.

3.3 Strategic Objective 3: Foundational, applied, and translational R&D

Major inflection points and existential threats are already impacting the advanced computing ecosystem; its future trajectory hinges on bold, urgent, and visionary action. There are three crucial trends in need of urgent response: the end of Dennard Scaling[13] and the slowing of Moore's Law[14]; the data and AI tsunami; and the shift from concentrated advanced computing resources (i.e., "supercomputers") towards distributed edge-to-cloud confederations of compute and data resources. (Please see Figure 7.)

Transitioning to a post-Moore and post-von Neumann era: The end of Dennard scaling and the slowing of improvement in Moore's Law are leading to a shift away from traditional, von Neumann-based models of computing[15] to alternative models such as neuromorphic, bioinspired, quantum, analog, hybrid, and probabilistic computing. Novel materials for compute and storage may emerge from a

[13] Dennard Scaling states that as transistors get smaller, their power density stays constant.

[14] Moore's Law is the observation that the number of transistors in a dense integrated circuit doubles about every two years.

[15] "von Neumann architecture" refers to any stored-program computer in which an instruction fetch and a data operation cannot occur at the same time because they share a common bus.

Figure 7. Examples of Current Activities Supporting Strategic Objective 3

- **DOE, NSF:** Supporting center-scale activities on emerging AI and quantum applications and technologies, and advancing full-stack efforts for both use-inspired and foundational research towards the future computing ecosystem.
- **DARPA, NIST:** Exploring novel advanced computing architectures that support research in evaluation of characteristics of trustworthy AI, including "explainability" and identification of bias.
- **DARPA, NASA:** Assessing advances in advanced computing systems with a focus on optimizing legacy codes and developing new efficient applications, including incorporating data analytics and AI approaches to solve mission challenges.
- **FBI:** Exploring accelerators specialized for agency mission capabilities, and researching applications inspired by specific computing approaches or technologies.
- **DHS:** Addressing advanced computing challenges based on current operational gaps and driven by emerging risk assessments for homeland security missions.
- **DoD:** Supporting dynamic and multiresolution AI approaches to improve scalability and tractability.
- **NOAA:** Supporting and conducting R&D on systems to advance reliability, fault tolerance, memory, and storage efficiency at scale.
- **DoD, DOE:** Investing in a grand challenge benchmarking suite to guide development of future architectures and software ecosystems.
- **DARPA, NSF:** Conducting crosscutting research programs such as the Principles and Practice of Scalable Systems, Electronics Resurgence Initiative, and others to attack broad post-Moore challenges; developing heterogeneous computing architectures and the associated theory to simultaneously leverage classical and non-von Neumann computing elements; exploring new and existing techniques to advance a secure and trustworthy cyberspace; mathematical and computational algorithms, models, and mechanisms in support of long-term national security objectives; and investigating technical feasibility of the real-time detection of hardware Trojans.

range of possible sources (e.g., DNA storage, and quantum technologies). These materials and device innovations are needed for compute elements, but also for system memory, storage, and communication functionalities as systems continue to get more specialized, heterogeneous, and complex over time.

Novel technology building blocks will likely be partnered with the use of new design approaches, such as non-von Neumann computing elements. Challenges lie in architecting elements that make good use of these approaches and methods for integrating them into designs for heterogeneous processors and accelerators, heterogeneous memories and models, new interconnect technologies, and special-purpose and energy-efficient architectures. To sustain computing leadership in the coming decade, the United States must lead in the use of these alternative computing approaches. These alternatives include the exploration and design of novel systems based on these approaches, the complex systems that utilize these approaches, software infrastructure to support such systems, and applications that effectively exploit the systems.

Data and AI tsunami: Simultaneously, the scalability of data analysis and storage is emerging as an additional trend that needs more research attention. The past decade has seen the unprecedented growth of massive distributed data and information bases, and this expansion is expected to accelerate. Maintaining U.S. leadership in extracting insights from this flood of data requires improvements in both hardware and computing techniques that address the storage, transmission, and processing of data. Continued advances depend on the capability to reduce latencies while improving insights extracted from data. Furthermore, clear application drivers and benchmark suites are essential to enable the broader research community to coordinate on important application categories such as weather forecasting, national security, and physical/chemical/materials modeling.

The recent explosion in available data combined with the availability of inexpensive mass storage and computation on demand has brought AI/ML to the forefront. However, ML has shown some limitations, including its inability to handle inputs outside of its training set and susceptibility to manipulation by malicious adversaries. Successful integration of AI/ML technology into transformative applications requires a rigorous understanding of the ability of AI/ML systems to handle uncertainty, reduce vulnerabilities, and improve robustness. Likewise, the design of the future advanced computing ecosystem will need to accommodate the needs of this increasingly important application domain.

Edge to HPC and cloud, at scale: Advanced computing is seeing a shift from geographically concentrated resources to the global, at-scale distribution of data and compute resources. Future advanced computing systems will be highly multicore and heterogeneous within a single node or cabinet. They will increasingly involve distributed and cloud-connected storage and compute facilities that are co-located with scientific facilities. Research efforts must explore the broad use of such edge-to-extreme-scales and cloud-compute approaches, as well as the large-scale storage resources they make possible. Software and tools are needed to scale performance and facilitate the integration of heterogeneous platforms, including those within a given architecture, that are network-centric, or that are at the edge.

Testbeds and prototypes: Computational resources at scale are also critical in research testbeds and prototypes, as they lead to broad operational adoption. Such testbeds will need to federate resources across different layers of the hardware/software stack and be able to employ large-scale, cloud-connected resources from different sites. As ideas translate from testbeds to operational systems, sustainability and reliability become major issues; testbeds should support resources in a way that distinguishes mature ones from more nascent efforts.

As discussed below, meeting these needs will require sweeping and innovative explorations into all levels of the ecosystem—from hardware devices to system architectures and software stacks—and the abstractions and workflows by which they interconnect.

3.3.1 Objectives

- **Ensure hardware leadership** in a post-Moore/von Neumann era ensuring broad investments across diverse candidate technologies.
- **Advance software and software-hardware research** to enhance the scale and resolution of important problems that are tractable.
- **Address challenges and opportunities related to growing data volumes and successful translation of data into insights**.
- **Enhance AI capabilities** including real-time, at-scale, and with attributes of fairness and explainability.
- **Expand availability of and access to testbeds, prototyping, and research infrastructure** to encourage research, development, and sustainment of software tools needed to deploy applications onto increasingly complex systems.
- **Address the need for technologies that ensure hardware supply chain security** for the manufacturing, packaging, and integration of advanced and trusted computing ecosystem electronics.

3.3.2 Execution Plan

- **Ensure hardware leadership** through technology-driven innovations (e.g., neuromorphic, bioinspired, quantum, analog, hybrid, and probabilistic computing), for compute, memory, and interconnect elements:
    - Support research on novel, beyond-CMOS device and/or /hardware topics to support implementations of these computing paradigms; and support R&D in beyond-CMOS devices/logics, memories/storages, and interconnects, and their integration in, for example, superconductors, biomolecules, spin electronics, and photonics, to establish advanced hardware platforms for implementation of emerging computing paradigms.
    - Provide early and broad access to novel computing systems and associated software infrastructure in order to encourage further research and speed deployment.
    - Cultivate portfolios that are diverse in hardware technology type and in likely near-term versus long-term payoffs and establish a national priority list of focused investments in new/emerging/future technologies with periodic updates, while maintaining reasonable resources for blue-sky research.
- **Advance software and software-hardware research**:
    - Encourage full-stack research that spans across devices, systems, software, and applications.
    - Advance software programmability, performance, and scalability through investments in application expertise, hardware-informed approaches, and tools that support reliability and portability across diverse platforms.
    - Create new techniques, technologies, and tools that can lead to improvements in engineering practice for software-based systems.
    - Research end-to-end software, algorithms, and workflow approaches with distinct compiler targets to demonstrate these capabilities.
- **Address challenges and opportunities related to growing data volumes and translation of data into insights**:
    - Integrate tools and enable continuous validation of large-scale modeling and simulation with literature, code, and data.
    - Catalyze the development of automated model discovery techniques and tools that enable non-data scientists to create empirical models of real and complex processes and phenomena.
    - Develop the technical means to protect private and proprietary information. This will relieve the tension between maintaining privacy and being able to tap into the huge value of data.
- **Enhance AI capabilities,** including real-time, at-scale, and with attributes of fairness and explainability:
    - Expand platforms that support hybrid approaches by integrating data-intensive AI/ML into modeling and simulation efforts to enable validation of models using real-time data feeds, and support novel hybrid approaches that combine physics-based simulation with deep learning modules to generate more-accurate predictions.
    - Support integrated efforts requiring collaboration technologies; a mixture of HPC, cloud, and accelerator (e.g., GPU) provisioning; and shared access to data across different types of computational platforms.
    - Incentivize active research collaborations (e.g., attract the AI community to computing centers by facilitating access to large datasets).

    - Maintain AI's strengths beyond deep learning and neural networks, including symbolic learning and other approaches.
- **Expand availability of and access to testbeds, prototyping, and research infrastructure**:
    - Provide early and broad access to novel computing systems and associated software infrastructure.
    - Develop edge-to-extreme-scale and cloud capabilities that offer programmability across distributed compute and storage elements.
    - Address challenges to post-Moore/post-von Neumann-era development by supporting access to and the translation of emerging technologies along with the necessary PPP models to sustain them.
    - Develop and deploy secure testbeds for distinct, novel technologies in at-scale configurations, and ensure their availability to the broad research community (i.e., academia and industry).
    - Meet the burgeoning demand of AI/ML disciplinary, application-focused, and translational research by provisioning university and government R&D computing centers and testbeds with data storage, hardware, and secure integration platforms (including GPUs).
    - Provide agency-facilitated pathways for the fabrication of research prototypes, including both semiconductor fabrication and the broader space of novel materials and computational substrates at the research frontier.
- **Address the need for technologies that ensure hardware supply chain security**:
    - Develop technologies and techniques that holistically ensure full-stack computer security.
    - Incorporate innovative data protection techniques such as differential privacy, data security standards, and flexible data access controls.
    - Invest in research that goes beyond performance, performance-per-watt, and accuracy, to prioritize other crucial characteristics including security (e.g., network operation analytics, attribution of attacks, applied cryptography such as secure multiparty computation, graceful degradation and recovery from attacks, and social engineering defenses).
- **Address the sociotechnical dimensions related to security, trust, privacy, and fairness**:
    - Support research on techniques and tools for verifying the integrity of data and information, including data received via media and other open sources as well as data from trusted sources that might have been compromised.
    - Support research on hardware-based computing security technologies (e.g., devices, circuits, systems, and combinations).
    - Prioritize research in sociotechnical constructs for computer security, programmability, and other dimensions.
    - Maintain and expand PPPs to share technological foresight and help catalyze the translation of these technologies.

3.4 Strategic Objective 4: Fostering a diverse, capable, and flexible workforce

Effectively leveraging the advanced computing ecosystem requires both turning these resources and capabilities into practical and usable forms (e.g., making the advanced computing, software, and data systems and services more intuitive and easier to use) and developing the skilled workforce that can build the tools, operate the systems, and support a broad range of users. This is especially challenging as new technologies are introduced into complex institutional settings where legacy systems cannot be easily displaced or reconfigured.

The essential workforce should be trained in the current state of the art, existing capabilities, and the ability to anticipate, exploit, and integrate future technologies and solutions. The new computing professional must communicate effectively with a variety of stakeholders and end users to understand their concerns and help them migrate to new, agile, and ultimately more effective environments. Such professionals must be able to adapt quickly and delve deeply as needs change and challenges arise. Training and skills range from those needed to stand up a computing or data center to targeted R&D, organizational management, and marketing and communication. Recruiting advanced computing talent including system architects, designers, integrators, administrators, and application developers can be challenging. Moreover, developing and sustaining a diverse, capable, and flexible workforce involves both training within educational institutions and on the job as technologies, platforms, and applications evolve. Developing the necessary curricula and tools for workforce training, reskilling, productivity, and collaboration is equally important. Finally, it will be essential to establish synergies across government, academic, nonprofit, and industry stakeholders and to develop creative recognition, incentive, and reward mechanisms.

The needs of the advanced computing ecosystem align with, and are embedded in, the societal challenges of encouraging and incentivizing U.S. citizens to progress through the formal educational process to the more highly skilled levels. The importance of expanding the pathways in the formal educational process to accommodate underrepresented U.S. populations has been well documented in numerous studies.[16] The challenges begin in the educational system itself, where greater specialization is required as students progress from undergraduate through postgraduate education and training. In contrast, the rapidly changing advanced computing ecosystem requires highly skilled individuals with breadth and depth in many fields (e.g., data science, AI/ML, and privacy-preserving computing) as well as interdisciplinary training and experience. Recruitment, retention, and cultivation of a highly capable, adaptive, and agile workforce is also difficult. Finally, while the formal educational pathway is critical, on-ramps for nontraditional students (e.g., those who are older and/or seeking reskilling or upskilling opportunities) requires thoughtful educational practices, mentoring, and support that provides them the means to advance and be successful.

Institutions of higher education have been the backbone for the education and training of the advanced computing workforce. However, Federal Government needs to provide the highly specialized training appropriate for agency missions. (Please see Figure 8.)

Within the ecosystem of higher education, increased engagement with institutions that serve underrepresented groups (e.g., Historically Black Colleges and Universities, Tribal Colleges and Universities, Hispanic-Serving Institutions, community colleges, and undergraduate institutions) will expand and broaden the talent pool. Government, academia, nonprofits, and industry play vital roles and benefit from migration of talent across sectors They can engage with higher education through fellowships; scholarships; internships; programs that support continuous lifelong learning, skilling, and reskilling; and alternative pathways for nontraditional learners.

[16] See, for example, *Assessing and Responding to the Growth of Computer Science Undergraduate Enrollments,* https://doi.org/10.17226/24926; *Retention in Computer Science Undergraduate Programs in the U.S.: Data Challenges and Promising Interventions*, https://www.acm.org/binaries/content/assets/education/retention-in-cs-undergrad-programs-in-the-us.pdf; *Barriers and Opportunities in 2-Year and 4-Year STEM Degrees*, https://doi.org/10.17226/21739.

Figure 8. Examples of Current Activities Supporting Strategic Objective 4

Fellowship Opportunities in Government, Academia, Nonprofits, and Industry

- **NASA, DoD, NSF** Pathways, an intern employment and recent graduates' program, offers employment, internship, fellowship, and scholarship opportunities for high school, undergraduate, and graduate students as well as educators.
- **NIH** Partnership with Civic Digital Fellows brings together student software engineers, data scientists, product managers, and designers who innovate at the intersection of technology and public service. Data and Technology Fellows recruits talented individuals from computer science and related fields to work in biomedical research.
- **NIST** Summer Undergraduate Research Fellowship Program and NIST/NRC Postdoctoral Associateship Program offer training to young researchers; partnerships in the Joint Quantum Institute and the Joint Center for Quantum Information and Computer Science at the University of Maryland provide training for graduate students and postdoctoral researchers.
- **NSF** INTERN program provides supplemental funding to support graduate student research in nonacademic/industry settings; the Advanced Technological Education (ATE) program focuses on two-year institutions of higher education, partnerships between academic institutions, and industry to promote improvement in the education of science and engineering technicians at the undergraduate and secondary institution school levels.
- **DoD** Future Scholars for STEM Workforce Development programs encourage formal and informal STEM education activities that generate impactful STEM educational experiences for both students and teachers.
- **DOE** Computational Graduate Fellowship as well as the Visiting Faculty Program and Community College Internships help ensure a supply of scientists and engineers trained to meet workforce needs in computational science, applied mathematics and computer science.
- **DARPA** Young Faculty Award Program engages rising stars in junior faculty positions in academia and equivalent positions at non-profit research institutions and introduces them to DoD and National Security challenges and needs.

Curriculum, Pedagogy, Training, etc.

- **DARPA** Joint University Microelectronics Program collaborates with the defense, commercial, and academic research communities to drive microelectronics path-finding research at scale for computing, artificial intelligence, and communication in the post-Moore's Law era.
- **DoD** HPC Internship and Mentoring Program provides training to high school, undergraduate, and graduate students, and Service cadets in key disciplines required for a wide class of DoD applications.
- **DOE/NNSA** programs engage with U.S. universities through mentoring, training, recruiting, and working with top researchers at the NNSA national laboratories in key disciplines required by stockpile stewardship.
- **NSF** Training-based Workforce Development for Advanced Cyberinfrastructure provides funding to programs to support learning and workforce development for cyberinfrastructure users, developers, and professionals.

Fostering a diverse, capable, and flexible workforce is a multidimensional challenge. Responding to it begins in primary school with ongoing STEM education efforts[17,18] and continues through the educational pipeline (undergraduate, graduate, and postgraduate) to career paths that provide professional growth and opportunities that attract and retain a talented and adaptive workforce. It also allows for nontraditional entrants to bring diverse skills and expertise as the ecosystem itself evolves.

[17] https://www.whitehouse.gov/wp-content/uploads/2018/12/STEM-Education-Strategic-Plan-2018.pdf

[18] DoD solicits proposals for Historically Black Colleges and Universities/Minority Institutions (HBCU/MI) research through various Broad Agency Announcement as well as other programs (e.g., DoD HPC program office and Army HPC Research Centers, https://basicresearch.defense.gov/Programs/HBCU-MI-Program). DoD HBCU/MI Intern Programs from high school to post-doctoral levels in HPC-ecosystem and cyber-related activities demonstrated great promise.

3.4.1 Objectives

- **Create the diverse workforce necessary** to achieve the goals of the future advanced computing ecosystem, support U.S. innovation, and push the leading edge of computation.
- **Develop training, upskilling, and reskilling strategies** that exploit the use of state-of-the-art technologies and anticipate future technologies and solutions.
- **Provide the necessary incentives, career paths, and reward structures** for retaining computing (hardware, software, data, security) professionals, technologists and practitioners.
- **Build synergies across government, academic, nonprofit, and industry stakeholders** focused on workforce development and training.
- **Foster relevant, mission-focused, on-the-job training** in the form of fellowships, academic programs, internships, and sabbaticals, in both intramural and extramural agency programs, federally funded R&D centers, and National Laboratories.

3.4.2 Execution Plan

- **Create a diverse workforce**:
  - Support, create, establish, and maintain K–12 STEM-education programs at the Federal, State, Tribal, and local levels. These programs could include informal mechanisms like boot camps, software factories, summer, and after-school programs that increase computational literacy in K–12 populations and expose diverse groups to advanced computing as a potential career.
  - Expand the formal educational pipeline to include under-resourced institutions and communities and broaden the talent pool.
  - Establish policies and programs to ensure that government organizations associated with the future advanced computing ecosystem include a diverse and agile workforce as a key component of their ecosystem.
  - Create career pathways in government, academia, nonprofits, and industry to attract, enable, and empower diverse talent.
- **Develop training, upskilling, and reskilling strategies**:
  - Develop and expand programs, pedagogy, and curricula through institutions of higher education, including community colleges and minority-serving institutions, appropriate to creating a diverse, multidisciplinary, and adaptive workforce. This training should include how to: conduct research to advance the technologies themselves (including AI/ML training) and to maintain the Nation's competitive advantage; deploy, manage, and operate the systems; and support outreach to future users of these systems to enable the broadest possible engagement by the R&D communities.
  - Create initiatives and programs to support assessment and evaluation, including identifying successful initiatives in the public and private sectors, at all levels (Federal, State and local) and their success factors; relevant information for monitoring and tracking metrics and evaluation strategies; and strategies for intervening to maximize programmatic effectiveness.
- **Provide the necessary incentives, career paths, and reward structures**:
  - Improve mechanisms that enable the transition of highly qualified graduates at all levels into the advanced computing and data science workforce.
  - Engage with a broad set of stakeholders to establish professional criteria and institutional and governmental job classifications together with broadly recognized certification programs that create career paths, provide incentives, and enable the transfer of newly acquired skills to take advantage of emerging employment opportunities.
  - Redesign contracts or other vehicles to enhance development of more capable software engineers.

- **Build synergies across government, academic, nonprofit, and industry stakeholders**:
  - Foster exchange programs within and across agencies and governments to address key collaborative issues, gaps, and concerns.[19]
  - Develop PPP programs that enable continuous learning, skilling, upskilling, and reskilling strategies targeting the use of the state-of-the-art computing technologies as well as the ability to anticipate and exploit future technologies and solutions together with on-ramps for nontraditional and mid-career learners.
- **Foster relevant, mission-focused, on-the-job training** through establishing programs to provide on-the-job training for undergraduate, graduate, and postdoctoral students through advanced computing developmental assignments (e.g., summer programs, internships, and fellowships) in academia, nonprofits, and industry.

4 Execution and Coordination

The FACE Subcommittee will lead the execution of the objectives in this strategic plan in coordination with other relevant NSTC subcommittees as appropriate and leveraging existing interagency working groups to foster coordination and collaboration at both strategic and tactical levels.

4.1 Agency Roles and Responsibilities

This strategic plan leverages three broad categories of agencies currently engaged in related research, development, and deployment that were identified in EO 13702 and presented in the 2016 *National Strategic Computing Initiative Strategic Plan*.[20] These categories are lead agencies, foundational R&D agencies, and deployment agencies, as described below. (Please see Table 1.) These groups may evolve over time as related mission needs emerge.

Table 1. **Summary of Agency Roles**

Agency	Lead Agencies	Foundational R&D Agencies	Deployment Agencies
Defense Advanced Research Projects Agency (DARPA)		X	
Department of Defense (DoD)	X		
Department of Energy (DOE)	X		
Department of Homeland Security (DHS)			X
Federal Bureau of Investigation (FBI)			X
ODNI/Intelligence Advanced Research Projects Activity (IARPA)		X	
National Aeronautics and Space Administration (NASA)			X
National Institutes of Health (NIH)			X
National Institute of Standards and Technology (NIST)		X	
National Oceanic and Atmospheric Administration (NOAA)			X
National Science Foundation (NSF)	X		

[19] Among the candidate topics for such exchanges are development of application software; generation and curation of data; and development and refinement of physics-informed machine/deep learning surrogate models for example, for vehicle design based on validated/verified physics-based application software and vast amounts of field data.

[20] https://www.whitehouse.gov/sites/whitehouse.gov/files/images/NSCI%20Strategic%20Plan.pdf

- ***Lead agencies*** will engage in mutually supportive R&D within their respective missions in areas related to objectives of the future advanced computing ecosystem; collaborate with deployment and foundational R&D agencies to accelerate deployment and integration; develop the workforce to support the objectives of the strategic plan for pioneering the future advance computing ecosystem, and for developing and delivering the next generation of integrated advanced computing capabilities.
- ***Foundational R&D agencies*** will enable fundamental scientific discovery and the associated advances in engineering necessary to support the objectives for the future advanced computing ecosystem strategic plan. The foundational R&D agencies will coordinate with academia and industry in conducting this R&D, and with industry and deployment agencies to enable effective transition of R&D efforts into commercial offerings that support the wide variety of requirements across the Federal Government.
- ***Deployment agencies*** will develop mission-based advanced computing requirements to influence the early stages of the design of new advanced computing systems, software, and applications; collaborate with lead and foundational R&D agencies to accelerate deployment and integration; integrate the special requirements of their respective missions; and engage with academia and industry to develop corresponding target requirements through bilateral engagements and Requests for Information.

The lead agencies will work closely with the foundational R&D agencies to develop and deploy novel future advanced computing technologies. The lead agencies will also work with the foundational R&D agencies and the deployment agencies to increase technology coherence, promote collaboration with academia and industry, and ensure the future advanced computing ecosystem addresses agency mission requirements while considering a wide variety of national needs.

4.2 Partnerships across U.S. Federal Agencies, Academia, and Industry

As outlined above, the Nation's advanced computing ecosystem relies upon advances in foundational, applied, and translational research leading to future generations of advanced computing systems and services; innovations in software and data infrastructure; and expertise to appropriately leverage these resources to support a broad range of application classes including large-scale data analytics, end-to-end workflows, as well as modeling and simulation. These requirements necessitate a holistic approach, spanning both technical and human dimensions, which in turn requires forging and expanding partnerships among Federal agencies, industry, nonprofits, and academia to appropriately leverage the ecosystem in support of the national computational infrastructure for science, engineering, and society.

Existing interagency partnerships take on many forms, with the overarching goal of coordinating and leveraging complementary agency strengths and mission foci. Likewise, existing partnerships between the Government, academia, nonprofits, and industry are multifold in their nature and foci, ranging from information exchanges, participation in consortia, and community events, to deeper technical R&D engagements. Federal agency partnerships with academic institutions typically focus on fundamental S&E, exploring next-generation technologies and solutions, clean-sheet ideation, and prototyping. These partnerships also address issues of workforce training and development. Complementing these partnerships, Federal agency partnerships with industry also focus on technology transfer and commercialization, system integration and large-scale system delivery, and maturation of components and their translation for production use. Federal agencies also have R&D agreements with industry and may leverage services provided by industry (e.g., cloud computing services) to support their missions. The COVID-19 HPC Consortium described in Figure 9 illustrates the value proposition for such multisector collaboration.

Figure 9. COVID-19 HPC Consortium

The White House, in partnership with the DOE, NSF, and IBM, launched the COVID-19 HPC Consortium (https://covid19-hpc-consortium.org/) on March 22, 2020, to help accelerate the pace of scientific discovery in the fight to stop the coronavirus causing the global COVID-19 pandemic. The consortium represents a unique PPP aimed at providing rapid access to the world's most powerful highly advanced computing resources to researchers worldwide spanning government, academia, nonprofits, and industry to support rapid COVID-19-related research. This PPP went from concept to practice within days, supporting breakthrough scientific research in epidemiology, virology, and microbiology, among other topics related to the pandemic.

The COVID-19 HPC Consortium has demonstrated the national and global importance and impact of U.S. leadership in the advanced computing ecosystem, the power of assembling complementary resources and expertise that can transcend organizational and geographic borders to address shared challenges, and the value proposition for multisector collaboration.

Long-term Federal agency engagement with academia and industry is necessary to inspire use cases; explore, develop, and produce novel technologies and capabilities; pilot solutions in specific mission contexts; and improve overall efficiency and productivity. Such a synergistic and collaborative approach leverages the Nation's unique innovation ecosystem that has pioneered innovative advances for decades and is essential to ensure U.S. leadership in science, technology, and innovation. Partnerships are also an important vehicle for recruitment and hiring. Finally, they increase awareness of global efforts related to the future advanced computing ecosystem.

Recognizing that the future of computing exists in a competitive global landscape, coordination of R&D should ensure that resources are available and efficiently used for both individual agency efforts and to enable cross-agency opportunities and collaborations. In addition, coordination activities should provide an appropriate venue for unclassified and classified discussions of both R&D and strategic considerations for the Nation's approach to the future advanced computing ecosystem.

5 Appendices

5.1 American Artificial Intelligence Initiative

On February 11, 2019, President Trump signed EO 13859 announcing the American Artificial Intelligence Initiative (AAII).[21] This strategy is a concerted effort to promote and protect national AI technology and innovation and to ensure that the United States maintains global leadership in AI. The Initiative implements a whole-of-government strategy in collaboration and engagement with the private sector, academia, the public, and like-minded international partners. It directs the Federal Government to pursue five pillars for advancing AI: invest in AI R&D, unleash AI resources, remove barriers to AI innovation, train an AI-ready workforce, and promote an international environment that is supportive of American AI innovation and its responsible use. The United States is also actively leveraging AI to help the Federal government work its own services and missions in smarter and more trustworthy ways.

Since the AI EO was signed, the Administration has called for record amounts of AI R&D investment and committed to doubling these investments over two years, led the development of the first international statement on AI Principles,[22] refreshed the national strategy for AI R&D,[23] issued the first-ever strategy for engagement in AI technical standards,[24] published the first-ever reporting of government-wide nondefense AI R&D spending,[25] released the first-ever AI regulatory document to foster confidence and trust in the use of AI technologies in the private sector,[26] and released the AAII year-one annual report summarizing progress and the continued long-term vision.[27] Under the AAII, the United States is embracing AI as a key industry of the future that will help Americans live longer, more productive, healthier, and safer lives.

5.2 National Quantum Initiative

Under President Trump's leadership, the United States is committed to strengthening its leadership in QIS. The President signed the National Quantum Initiative Act[28] in 2018, which greatly accelerated quantum R&D and established the National Quantum Coordination Office (NQCO) within OSTP to serve as a hub for national R&D and policy coordination. QIS as described in the *National Strategic Overview for Quantum Information Science*[29] arises from a synthesis of quantum mechanics and information theory and includes concepts and technology that support revolutionary advances in computing, communications, and metrology. QIS examines uniquely quantum phenomena that can be harnessed to advance information processing, transmission, measurement, and fundamental understanding in ways that classical approaches can only do less efficiently or not at all.

---

21 https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/

22 https://www.oecd.org/going-digital/ai/principles/

23 https://www.whitehouse.gov/wp-content/uploads/2019/06/National-AI-Research-and-Development-Strategic -Plan-2019-Update-June-2019.pdf

24 https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf

25 https://www.whitehouse.gov/wp-content/uploads/2019/09/FY2020-NITRD-AI-RD-Budget-September-2019.pdf

26 https://www.whitehouse.gov/ai/

27 https://www.whitehouse.gov/wp-content/uploads/2020/02/American-AI-Initiative-One-Year-Annual-Report.pdf

28 https://www.congress.gov/bill/115th-congress/house-bill/6227

29 https://www.whitehouse.gov/wp-content/uploads/2018/09/National-Strategic-Overview-for-Quantum-Information-Science.pdf

The NQI Act was signed into law "to accelerate quantum R&D for the economic and national security of the United States." The Act authorizes NSF, DOE, and NIST to strengthen QIS programs, centers, and consortia. The Act also calls for a coordinated approach to QIS R&D efforts across the Federal Government, including the civilian, defense, and intelligence sectors. To guide these actions, the NQI Act legislates responsibilities for the NSTC Subcommittee on QIS, the NQCO, and the National Quantum Initiative Advisory Committee. Recognizing that QIS technologies have commercial and defense applications, additional authorization for QIS R&D is legislated by the National Defense Authorization Act.[30] Civilian, defense, and intelligence agencies all have a long history of investments in QIS and have a critical stake in future QIS discoveries and technology development. NQI now provides an overarching framework to strengthen and coordinate Federal QIS R&D activities across departments and agencies, with industry, and with the academic community.

The National Strategic Overview for QIS provides recommendations to strengthen the U.S. approach to QIS R&D, focusing on six areas: prioritizing fundamental science research, growing the QIS workforce, deepening engagements with quantum industry, providing critical infrastructure, maintaining security, and encouraging international cooperation. These policy areas build on earlier interagency QIS policy discussed in the 2016 NSTC Report, *Advancing Quantum Information Science*,[31] and the 2009 NSTC Report, *A Federal Vision for Quantum Information Science*.[32] In recent years, the United States has greatly prioritized investment in QIS across the board. The NQI Act authorized $1.2 billion over five years for quantum activities, and the Administration is on track to surpass that. In February 2020, the President's FY2021 budget request included a call to double Federal QIS spending over two years.

30 https://www.congress.gov/bill/115th-congress/house-bill/5515/text

31 https://www.whitehouse.gov/sites/whitehouse.gov/files/images/Quantum_Info_Sci_Report_2016_07_22%20final.pdf

32 https://science.osti.gov/-/media/_/pdf/initiatives/qis/FederalVisionQIS.pdf?la=en&hash=31D3A8C29757D3E71177315516EE77A6CDB18BD5