

The National Strategic Computing Initiative

Saul Gonzalez Martirena
National Science Foundation

April 13, 2016



Executive Order 13702

July 29, 2015

EXECUTIVE ORDER

CREATING A NATIONAL STRATEGIC COMPUTING INITIATIVE

By the authority vested in me as President by the Constitution and the laws of the United States of America, and to maximize benefits of high-performance computing (HPC) research, development, and deployment, it is hereby ordered as follows:

The NSCI is a whole-of-government effort designed to create a cohesive, multi-agency strategic vision and Federal investment strategy, executed in collaboration with industry and academia, to maximize the benefits of HPC for the United States.



How did we get here?

- ***Initiatives and Investments in NIT R&D to Achieve America's Priorities*** (2010 PCAST Recommendations)
 - *“Recommendation 7-10: NSF, DARPA, and DoE should invest in a **coordinated program of basic research on architectures, algorithms, and software for next generation HPC systems**. Such research should not be limited to the acceleration of traditional applications, but should include work on systems capable of (a) **efficiently analyzing vast quantities of both numerical and non-numerical data**, (b) handling problems requiring real-time response, and (c) accelerating new applications...”*
- ***Initiatives and Investments in NIT R&D to Achieve America's Priorities and Advance Key NIT Research Frontiers*** (2012 PCAST Recommendations)
 - *“Recommendation 8: NSTC should lead an effort by NSF, DoE, DOD, member agencies of the Intelligence Community, and other relevant Federal agencies to **design and implement a joint initiative for long-term, basic research** aimed at developing fundamentally new approaches to high performance computing.”*



NSCI Intent



- **National**
 - “Whole-of-government” and “whole-of-Nation” approach
 - Public/private partnership with industry and academia
- **Strategic**
 - Leverage beyond individual programs (a key “platform” technology)
 - Long time horizon (decade or more)
 - Deploy broadly – lower barrier to entry; bring new communities
- **Computing**
 - HPC = most advanced, capable computing technology available in a given era
 - Multiple styles of computing and all necessary infrastructure
 - Scope includes everything necessary for a fully integrated capability
 - Theory and practice, software and hardware
- **Initiative**
 - Above baseline effort
 - Link and lift efforts

Enhance U.S. strategic advantage in HPC for economic competitiveness, scientific discovery, and national security



NSCI Policy Principles

It is the policy of the United States to sustain and enhance its scientific, technological, and economic leadership position in HPC research, development, and deployment through a coordinated Federal strategy guided by four principles:

- (1) The United States must deploy and apply new HPC technologies broadly for economic competitiveness and scientific discovery.
- (2) The United States must foster public-private collaboration, relying on the respective strengths of government, industry, and academia to maximize the benefits of HPC.
- (3) The United States must adopt a "whole-of government" approach that draws upon the strengths of and seek cooperation among all Federal departments and agencies with significant expertise or equities in HPC in concert with industry.
- (4) The United States must develop a comprehensive technical and scientific approach to efficiently transition HPC research on hardware, system software, development tools, and applications into development and, ultimately, operations.

This order establishes the NSCI to implement this whole-of-government strategy, in collaboration with industry and academia, for HPC research, development, and deployment.



Strategic Objectives

- (1) Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.
- (2) Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.
- (3) Establishing, over the next 15 years, a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached (the "post-Moore's Law era").
- (4) Increasing the capacity and capability of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development.
- (5) Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the United States Government and industrial and academic sectors.



NSCI "Objective 2"


Desire for Convergence

Data Intensity (Petabytes)



Internet-Scale Computing

Sophisticated data analysis
E.g., deep learning



Mixing simulation
with real-world data

Real-time analysis of
simulation results



Modeling &
Simulation-Driven
Science &
Engineering



Saul Gonzalez Martirena / NSCI


Computational Intensity (Petaflops)



Future Hardware Technology Challenges

NSCI “Objective 3”

- **Extending CMOS**
 - Serious problems as transistor lengths approach 5nm (Feature size 1.5nm)
 - International Technology Roadmap for Semiconductors predicts 2028
- **Other Possibilities**
 - Memristors, carbon nanotubes, cryogenic operation, quantum computing
 - All show promise, but nowhere near commercial deployment
 - HPC technology must fit within much larger context of consumer products
- **Other Impacts**
 - New architectures
 - New programming models based on localized data
 - Need for fault tolerance at higher hardware and software layers
 - New algorithms, languages, software



Federal Agency Roles

The order identifies [lead](#) agencies, [foundational](#) R&D agencies, and [deployment](#) agencies:

- [Lead](#) agencies (DOE, DOD, NSF) are charged with developing and delivering the next generation of integrated HPC capability and will engage in mutually supportive research and development in hardware, software, and workforce to support the objectives of the NSCI.
- [Foundational](#) R&D agencies (IARPA, NIST) are charged with fundamental scientific discovery work and associated advances in engineering necessary to support the NSCI objectives.
- [Deployment](#) agencies (NASA, FBI, NIH, DHS, NOAA) will develop mission-based HPC requirements to influence the early stages of design of new HPC systems and will seek viewpoints from the private sector and academia on target HPC requirements.



Lead Agency Responsibilities

- The DOE Office of Science and DOE National Nuclear Security Administration will execute a joint program focused on advanced simulation through a capable exascale computing program emphasizing sustained performance on mission relevant applications and on analytic computing to support its missions and post-Moore's Law HPC capability.
- NSF will play a central role in scientific discovery advances, the broader HPC ecosystem for scientific discovery, and workforce development.
- DOD will focus on data analytic computing to support its mission.

These responsibilities leverage the historical roles each of the lead agencies have played in pushing the frontiers of high-performance computing, and will keep the nation on the forefront of this strategically important field.

The lead agencies will also work with the foundational R&D agencies and the deployment agencies to support the objectives of the NSCI to address the wide variety of needs across the Federal Government.



Foundational R&D Agency Responsibilities

- The Intelligence Advanced Research Projects Activity (IARPA) will focus on future computing paradigms offering an alternative to standard semiconductor computing technologies.
- NIST will focus on measurement science to support future computing technologies.

The foundational research and development agencies will coordinate with deployment agencies to enable effective transition of research and development efforts that support the wide variety of requirements across the Federal Government.



Deployment Agency Responsibilities

There are five deployment agencies for the NSCI:

- National Aeronautics and Space Administration
- Federal Bureau of Investigation
- National Institutes of Health
- Department of Homeland Security
- National Oceanic and Atmospheric Administration

These agencies may participate in the co-design process to integrate the special requirements of their respective missions and influence the early stages of design of new HPC systems, software, and applications.

Agencies will also have the opportunity to participate in testing, supporting workforce development activities, and ensuring effective deployment within their mission contexts.



NSCI Timeline

- July 29, 2015 – Executive 13702 Order Issued
 - “To ensure accountability for and coordination of research, development, and deployment...there is established an NSCI Executive Council”
 - “The Executive Council shall, within 90 days of this Order, establish an implementation plan...”
- July 30, 2015 – NSCI Roundtable (Academia, Private Sector, Government) at the White House
- Inaugural meeting of the NSCI Executive Council
- September 15, 2015 – RFI on Science Drivers for Capable Exascale issued
- October 20-21, 2015 – White House NSCI Workshop (Academia, Private Sector, Government)
- October 27, 2015 – Executive Council delivers Implementation Plan
- February 9, 2016 – President’s FY 2017 budget
- 2nd Executive Council Meeting



NSCI Exascale Request for Information

RFI NOT-GM-15-122

- Specific scientific and research challenges that need 100-fold increase in performance
- Potential impact of the research to the scientific community, national economy, and society.
- Specific limitations/barriers HPC systems must overcome
- Related research areas that would benefit from this level of augmented computational capability.
- Important computational and technical parameters of computational problems in 10 years
- Alternative models of deployment and resource accessibility for exascale computing
- Capabilities needed by the end-to-end system, including data analytics and visualization tools, shared data capabilities, and data services
- Other issues, e.g., training, workforce development and collaboration environments.



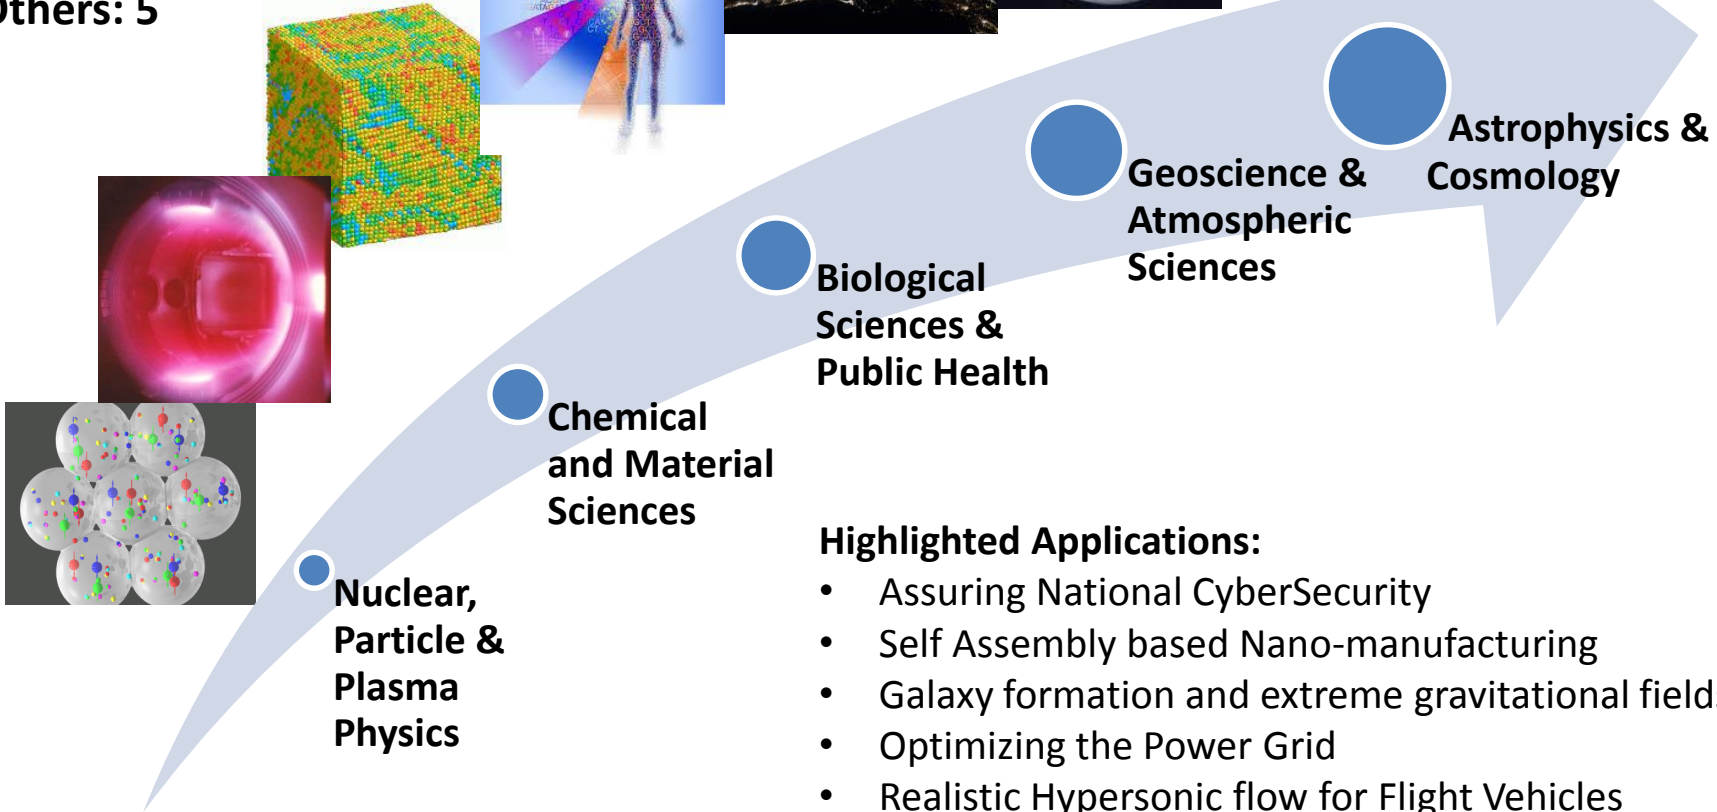
DOE National Laboratories Responses: 135

Academic Responses: 94

Industry Responses: 8

Foreign Responses: 2

Others: 5



Highlighted Applications:

- Assuring National CyberSecurity
- Self Assembly based Nano-manufacturing
- Galaxy formation and extreme gravitational fields
- Optimizing the Power Grid
- Realistic Hypersonic flow for Flight Vehicles
- Regional Scale Seismic predictions
- High Resolution Atmospheric & Climate Models
- N-by-N Comparison of All Patients in US





NSCI White House Workshop

A *White House Workshop on the National Strategic Computing Initiative* was held October 20-21, 2015, with around 250 participants representing industry, government, academia, and other organizations. During the workshop, many individual opinions were expressed. Several ideas emerged that are informing the NSCI implementation:

- The evolutionary path for HPC is more uncertain than during the previous decades. There will be different ways of coupling simulation with data analytics. There is reason to be optimistic for convergence of analytics and HPC in the long term – but diversity of approach is key in the short term. Potential to advance discovery science, engineering.
- A number of hardware technology and architectural innovations (neuromorphic, superconducting, quantum...) will be attempted to overcome physical limitations for charge-based CMOS. However, more can still be squeezed out of CMOS (migrating app function to hardware, optimizing data movement). NSCI must accommodate this breadth of choice and avoid any premature down select of technology.
- NSCI must support the exploration of different deployment models. Although existing clouds may not have the performance required to satisfy all of the most demanding HPC applications, they have already proved suitable for a number of scientific workloads.
- Deeper engagement with the industrial (non-computing) sector will be key to achieve broad deployment and advance the Nation's economic competitiveness.
- NSCI must support efforts to increase productivity of HPC software development process and to broaden the HPC workforce.



NSCI Governance: Executive Council

- a) Lead by the White House: *“...to be co-chaired by the Director of the Office of Science and Technology Policy (OSTP) and the Director of the Office of Management and Budget (OMB). The Director of OSTP shall designate members of the Executive Council from within the executive branch.”*
- b) Use new and existing Federal coordination mechanisms: *“...shall coordinate and collaborate with the National Science and Technology Council ... to ensure that HPC efforts across the Federal Government are aligned with the NSCI ... The Executive Council may create additional task forces as needed to ensure accountability and coordination.”*
- c) Accountability: *“The Executive Council shall meet regularly to assess the status of efforts to implement this order. The Executive Council shall meet no less often than twice yearly in the first year after issuance of this order.”*
- d) Public-private collaboration: *“...will encourage agencies to collaborate with the private sector as appropriate ... may seek advice from the President's Council of Advisors on Science and Technology ... and may interact with other private sector groups consistent with the Federal Advisory Committee Act.”*

The Executive Council has met two times



NSCI Implementation

- a) The Executive Council shall, within 90 days of the date of this order,* establish an implementation plan to support and align efforts across agencies in support of the NSCI objectives. Annually thereafter for 5 years, the Executive Council shall update the implementation plan as required and document the progress made in implementing the plan, engaging with the private sector, and taking actions to implement this order. After 5 years, updates to the implementation plan may be requested at the discretion of the Co-Chairs.
- b) The Co-Chairs shall prepare a report each year until 5 years from the date of this order on the status of the NSCI for the President. After 5 years, reports may be prepared at the discretion of the Co-Chairs.

*First internal implementation plan delivered October 26, 2015 and informed the FY 2017 President's Budget Request.



NSCI in the FY 2017 President's Budget

The 2017 President's Budget Request was sent to Congress on February 9, 2016. The total funding *explicitly* allocated to the NSCI was \$318 million.

- DOE National Nuclear Security Administration, (\$95 million)



- *Advanced Simulation and Computing*: Activities and research leading to deployment of exascale capability for national security applications in the early 2020's

- DOE Office of Science (\$190 million)



- *Advanced Scientific Computing Research* (\$154 million) R&D and design to ultimately achieve capable exascale systems with 1000x performance of current HPC
- *Basic Energy Sciences* (\$26 million): basic research resulting in codes to predictively design functional materials and chemical processes
- *Biological and Environmental Research* (\$10 million): develop science base for increasingly complex climate modeling and data analytic applications

- National Science Foundation (\$33 million):

- *Multi-directorate*: Scientific discovery for HPC, HPC for scientific discovery, the broader HPC ecosystem, and workforce development

Above *in addition to* existing core programs aligned with NSCI goals



FY 2017 Highlight: Exascale at DOE

- A partnership between the DOE Office of Science and the NNSA
- “All-in” approach: hardware, software, applications, large data, underpinning applied math and computer science
- Supports DOE’s missions in national security and science:
 - Stockpile stewardship – support annual assessment cycle
 - Discovery science – next-generation materials
 - Mission-focused basic science in energy – next-generation climate software
- The next generation of advancements will require Extreme Scale Computing
 - 100-1,000X capabilities of today’s computers with a similar physical size and power footprint
 - Significant challenges are power consumption, high parallelism, reliability
- Extreme Scale Computing, cannot be achieved by a “business-as-usual,” evolutionary approach
 - Initiate partnerships with U.S. computer vendors to perform the required engineering, research and development for system architectures for capable exascale computing
 - Exascale systems will be based on marketable technology – Not a “one off” system
 - Productive system – Usable by scientists and engineers



Exascale at DOE / continued

- Exascale Computing Initiative (ECI)
 - The ECI was initiated in FY 2016 to support research, development, and computer-system procurements to deliver an exascale (10^{18} ops/sec) computing capability by the mid-2020s.
 - The Exascale Crosscut includes primary investments by SC/ASCR and NNSA/ASC and software application developments in both SC (BES and BER) and NNSA.
- Exascale Computing Project (ECP)
 - In FY 2017, the ASCR ECI funding will be transitioned to the DOE Exascale Computing project (the ECP), consistent with Objective 1 of the NSCI
 - An ECP Project Office has been established at Oak Ridge National Laboratory to manage delivery of the exascale computers according to the principles of DOE Order 413.3b

Exascale Target System Characteristics:

- 20 pJ per average operation
- Billion-way concurrency
- High reliability and resilience
- Programming environments that increase productivity, reduce barrier to entry
- Ecosystem to support new application development, collaboration, portability, legacy apps



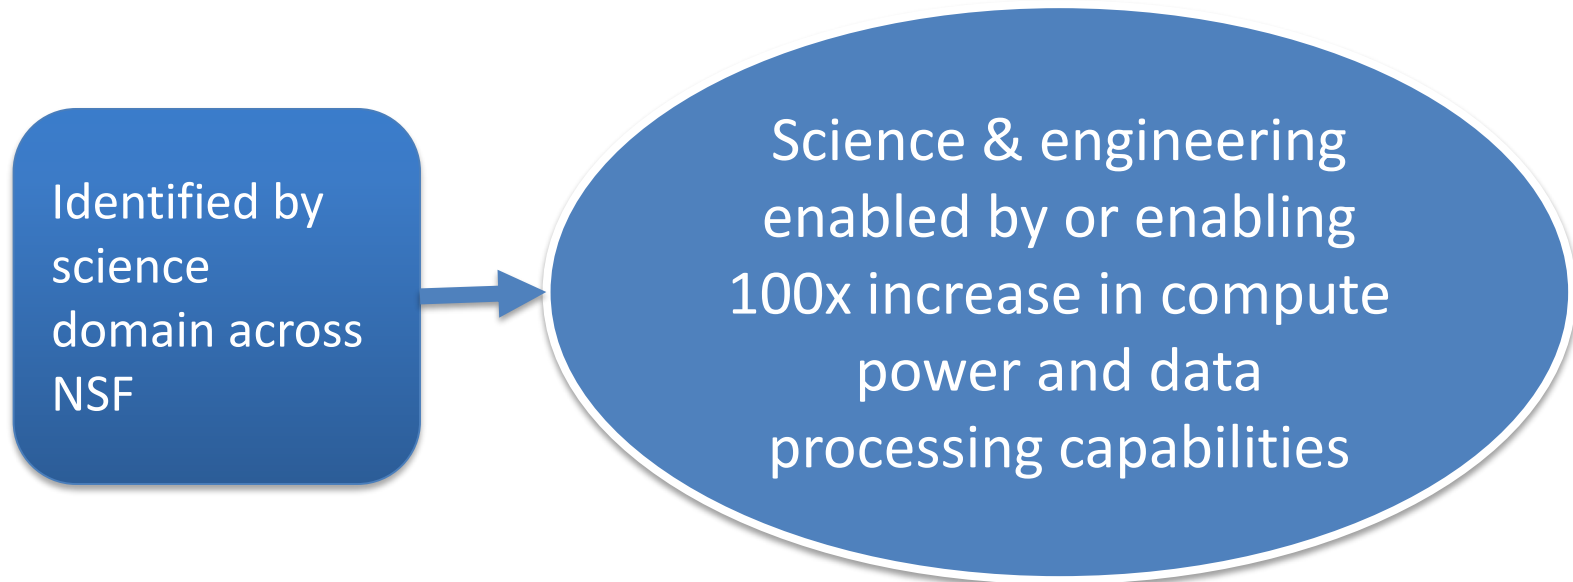
U.S. DEPARTMENT OF
ENERGY

Office of
Science



FY 2017 Highlight: NSCI at NSF

Driven by Science Frontiers



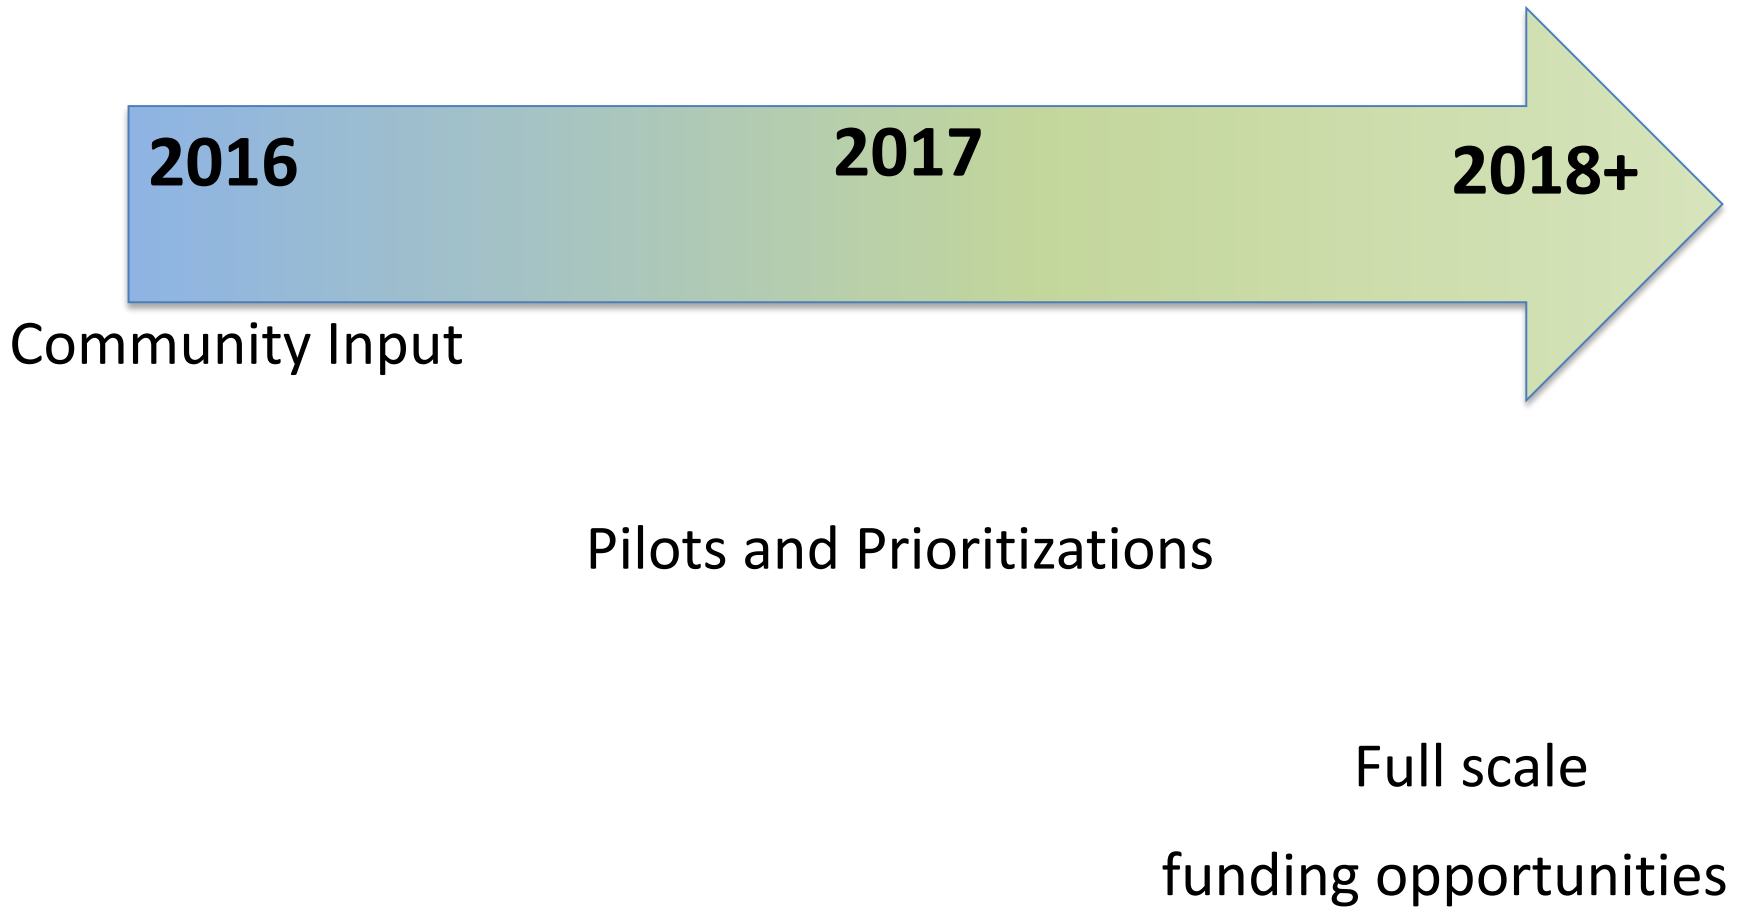
Example of science frontiers:

- Mapping the human brain with synaptic resolution
- Optimal decision making in complex high-dimensional problems
- Understanding multi-phase turbulent flow
- Device and systems technology enabling new compute paradigms
- Extreme computing and data analysis for experimental and observational facilities
- Analyzing social behavior from large-scale data collections
- Understanding high complexity in biological systems

Saul Gonzalez Martirena / NSCI



NSF Timeline for NSCI



Examples of NSF Activities in 2016

- Material Genome Initiative workshops have already taken place, and many of them are relevant to NSCI
- Workshop on topological insulators (<http://scholar.princeton.edu/nsfcmp/home>)
- Workshop on uncertainty quantification for data assimilation at Penn State (<http://adapt.psu.edu/2016EnKFWorkshop>)
- Workshop on “Intelligent Cognitive Assistants” in collaboration with SIA and SRC (May 12-13, 2016, San Jose, CA)
- National Academy of Science Study on HPC
- RFI with DOE and NIH on HPC
- White House NSCI workshop
- Joint NSF (CISE/ENG) solicitation with SRC on “Energy-efficient computing: from Devices to Architectures” (NSF 15-526)




NSF Investment Areas for 2017

- Low-power computing and future HPC systems in the post-Moore Law device and hardware systems era (28%)
- HPC algorithms and architectures for massive concurrency, energy-efficient computing and system resilience at extreme scale (13%)
- Novel scientific software architectures that are resilient, re-usable, and enduring yet agile (22%)
- Provision of shared infrastructure, workforce development (28%)
- Analyze large complex data sets and assimilate real-time data into models and forecasts (9%)



Pilots - Initial, exploratory projects into NSCI-relevant topics
Examples: Software, “Beyond Moore’s Law,” dynamic data workflows, hardware, industry engagement



Enhance existing programs and new programs based on experiences with pilots



Emphasis on cross-disciplinary

Saul Gonzalez Martirena / NSCI



Related Administration Priorities – an NSCI Ecosystem

- Materials Genome Initiative
- Advanced Manufacturing Initiative
- Precision Medicine Initiative
- The National Big Data R&D Initiative
- National Photonics Initiative
- The National Nanotechnology Initiative
- The BRAIN Initiative
- Grand challenge: *“Create a new type of computer that can proactively interpret and learn from data, solve unfamiliar problems using what it has learned, and operate with the energy efficiency of the human brain.”*



Outlook: A Lasting Initiative

- We have achieved a high level of coordination and collaboration across Federal agencies
 - It took more than two years to get here
 - There is a concrete plan to implement the NSCI – and it will evolve
- We are in the process of establishing similar level of coordination and collaboration with academia and private sector
 - Have held one industry roundtable, and one NSCI workshop, in addition to numerous NSCI-related workshops and meetings
 - A Request for Information closed on November 13
 - Have asked industry and academia to help us shape and promote follow-on activities
- The 2017 President's budget is the beginning of the implementation of a long-term strategy for High Performance Computing
 - Sustaining the NSCI across administration boundaries will require demonstrated progress and a continued Whole-of-Nation commitment to leadership in HPC



Thank You



And Thank You to my NSCI colleagues at the National Science Foundation,
the U.S. Department of Energy, the Department of Defense, and the
White House Office of Science and Technology Policy