Transforming Science Through Cyberinfrastructure: Coordination Services

NSF's Blueprint for National Cyberinfrastructure Coordination Services for Accelerating Science and Engineering in the 21st Century

Executive Summary

The national research cyberinfrastructure (CI) has become critical to computational and data intensive research across all of science and engineering (S&E) in the 21st century. The National Science Foundation (NSF) recently shared a vision, developed by its Office of Advanced Cyberinfrastructure (OAC), for ***an agile, integrated, robust, trustworthy and sustainable CI ecosystem that drives new thinking and transformative discoveries in all areas of science and engineering (S&E) research and education***. The envisioned CI ecosystem integrates advanced CI resources, services and expertise towards collectively enabling new, transformative discoveries across all of S&E. This document is the second in this series of blueprint documents that outline NSF's plan for realizing this vision and presents a forward-looking blueprint for national CI coordination services. These crosscutting services support critical aspects such as allocation, measurement and user support that are essential for the effective and efficient use of the CI ecosystem and for ensuring overall user productivity.

Table of Contents

1 Introduction

The national research cyberinfrastructure (CI) has become critical to computational and data intensive research across all of science and engineering (S&E) in the 21st century. It is a key catalyst for discovery and innovation and plays a critical role in ensuring US leadership in S&E, economic competitiveness and national security, consistent with the National Science Foundation's (NSF's) mission. NSF, through the Office of Advanced Cyberinfrastructure (OAC), has shared a vision[1] that calls for the broad availability and innovative use of ***an agile, integrated, robust, trustworthy and sustainable CI ecosystem*** that can ***drive new thinking and transformative discoveries in all areas of S&E research and education***.

NSF's vision for an agile, integrated, robust, trustworthy and sustainable CI ecosystem that drives new thinking and transformative discoveries in all areas of S&E research and education

- *View CI more holistically: CI continuum seamlessly integrating a spectrum of resources, tools, services, and expertise to enable transformative discoveries.*
- *Support translational research: Core innovations → development of community tools and frameworks → deployment and operation of sustainable production CI.*
- *Balance innovation with stability: Ensure continuity in production computational capacity while fostering innovation and transition to production.*
- *Couple discovery and CI innovation cycles : Rapidly address new challenges and opportunities in an era of disruptive technologies and evolving science needs.*
- *Improve usability: Ease pathways for discovering, accessing, understanding and using powerful CI capabilities and services to enhance researcher productivity and scientific impact.*

This document is the second in this series of blueprint documents that outline NSF's plan for realizing this vision. It presents a forward-looking blueprint for national CI coordination services. The crosscutting coordination services discussed in this document support critical aspects such as allocation, measurement and user support that are essential for the effective and efficient use of the CI ecosystem and for ensuring overall user productivity in the context of rapidly changing application and CI landscapes. This blueprint is informed by the community through advisory bodies, requests for information (RFIs), workshops and conferences, and national initiatives (including those listed in the vision document), and in particular, the ***2019 NSF-funded National Cyberinfrastructure Coordination Service Conference***[2] that specifically focused on the CI service ecosystem. This workshop resulted in the following observations:

- **Transforming Science:** The growth in CI capabilities is significantly transforming all aspects of our lives, and we are rapidly approaching a new inflection point on how research and education are being conducted.

[1] "Transforming Science Through Cyberinfrastructure: NSF's Blueprint for a National Cyberinfrastructure Ecosystem for Science and Engineering in the 21st Century," https://www.nsf.gov/cise/oac/vision/blueprint-2019/.

[2] "National Cyberinfrastructure Coordination Service Conference," https://www.rti.org/publication/national-cyberinfrastructure-coordination-service-conference.

- **It's an Ecosystem:** The core of CI is no longer just about HPC or data writ large. The evolving CI is an ecosystem that includes people, software, networks, data, and compute capabilities.
- **New Domains Served:** The CI ecosystem underpins today almost all scientific domains and frequently emerges as the pacesetter for advances in these domains.
- **Spanning all Campuses:** This CI Ecosystem is becoming a fundamental requirement and capability for all campuses and has evolved to become part and parcel of the university structure supporting education and research for undergraduates, graduates, and faculty.
- **Integrative Services:** The rapidly evolving CI Ecosystem requires new approaches to offering effective services to sciences that covers all the CI components.
- **Service and Resource Allocation:** The emerging CI Ecosystem requires new approaches on how to allocate CI services, which includes computer, data, support, and other resources.

1.1 NSF's Current CI Services Landscape

CI coordination services are integral to the NSF's CI vision and ecosystem. They enable researchers across every field of NSF-supported science and engineering to effectively and efficiently use the CI ecosystem and ensure overall user productivity. These services:

- Provide key capabilities including the discovery of and access to available resources; access to relevant and timely expert technical guidance and focused instruction on effective uses of these cyberinfrastructure capabilities; and support for developmental efforts to support the evolving portfolio of NSF supported science.
- Ensure the effective management, operation, monitoring and measurement, and overall use of computing resources, and integrating these resources into a coherent, coordinated national cyberinfrastructure ecosystem.
- Increase user accessibility; enable collaboration; simplify use of CI in dynamic, system-of-systems scenarios; support access to relevant data; and enable timely access to novel technologies and solutions.

Current NSF-supported CI coordination services landscape are comprised of three major tranches of activity: Extreme Digital (XD), Open Science Grid (OSG), and monitoring and measurement services. Each of these is explained in more detail in Section 3.2.1.3 of the initial vision document, *Transforming Science Through Cyberinfrastructure: NSF's Blueprint for a National Cyberinfrastructure Ecosystem for Science and Engineering in the 21st Century*[1]. Brief summaries are provided in the following paragraphs.

Extreme Digital (XD): XD provides crosscutting coordination, allocation, measurement, and user support services that add value to the Innovative HPC investments. The XD program's shared services model coherently and efficiently delivers researchers both access to and support for diverse, dynamic, and distributed resources and enables the connection between individual campuses and national resources. These shared services include allocation of resources to computational and data research projects; advanced user assistance; training, education, and outreach; architecture and operation of an integrated digital services infrastructure; metrics services; and overall coordination. Two awards are currently active within the XD program: XD

Metrics Service (XMS)[3] and the eXtreme Science and Engineering Discovery Environment (XSEDE)[4].

Open Science Grid (OSG): OSG[5] is jointly funded by NSF (by OAC and the Directorate of Mathematical and Physical Sciences (MPS)) and the Department of Energy and facilitates access to distributed high throughput computing for research in the U.S. It provides common service and support for resource providers and scientific institutions using a distributed fabric of high throughput computational services. OSG provides software and services to users and resource providers alike to enable the opportunistic usage and sharing of resources. It is primarily used as a high-throughput grid where scientific problems are solved by breaking them down into a very large number of individual jobs that can run independently. The resources accessible through the OSG are contributed by the community, organized by the OSG, and governed by the OSG consortium. OSG consists of computing and storage elements at over 100 individual sites spanning the United States. These sites, primarily at universities and national labs, range in size from a few hundred to tens of thousands of CPU cores. During 2018, OSG provided more than 1.2 billion CPU hours to researchers across a wide variety of projects. OSG services are funded through 2019.

Monitoring and measurements services: Deep instrumentation, monitoring, measurement, and reporting across all layers of the systems and services making up the CI ecosystem are essential to providing the situational awareness necessary for achieving desired levels of efficiency, understanding, autonomous operations, robustness, and performance. Complementing the services for monitoring usage of computational resources (i.e., XSEDE's XDMoD[3]) listed above, OAC has invested in a portfolio of award activities in monitoring and measurement research and development over the years, including networking-based monitoring and measurement (e.g., CAIDA[6], Route Views[7], NetSage[8] and perfSONAR[9]).

[3] XD Metrics Service (XMS), https://open.xdmod.org/7.5/index.html.
[4] eXtreme Science and Engineering Discovery Environment (XSEDE), https://www.xsede.org/.
[5] Open Science Grid (OSG), https://opensciencegrid.org/.
[6] "Center for Applied Internet Data Analysis," http://www.caida.org/home/.
[7] "University of Oregon Route Views Project," http://www.routeviews.org/routeviews/.
[8] "NetSage," http://www.netsage.global/.
[9] "perfSONAR," https://www.perfsonar.net/.

2 A Blueprint for National CI Coordination Services for Accelerating Science and Engineering in the 21st Century

This section outlines NSF's forward-looking blueprint for national coordination CI services and outlines a plan for translating this blueprint into action. The fabric of envisioned CI coordination services and is composed of five interrelated services (illustrated in Figure 1): Allocations Services, Operations Services, User Engagement Services, Monitoring and Measurement Services, and Technology Translation Services. NSF aims to conceptualize and implement this fabric of interrelated services that can effectively support a broad and diverse set of requirements, users, and usage modes from all areas of S&E research and education. It is also essential that this fabric of services is agile and can evolve and adapt to respond to emerging requirements and technology landscapes. Note that the CI services are complemented by the other components of the CI ecosystem as presented in NSF's CI vision[Error! Bookmark not defined.]. The five services are described in more detail in the following sections.

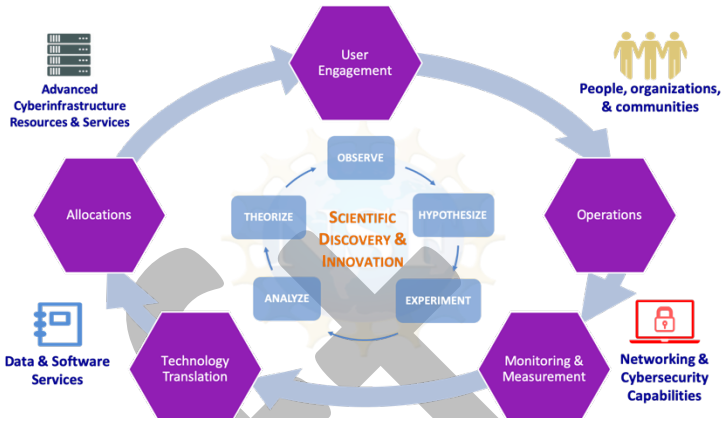

Figure 1. Overview of the National CI Coordination Services Ecosystem

2.1 Allocation Services

Allocation services are responsible for allocating the range NSF-funded CI resources to the user community to support their S&E research and education needs. However, NSF also expects that new allocation approaches that go beyond the current practices targeting more traditional batch-based computational resources and are better integrated with the S&E research goals of NSF-funded projects as they emerge. Specifically, NSF expects allocation services will evolve to respond to emerging technologies, resources classes and usage modes (e.g., cloud services, federated resources, testbeds and prototype systems, etc.), new and diverse application classes (machine learning based, data-driven, online/real-time, elastic, long-running, high-throughput, etc.), and to enable the potential allocation other elements of the CI ecosystem (e.g., network resources, storage, etc.) that are relevant to end-to-end application workflows.

2.2 Operations Services

Operations services provide essential integrative and crosscutting functions to support the coordinated, secure, robust operations of NSF-funded resources that are part of the NSF CI ecosystem. The NSF CI ecosystem is composed of a wide variety of independent geographically distributed resource providers, each provisioning unique capabilities, services, and expertise to the wider S&E research community. NSF expects that this service will provide coordinating functions to the resource providers to enable the different elements of the CI ecosystem to work together effectively and securely. Some of the principal goals of the coordinating activities will

be to enhance the cybersecurity of the ecosystem, enable data-sharing and information dissemination where appropriate, as well as provide technical support and best practices to assist in resource provider integration and operations.

In this model, operations services are expected to have clearly defined relationships and divisions of responsibilities with resource providers. Furthermore, operations services are also expected to facilitate and support synergies and collaborations and to ensure that the resource providers collectively support the needs of the broad S&E research and education community using NSF-supported CI.

2.3 User Engagement Services

User engagement services are coordinated services and activities that ensure a high-quality productive experience of prospective and current users when engaging with the NSF CI ecosystem at any stage. The array of services targets (a) end users with planned, pending, current, or recent allocations on the NSF-supported resources, or who are otherwise utilizing associated services such as for data and software and (b) the wider community of prospective and current end users whose computational science interests intersect with the capabilities provided by the national CI ecosystem. End users may be individuals, groups or organizations such as Science Gateways and others. Key services include: *general informational assistance*, primarily targeted to the pre-allocation stage; *allocation and utilization assistance*, assisting and advising end users with their allocations and utilization strategy; *end user training* focused on how to access, utilize and successfully accomplish science on the computational ecosystem; and *computational science advisory network* focusing organizing and coordinating campus, regional and national expertise into a strong community of experts.

User engagement services will assimilate and coordinate the human capital that is separately funded by NSF at the national, regional and campus levels, and engage with existing community organizations and structures, to collectively ensure the effective and efficient use of the NSF support CI ecosystem and promote high user productivity and satisfaction. Note that in this model, user engagement services do not include general CI community learning and workforce development; these are supported separately through OAC's Learning and Workforce Development (LWD) programmatics.

2.4 Monitoring and Measurements Services

Monitoring and measurement services will provide an integrated and open data analytics platform to support optimal performance and usage of NSF-funded resources, as well as facilitate timely decision making for a broad range of stakeholders. The service is expected to provide an extensible framework to support the assimilation of instrumentation data from the broad range and diversity of resource providers constituting the NSF CI ecosystem. Some of the functions provided by this service will include customizable multi-view reporting of operational and higher-level impact metrics across time, as well as other dimensions, to garner insights that would not be available otherwise to S&E researchers, CI software developers, resource administrators, center directors, and other stakeholders using the NSF CI ecosystem. Machine-readable interfaces should also be provided to enable automation of S&E workflows and the incorporation of forward-looking predictive data analytics techniques such as machine learning.

NSF expects that the monitoring and measurement services will leverage a broad range of current capabilities, as well as integrate/interoperate with monitoring and measurement services and instrument data across all aspects and layers of the CI ecosystem including computation, network and other resources provisioned by other CI projects.

2.5 Technology Translation Services

Technology translation services are responsible for providing the processes and mechanisms for transitioning innovative CI solutions developed by the CI community, possibly through awards from other OAC (and CISE) programs, such as OAC Core Research[10], Cyberinfrastructure for Sustained Scientific Innovation (CSSI)[11], Cybersecurity Innovation for Cyberinfrastructure (CICI)[12], Campus Cyberinfrastructure (CC*)[13], etc., to production operations as part of NSF's production CI ecosystem. Specifically, providers of technology translation services will develop, document, and execute processes for determining (using established criteria and metrics, and with community inputs) which solutions to move to production and for how long. Furthermore, they will work with resource providers to sustain these solutions in production operations so that CI innovations can be leveraged by user community for S&E research and education. The technology translation service providers are also expected to maintain well defined and documented interfaces with CI research and developer community as well as the resources providers to ensure the translation of CI innovations to the production and their potential impact on S&E.

2.6 Putting the Plan into Action: Programs, Projects and Opportunities

NSF plans to move ahead aggressively to put this blueprint for national CI coordination services into action through programs and projects to ensure continuity of operations while also ensuring that the resulting fabric of services are responsive to current and emerging S&E needs. Furthermore, and as illustrated in Figure 2, NSF is committed to ensuring the availability of such CI coordination services (contingent as always on budgets and national priorities) over the lifetimes of the NSF-funded computational ecosystem to ensure its effective and efficient use.

Figure 2. Notional execution timeline for the national CI coordination services blueprint.

[10] "Office of Advanced Cyberinfrastructure (OAC): Research Core Program," https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=505571.
[11] "Cyberinfrastructure for Sustained Scientific Innovation (CSSI)," https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=505505.
[12] "Cybersecurity Innovation for Cyberinfrastructure (CICI)," https://www.nsf.gov/pubs/2019/nsf19514/nsf19514.htm.
[13] "Campus Cyberinfrastructure (CC*)," https://www.nsf.gov/pubs/2019/nsf19533/nsf19533.htm.

3 Ongoing Strategic Planning and Community Engagement

NSF will continue its strategic planning and community engagement activities as it explores other elements of computational ecosystem, including software and data systems and services, networking, cybersecurity, and learning and workforce development. NSF notes that this blueprint is the second of a series of blueprints focused on different elements of the CI ecosystem to be developed in partnership with the community.

Looking to the future, the road ahead for research CI promises to be exciting with many new opportunities. In the near term, NSF Big Ideas are moving into full gear with multiple new solicitations and are complemented by recent relevant national initiatives in areas such as Quantum Information Science (including a National Quantum Initiative Act signed into law in December 2018[14]) and Artificial Intelligence (including an Executive Order on *Maintaining American Leadership in Artificial Intelligence*[15]). NSF's investments, for example, as part of the Quantum Leap Big Idea[16], including the Quantum Computing & Information Science Faculty Fellows (QCIS-FF) program (NSF 19-507)[17], as well as in foundational and translational Artificial Intelligence and Machine Learning[18] research [see, for example, the recently-launched National Artificial Intelligence (AI) Research Institutes program (NSF 20-503)[19]], will help define the nature and structure of the CI ecosystem over the longer term. NSF looks forward to continuing to work with the community to define the future of cyberinfrastructure research and research cyberinfrastructure, with the overarching goal of realizing and integrated CI ecosystem that transforms all of S&E research and education.

4 Conclusion

The NSF-funded CI ecosystem is playing an increasingly critical role across all of S&E research and education, enabling discoveries and driving innovation. It is an important part of the national CI ecosystem that is critical for ensuring US leadership in S&E, economic competitiveness and national security. As a result, it is essential that NSF strategically rethink and evolve this CI ecosystem in response to the changing nature of needs of S&E, driven by the changing technology landscape, and informed by community inputs. This document builds on the NSF's recently articulated vision for a national CI ecosystem that integrates computational, data, software, networking, and security resources, tool and services, and computational and data skills and expertise towards collectively enabling new, transformative discoveries across all of S&E, and presented NSF's blueprint for CI coordination services to support science and engineering research and education in the 21st century. It also outlined a plan to implement this blueprint. The vision and blueprint presented in this document have been informed by the community

[14] "National Quantum Initiative," https://www.congress.gov/115/bills/hr6227/BILLS-115hr6227enr.pdf.
[15] "Executive Order on Maintaining American Leadership in Artificial Intelligence," https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/.
[16] "NSF Quantum Leap Big Idea," https://www.nsf.gov/news/special_reports/big_ideas/quantum.jsp.
[17] "Quantum Computing & Information Science Faculty Fellows (QCIS-FF)," https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=505535.
[18] "Artificial Intelligence (AI) at NSF," https://nsf.gov/cise/ai.jsp.
[19] "National Artificial Intelligence (AI) Research Institutes," https://www.nsf.gov/pubs/2020/nsf20503/nsf20503.htm.

through advisory bodies, requests for information (RFIs), workshops and conferences, and national initiatives. NSF intends to continue to work with the community to evolve and implement the vision and blueprint presented in this document, as well as to develop complementary blueprints for other CI elements.