

NSF/VMware Partnership on the Next Generation of Sustainable Digital Infrastructure (NGSDI)

August 31, 2020 Webinar



Agenda

- **Welcome and CISE Context**

Margaret Martonosi
Assistant Director, CISE

- **NGSDI Program--VMware Perspective**

David Tennenhouse
Chief Research Officer, VMWare

- **NGSDI Research Overview**

Gurdip Singh
Division Director, CNS

- **NGSDI Proposals, Review, & Project Management**

Matt Mutka
Program Director, CNS

- **Questions**

NSF-VMWare Team 

NGSDI: NSF – VMware Partnership

- **NGSDI: Continuing our NSF-VMware Partnership**
 - 2016: Software Defined Infrastructure as a Foundation for Clean-Slate Computing Security (SDI-CSCS)
 - 2018: Edge Computing Data Infrastructure (ECDI)
- **VMware's long-term commitment to Sustainability**
 - VMware's virtualization and resource management resulted in customer server consolidation, reducing power consumption by 120 million MWh and saving 67 million Metric Tons of CO₂ in 2015 alone
- **NGSDI Results: Dedicated to the Public**



NGSDI: Background

- LBNL US Data Center Energy Report found energy consumption in data centers grew 90% annually (2000-2005)
- Since 2005, growth rate has plateaued at 4%, due to server virtualization and hardware improvements
- First generation of innovation
 - increased density of workloads per computing unit via virtualization
 - redesign of computing systems power management
 - redesign of data center electrical and mechanical architectures



NGSDI: Program Goals

- Foster transformative research in fundamental and systemic approaches to bring dramatic increases in environmental sustainability of the Digital Infrastructure leading to practical methodologies and tools
 - The Digital Infrastructure is broadly defined as the totality of software, hardware, and the methods for managing them for efficient computation
- Progress to carbon-neutrality or carbon freedom requires efficiency in entire computation chain



NGSDI: New Considerations

- Metrics and benchmarks for systemic computational efficiency
- Software bloat and inefficiency
- Distributed resource allocation
- Capacity planning and provisioning
- Compute-storage-networking tradeoffs and placement
- Hardware-software co-optimization
- Service level agreement (SLA) trade space
- Renewable energy source optimization
- Hardware lifetime optimization
- Renewable-energy driven workload shifting



NGSDI: Central Question

- What fundamental and systematic approaches in *measurement, design, development* and *management* of the Digital Infrastructure resources and workloads will enable significant progress toward maximizing sustainability of the Digital Infrastructure *with minimal impact on traditional concerns* such as programmer *productivity* and Digital Infrastructure *performance* and *scalability*?



NGSDI: Research Objectives

- Consider fundamental nature of the problem
- Multiple aspects of sustainability
- Multitude of competing goals of Digital Infrastructure Management
 - Economics, performance, efficiency, sustainability



NGSDI: Example Research Vectors

■ Metrics

- E.g. Metrics and Instrumentation, Benchmarks, Service-Level Objectives (SLO)

■ Workload Design and Development

- E.g. DevOps Divide, Full Stack Visibility and Optimization, Migration of Applications to the Cloud

■ Workload and Digital Infrastructure Management

- E.g. IT/OT Divide, Automation, Full Digital Infrastructure Optimization, Power-sensitive execution, Renewable power sources



NGSDI: Possible Techniques

- Passive and active measurements at systems, local and wide area levels
- Metrics aggregation services, statistical analysis and inference
- Modeling and learning techniques to assist automated control, complex resource management and optimization
- Process and system isolation (virtual machines, containers, functions-as-a-service) to enable control at a wide range of parameters and scale to include sustainability objectives
- Technologies for agile development and convergence of development and production environments



NGSDI: Expectations

- Disruptive innovative approaches
- Focus on software-layer sustainability
 - Narrow focus on non-software components and disciplines are **out of scope** (e.g., strictly hardware architectures, water, power, cooling)



NGSDI: Develop Prototypes

Purpose of prototypes

- Explore implementation aspects of designs
- Empirical demonstration of effectiveness of approaches

Prototypes should leverage existing software, tools, frameworks, testbeds if possible



NGSDI: Broader Impact Results

- Strive to achieve broader industry impacts with any foundational results.
- Interest in dissemination results including open source software, production and publication of datasets, activities leading to real-world experimentation, measurements, and deployments.
- When appropriate, proposers should be clear how to navigate broader policy, economic and social considerations.



NGSDI: Solicitation and Review

- Solicitation Requirements
- Review Process
 - Solicitation-Specific Review Criteria
- Award Selection Process
- Management of the Projects
- Q & A



NGSDI: Key Numbers

NSF 20-594

- Proposals due: Nov. 4, 2020
- Approximately 2 project awards
 - Up to \$3,000,000 per project
 - Over 3 years
- NSF funds from FY2021
- Awards Early 2021



NGSDI: Who Can Submit

- Institutes of Higher Education (IHEs)
Universities and two- and four-year
Colleges (including community colleges)
 - See special instructions for
International Branch Campuses of IHEs
 - Sub-awardee requirements: same as
submitting institutions



NGSDI: Proposal Requirements

- Personnel:
 - 1 proposal submission per person as PI, co-PI, or senior personnel in response to this solicitation.
 - Inclusion of each member needs to be justified with respect to the goals of the project
 - Some number of graduate students expected
 - Some number software engineers or programmers may be submitted as needed
- Proposal Sections
- 20-pages for the Project Description



NGSDI: Proposal Requirements

- 1-page Postdoc Mentoring Plan (if applicable)
- 2-page (max) Collaboration Plan (if applicable)
 - Appropriateness of team participants and expertise
 - Role of each team member
 - Management and Coordination mechanisms
 - Interdependencies among tasks
 - Reference of budget lines to support collaboration
- 2-page (max) Data Management Plan
 - See http://www.nsf.gov/cise/cise_dmp.jsp for guidance.
- 0 general letters of support



NGSDI: Intellectual Property

- NSF/VMware Partnership awardees will agree to dedicate to the public all intellectual property resulting from the research funded as part of this program, and further:
- The awardees will offer its software through an open source license under an Apache 2.0 license found at:
 - <http://www.opensource.org/licenses/apache2.0.php> or other similar open source license; in the event the software already contains code licensed under GNU's General Public License (GPL), then the open source shall be through GPL version 3 found at <http://www.gnu.org/licenses/gpl.html>;
- The awardees will submit for publication in openly available literature any results of this research; and
- The awardees will deposit all published manuscripts and juried conference papers in a public access-compliant repository in accordance with the guidelines set forth in NSF's Public Access Policy (see NSF Public Access Frequently Asked Questions at:
 - https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf18041)
no later than 12 months after initial publication.



NGSDI: Review Process

NSF: Panel with ad hoc reviews as appropriate:

- Intellectual Merit & Broader Impacts
- See NSF 20-1; Proposal and Award Policies and Procedures Guide (PAPPG) for more information
- Additional Review Criteria—see next slide
- VMware team members participate as observers

Joint NSF-VMware reverse site visits as needed

Joint NSF-VMware decisions on awards based on NSF Merit Review process



NGSDI: Solicitation-specific Review Criteria

In addition to Intellectual Merit and Broader Impact, the proposal will be evaluated on the degree to which:

- Project pursues a systems perspective and the creation, deployment, and evaluation of demonstrations or prototypes at the component and eventually the system levels;
- Project features a lean, well-integrated team of researchers with expertise area(s) necessary to conduct the proposed work;
- Convincingly frames meaningful system-level sustainability metrics and argues that successful results will have a meaningful impact in terms of those metrics;
- Projects demonstrate concrete plans to impact the broader industry;
- Researchers leverage existing components and infrastructure such as the NSFFutureCloud projects, Chameleon and CloudLab, and NSF-funded CloudBank; if proposing to build a new infrastructure, justification is needed for why the existing infrastructures do not suffice.



NGSDI: Funding Model

Projects will be jointly funded by NSF and VMware through separate NSF and VMware funding instruments.

- NSF awards will be made as grants.
- VMware awards will be made as VMware agreements (Contracts or Grants) through VMware or its Vanguard-managed University Research Fund.
- NSF and VMware will manage their respective awards/agreements in accordance with their own guidelines and regulations.
- Either organization may supplement a project without requiring the other party to provide any additional funds.



NGSDI: Program Management

- NSF and VMware will each designate a Program Director for each NSF/VMware Partnership award who will jointly oversee the execution of the project
- The VMware Program Director may become a member of the NSF/VMware Partnership Project Management Team.
- Annual on-site reviews may be conducted jointly by NSF and VMware.
- Institutions may request site visits to VMware or invite site visits from VMware.
- VMware may invite academic faculty and students to visit VMware and may visit research institutions upon request.



NGSDI: Proposal Summary

- Title: NGSDI: <title>
 - For Collabs: Collaborative Research:NGSDI: <title>
- Project Description: 20 pages
- Supplementary Documents
 - A list of Project Personnel and Partner Institutions
 - Collaboration Plan
 - Data Management Plan
 - Post-Doctoral Mentoring Plan (if applicable)
- Single Copy Documents
 - list of collaborators



NGSDI: Full Proposals

Deadline

5:00 pm submitter's time on Nov. 4, 2020



Questions?

