

Campus Cyberinfrastructure – Network Infrastructure and Engineering (CC-NIE)

Kevin Thompson
NSF Office of CyberInfrastructure
December, 2012



(Post NSFnet) Brief History of NSF Investments in Network Infrastructure

- ❖ vBNS and High Performance Connections Program (HPNC) – 1995-2003
 - National backbone and connections
- ❖ International Networking (IRNC) – 1997 – present
 - Connecting US to the world
- ❖ Experimental Infrastructure Networking (EIN) - 2003
- ❖ “Academic Research Infrastructure Program – Recovery and Reinvestment” - 2009
 - Subset: Optical exchange, regional networking upgrades
- ❖ EPScOR – Research Infrastructure Improvement (RII) – 2011
 - Inter-campus, intra-campus connectivity
- ❖ STCI program (2011 – “100G Connectivity for Data-Intensive Computing at JHU”, Lead PI: Alex Szalay)



ACCI Task Force on Campus Bridging

- ❖ *Strategic Recommendation to the NSF #3: The National Science Foundation should create a new program funding high-speed (currently 10 Gbps) connections from campuses to the nearest landing point for a national network backbone. The design of these connections must include support for dynamic network provisioning services and must be engineered to support rapid movement of large scientific data sets."* - pg. 6, National Science Foundation Advisory Committee for Cyberinfrastructure Task Force on Campus Bridging, Final Report, March 2011
- ❖ www.nsf.gov/od/oci/taskforces/TaskForceReport_CampusBridging.pdf
- ❖ Also see Campus Bridging Technologies Workshop: Data and Networking Issues Workshop Report. G.T. Almes, D. Jent and C.A. Stewart, eds., 2011, <http://hdl.handle.net/2022/13200>



Campus Cyberinfrastructure – Network Infrastructure and Engineering (CC-NIE)

- ❖ NSF 12-541 – solicitation released March 1, 2012
- ❖ http://www.nsf.gov/publications/pub_summ.jsp?WT.z_pims_id=504748&ods_key=nsf12541
- ❖ 1st area: Data Driven Networking Infrastructure for the Campus and Researcher
- ❖ 2nd area: Network Integration and Applied Innovation
- ❖ Proposals were due May 30, 2012



CC-NIE

- ❖ **Estimated Number of Awards:** 10 to 20
- ❖ **Anticipated Funding Amount:**
 - \$12,000,000 to \$15,000,000 will be available for this competition in FY 2012.
 - Data Driven Networking Infrastructure for the Campus and Researcher awards will be supported at up to \$500,000 total for up to 2 years.
 - Network Integration and Applied Innovation awards will be supported at up to \$1,000,000 total for up to 2 years.
- ❖ Proposals may only be submitted by Universities and Colleges - Universities and two- and four-year colleges (including community colleges) accredited in, and having a campus located in the US, acting on behalf of their faculty members. Such organizations also are referred to as academic institutions.



CC-NIE Area#1 - Data Driven Networking Infrastructure for the Campus and Researcher

- ❖ network infrastructure improvements at the campus level
- ❖ network improvements include:
 - network upgrades within a campus network to support a wide range of science data flows
 - re-architecting a campus network to support large science data flows, for example by designing and building a "science DMZ" (see <http://fasterdata.es.net/science-dmz/> for more information on the "science DMZ" approach)
 - campus network upgrades addressing sustainable infrastructure through improvements in energy efficient networking.
 - campus network upgrades addressing the growing needs in mobile networking.
 - Network connection upgrade for the campus connection to a regional optical exchange or point-of-presence that connects to Internet2 or National Lambda Rail.



Other Notes on Area#1

- ❖ Must address scientific and engineering project and application drivers
- ❖ **Must present project-specific end-to-end scenarios for data movement**, distributed computing, and other end-to-end services driving the networking upgrade.
- ❖ Data movement scenarios are encouraged to describe end-to-end data transfers that include access to and use of wide area dynamic circuit networking services
- ❖ **Proposals must include a Campus Cyberinfrastructure plan** within which the proposed network infrastructure improvements are conceived, designed, and implemented in the context of a coherent campus-wide strategy and approach to CI.
- ❖ This Campus CI plan must be included as a supplementary document and is limited to no more than 5 pages. The plan should also address campus IPv6 deployment and use of the InCommon Federation global federated system.



Other Notes on Area#1

- ❖ **Must document explicit partnerships or collaborations with the campus IT/networking organization, as well as one or more domain scientists, research groups, and educators in need of the new network capabilities.**
- ❖ Partnership documentation from personnel not included in the proposal as PI, Co-PI, or Senior Personnel should be in the form of a letter of commitment located in the supplementary documents section of the proposal.
- ❖ Should describe an approach to end-to-end network performance measurement based on the perfSonar framework with associated tool installation and use; proposals may describe an alternative approach to perfSonar with sufficient justification.
- ❖ Title should start with:” CC-NIE Network Infrastructure:”
- ❖ Funding request not to exceed \$500k for up to 2 years



CC-NIE Area#2 – Network Integration and Applied Innovation

- ❖ end-to-end network CI through integration of existing and new technologies and applied innovation
- ❖ Applying network research results, prototypes, and emerging innovations to enable (identified) research and education
- ❖ May leverage new and existing investments in network infrastructure, services, and tools by combining or extending capabilities to work as part of the CI environment used by scientific applications and users



Area#2 Examples of Relevant Activities

- ❖ Integration of networking protocols/technologies with application layer
- ❖ Tool development supporting native IPv6 environments
- ❖ **Transitioning successful research prototypes in SDN, and activities supported by GENI and FIA programs, to distributed scientific environments and campus infrastructure**
- ❖ Development/deployment of mobile communication protocols/capabilities to support remote sensor and instrument access, real-time and adaptive steering, and energy efficiency.



Area#2 Examples of Relevant Activities

- ❖ Apply new energy efficient algorithms/techniques to production net environments
- ❖ Network engineering support through expert teams, common tools and processes for solving end-to-end network performance issues
- ❖ Applying network test, monitoring, measurement, or security tools and capabilities, including PerfSonar, Bro, and Web10G, to scientific domains, project end systems, and distributed environments.
- ❖ Experimental deployment of new networking protocols and technologies



Other Notes on Area#2

- ❖ **Must identify one or more supported science or engineering research projects or applications and describe how the proposed network integration activities will support those projects**, particularly in the context of addressing data movement, throughput, and predictable performance end-to-end.
- ❖ Must include clear project goals and milestones.
- ❖ Any software development must be made available under an open source license.
- ❖ Title should start with “CC-NIE Integration:”
- ❖ Funding request not to exceed \$1M total for up to 2 years



Additional Review Criteria for CC-NIE proposals

- ❖ expected impact on the deployed environment described in the proposal.
- ❖ extent to which the value of the work is described in the context of a needed capability required by science and engineering, and potential impact across a broader segment of the NSF community.
- ❖ A project plan that addresses in its goals and milestones the end result of a working system in the target environment.
- ❖ Where applicable, how resource access control, federated identity management, and other cybersecurity related issues and community best practices are addressed.
- ❖ Also for CC-NIE Network Infrastructure projects: Cyberinfrastructure Plan - How well does the cyberinfrastructure plan support and integrate with the institutions' science and technology plan? To what extent is the cyberinfrastructure plan likely to enhance capacity for discovery, innovation, and education in science and engineering? How well does the plan as presented position the proposing institution(s) for future cyberinfrastructure development? Are IPv6 deployment and InCommon federation addressed?
- ❖ Also for CC-NIE Integration projects: Tangible metrics described to measure the success of the integrated systems and any associated software developed, and the steps necessary to take the systems from prototype status to production use.



CC-NIE 2012 Stats

- ❖ Proposals were due May 30
- ❖ 89 proposals received (\$52M+ requested)
- ❖ 39 awards made (34 projects total)
 - 34 different institutions
 - 23 states
 - Total funding: \$21.6M (that includes \$3M in co-funding from CISE/CNS)
 - Area#1: \$9.7M, 21 awards
 - Area#2: \$11.9M, 18 awards



Award List Area#1(unordered)

Institution	PI	Title
Colorado State U	Burns, Patrick	CC-NIE Data-Driven Network Infrastructure Upgrade for Colorado State University
U of Washington	Lazowska, Edward	CC-NIE Network Infrastructure: Enhancements to Support Data-Driven Discovery at the University of Washington
Virginia Tech	Gardner, Mark	CC-NIE Network Infrastructure: ASCED -- An Advanced Scientific Collaboration Environment and DMZ
U of Chicago	Jelinkova, Klara	CC-NIE Network Infrastructure: High Performance Research Networking (HiPerNet)
Penn State	Agarwala, Vijay	CC-NIE Network Infrastructure: Accelerating the Build-out of a Dedicated Network for Education and Research in Big Data Science and Engineering
Duke	Futhey, Tracy	CC-NIE Network Infrastructure: Using Software-Defined Networking to Facilitate Data Transfer
U of Florida	Deumens, Erik	CC-NIE Network Infrastructure: 100Gig Connection to FLR
U of Wisconsin	Maas, Bruce	CC-NIE Network Infrastructure: Advancing Network Capacity, Efficiency, and Security for Wisconsin Big Data Research Through Improvement of campus research DMZ
U of Oregon	Rejaie, Reza	CC-NIE Network Infrastructure: Bridging Open Networks for Scientific Applications and Innovation (BONSAI)
Florida International U	Ibarra, Julio	CC-NIE Network Infrastructure: FlowSurge: Supporting Science Data Flows towards discovery, innovation and education



Award List Area#1

Institution	PI	Title
UT Knoxville	Hazelwood, Victor G.	CC-NIE Network Infrastructure: Bandwidth for Leadership in Advancing Science and Technology (BLAST)
UC San Diego	Papadopoulos, Philip	CC-NIE Network Infrastructure: PRISM@UCSD: A Researcher Defined 10 and 40Gbit/s Campus Scale Data Carrier
San Diego State	Castillo, Jose	CC-NIE Network Infrastructure: Implementation of a Science DMZ at San Diego State University to Facilitate High-Performance Data Transfer for Scientific Applications
U of North Carolina	Aikat, Jay	CC-NIE Network Infrastructure: Enabling data-driven research
Florida State U	Barret, Michael	CC-NIE Network Infrastructure: NoleNet Express Lane -- a private network path for research data transmission at Florida State University and beyond
U of Michigan	Noble, Brian	CC-NIE Network Infrastructure: Expanding Connectivity to Campus-Wide Resources for Computational Discovery
Wayne State	Cinabro, David	CC-NIE Network Infrastructure: Wayne State University
Yale	Sherman, Andrew	CC-NIE Network Infrastructure: The Future of Research & Collaboration: The Dedicated Science Network
Louisiana State U	Tohline, Joel	CC-NIE Network Infrastructure: CADIS -- Cyberinfrastructure Advancing Data-Interactive Sciences
U of Colorado	Hauser, Thomas	CC-NIE Network Infrastructure: Improving an existing science DMZ
Texas A&M	Cantrell, Pierce	CC-NIE Network Infrastructure: Advanced Connectivity for Texas A&M University



Award List Area#2

Institution	PI	Title
Indiana U	Swany, Douglas	Collaborative Research: CC-NIE Integration: An Open Cloud Infrastructure for Scalable Data Intensive Collaboration
UT Knoxville	Beck, Micah	Collaborative Research: CC-NIE Integration: An Open Cloud Infrastructure for Scalable Data Intensive Collaboration
Vanderbilt	Sheldon, Paul	Collaborative Research: CC-NIE Integration: An Open Cloud Infrastructure for Scalable Data Intensive Collaboration
U of Chicago	Tuecke, Steven	Collaborative Research: CC-NIE Integration: A Data Movement Solution for Next-Generation Campus Cyberinfrastructure
Indiana U	Swany, Douglas	Collaborative Research: CC-NIE Integration: A Data Movement Solution for Next-Generation Campus Cyberinfrastructure
U of Maryland	Voss, Brian	CC-NIE Integration: SDNX - Enabling End-to-End Dynamic Science DMZ
Ohio State	Whitacre, Caroline	CC-NIE Integration: Innovations to Transition a Campus Core Cyberinfrastructure to Serve Diverse and Emerging Researcher Needs
UMass Amherst	Dubach, John	CC-NIE Integration: Multi-Wave - a Dedicated Data Transport Ring to Support 21st Century Computational Research
Clemson	Wang, Kuang-Ching	CC-NIE Integration: Clemson-NextNet
U of Kentucky	Kellen, Vincent	CC-NIE Integration: Advancing Science through Next Generation SDN Networks



Award List Area#2

Institution	PI	Title
Stanford	McKeown, Nick	CC-NIE Integration: Bringing SDN based Private Cloud to University Research
U of North Carolina	Baldine, Ilia	Collaborative Research: CC-NIE Integration: Transforming Computational Science with ADAMANT (Adaptive Data-Aware Multi-domain Application Network Topologies)
U of Southern California	Deelman, Eva	Collaborative Research: CC-NIE Integration: Transforming Computational Science with ADAMANT (Adaptive Data-Aware Multi-domain Application Network Topologies)
Duke	Chase, Jeffrey	Collaborative Research: CC-NIE Integration: Transforming Computational Science with ADAMANT (Adaptive Data-Aware Multi-domain Application Network Topologies)
U of Nebraska	Bockelman, Brian	CC-NIE Integration: Bringing Distributed High Throughput Computing to the Network with Lark
Caltech	Newman, Harvey	CC-NIE Integration: ANSE (Advanced Network Services for Experiments)
Missouri U - Columbia	Springer, Gordon	CC-NIE Integration: Creation of an Institutional Cyberinfrastructure to Enable Researcher-Oriented, Federated Environment for Large, Collaborative Science Projects
UC Davis	Bishop, Matt	CC-NIE Integration: Improved Infrastructure for Data Movement and Monitoring



Selected quotes from Award Abstracts

- ❖ *"The Multi-Service Exchange (MSX) platform facilitates an environment where the advanced network functions are extended throughout campus over a Software Defined Networking (SDN) solution...the main focus of this effort is to integrate the advanced network functions with individual science applications, algorithms and workflows."* - **#1246386 (PI: Voss, UMd, "CC-NIE Integration: SDNX - Enabling End-to-End Dynamic Science DMZ")**



Selected quotes from Award Abstracts

- ❖ *"...scientists lack the tools that integrate the operation of workflow-driven science applications on top of dynamic infrastructures that link campus, institutional and national resources into connected arrangements targeted at solving a specific problem...ADAMANT leverages ExoGENI - an NSF-funded GENI testbed, as well as national providers of on-demand bandwidth services (NLR, I2, ESnet) and existing OSG computational resources to create elastic, isolated environments to execute complex distributed tasks."* – #1245926 (**Lead proposal – PI: Baldine, UNC, collabs: PI: Deelman, USC and PI: Chase, Duke – “Collaborative Research: CC-NIE Integration: Transforming Computational Science with ADAMANT (Adaptive Data-Aware Multi-domain Application Network Topologies)”**)



Selected quotes from Award Abstracts

- ❖ *"The resulting SDN network will have a lasting impact on research projects spanning a wide range of areas including astrophysics, bio-medical, computer vision, visualization, and networking research...**It will also enable end-to-end user-defined provisioning of network access and capacity** so that each research project can obtain precisely the performance it requires of the network. Finally, being integrated with the GENI network will enable researchers to access additional resources all across the GENI network."* - **#1246332 (PI: Kellen, U of Kentucky, "CC-NIE Integration: Advancing Science through Next Generation SDN Networks")**



Selected quotes from Award Abstracts

- ❖ *"...Specifically, Lark will integrate the capabilities of perfSONAR, a network performance monitoring tool, and DYNES, a cyber-instrument for allocation of bandwidth guarantees, into distributed computing grid middleware such as Condor and glideinWMS. By using a multi-site IPv6 test-bed and collaborating with a Nebraska project to provide small starter clusters to local colleges, the project aims to produce production-quality technology that could be deployed into the **Open Science Grid (OSG)**...This work will enable better trade-offs between selection of a site to execute jobs, data location decisions, performance, and security in large distributed computing environments, hence providing more effective use of limited resources. **This project addresses a clear and immediate need in the Large Hadron Collider (LHC) computing environment, as network-aware scheduling will increase the amount of analysis the LHC experiment collaborations can perform.** By integrating this work into the widely-used Condor high throughput computing software and into the OSG stack, the benefits will become available not only to LHC researchers, but to a number of science and industry projects that rely upon distributed high throughput computing." - #1245864, PI: **Bockelman, U of Nebraska, "CC-NIE Integration: Bringing Distributed High Throughput Computing to the Network with Lark"***



Selected quotes from Award Abstracts

- ❖ *"...a research infrastructure component is added to the campus network, encompassing 40 gigabit per second connections to research projects on campus with large data transport needs...Additional network capacity is provided to the Front Range GigaPoP...Specific projects with large data transport needs that benefit from the project are: 1)... Extreme Ultraviolet Lasers at nanoscale (huge data sets, transported regularly to and from HPC centers), 3) innovative data collection and dissemination of Internet traffic exhibiting malware patterns (huge, shared data sets, shared internationally), 4) two complementary next-generation climate modeling projects involving comprehensive, regional ecosystem analysis (huge, multi-faceted data sets, shared regionally and worldwide), and 5) state-of-the-art genomic analysis for pre-clinical drug design (huge and numerous data sets). However, the entire community of users benefits from added network capacity for dynamic and other research traffic that frees up capacity on the generic network."* - **#1245428, PI: Burns, CSU, "CC-NIE Data-Driven Network Infrastructure Upgrade for Colorado State University"**



Selected quotes from Award Abstracts

- ❖ *"The availability and growing use of high-performance heterogeneous computing and storage components by scientists across university campuses has led to realization that the general-purpose campus network offers neither the capacity nor the capabilities necessary for data-intensive research project. This project designs, builds and maintains a new network at the University of Oregon (UO) campus called Bridging Open Networks for Scientific Applications and Innovation (BONSAI). BONSAI is designed as a high-performance science network providing high end-to-end throughput and unique capabilities between five interconnected UO facilities, as well as computing resources located at other institutions throughout Internet2.*
- ❖ *This project primarily focuses on the following five major tasks:(1) Creating a Science DMZ platform among major computing facilities across UO; (2) Adding a new 10Gbps network circuit between the UO and Internet2; (3) Implementing and operating Software-Defined Networking(SDN) technologies throughout the network; (4) Promoting the development of IPv6- and service-aware scientific applications, and (5) Socializing the use of the UO's membership to the InCommon federation." - #1246136, PI: Rejaie, U of Oregon, "CC-NIE Network Infrastructure: Bridging Open Networks for Scientific Applications and Innovation (BONSAI)"*



Selected quotes from Award Abstracts

- ❖ *"...The significant expansion of the network being undertaken in this project will make several hundred 10 Gbps ports available in laboratories and offices across two major campuses of Penn State. This expansion is accomplished by deploying 48-port 10 Gbps Ethernet switches in 12 separate buildings with highest concentration of faculty who rely on large computational and data resources for advanced teaching and research. Each of these 12 switches is slated to have two 10 Gbps uplinks, one to each core router for redundancy. Both core routers will have 10 Gbps and 100 Gbps ports for external connectivity. The overall goal of the project is to provide at least a ten-fold increase in end-to-end network connectivity in labs and offices, making sustained 10 Gbps connections ubiquitous between faculty laboratories, data-generating instruments, select classrooms, and local and national computational and data resources. "* - **#1245980, PI: Agarwala, PSU, "CC-NIE Network Infrastructure: Accelerating the Build-out of a Dedicated Network for Education and Research in Big Data Science and Engineering"**



Selected quotes from Award Abstracts

- ❖ *"...These applications and datasets support research projects in the numerical simulation of earthquake rupture and wave propagation, coastal ocean modeling, pulse detonation engines, vortex rings in Bose-Einstein condensates, and large-scale data for proteomics, gene promoter bioinformatics, and microbial metagenomics. The project will directly impact the faculty and students who encompass the Computational Science Research Center which draws participation from all science and engineering departments at SDSU. **The Science DMZ** provides researchers with the capability to rapidly exchange large datasets and deploy Web-based science applications that are accessible through a dedicated network ...#1245312, PI: Castillo, San Diego State, "CC-NIE Network Infrastructure: Implementation of a Science DMZ at San Diego State University to Facilitate High-Performance Data Transfer for Scientific Applications"*



Selected quotes from Award Abstracts

- ❖ *"Prism@UCSD creates a campus-wide "Big Data freeway" composed of high-bandwidth end-to-end optical connections routed by a next-generation Arista switch. This creates an optical fabric capable of more than 10 Terabit/s of aggregate bandwidth, has full bisection similar to in-machine room clusters but is deployed at a campus scale. This researcher-defined network unites users of in-lab scientific instruments such as genome sequencers and microscopes with remote compute, visualization, data-storage and analysis systems...Prism builds upon and upgrades the Quartzite "campus-scale network laboratory" NSF MRI (awarded 2006) that was motivated by applications with extreme-scale bandwidth requirements. Compared to Quartzite, Prism not only adds IPv6 capability and support for software defined networks via OpenFlow, but also increases port capacity by 4x, lowers power consumption by 3x, and removes all card-to-switch-backplane over-subscription at the core switch..."*

#1246396, PI: Philip Papadopoulos, UCSD, "CC-NIE Network Infrastructure: PRISM@UCSD: A Researcher Defined 10 and 40Gbit/s Campus Scale Data Carrier"



Wrap up

- ❖ Award abstracts available on fastlane.nsf.gov (try searching on the term "CC-NIE")
- ❖ CC-NIE program activities ahead
 - Jan 7/8 GENI/CC-NIE 2 day workshop at NSF
 - Multiple sessions at I2 Joint Techs in Jan.
 - ScienceDMZ tutorial
 - CC-NIE talk and separate BOF
- ❖ Future CC-NIE funding opportunities
 - CC-NIE saw significant community interest in 2012
 - **FY13 Solicitation under joint development with CISE/CNS**
- ❖ Any comments/questions – kthompso@nsf.gov