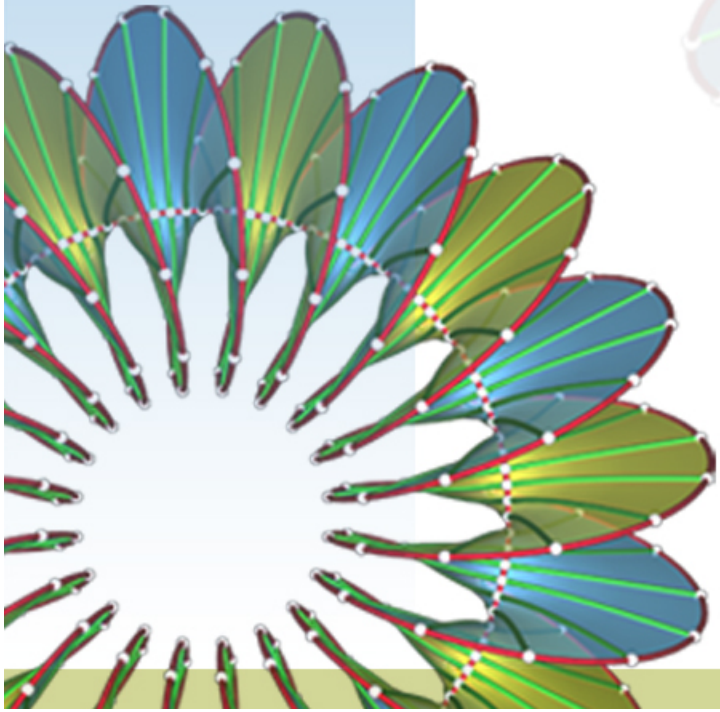

Distributed Computing at NSF

Daniel S. Katz
NSF/ACI

Research vs. Infrastructure

- Research: projects that support research in distributed computing
- Includes many CISE core programs, such as Computer Systems Research (CSR) and Networking Technology and Systems (NeTS), as well as Exploiting Parallelism and Scalability (XPS)
  - ~17% of new FY14 CISE projects support research into or are closely related to cloud computing, across:
    - Computer systems
    - Computer networks
    - Security and privacy
    - Data Management
    - Applications and Software Engineering
- Many other programs from across NSF, including BIGDATA, Computational and Data-Enabled Science & Engineering (CDS&E)

Research Partnerships

- With SRC – Secure and Trustworthy Cyberspace: Secure, Trustworthy, Assured and Resilient Semiconductors and Systems (SaTC: STARSS)
  - Hardware design for assured function
  - 9 FY14 projects funded jointly by NSF & SRC
- With Intel – Cyber-Physical Systems Security (CPS-Security)
  - Ensuring that cyber-physical systems are secure and preserve privacy
  - Small number of FY14 projects funded jointly by NSF & Intel
  - 2 Larger-scale, interdisciplinary projects in future
    - ~$3M/project, with multiple PIs

Research vs. **General Infrastructure**

- Infrastructure: projects that provide distributed computing for use in general research
  - XSEDE and OSG
  - Campus Cyberinfrastructure – Data, Network, and Innovation (CC*DNI)
  - CI platforms, e.g. iPlant, nanoHUB, OOI, ...
  - Major Research Instrumentation (MRI)

XSEDE (NSF-funded)

- Vision: A world of digitally enabled researchers, engineers, and scholars participating in multidisciplinary collaborations to tackle society's grand challenges
- Mission: To substantially enhance the productivity of a growing community of researchers, engineers, and scholars through access to advanced digital services that support open research
  - Leads to more science
  - Makes a previous impractical project feasible
  - Lowers barriers to adoption
- An ecosystem of advanced digital services accelerating scientific discovery
  - Supports a growing portfolio of interconnected resources and services
    - Advanced computing, high-end visualization, data analysis, and other resources and services
    - Interoperability with other infrastructures
- Intended for large number of projects per year (production class)
- Allocated through XSEDE-run national peer review process
- Open to US researchers and their collaborators in any field

Credit: John Towns (modified)

ACI-Supported Resources

Increasingly diverse yet collaborative

UCSD
Trestles
IO-intensive
10k cores
160 GB SSD/Flash

Gordon
Data-intensive
300 TB Flash Mem

❖ **Comet**
Long Tail Science
47k cores/2 PF
High Throughput

UIUC/Blue Waters
Leadership Class

Open Science Grid (OSG)
High throughput

PSC
Blacklight/Supercell
Large Shared Memory

❑ **Bridges**
Data-intensive

Indiana U.
❑ **Jetstream**
Cloud-based

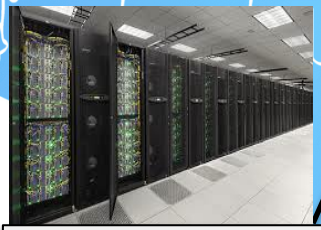

UT- Austin

Stampede
460K cores

Maverick
Visualization, Data Analytics

❖ ***Wrangler***
Novel Data-intensive

LSU SuperMIC
380 nodes – 1PF

UT-Knoxville/ORNL
Darter
250 TFlops

Nautilus
Visualization

Keeneland
CPU/GPGPU

Distributed resource

❖ ***To be deployed in 2015***
❑ ***To be deployed in 2016***

OSG (NSF- and DOE-funded)

- Mission: The Open Science Grid aims to promote discovery and collaboration in data-intensive research by providing a computing facility and services that integrate distributed, reliable and shared resources to support computation at all scales.
- OSG is a Consortium
  - Resource owners and campuses, scientist and research users, computer scientists and software providers, national and International partners
- OSG is a Project
  - Provides a fabric of services across contributed resources
    - E.g., Operations, user support, software packaging, testing, patching
- OSG is an Ecosystem
  - Provides a framework for exploring ways of scientific discovery through the use of distributed high throughput computing
  - Domain and computer scientists collaborating for more than decade
  - Contributing to state of the art through innovation and collaboration
- Connected to campus grids through BOSCO & OSG Connect

Credit: Lothar Bauerdick (modified)

OSG in Numbers

- OSG Delivers up to 2 Million CPU hours every day
- Almost 700M hours of Distributed High-Throughput Computing per year, of which ~90M were provided as “opportunistic resources”
- About 60% go to LHC, 20% to other HEP, 20% to many other sciences
- OSG has a footprint on ~120 campuses and labs in the U.S.
- OSG transfers ~1 PetaByte of data every day
- Supports active community of 20+ multi-disciplinary research groups

Credit: Lothar Bauerdick (March 2014)

CC*DNI

- Campus networking program, investing in improvements and re-engineering at the campus level to support a range of data transfers supporting computational science and computer networks and systems research
- Data infrastructure program, integrates new data-focused services, capabilities, and resources beyond single institutions; incorporates innovative developments in networking
- Also supports Network integration activities tied to achieving higher levels of network performance, reliability and predictability for science applications and distributed research projects

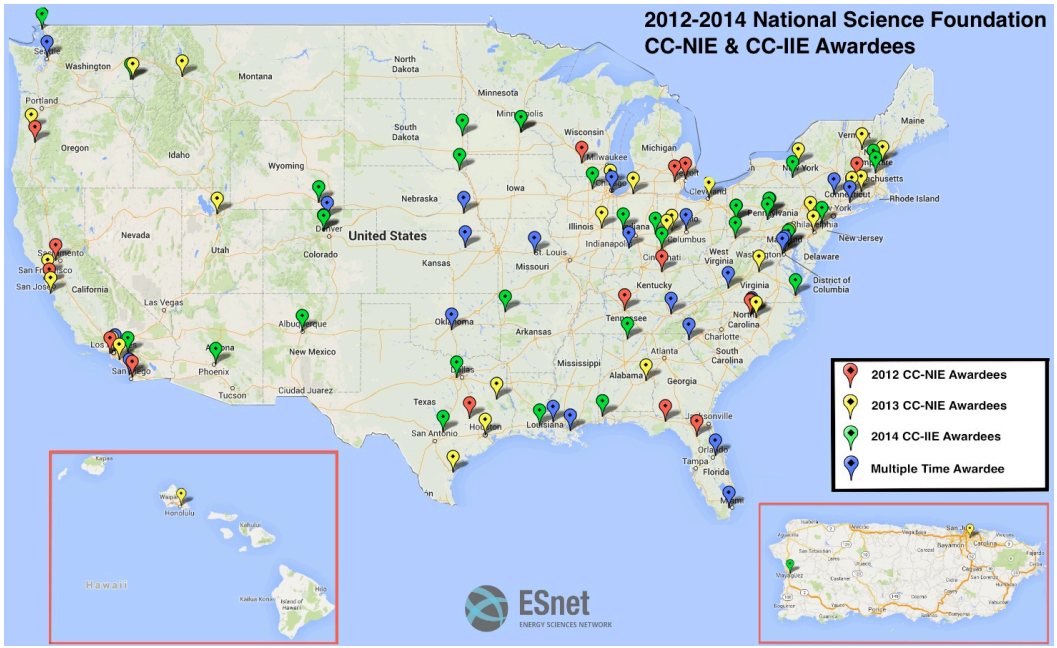

Research vs. **Computer Science Infrastructure**

- Infrastructure: projects that provide distributed computing for use in CS research
- Intended to give researchers systems on which they can do research at scale/level appropriate for impacting commercial systems
  - NSF FutureCloud
  - Parallel Reconfigurable Observational Environment (PRObE)
  - Global Environment for Network Innovations (GENI)
  - CISE Research Infrastructure (CRI)
  - Major Research Instrumentation (MRI)

an NSF FutureCloud

- A Large-Scale, Reconfigurable Experimental Environment for Cloud Research
- Kate Keahey (U Chicago), J. Mambretti (Northwestern), D.K. Panda (Ohio State), P. Rad (UT San Antonio), W. Smith, D. Stanzione (UT Austin)
- Large-scale, responsive experimental testbed
  - Targeting critical research problems at scale
  - Evolve with the community input
- Reconfigurable environment
  - Support use cases from bare metal to production clouds
  - Support for repeatable and reproducible experiments
- One-stop shopping for experimental needs
  - Trace and Workload Archive, user contributions, requirement discussions
- Intended to engage the community
  - Network of partnerships and connections with scientific production testbeds and industry
  - Partnerships with existing testbeds
  - Outreach activities
- http://www.chameleoncloud.org

CloudLab an NSF FutureCloud

- "With CloudLab, it will be as easy to get a cloud tomorrow as it is to get a VM today"
- Robert Ricci (U Utah), A. Akella (U Wisconsin), KC Wang (Clemson), C. Elliott (Raytheon/BBN), M. Zink (UMass), G. Ricart (US Ignite)
- Built on Emulab and GENI
- Exploring emerging and extreme cloud architectures
- Evaluating design choices that exercise hardware and software capabilities
- Studying geo-distributed data centers for low-latency applications
- Developing different isolation models among tenants
- Quantifying resilience properties of architectures
- Developing new diagnostic frameworks
- Exploring cloud architectures for cyber-physical systems
- Enabling realtime and near-realtime compute services
- Enabling data-intensive computing ("big data") at high performance in the cloud
- http://www.cloudlab.us

Funding Platforms

- Computer science research, supported by CISE research programs, NSF FutureCloud, CRI, some MRI
  - Platforms are central
- Domain science research, supported by other NSF research programs, general infrastructure including some MRI
  - Platforms are essential, but not central

- Need to continue to support both

Going Forward

- Emphasize research to infrastructure process (aka transition to practice)
- Emphasize federation (for both compute and data)
  - Using existing work from XSEDE, OSG, OGF, GENI, IETF, EGI, etc.
  - Common authentication, authorization, accounting (AAA) services
- Continue to make sure infrastructure needed for distributed computing research is available
- Continue to make sure distributed infrastructure needed for general research is available
- Determine needs, e.g.
  - National Academies Activity: Future Directions for NSF Advanced Computing Infrastructure to Support US Science in 2017-2020
  - Upcoming CCC Extensible Distributed Systems Workshop Jan 21-22 Arlington - http://www.cra.org/ccc/component/content/article/437-extensible-distributed-systems-workshop

Looking Farther Forward

- Start to think of infrastructure in wider sense

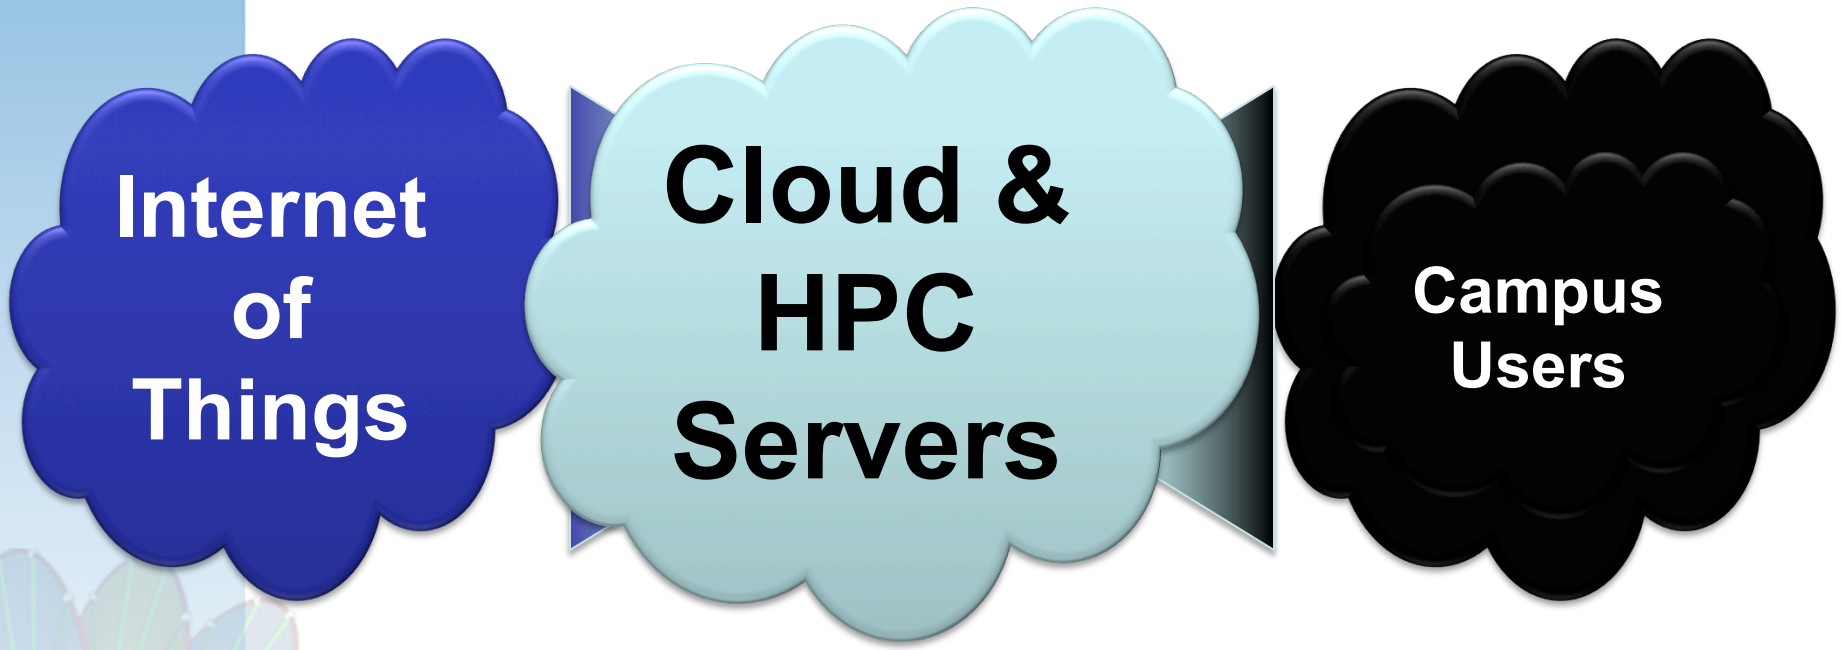

- What are the abstractions & tools for using this entire infrastructure?
- Cross-cutting concerns: security, programmability, reliability, verifiability

Credit: Snir et al., "21st Century Computer Architecture," 2012