Cyberinfrastructure Uptake
Some observations over 20 years

Nancy Wilkins-Diehr

San Diego Supercomputer Center

wilkinsn@sdsc.edu

Stated goals of this workshop
I hope to touch on two

- Serve as a focused forum for PIs to share technical information.
- **Explore innovative topics emerging within software communities.**
- **Discuss emerging best practices across the software projects.**
- Stimulate thinking on new ways of achieving software sustainability.
- Gather the shared experiences in an online web portal.

What do most developers have in common?
They'd like to see their software used

- Why does software get used?
  - Because it serves a need it's dead easy to use
  - Because of force
  - Because of desperation
  - Because of trust
  - Through science gateways

My experiences
Early 1990s

- Cray supercomputers
  - XMP, YMP, C90, T90
- Easy to convince people to to use CI
  - There weren't many people and there wasn't much CI
- Not a lot of opportunity cost
  - Great libraries and compilers
  - Codes typically ran much faster with a simple recompile on a vector supercomputer

My experiences
Late 1990s

- Parallel computing, grid computing
  - CI becomes harder to use, but benefits increase
- Rise of information technology
- Application/technology partnerships (aka force)
  - NSF Partnerships for Advanced Computational Infrastructure (NPACI for me)
    - **As project manager, I sometimes felt as though I were holding the shotgun**
      - Electron microscopes linked to supercomputers for data refinement, data stored in SRB, images retrieved for multiscale brain mapping project.
      - Large database scans and visualization including Protein Data Bank (PDB), Molecular Trajectories database, and Protein Sequence database using Legion and SRB
      - Scalable vis (MPIRE) as applied to astrophysics
      - Groundwater simulators (GWM, IPARS, SWM) paired with computing and data handling tools Active Data Repository (ADR), KeLP, MetaChaos, and Globus
      - CHARMM and Amber paired with Legion
      - AppLeS scheduler and Network Weather Services paired with Globus and Legion
      - immersed Boundary Method model of the heart and Titanium
      - Grid-enabled MCell
      - Strategic Applications Collaborations
- But these early partnerships really did pave the way for more to come later on

My experiences
Early 2000s

- Rise of the Web as truly functional and ubiquitous
- TeraGrid and the development of the Science Gateway program
  - Initial focus on NSF Information Technology Research (ITR) projects (aka more forced marriages)
    - Linked Environments for Atmospheric Discovery
    - National Virtual Observatory
    - RENCI Bioportal
    - SPRUCE (urgent coomputing)
    - nanoHUB
    - GISolve/CyberGIS
    - Open Life Sciences Gateway
    - Neutron Science Gateway
- As opportunities grew, we saw more and more gateways developed
- Explosion of software offerings
  - Apps, clouds, workflows
  - Many more options that all promise to solve scientists' problems
  - Age of disillusionment

My experiences
Today

- XSEDE Extended Collaborative Support program
  - Extensive program dedicated specifically to **in depth** engagements with the research community to improve CI uptake
    - ~70 ECSS staff across ~12 organizations
  - Expertise in supercomputing, but also workflows, gateways, visualization, databases and data analytics
    - Portion of the program dedicated to work with new communities
  - Extended, collaborative nature of the work helps make these projects a success

Open Invitation

XSEDE is actively looking to expand its software infrastructure

- Community Software Areas
  - Developer-supported software
  - Includes
    - Allocation on any XSEDE platforms
    - Backed up disk space
    - Listing in software catalog
  - XSEDE can provide usage statistics, promote training that you offer, highlight successes

http://www.xsede.org/software

Today, more users access XSEDE via gateways than by logging in

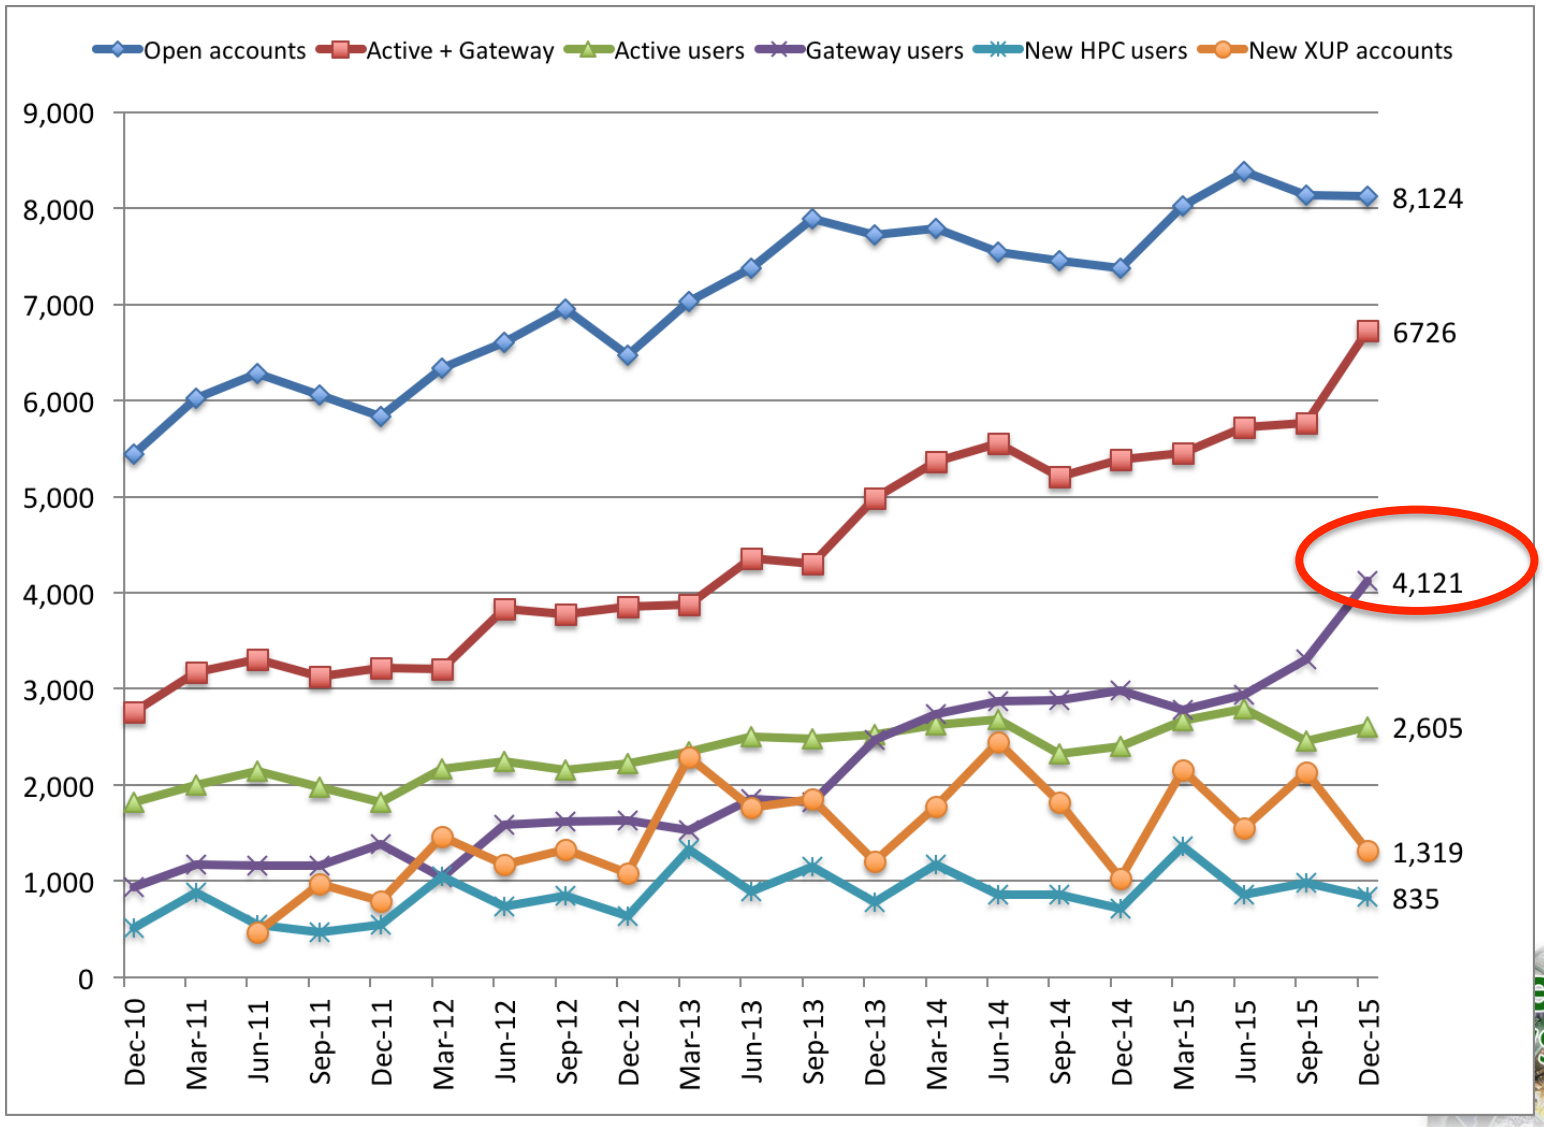

My experiences
Today

- S2I2 conceptualization phase award for a Science Gateway Institute
- 5000 respondent survey to understand
  - Use of science gateways
  - Development of science gateways
  - Results published at http://sciencegateways.org/resources/our-work/

88% indicate Web-based applications are important to their work

Specialized Resources	Percent
Data collections	75%
Data analysis tools, including visualization and mining	72%
Computational tools	72%
Tools for rapidly publishing and/or finding articles and data specific to my domain	69%
Educational tools	67%
Platforms for fostering group or community collaboration	63%
Simplified interfaces that eliminate the need to learn coding	62%
Citizen science and other public engagement resources	47%
Workflows that automate or capture tasks or processes	42%
Scientific instruments, such as telescopes, microscopes, or sensors	39%

n=4,004, or 88% of 4,538 researcher/educators. Percentage indicates these resources are “somewhat” or “very” important to their work.

57% played some role in gateway creation
and these gateways were used for a variety of purposes

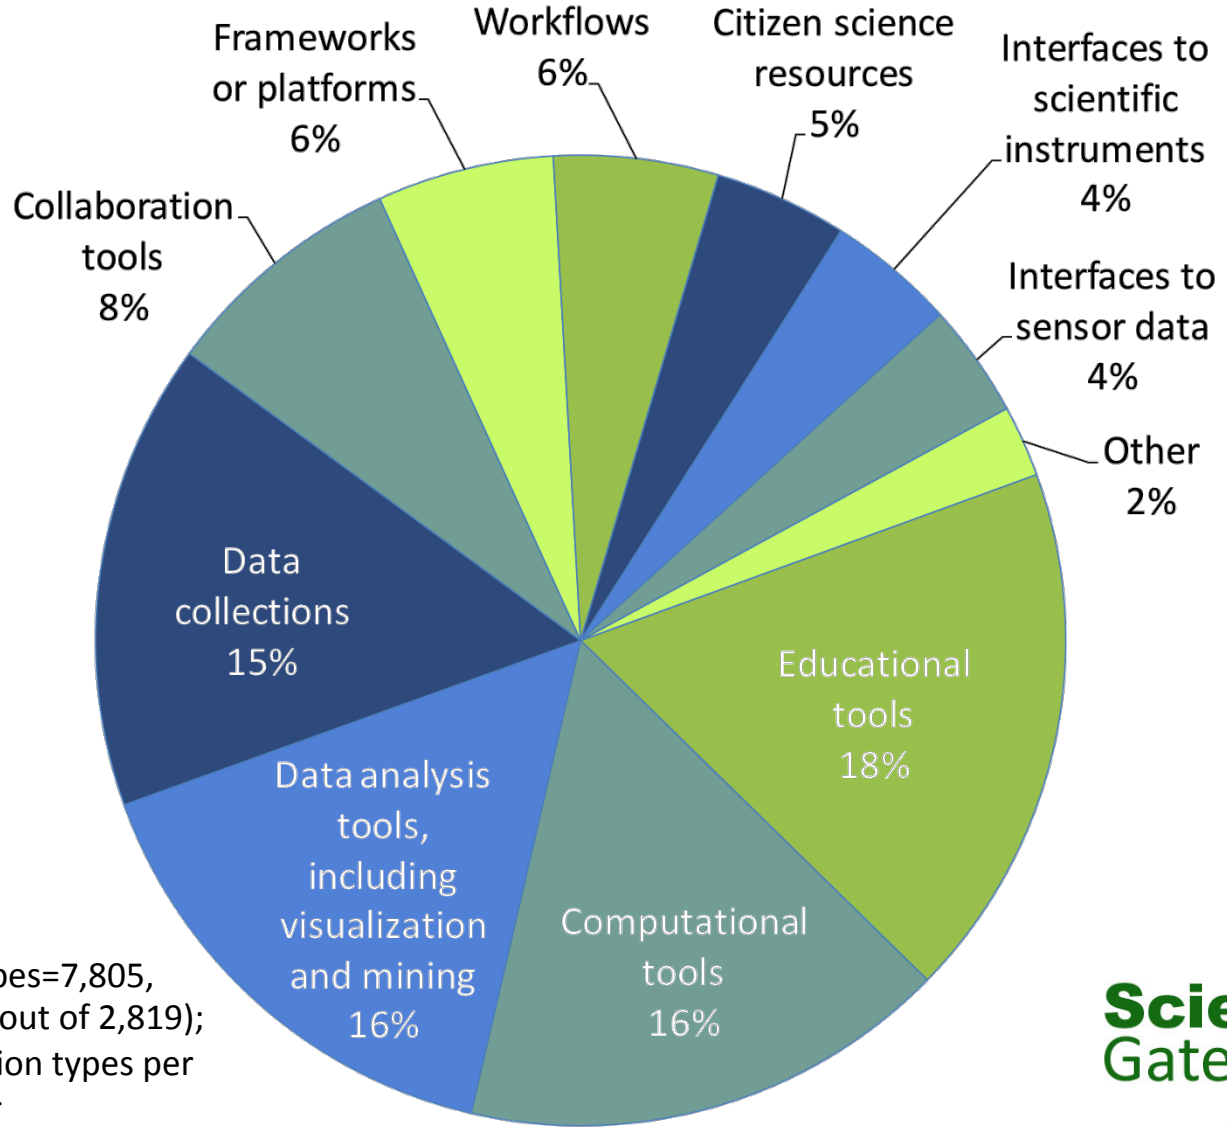

n of application types=7,805, by 2,756 creators (out of 2,819); mean=2.8 application types per application creator

Science Gateways

We also asked developers about new technologies

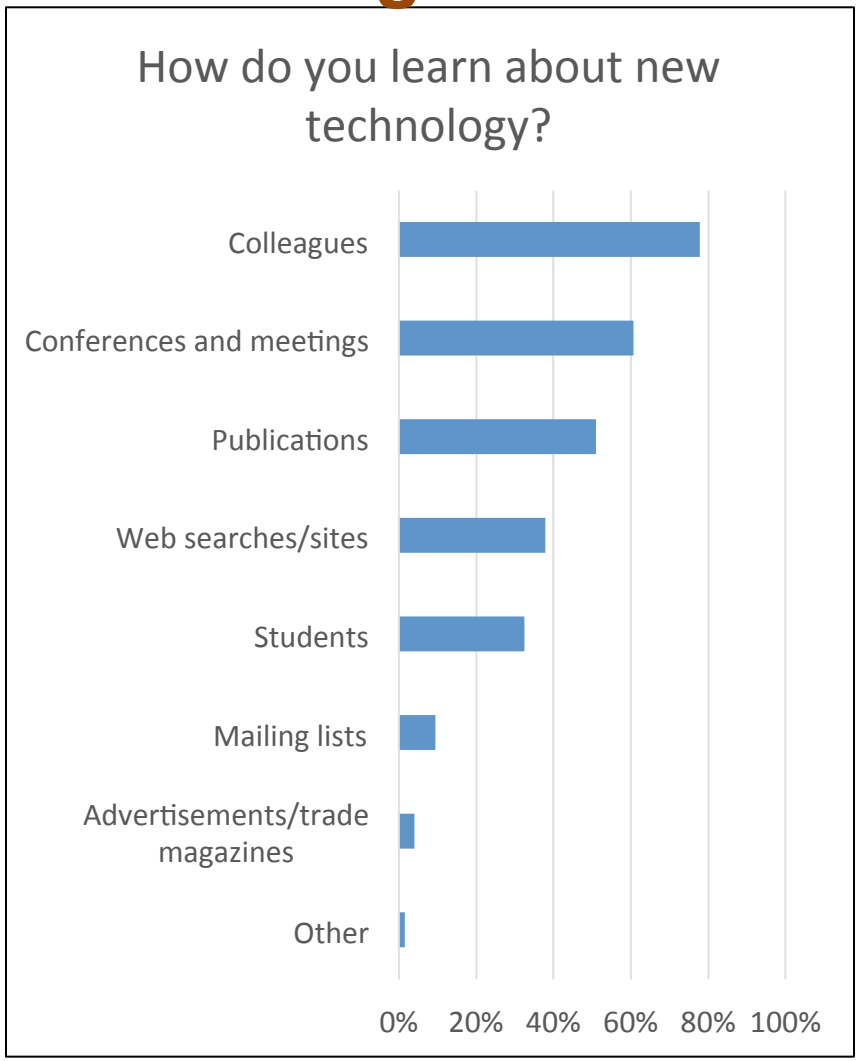

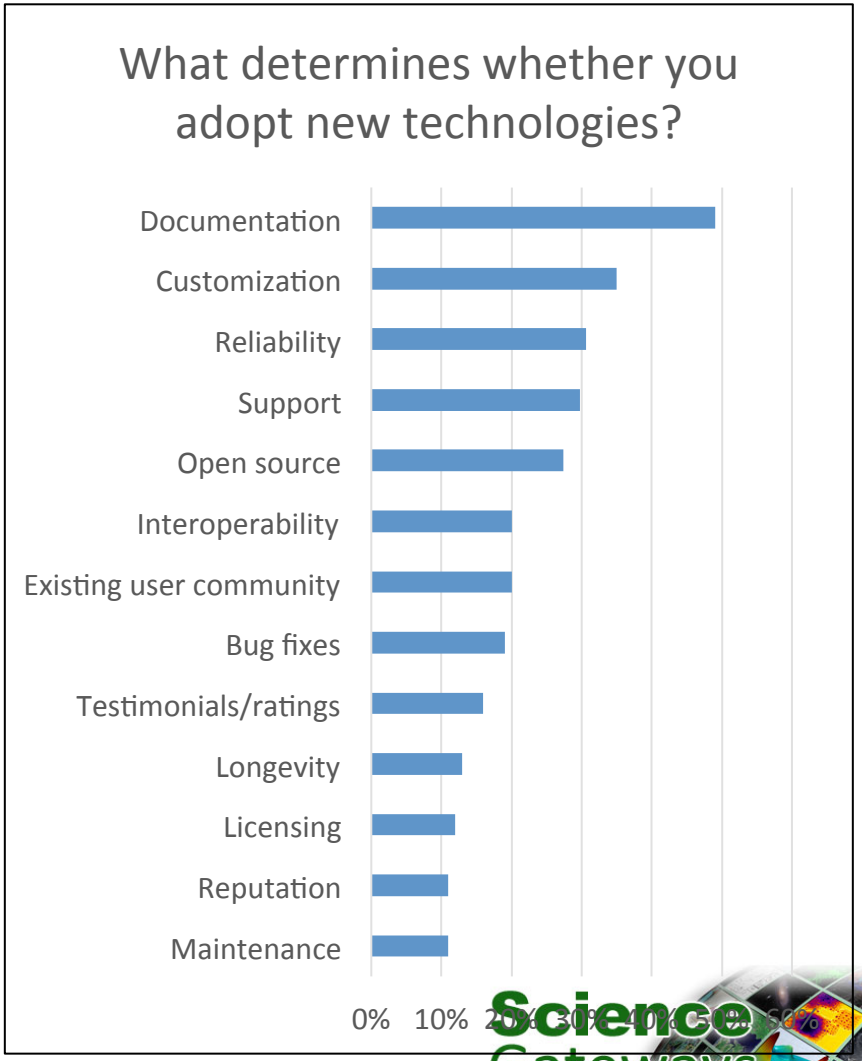

I've been trying to help scientists use technology for a very long time!

In December, I attended this NSF workshop
At this very same Westin

Cyberinfrastructure for NSF Large Facilities Workshop

December 1st, 10 a.m. – 6:00 p.m. & December 2nd, 8:30 a.m. to 2:30 p.m., 2015
Westin Arlington Gateway, Arlington, VA

Organizers

Workshop Chair:
Alex Szalay, Johns Hopkins University
Program Committee:
- Christine Borgman, University of California, Los Angeles
- Peter Couvares, Syracuse University
- Brian Glendenning, National Radio Astronomy Observatory
- Kerstin Lehnert, Columbia University
- Chuck Meertens, UNAVCO
- Manish Parashar, Rutgers University

Goal was to "create a forum for direct interaction between the NSF large facilities and CI developer community"

- Peter Couvares from the LIGO project shared observations on CI uptake by NSF large facilities
  - He nicely captured what I have observed over 20 years
- These observations can be applied to many things
  - When to use someone else's software
  - When to use a supercomputer
  - When to develop a science gateway

Insource/Outsource Tradeoffs

- CI interest and adoption is fundamentally a problem of desperation, credibility and risk
- “the best collaborators are the desperate ones” (Jim Gray)
- short of desperation, scientists and facilities are right to be skeptical and conservative when it comes to CI adoption

Source: Peter Couvares

Outsourcing: Risks/Costs

- Time and attention cost
- Uncertain benefit
- CI is usually research to help research — so **some home runs, some base hits, some strikeouts**.
- Unstable funding for turning research CI "hits" into sustainable, production CI infrastructure.

Source: Peter Couvares

Peter's Thoughts

- Why is a large facility like LIGO solving old computing problems itself (maybe poorly), when proven solutions developed by experts exist in the CI community (for job scheduling, data movement, etc.)?
  - Well, there's "proven" (toy) and there's "proven" (production at scale).
  - Sometimes CI was examined and rejected early before it was proven, or we had a best-of-breed in-house solution years ago but the world changed.
  - Sometimes our problems *are* unique.
  - Sometimes our inferior solution is "good enough" and there is little/no science benefit to improving it.
  - Sometimes our inferior solution *isn't* "good enough" but we don't have the expertise, judgement, or time to understand that, and pay a science cost — **How do we identify/differentiate THIS case**?
  - Sometimes we know or suspect we should, but don't have the internal resources to engage.
  - Sometimes we know or suspect we should, but the collaborators are difficult to work with (or visa versa) and the collaboration breaks down or never starts.

Source: Peter Couvares

How Do We Bridge the Gap?

- A facility needs to be desperate to make a serious investment of time or resources.
  - Desperation can come in many forms: e.g., big problem we know we can't solve ourselves (carrot), scary NSF mandate (stick).
- The CI provider needs to offer something concrete, and easy, and low-risk to entice engagement, and then manage the engagement skillfully from there
  - Solve one real problem for one person or group, and solve it well — earn allies and advocates, and work from that foothold.
- Sustainability: need a roadmap to continue CI support beyond the life of the original award: produce CI that can be turned into unfunded, truly community-supported OSS (rare), that can be neatly handed off to facilities to efficiently develop/maintain forever themselves (rare), that can find NSF funding as production infrastructure and not just research (rare), or self-fund by user fees (rare — Globus online is one example)
- Can we create more hybrid CI/facility roles? Rather than a provider/consumer relationship, put the right human in the middle. That becomes a hedge against many of the risks described earlier (unaligned priorities, abandoned CI, etc.).

Source: Peter Couvares

Thank you