



Computation Institute

Production Distributed Infrastructure (PDI)

Ian Foster

State of the art in PDI



- Recall the two-part grid vision
 - “On-demand access to computing power—and new modes of working enabled by same” (aka ‘cloud’ today)
 - “Federation of resources within virtual organizations”
- On-demand access
 - Limited success due to lack of economic model for service providers (supercomputer centers are not incentivized for this access modality)
 - Thus use of grid technologies in e.g. TeraGrid streamlines access, but does not enable fundamentally new ways of working
 - Commercial IaaS (cloud) finally delivers what we need! (at least in part)
- Federation/virtual organizations
 - Lots of wonderful technical progress, e.g., in security infrastructure
 - Numerous successes: LIGO, Earth System Grid, LHC Computing Grid, Open Science Grid, BIRN, ... : All operating 24/7, delivering production service that is vital to scientific progress in their disciplines

Open challenges (R&D, code, support, policy)



- What should future cyberinfrastructure look like, anyway?
 - Where does data sit, and where do researchers compute on it? (Isn't that the fundamental question?)
- If on-demand computing (conceived widely) is fundamental to the future of science, how do we deliver it at scale?
 - 10% of NSF budget for Amazon credit cards? Specialized science clouds?
- New delivery models for scientific software and discovery processes
 - Software delivery via tarballs is no longer the best method available
 - Maybe software as a service (SaaS) is the answer: e.g., see Globus Online
- Preparing for the next generation of cyberinfrastructure, circa 2017
 - Data deluge means we have >10x more cyberinfrastructure users
- No sustainability model for key grid infrastructure software
 - E.g., GridFTP: 20M transfers/day logged; who pays for it?
 - How does the research community want to pay for infrastructure software?

Path forward: I'll just present one idea ...



- Develop an alternative to current flawed system of support that perversely incentivizes replication while starving successful software of funding. E.g., maybe:
 - Fund research based on peer-reviewed innovation
 - Fund deployment and enhancement based on some measure of use
- Work out how to do this for
 - Software
 - Services
 - On-demand computing services
- How to do this? Let's discuss. Engage economists, perhaps? Study what works elsewhere?