



Computation Institute

Opinions on the State of Production Distributed Infrastructure (PDI)

Daniel S. Katz (dsk@ci.uchicago.edu)

Senior Fellow, Computation Institute (University
of Chicago & Argonne National Laboratory)

(all the following are personal opinions only)

State of the Art in PDI



- PDI - a distributed set of computational hardware and software that is intended to allow multiple people who are not the developers of the infrastructure to do something useful
- Three types of PDI exist:
 - Academic/Public Production (aimed at science)
 - TeraGrid/XSEDE, DEISA (HPC), OSG, EGI (HTC), Open Science Data Cloud (cloud), etc.
 - Academic/Public Research (aimed at computer science)
 - Grid5000, DAS, FutureGrid, PlanetLab, etc.
 - Open question – how to transition lessons to production
 - Commercial (aimed at whoever will pay)
 - Amazon (IaaS), Microsoft (Paas), perhaps SGI, Penguin (HPCaaS)
- In addition, lots of other components that may not be integrated together:
 - Campus clusters (HPC), campus Condor pools (HTC)
- All seem to do what they were designed to do reasonable well
- Note: this view is mostly by funding and purpose
- Note: Could also view by interface:
 - Command-line, grid (Globus, Condor, etc.), cloud

User View of PDI



- Need to answer:
 - Application developers:
 - How do I develop an applications?
 - What is the model?
 - How do I deploy an application?
 - Particularly hard for HPC
 - Applications users:
 - How do I find an application?
 - Application builders:
 - How do I compose an application?
 - First need to find components...
 - In all cases:
 - How do I execute an application?



Open Challenges (in R&D, code, support, and/or policy)



- Goal – deliver maximum science (at minimum cost?)
 - Note – sustainability seems to be a hot topic, but it seems to be defined as: useful work should continue to be done, with someone else paying for it
- 2 views of this?
 - What are the components of infrastructure?
 - Hardware (nodes, interconnect, storage, network), software (system software, middleware, tools/libraries, applications), training (material, people), support (people), integration into PDIs
 - What’s the vision of “the” PDI?
 - For those who build the infrastructure components, what is the architecture?
 - For applications people, what is the abstraction of the PDI?

Open Challenges (in R&D, code, support, and/or policy) (1)



- Issues

- How measure delivered science
 - We don't really know how to do this – we can measure papers (w/ a small time delay) or citations (w/ longer time delay)
- How to develop the infrastructure/tools that will best do this
 - I think we are doing reasonably well at this at an individual infrastructure/tool level—we identify things that users want to do, then provide tools/systems that let them do these things—but it's not clear that we ever solve the whole integrated problem
- How to integrate them
 - We can do this on a piece-by-piece basis, but don't really know how to do this in general (maybe because we don't have a single vision for where we are going)
 - Note that if we can do this, we can then allow the pieces to compete

Open Challenges (in R&D, code, support, and/or policy) (2)



- Issues

- How to represent users

- Large projects have the ability and opportunity to express their needs; small projects and individual users are often not represented

- Role of virtualization in HPC

- HPC executables are not portable now, but could be with virtualization
- But virtual HPC applications don't perform well (I/O, network issues)

- How to build the application programming system that matches the abstract PDI?

- We've had a nice run of ~20 years with the MPI model, which assumes a platform abstraction of interconnected nodes, each with CPUs and memory



- Greatest need is a single vision that defines the overall goal
 - NSF has CIF21 (<http://www.nsf.gov/pubs/2010/nsf10015/nsf10015.pdf>) which calls for such a vision, and 6 task forces that have created reports
 - DOE seems more focused on single large systems, and less on integrating them into an infrastructure
 - Some other US agencies looking at clouds, but not distributed computing in general
 - I haven't seen a European vision that includes integration of all PDIs, though the UMD implies one for HPC and Grids
- and is flexible enough to allow groups to work towards that goal in various ways
 - Should define metrics (to enable progress to be measured and see which work is most useful)
 - And an overall architecture with interfaces
 - OGF is trying to do some of this, but without sufficient US buy-in
- DARPA funded effort in understanding user productivity in HPCS – but no one is doing this for PDIs

Acknowledgements



- First slide (and probably some other thinking) is based on chapter 3 of upcoming book:
 - S. Jha, D. S. Katz, M. Parashar, O. Rana, and J. Weissman. Abstractions for Distributed Applications and Systems. Wiley.
- And technical report that's coming sooner:
 - D. S. Katz, S. Jha, M. Parashar, O. Rana, and J. Weissman. Distributed Cyberinfrastructures. Technical Report. Computation Institute, University of Chicago & Argonne National Laboratory, URL/number TBD
- And related eSI research themes (DPA and 3DPAS)