

# Volunteer Computing for Science Gateways

David P. Anderson

Space Sciences Lab  
University of California, Berkeley



11 October 2017



# Resources for scientific computing

## Organizational

supercomputers  
clusters  
grids  
clouds



## Consumer

desktop/laptop  
video game consoles  
smartphones, tablets  
appliances  
Internet

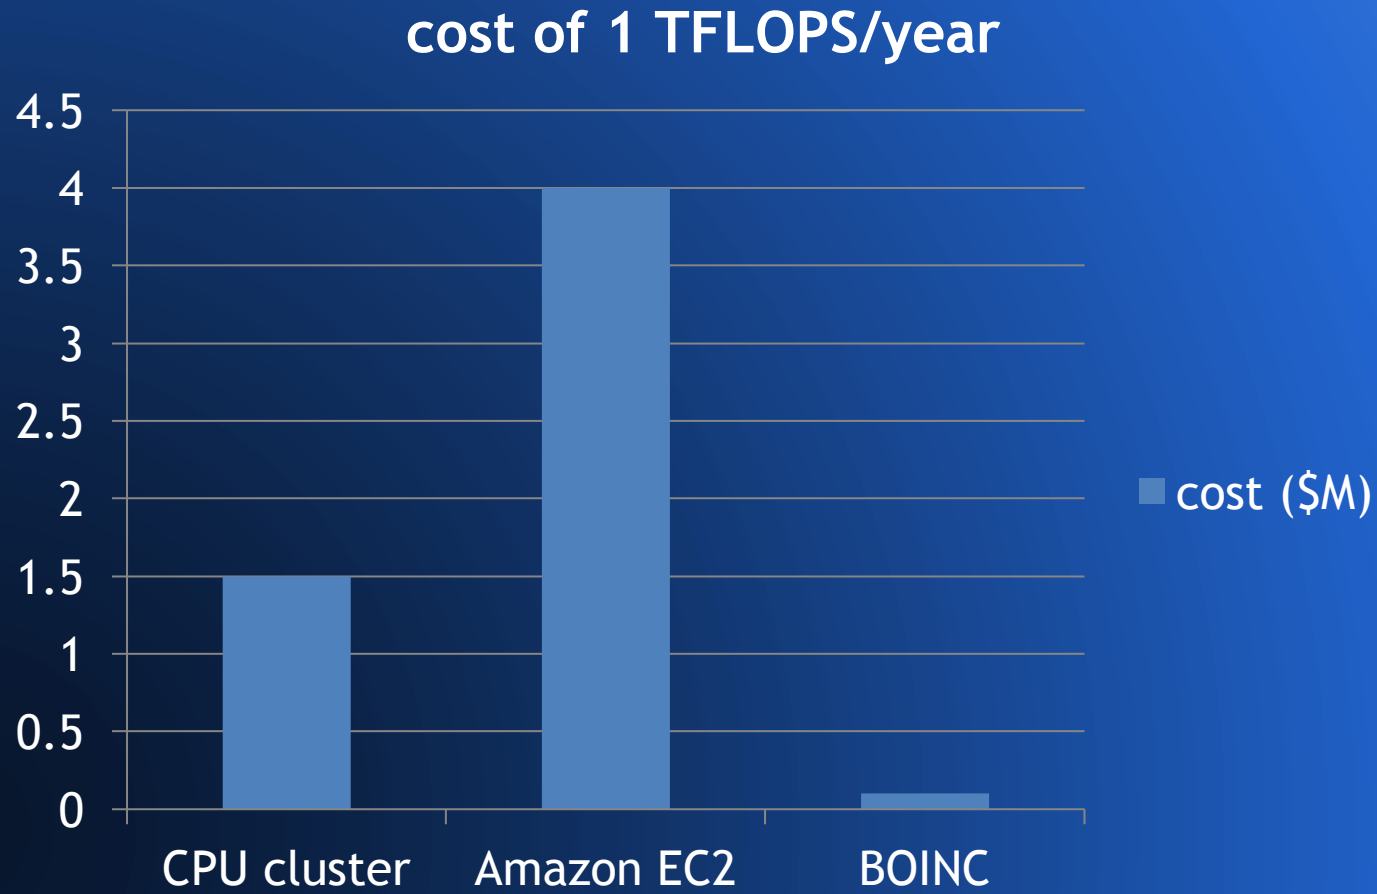
# Performance potential

- 1 billion desktop/laptop PCs
  - CPUs: 10 ExaFLOPS
  - GPUs: 1,000 ExaFLOPS
- 5 billion smartphones
  - CPUs: 20 ExaFLOPS
  - GPUs: 1,500 ExaFLOPS

# BOINC: middleware for volunteer computing

- Open-source, NSF-funded since 2002
- Server: distributes jobs, provides web site
- Client: runs jobs on consumer devices
- BOINC deals with
  - invisibility
  - device heterogeneity
  - recovery from random and malicious faults

# Cost

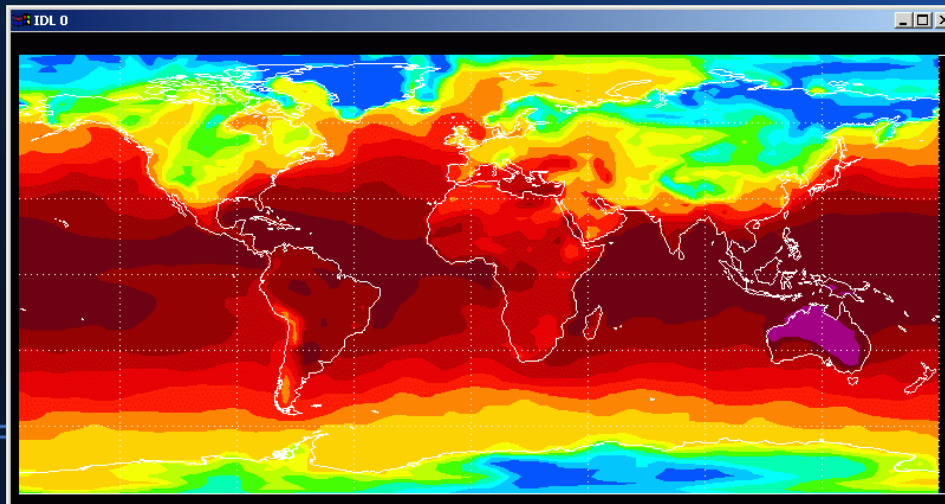



# The original BOINC model

- Scientists create BOINC projects, compete for volunteers
- Volunteers evaluate projects, participate in the “best” ones

# Projects using BOINC

- [Climateprediction.net](http://Climateprediction.net)
- [Rosetta@home](http://Rosetta@home)
- [Einstein@home](http://Einstein@home)
- [IBM World Community Grid](http://IBMWorldCommunityGrid)
- [CERN](http://CERN)
- . . .



# Volunteer computing today

- 500,000 active computers
  - 2.3 million CPU cores
  - 290K GPUs
- 40 projects
- 90 PetaFLOPS



# Problems and limitations

- The set of projects is small and static
  - creating a BOINC project is hard
  - risk: what if no volunteers?
  - not useful for sporadic or one-time workload
- The set of volunteers is declining
  - project lock-in
  - Marketing/branding

# A new “coordinated” model for VC

## Goals

- Serve more scientists
- Get more volunteers
- Move VC toward the mainstream of
  - HPC computational science
  - computer science

# The model

- Scientists
  - add VC to existing HPC facilities like science gateways and supercomputer centers
  - prototypes: TACC, nanoHUB
- The coordinator
  - an interface to VC based on science goals rather than projects
  - a unified brand
  - an allocator of computing power to projects

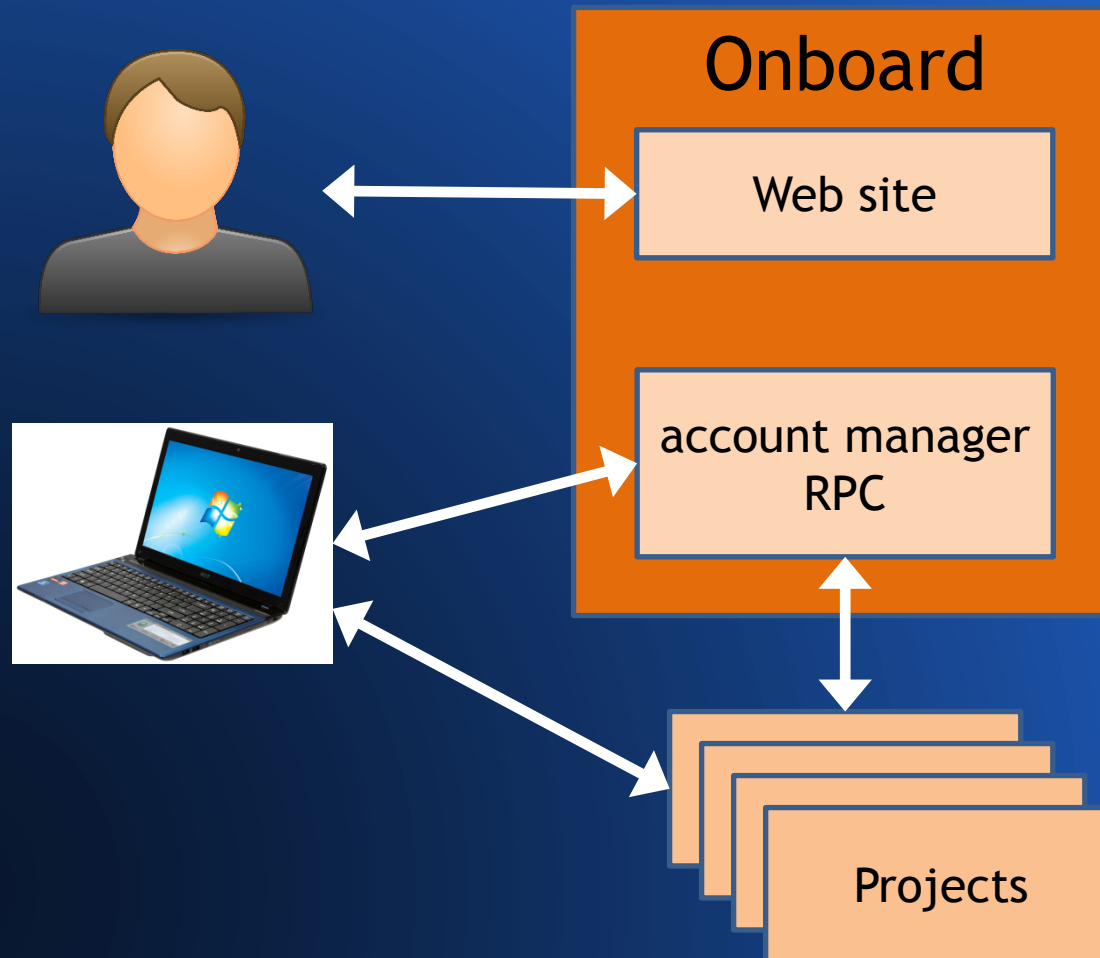
# Adding VC to existing HTC systems

- Application packaging
  - VirtualBox, Docker
- Remote job submission and file management
  - Web RPC interfaces

# nboard: a VC coordinator

- Volunteers pick science areas rather than projects.
- A resource allocation mechanism
  - XSEDE-based model
  - Projects no longer have to do PR
- A unified brand for marketing VC

# nboard Architecture



# Project status

- Funded 2016-2019
- TACC, nanoHUB
  - app packaging, job routing: done
  - public launch Q4
- Onboard
  - mostly done; Q4 launch

# Interested?

- If your gateway could use lots of computing
  - high throughput/latency
  - large compute-intensive jobs

... then let's talk!

[davea@ssl.berkeley.edu](mailto:davea@ssl.berkeley.edu)

<https://boinc.berkeley.edu>