

# Distributed Computing with the **B**erkeley **O**pen **I**nfrastructure for **N**etwork **C**omputing

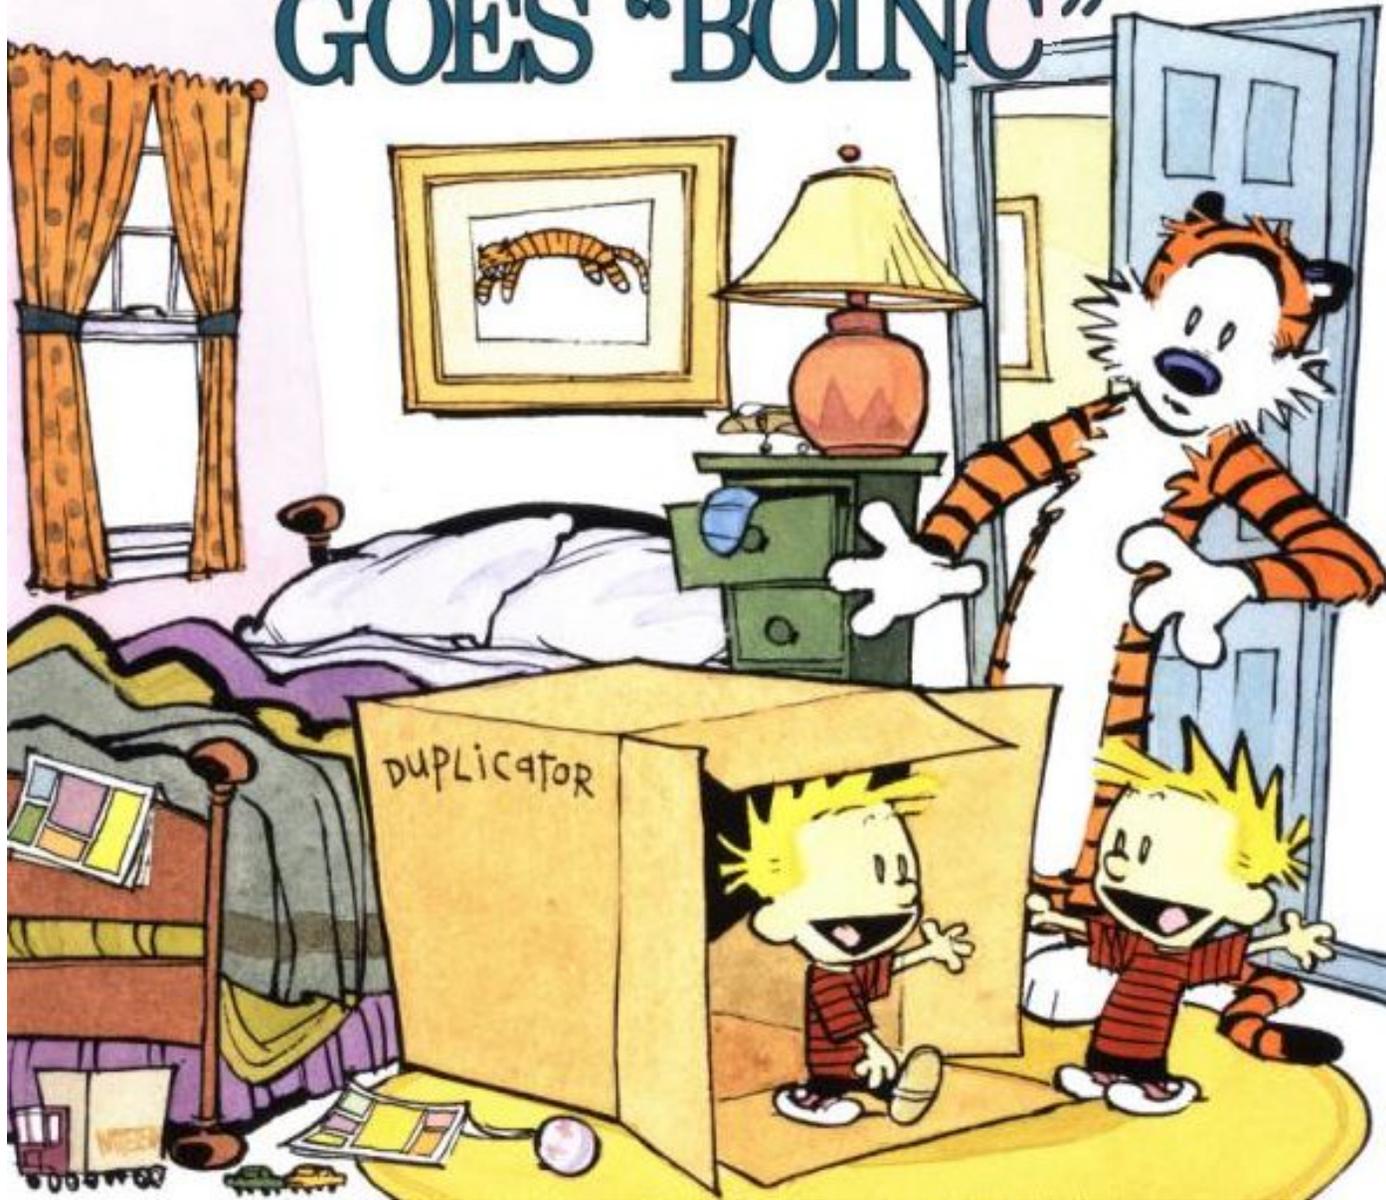


Eric Myers

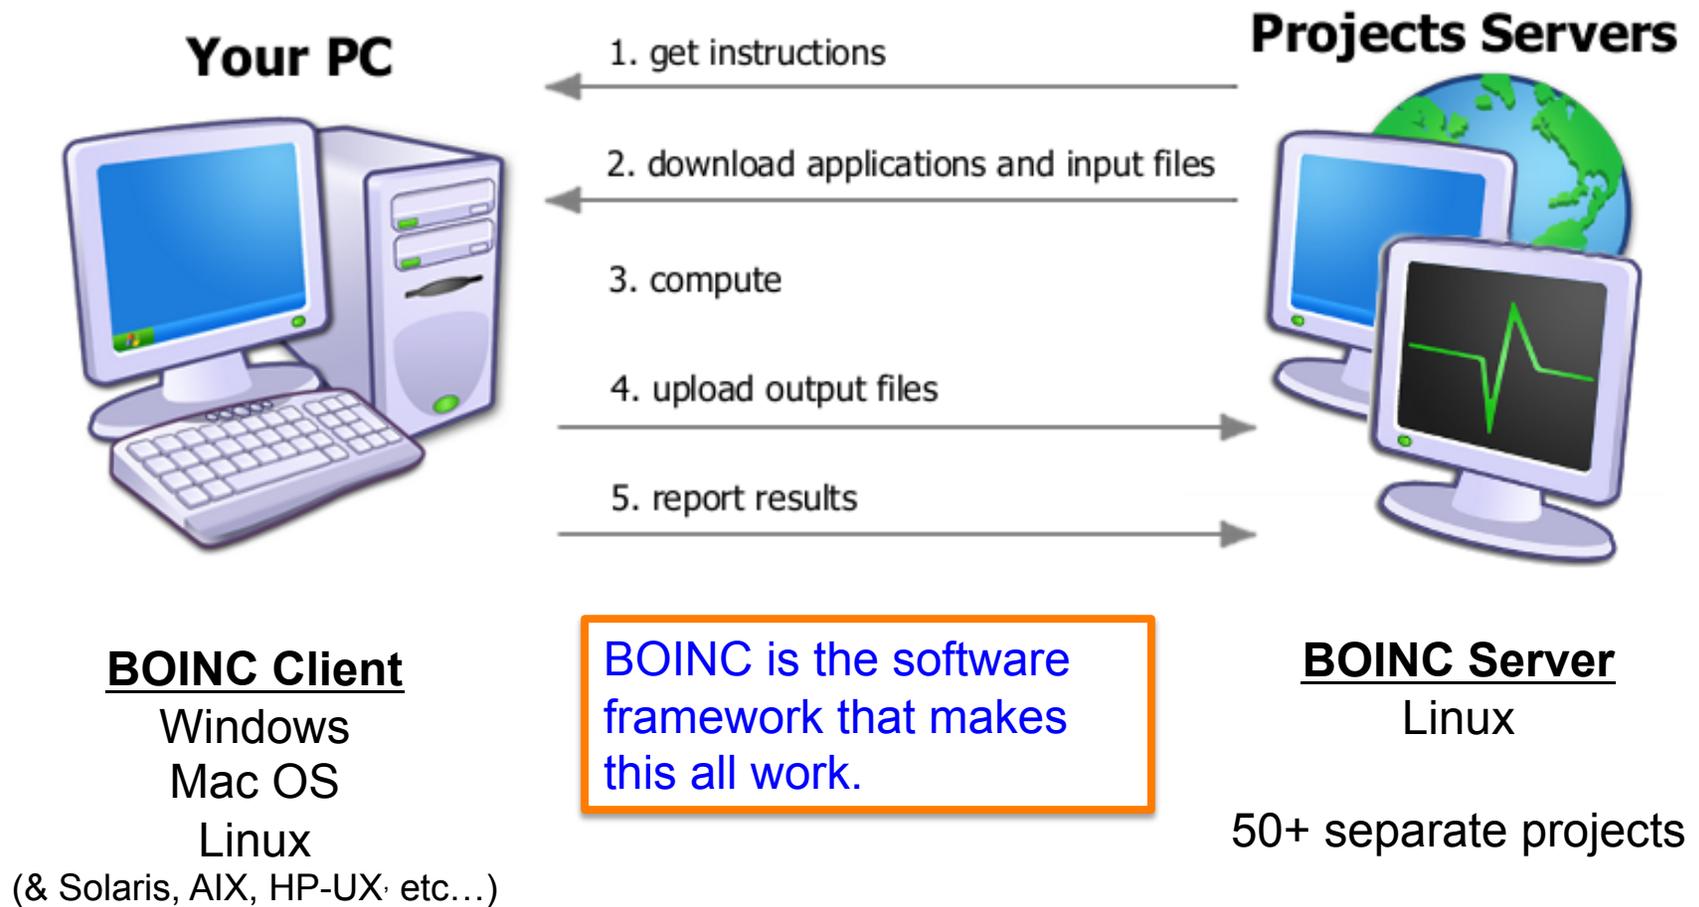
1 September 2010

*Mid-Hudson Linux Users Group*

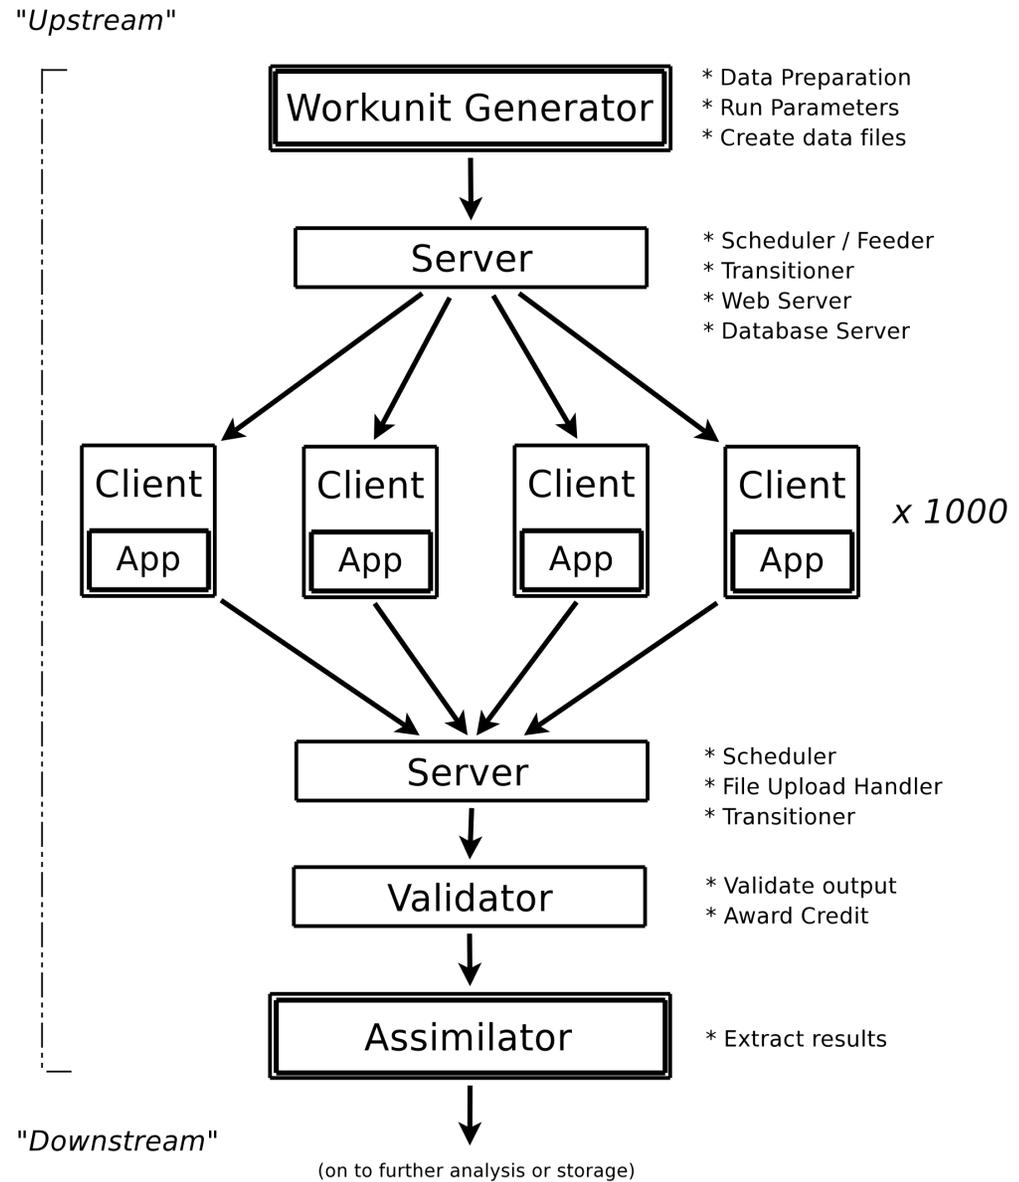
# SCIENTIFIC PROGRESS GOES "BOINC"



# How BOINC Works



# BOINC Dataflow



# Early History

## ***SETI@home***

- May 1999 to Dec 2005 (“Classic”)
- 2003 (BOINC)

## ***Climateprediction.net***

- Sept 2003 (“Classic”)
- August 2004 (BOINC)

## ***Predictor@Home***

- June 2004 - Scripps
- October 2008 – U. Michigan

## ***LHC@Home***

- Sept 2004 - CERN
- October 2007 - QMC

## ***World Community Grid*** (IBM)

- Nov 2004 to mid 2008 (Grid MP)
- Nov 2007 (BOINC)

## ***Einstein@Home***

- February 2005
- (*Pirates@Home* – June 2004 :-)

## ***Rosetta@Home*** (University of Washington)

- June 2005

## ***PrimeGrid*** (Lithuania)

- July 2005

With 50+ to follow...

# SETI@HOME Needs your Help

Donate to SETI@home

[Click Here for More Information](#)



## What is SETI@home?

SETI@home is a scientific experiment that uses Internet-connected computers in the Search for Extraterrestrial Intelligence (SETI). You can participate by running a free program that downloads and analyzes radio telescope data.

<http://setiathome.berkeley.edu>

### PARTICIPATE

Download  
Get help  
Tell a friend  
Donate  
Porting & optimization  
... more

### ABOUT

About SETI@home  
About Astropulse  
Science newsletters  
Technical news  
Server status  
Science status  
Sponsors  
... more

### COMMUNITY

Message boards  
Questions & answers  
Profiles  
User search  
Teams  
Web sites & IRC  
Pictures & music

### YOUR ACCOUNT

Your account  
Preferences  
Certificate

### STATISTICS

Top participants  
Top computers  
Top teams

Site search:



[Languages](#)

### Get started

**1** [Read our rules and policies](#)

**2** **Download, install and run** the BOINC software used by SETI@home. When prompted, enter the URL:  
**<http://setiathome.berkeley.edu>**

Have questions or need help? Contact a volunteer using [BOINC online help](#).

Special instructions:

- [For SETI@home Classic participants](#)
- [For users of command-line and pre-5.0 clients](#).

POWERED BY

Keep your computer busy when

### News

#### Weekly Outage and Initial Catch Up

Every Tuesday morning ([Pacific time](#)) we begin a 3 day data distribution outage to focus on science processing and development plus any needed systems maintenance. The upload/download servers will be offline during this time. The web site (including the forums) will only be offline during the database maintenance and backup portion of the outage. On Friday, you may experience connectivity issues as the servers catch up with demand. 6 Aug 2010 19:07:50 UTC

#### Don Backer has passed away.

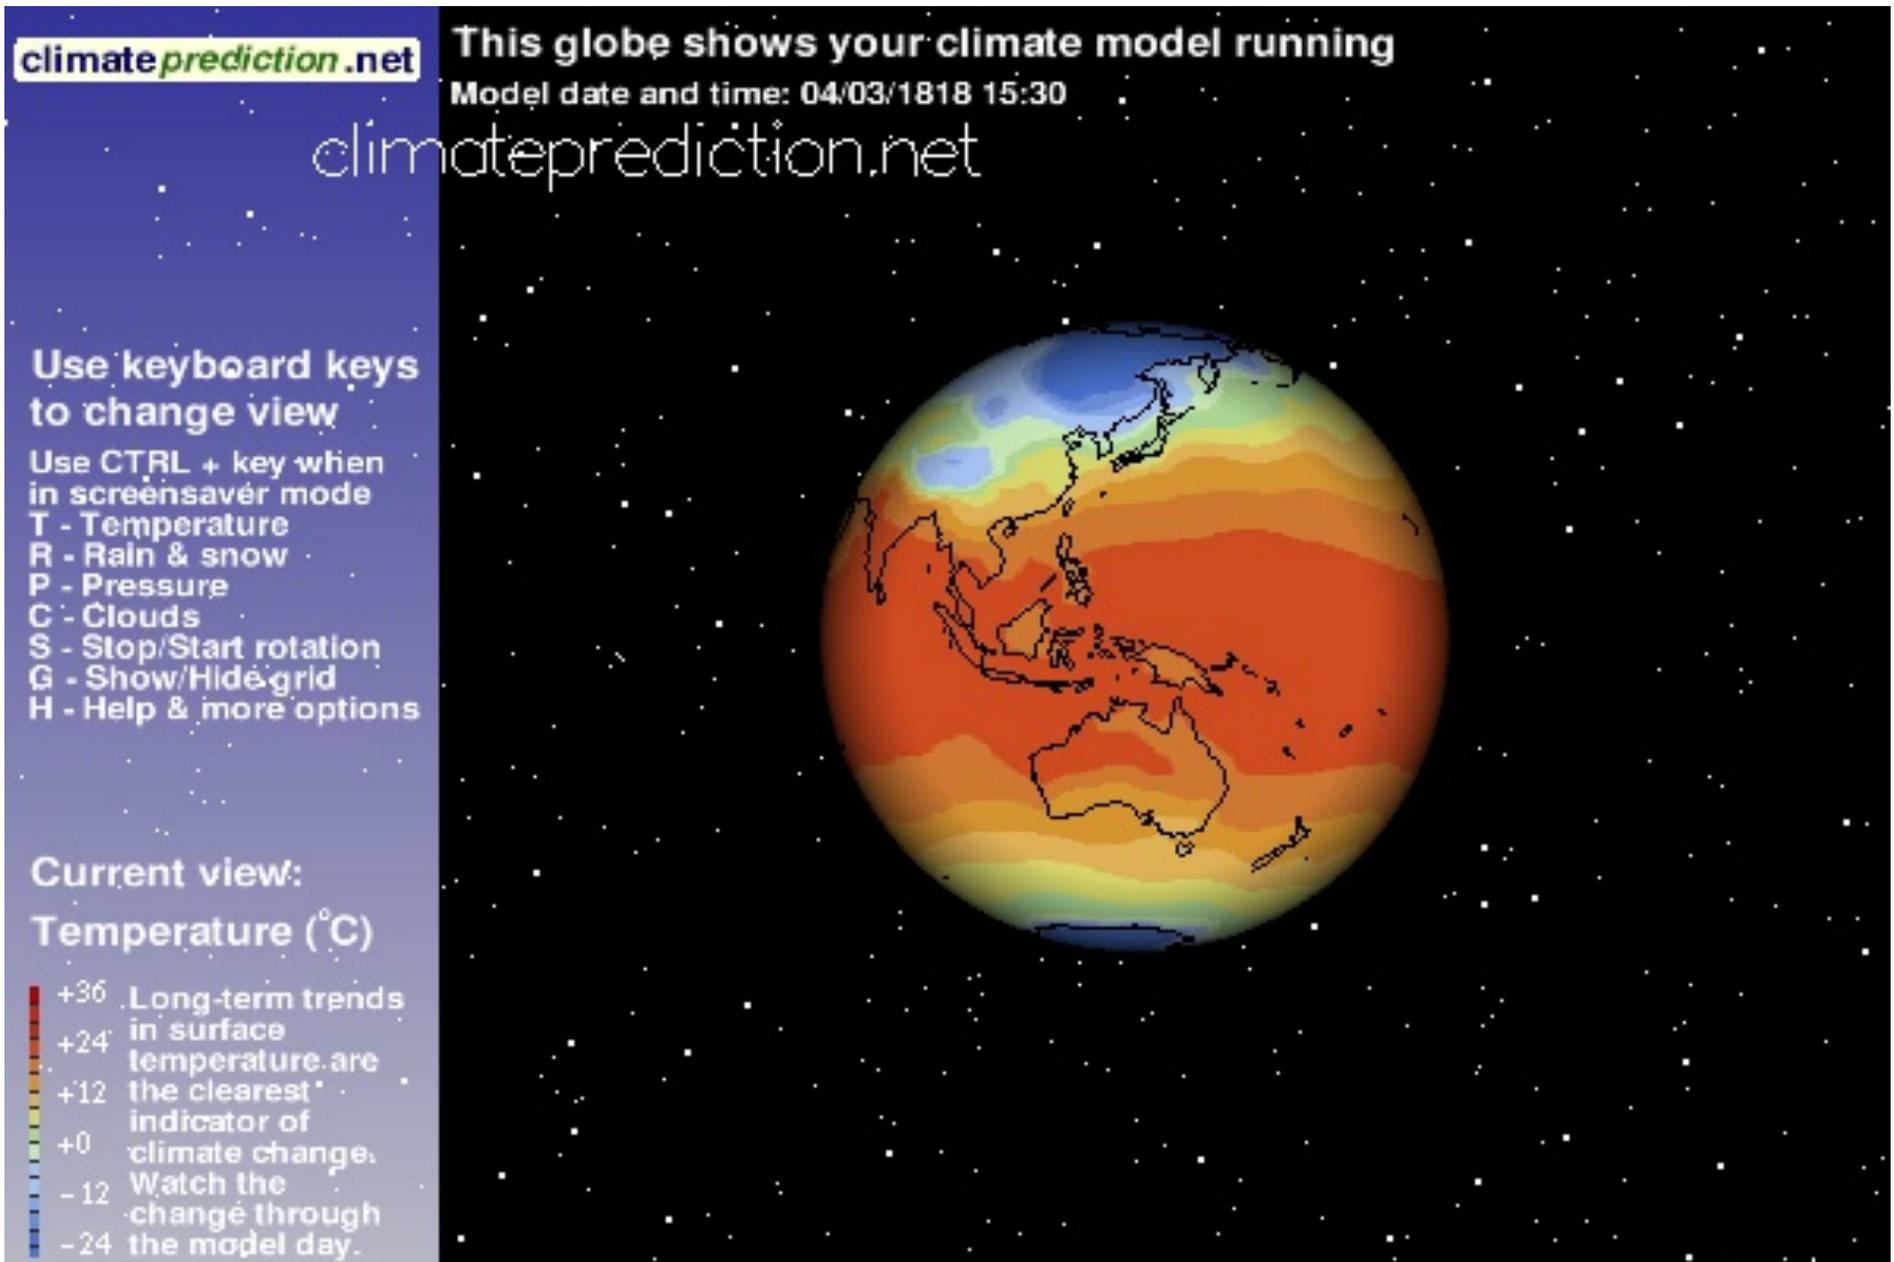
Don Backer, Berkeley professor and director of the Allen Telescope Array, passed away on Sunday, July 25th. Don was a very important contributor to radio astronomy and to the study of pulsars. He will be missed.

27 Jul 2010 0:36:12 UTC · [Comment](#)

#### Arecibo Observatory Repair Status

A structural failure at the telescope on February 3rd was partially repaired in March, and we've been observing with reduced motion since then. The next phase of repairs is slated to start July 12th and may take up to 6 weeks for full

<http://climateprediction.net/>



1 September 2010

Mid-Hudson Valley Linux Users Group

7

<http://www.worldcommunitygrid.org/>



**world community grid.**  
technology solving problems

## We're helping develop affordable solar energy.

**What are you doing?**

Why not donate your unused computer time to World Community Grid and The Clean Energy Project to help discover new materials that efficiently capture and store solar radiation as energy for later use?

> **Learn More**

**Join Today!**

The Clean Energy Project

### Active

- The Clean Energy Project - Phase 2
- Help Cure Muscular Dystrophy – Phase 2
- Help Fight Childhood Cancer
- Help Conquer Cancer
- Human Proteome Folding - Phase 2
- FightAIDS@Home

### Intermittent

- Discovering Dengue Drugs - Together - Phase 2
- Influenza Antiviral Drug Search
- The Clean Energy Project
- Discovering Dengue Drugs - Together

Funded and operated by IBM

### Completed

- Nutritious Rice for the World
- AfricanClimate@Home
- Help Cure Muscular Dystrophy
- Genome Comparison
- Help Defeat Cancer
- Human Proteome Folding

<http://einstein.phys.uwm.edu/>

or <http://einsteinathome.org>

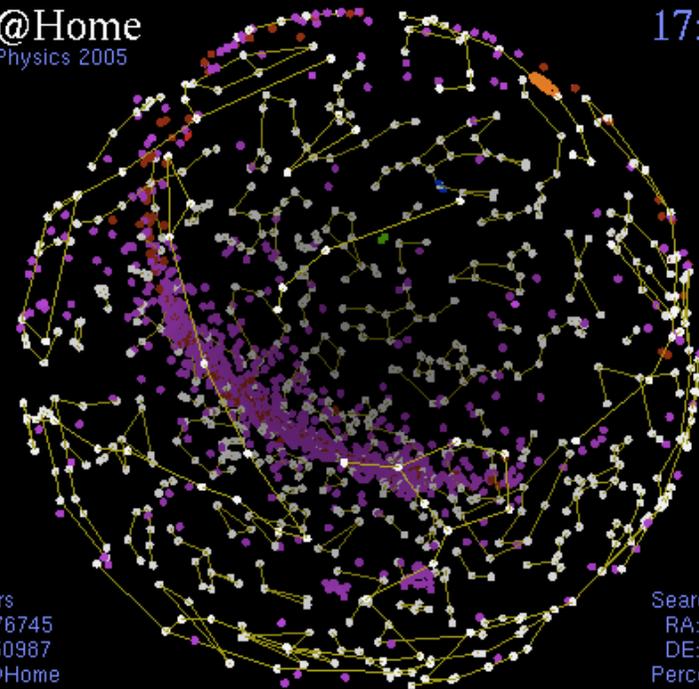


EINSTEIN@HOME

Einstein@Home  
World Year of Physics 2005

17:58:53

E  
ive



User: Eric Myers  
Total credit: 376745  
Host credit: 50987  
Team: Pirates@Home

Search information:  
RA: 58.67  
DE: 41.83  
Percent done: 0.200%

## About Einstein@Home

Thank you for your interest in Einstein@Home!

Einstein@Home is a program that uses your computer's idle time to search for gravitational waves from spinning neutron stars (also called pulsars) using data from the LIGO gravitational wave detector. Learn about this search at [einsteinathome.org](http://einsteinathome.org), [Einstein Online](#) and in our [S3 report](#).

Einstein@Home also searches for radio pulsars in binary systems, using data from the Arecibo Observatory in Puerto Rico. Read more about this search [here](#).

Einstein@Home is a World Year of Physics 2005 and an International Year of Astronomy 2009 project supported by the American Physical Society (APS) and by a number of international organizations.

If you would like to take part, please follow the "Join Einstein@Home" instructions to the left. Einstein@Home is available for Windows, Linux

## News

### First Einstein@Home Discovery!

**We are delighted to announce that Einstein@Home has made its first discovery: a radio pulsar, found in data from the Arecibo Observatory in Puerto Rico.**

Details are available in a paper published online by Science today. Science has given us permission to post a copy of the [abstract](#) and [paper](#) here. They are also on the Science website [here](#). The manuscript is also in the [arXiv preprint](#) archive (use the PDF link in the top right corner).

# Rules and Policies

## Run BOINC only on authorized computers

Run BOINC only on computers that you own, or for which you have obtained the owner's permission. Some companies and schools have policies that prohibit using their computers for BOINC.

## Is it safe to run BOINC?

Any time you download a program through the internet you are taking a chance: the program might have dangerous errors, or the download server might have been hacked. The well-known BOINC projects make efforts to minimize these risks. They test their applications and configure their servers to be secure.

The applications run by some BOINC projects may cause some computers to overheat. If that happens, stop running BOINC, or you can use the `cpulimit` utility to throttle CPU usage.

# Linux System Requirements

- Linux kernel 2.2.14 or higher
  - glibc 2.3.2 or higher
  - BOINC uses file locking. If you run BOINC in an NFS-mounted directory, make sure the NFS lock daemon is working.
- For BOINC Manager:
  - XFree86-3.3.6 or higher
  - gtk+2.0 or higher
- For application graphics:
  - Accelerated OpenGL support in your X server (version of OpenGL and graphics hardware required depends on the project).
- Hardware
  - Pentium 500 MHz or greater
  - 64 MB RAM
  - 50 MB disk space

Fedora 7+  
Debian 4.1+  
Ubuntu 7.10+

See project web sites for  
project-specific requirements!!

# Installing BOINC on Fedora

```
$ su root
Password:
# yum install boinc-client boinc-manager
# chkconfig boinc-client on
# service boinc-client start
# ^D
$ boincmgr
```

## **Where stuff ends up:**

```
/var/lib/boinc/...
/etc/rc.d/init.d/boinc-client
/etc/sysconfig/boinc-client
/etc/logrotate.d/boinc-client
/usr/bin/boinc_client , boinc , boinccmd
```

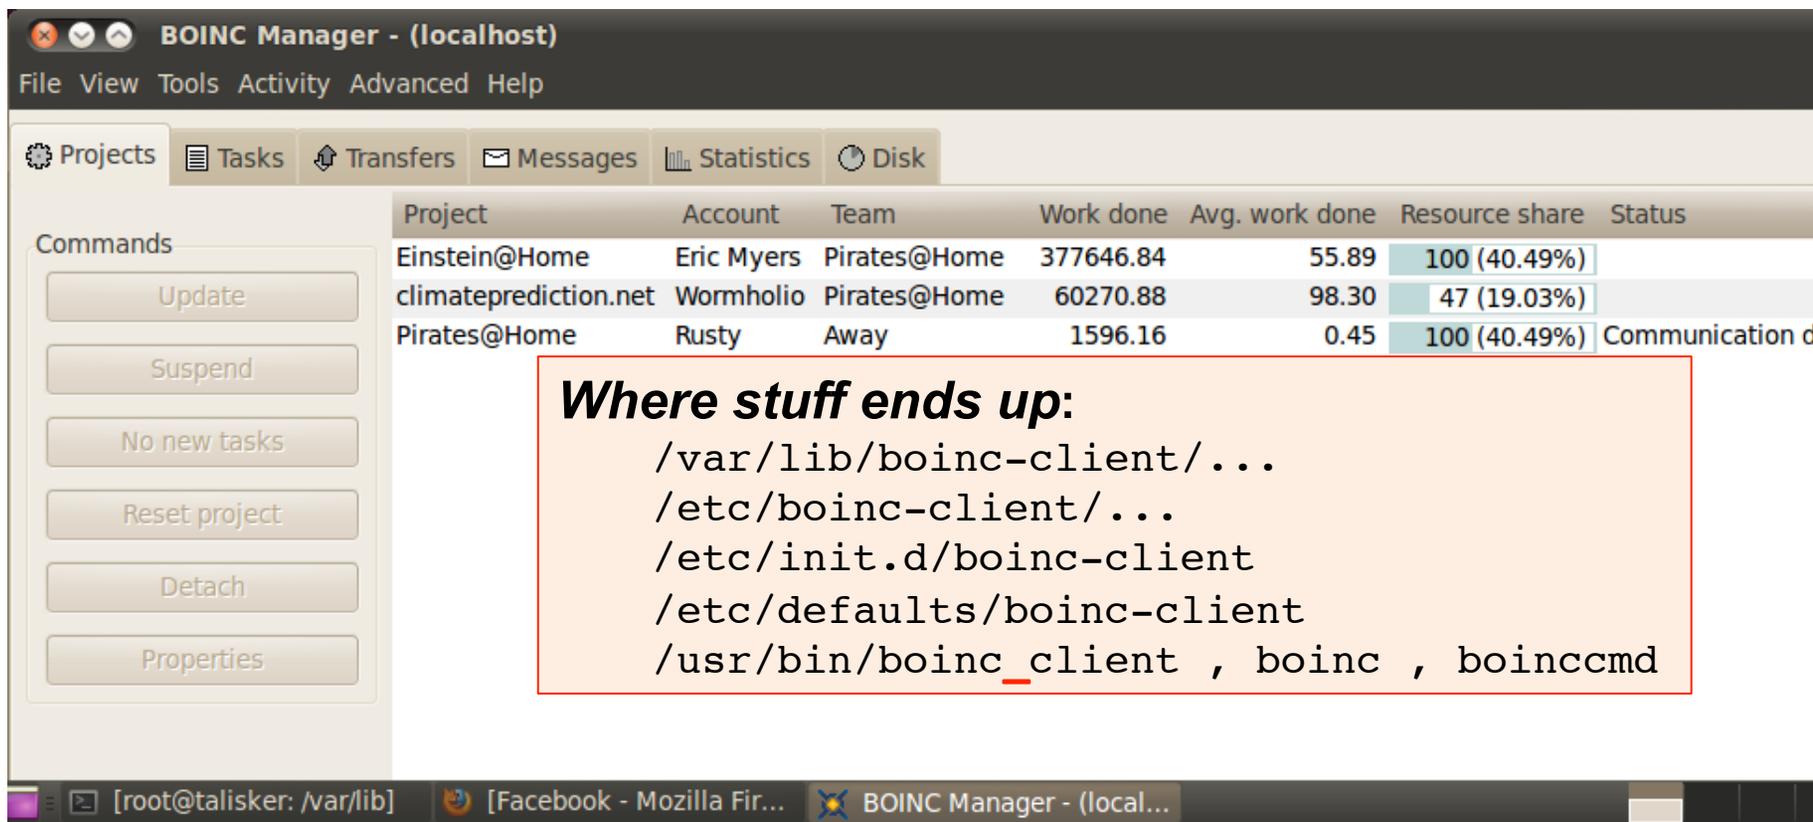
[http://boinc.berkeley.edu/wiki/Installing\\_BOINC\\_on\\_Fedora](http://boinc.berkeley.edu/wiki/Installing_BOINC_on_Fedora)

# Installing BOINC on Ubuntu

System > Administration > Synaptic Package Manager

- Search for “boinc”
  - For “boinc-client” and “boinc-manager” set “*Mark for installation*”
- ✓ Apply

Applications > System Tools > BOINC Manager



The screenshot shows the BOINC Manager application window. The title bar reads "BOINC Manager - (localhost)". The menu bar includes "File", "View", "Tools", "Activity", "Advanced", and "Help". The main window has a toolbar with "Projects", "Tasks", "Transfers", "Messages", "Statistics", and "Disk". On the left, there is a "Commands" panel with buttons for "Update", "Suspend", "No new tasks", "Reset project", "Detach", and "Properties". The main area displays a table of projects:

| Project               | Account    | Team         | Work done | Avg. work done | Resource share | Status          |
|-----------------------|------------|--------------|-----------|----------------|----------------|-----------------|
| Einstein@Home         | Eric Myers | Pirates@Home | 377646.84 | 55.89          | 100 (40.49%)   |                 |
| climateprediction.net | Wormholio  | Pirates@Home | 60270.88  | 98.30          | 47 (19.03%)    |                 |
| Pirates@Home          | Rusty      | Away         | 1596.16   | 0.45           | 100 (40.49%)   | Communication d |

Below the table, a red-bordered box contains the following text:

**Where stuff ends up:**  
/var/lib/boinc-client/...  
/etc/boinc-client/...  
/etc/init.d/boinc-client  
/etc/defaults/boinc-client  
/usr/bin/boinc\_client , boinc , boinccmd

The bottom of the screenshot shows the system tray with the terminal prompt [root@talisker: /var/lib], a Facebook browser window [Facebook - Mozilla Fir...], and the BOINC Manager window [BOINC Manager - (local...)].

# Installing BOINC from tarball

Use the *Berkeley Installer* for other Linux distributions, or to participate in **alpha testing**: <http://boinc.berkeley.edu/download.php>

```
$ mv ~/Desktop/boinc_6.10.56_i686-pc-linux-gnu.sh $HOME
$ cd
$ sh boinc_6.10.56_i686-pc-linux-gnu.sh
$ cd BOINC
$ ./run_client -daemon
$ ./run_manager
```

The disadvantages to using the Berkeley installer are:

- installs BOINC in a way that fails to take advantage of Linux security features.
- the BOINC binaries use more memory than they really need to because several shared libraries are statically linked.
- does not install BOINC as a daemon (though you can modify the installation)

# Building BOINC from source

```
$ cd /tmp
$ svn co http://boinc.berkeley.edu/svn/branches/boinc_core_release_6_10
$ cd boinc_core_release_6_10
$ ./_autosetup
Bootstrapping configure script and makefiles...

$ ./configure --prefix=/usr/local/boinc --disable-server
. . .
--- Configuring BOINC 6.10.58 (Release) ---
--- Build Components: (client manager libraries) ---
$ make
```

Building BOINC Manager  
requires wxWidgets 2.8.3  
and GTK 2

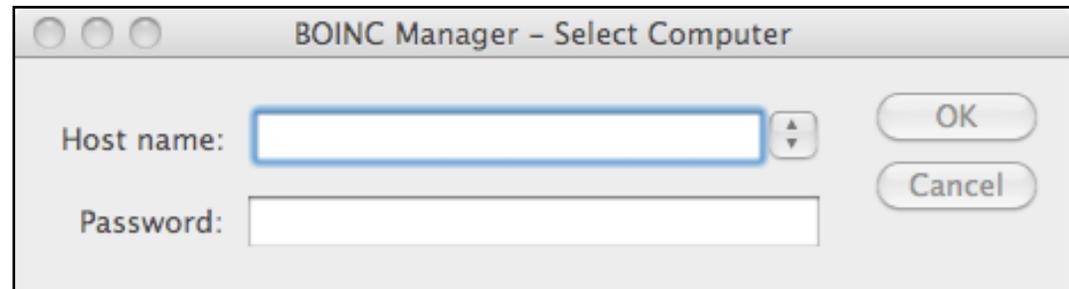
**You will find** boinc\_client **and** boinccmd **in the** client/ **subdirectory**

**You will find** boincmgr **in the** clientgui/ **subdirectory**

# Remote Control of Clients

The BOINC Manager on one host can connect to and control the BOINC core client (daemon) on another host. (RPC via port 31416)

BOINC Manager > Advanced > Select Computer ...



To enable this requires:

1. `boinc-client --allow_remote_gui_rpc`  
(or put selected hosts in the file `remote_hosts.cfg` )
2. Password in file `gui_rpc_auth.cfg`

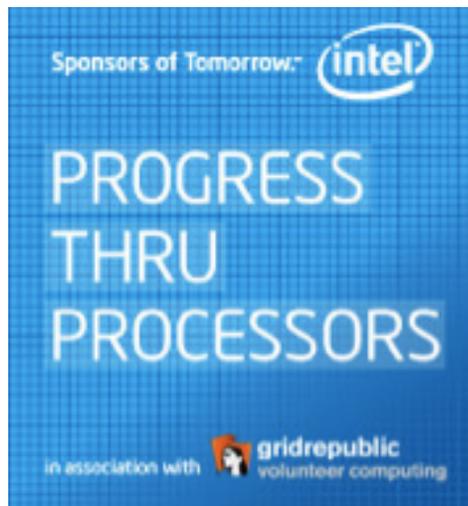
See [http://boinc.berkeley.edu/wiki/Controlling\\_BOINC\\_remotely](http://boinc.berkeley.edu/wiki/Controlling_BOINC_remotely)

# Account Managers

Select projects and control hosts, all from a single centralized web site.



<http://boincstats.com/bam/>



<http://www.facebook.com/progressthruprocessors>



<http://www.gridrepublic.org/>

# GPU Coprocessors

Some Graphics Processing Units (GPU's) can perform mathematical calculations at very high speed, and some BOINC applications can make use of the GPU as a co-processor. **These applications can run from 2X to 10X faster than the CPU-only version.**

- NVIDIA devices must support CUDA ([see list](#))
- ATI devices must be based on AMD R600 and AMD R700 or later.

Projects must provide a GPU version of their application.

## Projects with NVIDIA applications:

- GPUgrid.net
- SETI@home
- Milkyway@home
- AQUA@home
- Einstein@Home (GPU+CPU)
- Collatz Conjecture
- PrimeGrid
- DNETC@Home

Requires BOINC  
6.4.5 or newer

## Projects with ATI applications:

- Collatz Conjecture
- Milkyway@home
- DNETC@Home

Requires BOINC  
6.10.17 or newer

See [http://boinc.berkeley.edu/wiki/GPU\\_computing](http://boinc.berkeley.edu/wiki/GPU_computing)

# Credit and Teams



<http://boincstats.com/> (and see <http://boinc.berkeley.edu/links.php#stats>)

| Statistics             |  | Project list          |           |          |           |          |        |          |           |          |                 |              |                    |
|------------------------|--|-----------------------|-----------|----------|-----------|----------|--------|----------|-----------|----------|-----------------|--------------|--------------------|
| Active Projects:       |  | Projects              | Users     | last day | Hosts     | last day | Teams  | last day | Countries | last day | Total credit    | last day     | Last XML update    |
| ▶BOINC combined        |  | <b>BOINC combined</b> | 2,020,646 | +545     | 5,592,556 | +6,088   | 88,542 | +18      | 272       | 0        | 290,095,588,508 | +446,866,114 |                    |
| ▶ABC@home              |  | SETI@Home             | 1,118,704 | 0        | 2,711,035 | 0        | 58,114 | 0        | 234       | 0        | 81,065,992,099  | 0            | 08:11:23 old       |
| ▶AlmereGrid Boinc Grid |  | World Community Grid  | 304,149   | +133     | 1,014,303 | +760     | 17,858 | +8       | 218       | 0        | 29,387,500,724  | +41,331,723  | 04:57:05 old       |
| ▶AQUA@home             |  | Rosetta@Home          | 297,018   | +93      | 915,097   | +360     | 8,907  | +1       | 222       | 0        | 11,339,564,677  | +5,456,683   | 00:38:50 old       |
| ▶BOINC Alpha Test      |  | Einstein@Home         | 271,444   | +291     | 1,766,742 | +3,703   | 9,360  | +6       | 215       | +1       | 22,455,012,470  | +31,589,429  | 07:04:09 old       |
| ▶BURP                  |  | Climate Prediction    | 232,220   | 0        | 458,249   | +242     | 7,167  | 0        | 217       | 0        | 12,718,976,647  | 0            | 1 day 10:09:03 old |
| ▶CAS@HOME              |  | MilkyWay@home         | 70,379    | +101     | 137,166   | +114     | 2,134  | +3       | 181       | +1       | 44,962,756,186  | +140,505,413 | 01:09:16 old       |
| ▶Chess960@Home         |  | QMC@Home              | 43,105    | +9       | 105,005   | +42      | 1,983  | +1       | 172       | 0        | 3,108,551,824   | +2,305,091   | 00:59:24 old       |
| ▶Climate Prediction    |  | SIMAP                 | 34,342    | +18      | 104,457   | +60      | 1,942  | +2       | 174       | 0        | 1,138,935,839   | +3,754,430   | 03:34:23 old       |
| ▶Collatz Conjecture    |  | PrimeGrid             | 33,759    | +27      | 105,107   | +56      | 1,784  | +2       | 172       | 0        | 3,728,979,191   | +10,126,291  | 02:39:25 old       |
| ▶Cosmology@Home        |  | Docking@Home          | 19,140    | +23      | 49,290    | +46      | 736    | +3       | 127       | 0        | 1,139,636,526   | +2,482,224   | 01:14:19 old       |
| ▶DistributedDataMining |  | AQUA@home             | 15,408    | +38      | 29,743    | +81      | 766    | 0        | 138       | +1       | 6,832,649,340   | +34,309,707  | 05:19:25 old       |
| ▶DistrRTgen            |  | Collatz Conjecture    | 11,870    | +47      | 25,887    | +88      | 773    | +3       | 123       | +1       | 28,752,639,392  | +98,929,912  | 05:29:26 old       |
| ▶DNETC@HOME            |  | GPUGRID               | 9,666     | +12      | 17,005    | +16      | 713    | +1       | 110       | 0        | 10,494,044,074  | +20,920,748  | 03:34:25 old       |
| ▶Docking@Home          |  | FreeHAL               | 7,976     | +33      | 32,463    | +58      | 478    | +2       | 110       | 0        | 2,006,400,969   | +2,921,953   |                    |
| ▶DrugDiscovery@Home    |  | DNETC@HOME            | 2,335     | +3       | 5,738     | +20      | 316    | 0        | 91        | 0        | 14,844,139,352  | +38,963,151  |                    |
| ▶DynaPing              |  |                       |           |          |           |          |        |          |           |          |                 |              |                    |

This list only shows the 15 most active projects. Go [here](#) to see the full list.



**Jakob Nielsen's Alertbox, October 9, 2006:**

## Participation Inequality: Encouraging More Users to Contribute

### Summary:

In most online communities, 90% of users are lurkers who never contribute, 9% of users contribute a little, and 1% of users account for almost all the action.

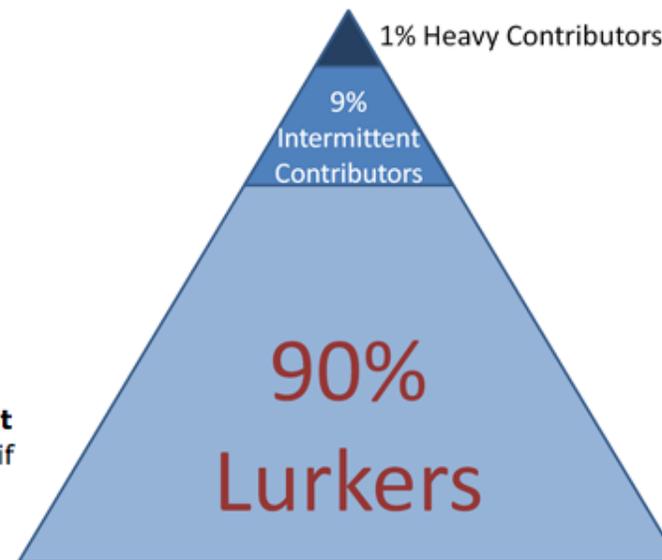
All large-scale, multi-user communities and online social networks that rely on users to contribute content or build services share one property: **most users don't participate** very much. Often, they simply **lurk** in the background.

In contrast, a tiny minority of users usually accounts for a disproportionately large amount of the content and other system activity. This phenomenon of **participation inequality** was first studied in depth by Will Hill in the early '90s, when he worked down the hall from me at Bell Communications Research (see references below).

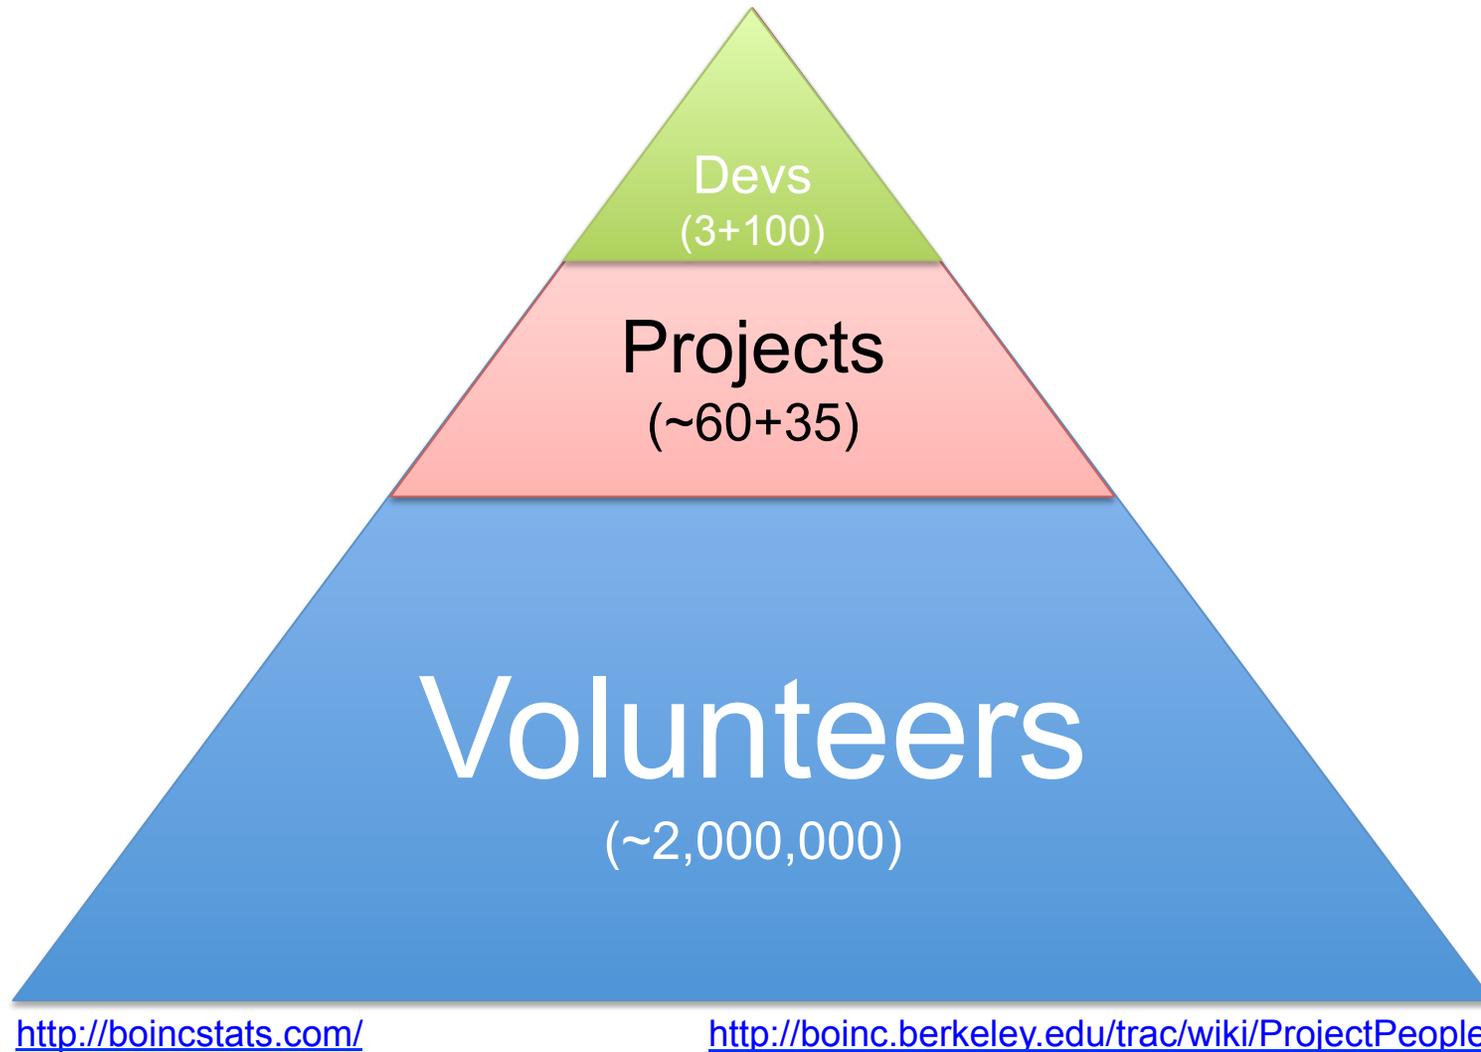
When you plot the amount of activity for each user, the result is a [Zipf curve](#), which shows as a straight line in a [log-log diagram](#).

User participation often more or less follows a **90-9-1 rule**:

- **90%** of users are **lurkers** (i.e., read or observe, but don't contribute).
- **9%** of users contribute **from time to time**, but other priorities dominate their time.
- **1%** of users participate a lot and **account for most contributions**: it can seem as if they don't have lives because they often post just minutes after whatever event they're commenting on occurs.



# Amplification of Effort





Open-source software for **volunteer computing** and **grid computing**.



-- language --

Search

### Volunteer

[Download](#) · [Help](#) · [Documentation](#)

Use the idle time on your computer (Windows, Mac, or Linux) to cure diseases, study global warming, discover pulsars, and do many other types of scientific research. It's safe, secure, and easy:

1. **Choose** projects
2. **Download** and run BOINC software
3. **Enter** an email address and password.

Or, if you run several projects, try an [account manager](#) such as [GridRepublic](#) or [BAM!](#).

### Compute with BOINC

[Documentation](#) · [Software updates](#)

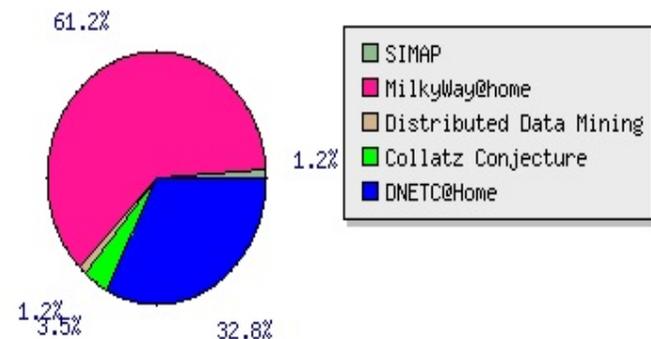
- **Scientists:** use BOINC to create a [volunteer computing project](#) giving you the computing power of thousands of CPUs.
- **Universities:** use BOINC to create a [Virtual Campus Supercomputing Center](#).
- **Companies:** use BOINC for [desktop Grid computing](#).

### Computing power

[Top 100 volunteers](#) · [Statistics](#)

Active: 294,457 volunteers, 526,048 computers.  
24-hour average: 2,943.76 TeraFLOPS.

[Silent M??bius](#) is contributing 5,703 GFLOPS.  
Country: Germany; Team: SETI.Germany



### News

#### Einstein@Home pulsar discovery

The recent discovery of a new pulsar by [Einstein@Home](#) is described in an [article in Science](#) and on the [National Science Foundation web site](#).

14 Aug 2010 14:01:15 UTC · [Comment](#)

# Documentation and Help

- BOINC User Manual: [http://boinc.berkeley.edu/wiki/User\\_manual](http://boinc.berkeley.edu/wiki/User_manual)
- BOINC FAQ Service: <http://boincfaq.mundayweb.com/>
- The Unofficial BOINC wiki: <http://www.boinc-wiki.info/>
- BOINC Forums: <http://boinc.berkeley.edu/dev/>
- Online Help via Skype (or text): <http://boinc.berkeley.edu/help.php>

...and every project has on-line forums, for both  
project help and BOINC help

Project Creation/Software Development: <http://boinc.berkeley.edu/trac/wiki>

# Closing points...

- **BOINC is the software that makes distributed volunteer computing work**
- BOINC projects are separate from The BOINC Project (though many projects cooperate together)
- **BOINC is now fairly easy to install and manage on Linux (still no screensaver!)**
- The “spare” cycles which volunteers donate can add up to real research results.

# Creating a BOINC project

1. Choose a server (Linux, but which distro?)
2. Obtain required Software (`server_stable` via SVN)
3. Set up database server (MySQL)
4. Build the BOINC server software
5. Run the `make_project` script
6. Set up the web server (Apache)
7. Adjust and check the configuration
  - a) File Permissions
  - b) Database passwords
  - c) `config.xml`
  - d) Scheduler
  - e) Feeder
  - f) periodic tasks
8. *Create and deploy your application...*
9. Customize your web site
  - a) Project Home Page
  - b) Discussion Forums
  - c) RSS News
  - d) Client web links (`gui_urls.xml`)
  - e) Simple GUI icons

*It's just so easy!*

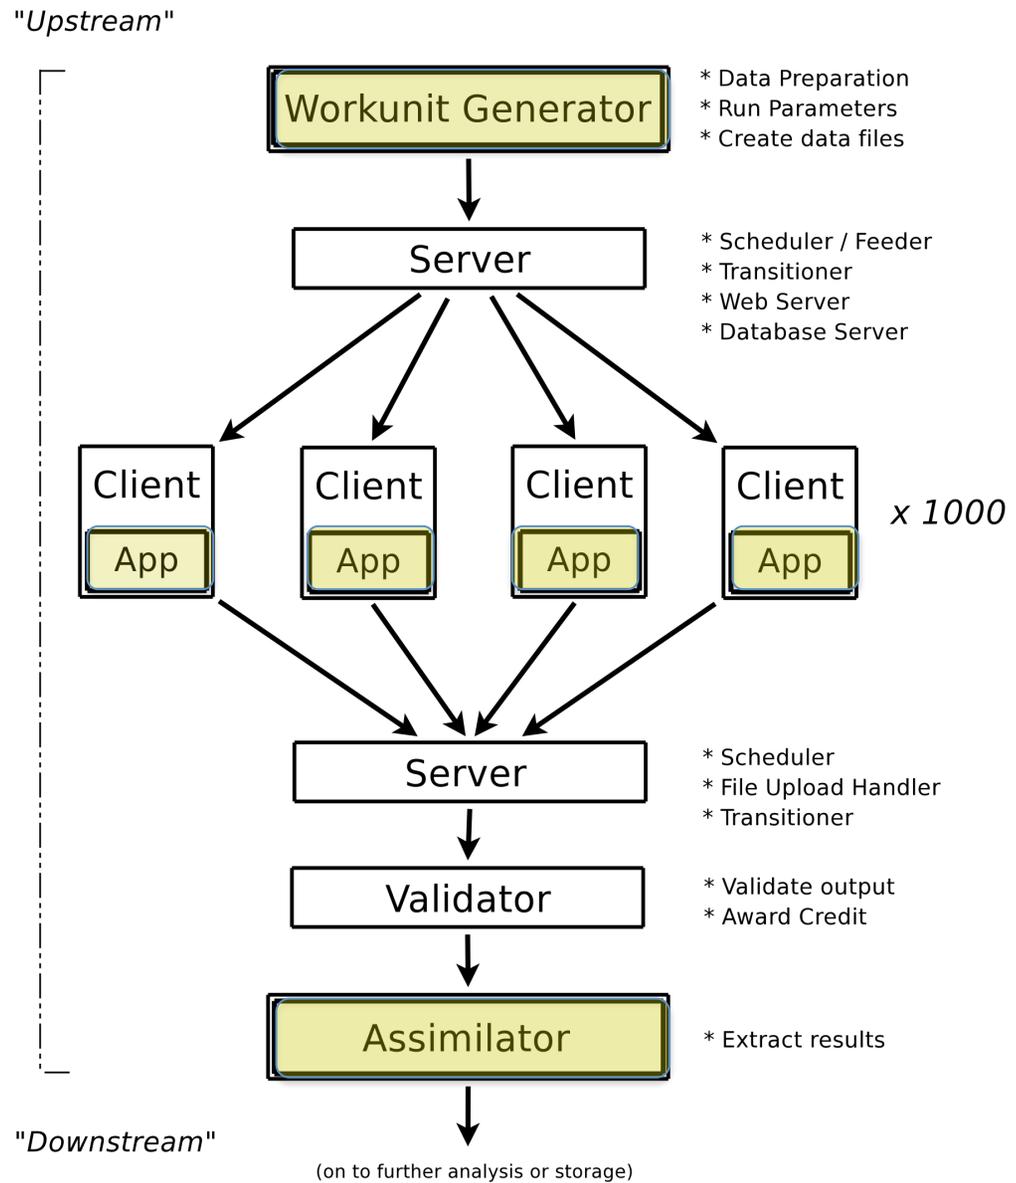


Quick-start alternative:  
[BOINC server Virtual Machine](#)

<http://www.spy-hill.net/~myers/help/boinc/>

<http://boinc.berkeley.edu/trac/wiki/ProjectMain>

# BOINC Dataflow



# Creating BOINC Applications

See <http://boinc.berkeley.edu/trac/wiki/ProjectMain>  
and [http://www.spy-hill.net/~myers/help/boinc/BOINC\\_Apps.html](http://www.spy-hill.net/~myers/help/boinc/BOINC_Apps.html)

## Sample Applications

- `hello` - Hello, World -- the simplest BOINC program
- `concat` - concatenates two or more input files into one output file (from BOINC)
- `yello` - Yello, World! -- the simplest graphics program for BOINC,
- `cube` - simplest non-trivial 3D graphics application for BOINC
- `uppercase` - reads text from a file, converting it to upper case. (from BOINC)  
(Do not confuse this with `boinc/apps/upper_case`, which is a bare-bones application which is used by the `test_uc.py` script.)
- `lalanne` - The Jack LaLanne program - an API Exerciser
- `sextant` - the Einstein@Home screensaver graphics with some fancy additions
- `starboard` - Starboard! - xscreensaver GL graphics suite