

# Release Management and New Capabilities in Chirp

Patrick Donnelly, University of Notre Dame  
CCL Workshop, June 2012

# Release Wrangler



# What is CCTools?

Collection of Distributed Computing Tools

- AllPairs/Wavefront
- Chirp/Parrot
- Makeflow/WorkQueue
- Sand

# Previous Release Methodology

- Tagged /trunk every few months.
- Contained features and bug fixes.
- Version number bumps mostly arbitrary.
- 3.4.2 --> MAJOR.MINOR.RELEASE

# New CTools Release Philosophy

- Versions are now MAJOR.MINOR.  
REVISION
- REVISION changes may not break  
API/ABI compatibility --> Bug Fix

# Autobuild

Latest builds from:

trunk

## Tarballs:

- [cctools-2116-i386-darwin-8.11.tar.gz](#)
- [cctools-2106-i686-redhat5.tar.gz](#)
- [cctools-2116-i686-redhat6.tar.gz](#)
- [cctools-2116-source.tar.gz](#)
- [cctools-2116-x86\\_64-redhat5.tar.gz](#)
- [cctools-2116-x86\\_64-redhat6.tar.gz](#)

## Revisions

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 Next

Id	Author	Message	Branch	Commit Time	Tarballs	Logs
2116	dpandiar	a. Fix typo. b. Add reference to Python API.	trunk	Thu, Jun 07, 2012 04:31 PM	<ul style="list-style-type: none"> <li>• <a href="#">i386-darwin-8.11</a></li> <li>• <a href="#">i686-redhat5 (timeout)</a></li> <li>• <a href="#">i686-redhat6</a></li> <li>• <a href="#">source</a></li> <li>• <a href="#">x86_64-redhat5</a></li> <li>• <a href="#">x86_64-redhat6</a></li> </ul>	<ul style="list-style-type: none"> <li>• <a href="#">summary</a></li> <li>• <a href="#">condor.submit</a></li> <li>• <a href="#">condor.log</a></li> <li>• <a href="#">i386-darwin-8.11</a> <ul style="list-style-type: none"> <li>◦ <a href="#">build.summary</a></li> <li>◦ <a href="#">build.log</a></li> </ul> </li> <li>• <a href="#">i686-redhat5</a> <ul style="list-style-type: none"> <li>◦ <a href="#">build.summary</a></li> </ul> </li> <li>• <a href="#">i686-redhat6</a> <ul style="list-style-type: none"> <li>◦ <a href="#">build.summary</a></li> <li>◦ <a href="#">build.log</a></li> </ul> </li> <li>• <a href="#">source</a> <ul style="list-style-type: none"> <li>◦ <a href="#">build.summary</a></li> <li>◦ <a href="#">build.log</a></li> </ul> </li> <li>• <a href="#">x86_64-redhat5</a> <ul style="list-style-type: none"> <li>◦ <a href="#">build.summary</a></li> <li>◦ <a href="#">build.log</a></li> </ul> </li> <li>• <a href="#">x86_64-redhat6</a> <ul style="list-style-type: none"> <li>◦ <a href="#">build.summary</a></li> <li>◦ <a href="#">build.log</a></li> </ul> </li> </ul>
2115	dpandiar	Update with the -j option that submits workers as a job array.	trunk	Thu, Jun 07, 2012 03:27 PM	<ul style="list-style-type: none"> <li>• <a href="#">i386-darwin-8.11</a></li> <li>• <a href="#">i686-redhat5 (timeout)</a></li> <li>• <a href="#">i686-redhat6</a></li> <li>• <a href="#">source</a></li> <li>• <a href="#">x86_64-redhat5</a></li> </ul>	<ul style="list-style-type: none"> <li>• <a href="#">summary</a></li> <li>• <a href="#">condor.submit</a></li> <li>• <a href="#">condor.log</a></li> <li>• <a href="#">i386-darwin-8.11</a> <ul style="list-style-type: none"> <li>◦ <a href="#">build.summary</a></li> <li>◦ <a href="#">build.log</a></li> </ul> </li> </ul>



Sign in

[Join Google+](#) Share the right things with just the right people.

## Patrick Donnelly



**Patrick Donnelly** Apr 17, 2012 - Public

My reign of terror as Release Wrangler is about to begin!

Thanks to [+Michael Albrecht](#), [+Dinesh Rajan](#), and [+Peter Bui](#) for this awesome Birthday Gift!



- Works at **University of Notre Dame**
- Attends **University of Notre Dame**
- Lived in **Henderson, Nevada**

[View full profile](#)

Report this profile



**Peter Bui** Apr 17, 2012  
I eagerly await autobuild integration. Good thing I rarely commit to trunk anymore...

+2

## [The Cooperative Computing Lab](#)

### Research

[Projects](#) [Papers](#)[People](#) [Jobs](#)[Events](#) [REU](#)[News](#) [Blog](#)

### Software

[Chirp](#) [Downloads](#)[Parrot](#) [Mailing List](#)[Makeflow](#) [Manuals](#)[Work Q](#) [HOWTO](#)[SAND](#) [API](#)[AllPairs](#) [SVN](#)[Wavefront](#)

### Operations

[Visual System Status](#)[Condor](#) [BXGrid](#)[Chirp](#) [Biocompute](#)[Hadoop](#) [DISC](#)[CertAuth](#) [CondorLog](#)[Internal Docs](#)  

## Downloads

The cctools package contains [Parrot](#), [Chirp](#), [Makeflow](#), [Work Queue](#), [SAND](#), and other software. This software is Copyright (C) 2004-2011 The University of Notre Dame and licensed via the [GNU General Public License V2](#).

[Report bugs or other problems to our mailing list.](#)

[Users at Notre Dame click here for the local installation.](#)

### Download Stable Version 3.5.0

<a href="#">cctools-3.5.0-src.tar.gz</a>	Source Code
<a href="#">cctools-3.5.0-x86_64-redhat6.tar.gz</a>	Red Hat Enterprise Linux 6, x86 CPU (64 bit)
<a href="#">cctools-3.5.0-x86_64-redhat5.tar.gz</a>	Red Hat Enterprise Linux 5, x86 CPU (64 bit)
<a href="#">cctools-3.5.0-i686-redhat5.tar.gz</a>	Red Hat Enterprise Linux 5, x86 CPU (32 bit)
<a href="#">cctools-3.5.0-i386-darwin-8.11.tar.gz</a>	MacOS 10.4 "Tiger", x86 CPU
<a href="#">cctools-3.5.0-i686-windows.tar.gz</a>	Windows 7/Vista/XP, x86 CPU (requires Cygwin installation)

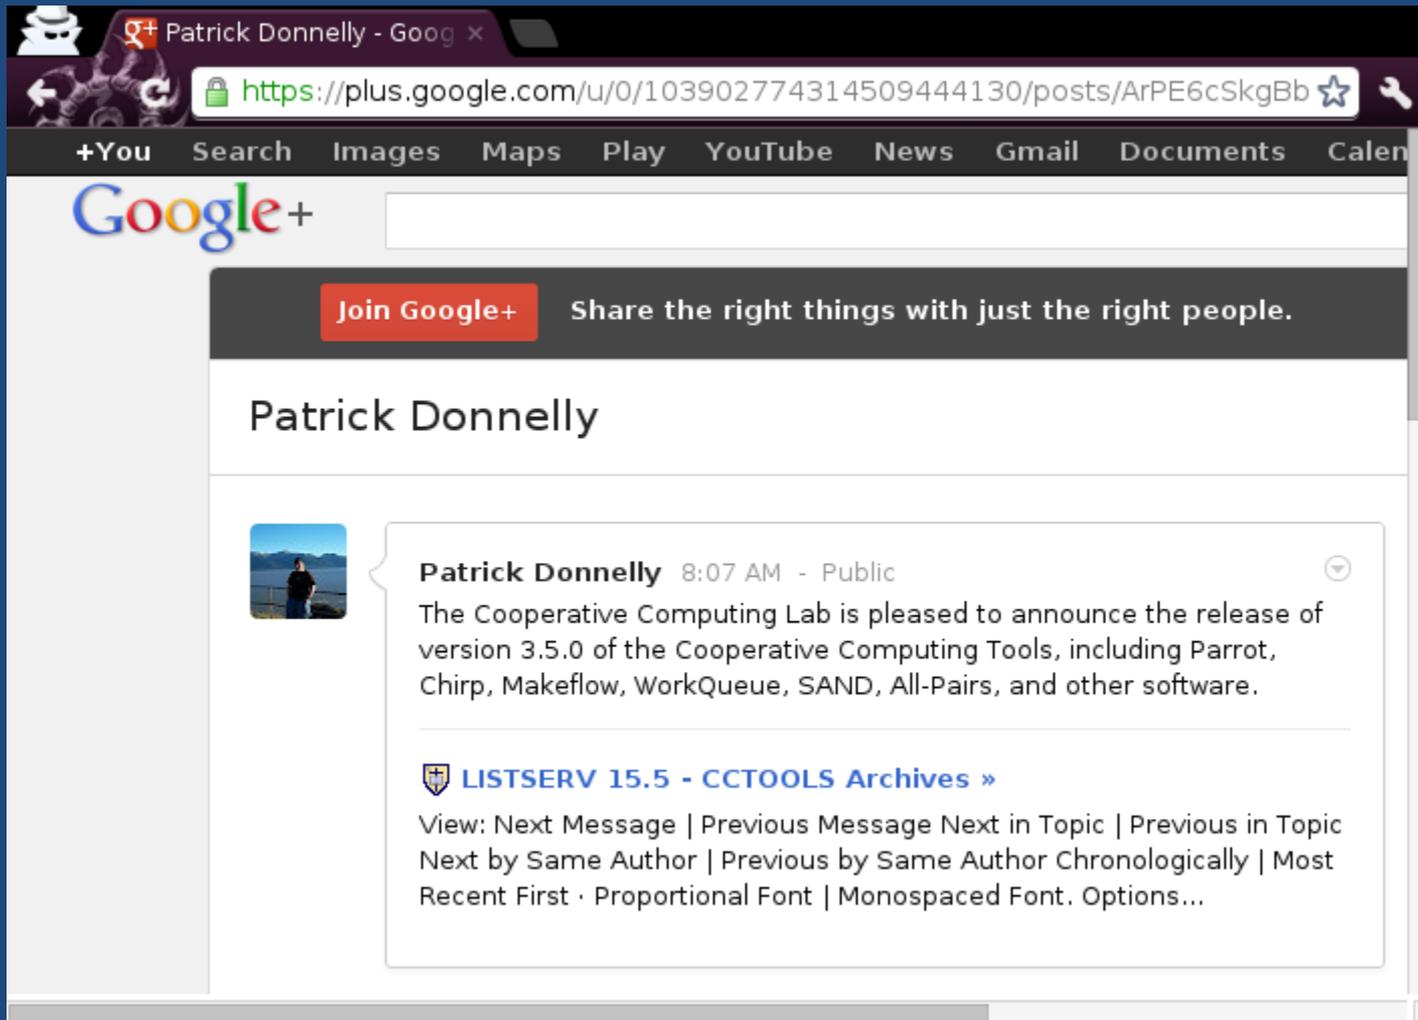
### Download Current Unstable Version:

<a href="#">cctools-current-src.tar.gz</a>	Source Code
<a href="#">cctools-current-x86_64-redhat5.tar.gz</a>	Red Hat Enterprise Linux 5, x86 CPU (64 bit)
<a href="#">cctools-current-i686-redhat5.tar.gz</a>	Red Hat Enterprise Linux 5, x86 CPU (32 bit)

### Software Engineering Tools

- [CCTools Autobuild System](#)
- [CCTools SVN Repository Browser](#)
- [Download Older Versions](#)

# Outreach



A screenshot of a web browser displaying a Google+ post. The browser's address bar shows the URL: <https://plus.google.com/u/0/103902774314509444130/posts/ArPE6cSkgBb>. The browser's navigation bar includes links for +You, Search, Images, Maps, Play, YouTube, News, Gmail, Documents, and Calendar. The Google+ logo is visible in the top left of the page content. Below the logo is a dark grey banner with a red button labeled "Join Google+" and the text "Share the right things with just the right people." The profile name "Patrick Donnelly" is displayed below the banner. The main content area shows a post by Patrick Donnelly, dated 8:07 AM, with a public visibility setting. The post text reads: "The Cooperative Computing Lab is pleased to announce the release of version 3.5.0 of the Cooperative Computing Tools, including Parrot, Chirp, Makeflow, WorkQueue, SAND, All-Pairs, and other software." Below the text is a link: "LISTSERV 15.5 - CCTOOLS Archives »". At the bottom of the post, there are navigation options: "View: Next Message | Previous Message Next in Topic | Previous in Topic Next by Same Author | Previous by Same Author Chronologically | Most Recent First · Proportional Font | Monospaced Font. Options..."

Patrick Donnelly - Goog x

<https://plus.google.com/u/0/103902774314509444130/posts/ArPE6cSkgBb>

+You Search Images Maps Play YouTube News Gmail Documents Calen

Google+

Join Google+ Share the right things with just the right people.

Patrick Donnelly

 **Patrick Donnelly** 8:07 AM - Public

The Cooperative Computing Lab is pleased to announce the release of version 3.5.0 of the Cooperative Computing Tools, including Parrot, Chirp, Makeflow, WorkQueue, SAND, All-Pairs, and other software.

 [LISTSERV 15.5 - CCTOOLS Archives »](#)

View: Next Message | Previous Message Next in Topic | Previous in Topic Next by Same Author | Previous by Same Author Chronologically | Most Recent First · Proportional Font | Monospaced Font. Options...



Projects / Makeflow

## Makeflow

Makeflow is a workflow engine for executing large complex applications on clusters, clouds, and grids. It can be used to drive several different distributed computing systems, including Condor, SGE, and the included Work Queue system. It does not require a distributed filesystem, so you can use it to harness whatever collection of machines you have available. It is typically used for scaling up data-intensive scientific applications to hundreds or thousands of cores.

- Tags: [Distributed Computing](#) [cloud](#) [Cluster](#) [GRID](#) [Workflow](#) [Scientific/Engineering](#) [Clustering/Distributed Networks](#)
- Licenses: [GPLv2](#)
- Operating Systems: [Linux](#) [Mac OS](#) [Solaris](#) [Windows \(Cygwin\)](#)
- Implementation: [C](#)
- Translations: [English](#)

[Website](#) [Download](#)

[Tweet](#) 2 [+1](#) 0

### Recent releases

[All releases](#) [Release tags](#)

3.5.0 11 Jun 2012 21:37

Minor [workqueue](#)

**Release Notes:** Numerous improvements have been made to WorkQueue to improve support for various batch job submission platforms. WorkQueue now allows extensive configuration and automation... [\(more\)](#)

### Links

[Download](#) [Website](#)



dthain

22 Jul 2010 15:34



- Dependencies
- Request ownership
- Report problem
- Graphs
- Submit a comment

[ignore](#)

[follow](#)

### Heartbeat



Popularity 231.15

Score

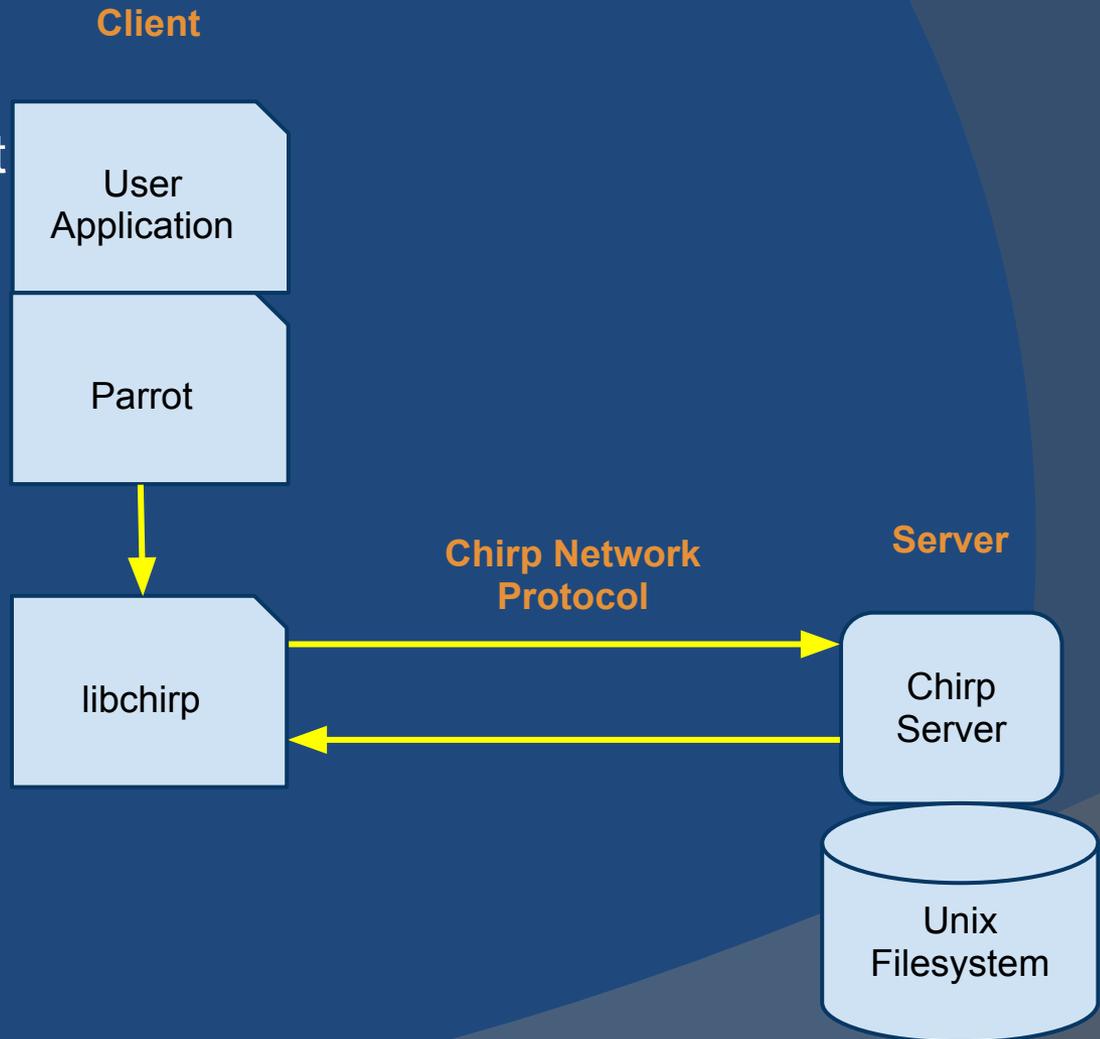
Vitality Score 117.54

Subscriptions 22

# New Capabilities in Chirp...

# Chirp

- Distributed File System for use on a Grid
- Exports file system on host
- User-level filesystem
- Authentication mechanisms
  - Grid Security Infrastructure
  - Kerberos
  - Hostnames
  - Unix
  - Tickets (new!)
- Secure Authorization through Access Control Lists



# Chirp with Backend Storage

- Supports a filesystem abstraction layer.
- Currently Available:
  - Hadoop
  - Locally mounted (Unix)
  - Chirp (Chirp mounting Chirp; how perverse)
- Why?

# Hadoop Distributed Filesystem

- Java open source implementation of the concepts in the Google File System.
- Offers very large file storage on the order of terabytes.
- Replicated file storage.
- Active Storage and Map-Reduce.
- Streaming data access.

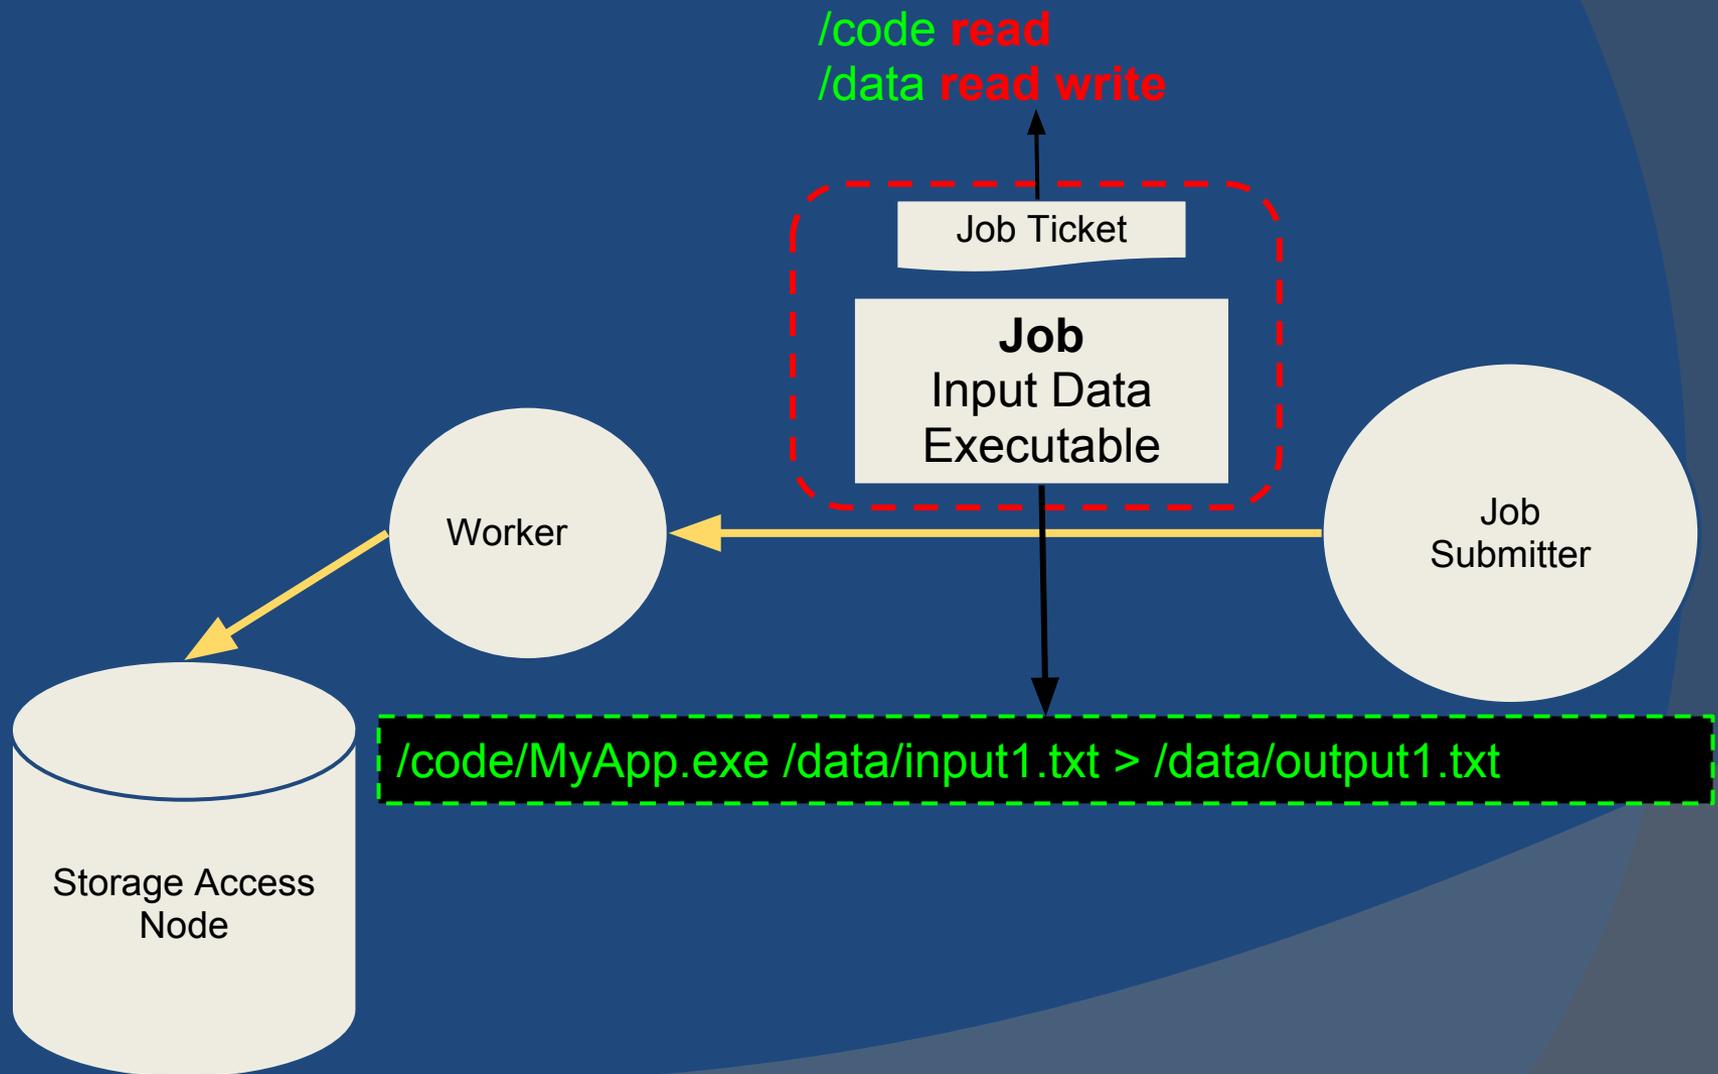
# Problems with using Hadoop on a Grid

- Java API/C API
- Carry around JVM +- FUSE
- No authenticated access
- Tight coupling of JVM with Hadoop versions

# Authentication in Chirp

- Globus
- Kerberos
- Hostname/IP
- Unix (local)
- **Tickets (new!)**

# What are Tickets for?



# What we want in a Ticket

- Available **secure** authentication mechanisms rely heavily on user interaction.
- **Desire:**
  - "Token" passed along with Job for authentication.
  - Temporary lifetime
  - Simple to setup
  - Reasonable security; minimal risk

# Authentication Complications

- Users and machines do not have dedicated credentials with which to sign a token.
- Chirp servers and workers are transient.

# Ticket Authentication

- Using Public Key Encryption, setup a ticket which is a private key credential for user Jobs.
- Authenticate similarly to SSH.
- Authenticated client *registers* a ticket for their current subject.
- Server maintains a list of registered tickets (public keys). Checks ACL with ticket ACL mask.

# Simple Ticket Setup Steps

1. Client creates a ticket (Public/Private key pair);
2. Client registers the ticket with a storage node (send Public key);
3. Client sets capabilities of the ticket (sets ACL masks).

# An Access Control List

For a directory, storage node maintains a list of tuples (ACL):

**<Subject, Rights>**

**Subject** is a tuple:

**<Means of Authentication, User>**

**Rights** is a list of primitive access methods  
(**read, write, etc.**)

# Access Control List Masks

Resolution of directory authorization: Logical AND of **ticket ACL mask** and the **rights of the subject** from storage-node ACL.

Key point: A ticket cannot exceed the rights of the subject's **current** rights. Conversely, the ticket is still limited by the ACL mask the user assigns.

# Typical Example

```
pdonnel3@cclws13 ~$ chrip -a unix disc01.crc.nd.edu:9090 ticket_create -o
my.ticket / 1 /pdonnel3/data rwl
connected to disc01.crc.nd.edu:9090 as unix:pdonnel3
Generated ticket my.ticket.
ticket 'my.ticket': successfully created with 1024 bits.
ticket 'my.ticket': successfully registered.
ticket 'my.ticket': directory '/' aclmask = 'l'.
ticket 'my.ticket': directory '/pdonnel3/data' aclmask = 'rwl'.
pdonnel3@cclws13 ~$
```

```
pdonnel3@cclws13 ~$ chrip -a unix disc01.crc.nd.edu:9090 listacl /pdonnel3
connected to disc01.crc.nd.edu:9090 as unix:pdonnel3
unix:pdonnel3 rwlda
hostname:*.nd.edu rl
* rl
pdonnel3@cclws13 ~$
```

```
pdonne13@cclws13 ~$ chirp -a ticket -i my.ticket disc01.crc.nd.edu:9090 ls /pdonne13
connected to disc01.crc.nd.edu:9090 as unix:pdonne13
bash
congress-records
data
derp
foo.txt
foobar.txt
fooo.txt
foooo.txt
lots
pdonne13@cclws13 ~$ █
```

```
pdonnel3@cclws13 ~$ chrip -d auth -a ticket -i my.ticket disc01.crc.nd.edu:9090
mkdir /pdonnel3/data/foo
2012/06/12 07:47:22.05 [20973] chrip: auth: ticket: registered
2012/06/12 07:47:22.05 [20973] chrip: auth: requesting 'ticket' authentication
2012/06/12 07:47:22.79 [20973] chrip: auth: server agrees to try 'ticket'
2012/06/12 07:47:22.80 [20973] chrip: auth: trying ticket 3272b849b53c62df9efd89
e7f2d4f737
2012/06/12 07:47:22.83 [20973] chrip: auth: receiving challenge of 64 bytes
2012/06/12 07:47:22.89 [20973] chrip: auth: succeeded challenge for 3272b849b53c
62df9efd89e7f2d4f737
2012/06/12 07:47:22.89 [20973] chrip: auth: successfully authenticated
2012/06/12 07:47:22.89 [20973] chrip: auth: reading back auth info from server
2012/06/12 07:47:22.89 [20973] chrip: auth: server thinks I am ticket:3272b849b5
3c62df9efd89e7f2d4f737
connected to disc01.crc.nd.edu:9090 as unix:pdonnel3
pdonnel3@cclws13 ~$ █
```

```
pdonnel3@cclws13 ~$ chirp -a ticket -i my.ticket disc01.crc.nd.edu:9090
rmdir /pdonnel3/data/foo
connected to disc01.crc.nd.edu:9090 as unix:pdonnel3
couldn't rmdir: Permission denied
pdonnel3@cclws13 ~$ █
```

# Future Work?

**Integrate Chirp URLs and Ticket Authentication with Makeflow/WorkQueue**

# Questions?

Website: <http://www.cse.nd.edu/~ccl>

Autobuild: <http://www.cse.nd.edu/~ccl/software/autobuild>

Chirp: <http://www.cse.nd.edu/~ccl/software/chirp>

Patrick Donnelly: [pdonnel3@nd.edu](mailto:pdonnel3@nd.edu)