

Flexible Data Placement Mechanisms in Condor

Zach Miller
Condor Project
Computer Sciences Department
University of Wisconsin-Madison



Overview

- > What existed in 2010
- > What's new in 2011
- > Problems, Solutions, and Future Work

File Transfer Basics

- Condor will transfer your input files for you if ask nicely in your submit file

```
universe = vanilla
```

```
should_transfer_files = ALWAYS
```

```
when_to_transfer_output = ON_EXIT_OR_EVICT
```

```
transfer_input_files = input_file_1, input_file_2
```

```
executable = myexe
```

```
arguments = 30
```

```
queue
```



File Transfer Basics

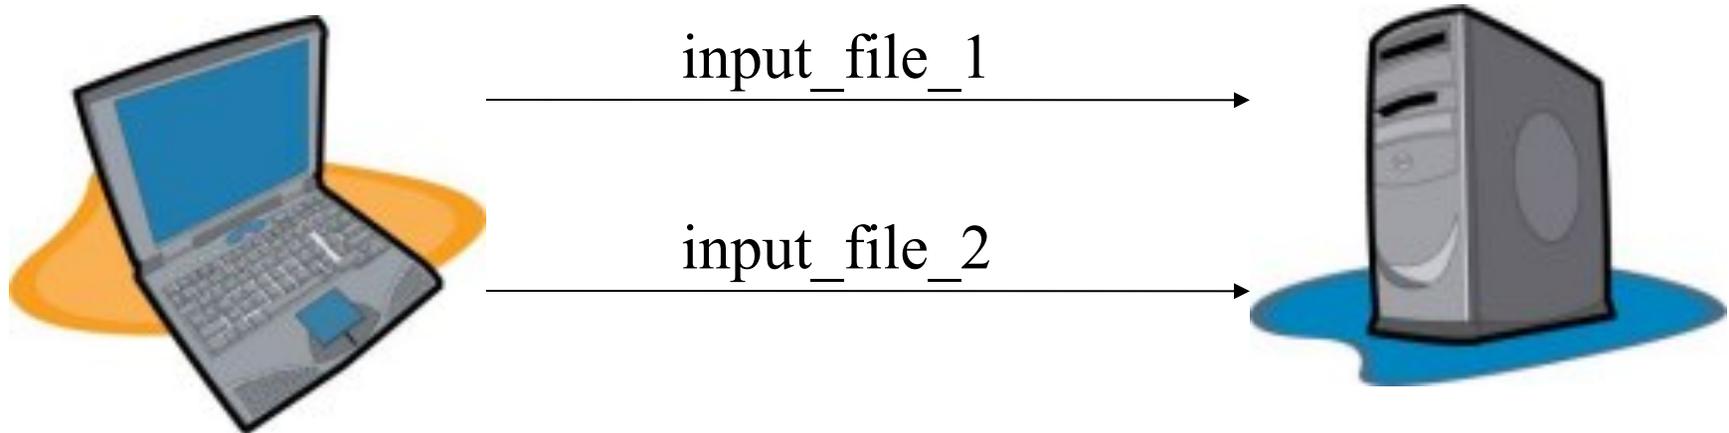
- These files are transferred from the submit machine to the execute machine when the job starts up
- Starting many jobs simultaneously can place a lot of demand for bandwidth on the submit machine
- Throttles were added to counteract

Third Party Transfer

- > Wouldn't it be nice if the data didn't have to come from the submit machine?
- > Especially if it's the same input file for each job, e.g. a reference database of genomes?
- > That would be great!

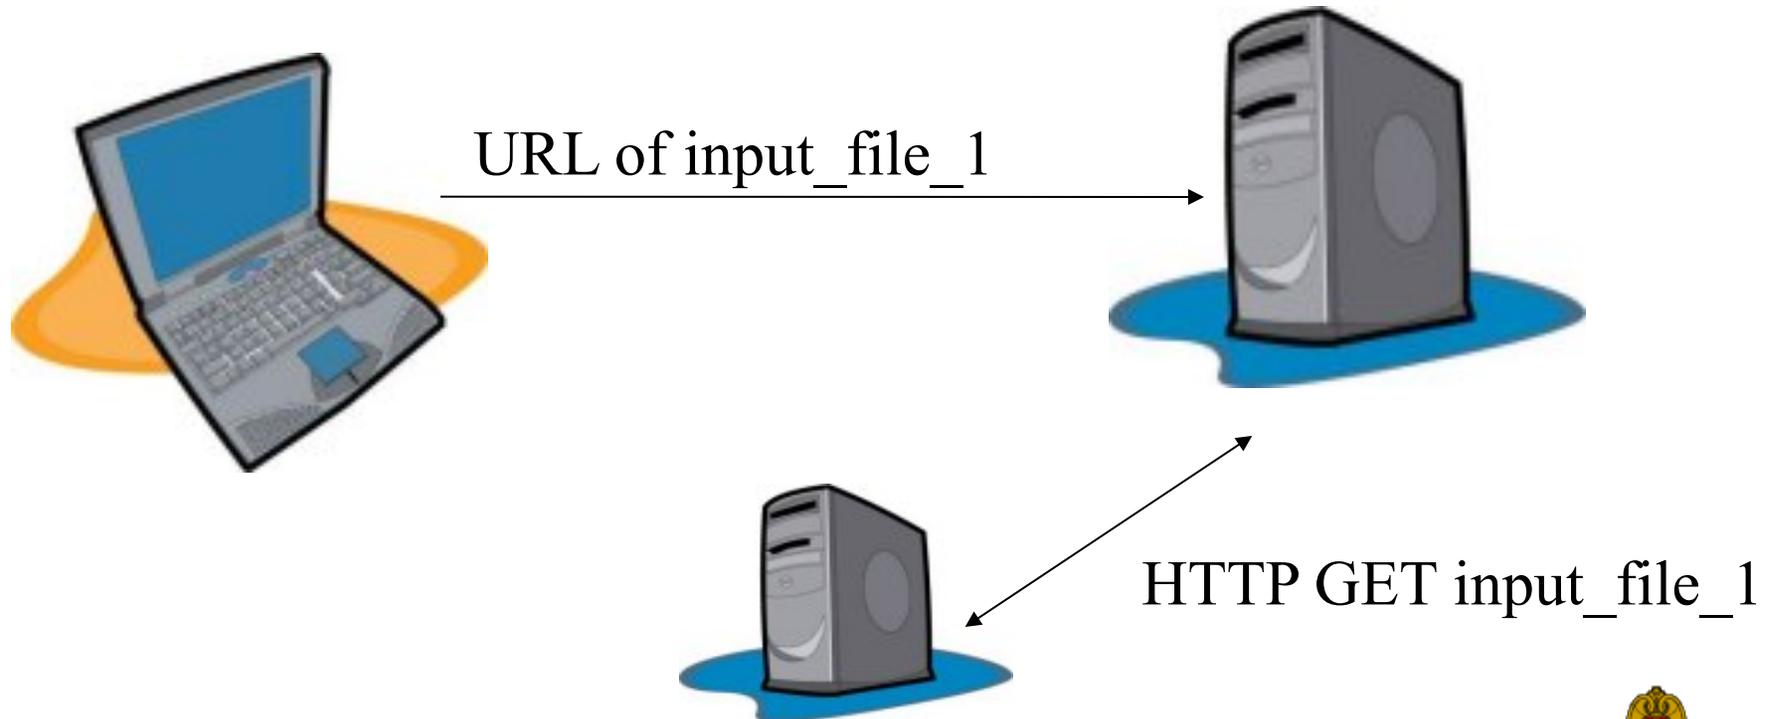
Normal File Transfer

- > File contents sent by submit machine



Third Party Transfer

- > Only URL sent by submit machine



URL Format Assumptions

- > urltype
- > ://
- > Anything else
- > Everything after the last slash is used as the downloaded filename
- > Example: `http://www.cs.wisc.edu/path/to/file`

Example Submit File

```
universe = vanilla
should_transfer_files = ALWAYS
when_to_transfer_output = ON_EXIT_OR_EVICT
transfer_input_files = input_file_1,
    http://www.cs.wisc.edu/~zmiller/input_file_2
executable = myexe
arguments = 30
queue
```

Submit File

- If an input file name matches the format of the URL described, Condor assumes it should be transferred with a plugin
- Normal files and third party transfers can be mixed, as well as different types of plugins

Another Submit File

```
universe = vanilla
should_transfer_files = ALWAYS
when_to_transfer_output = ON_EXIT_OR_EVICT
transfer_input_files = input_file_1,
    http://www.cs.wisc.edu/~zmiller/input_file_2
    ftp://zmiller:pass@ftp.cs.wisc.edu/data/input_file_3
    data://;base64,iVBORw0KGgoAAAANSBlahBlah/input_file_4
executable = myexe
arguments = 30
queue
```

How Does It Work?

- Condor uses a general plugin architecture to support additional methods of file transfer
- This allows Condor to be customized with whatever mechanisms you like
- Also allows plugins to be created and updated out-of-band of Condor releases and stable series

condor_config Settings

- It's easy for users to add these URLs to their submit files
- However, the URLs are meaningless unless they are interpreted by a plugin
- Plugins must be installed by the administrator of the execute machine

condor_config Settings

- > Plugins are executables and can be written in C, perl, bash, whatever.
- > Fairly straightforward to configure:

```
FILETRANSFER_PLUGINS = \  
    /path/to/curl_plugin, \  
    /path/to/data_plugin, \  
    /path/to/custom_plugin
```

Plugin Behavior

- > How do the plugins work?
- > Query Mode
 - Describe what the plugin is capable of doing
- > Action Mode
 - Perform an actual transfer

Plugin Behavior

> Query Mode

- Condor invokes the plugin with the `-classad` option to ask for information

```
% /unsup/condor/libexec/curl_plugin -classad
PluginVersion = "0.1"
PluginType = "FileTransfer"
SupportedMethods = "http,ftp,file"
```

Plugin Behavior

> Action Mode

- Condor invokes the plugin with two parameters, source and destination, to actually perform the transfer.

```
% /unsup/condor/libexec/curl_plugin \  
http://www.cs.wisc.edu/~zmiller/file_1 \  
/scratch/condor/exec.22356/file_1
```

Support in Condor

- > All Condor functionality previously mentioned was available in Condor 7.4.0 and later
- > Two major shortcomings
 - No ability to transfer job output
 - X.509 credentials are not available to the plugins

Support in Condor

- > In 7.6.0, better X.509 support
- > If your submit file specifies an X.509 credential, Condor will
 - Make sure the proxy is the first thing transferred to the job sandbox
 - Set the X509_USER_PROXY environment variable before invoking the transfer plugins

Support in Condor

- > This allows plugins to use your Globus credentials to do third party transfers
 - gridftp
 - globusonline
 - Your custom plugins

Example Submit File

```
universe = vanilla
should_transfer_files = ALWAYS
when_to_transfer_output = ON_EXIT_OR_EVICT
x509userproxy = /tmp/x509up_u24842
transfer_input_files = input_file_1,
    gridftp://newbio.cs.wisc.edu/~zmiller/giant_protein_db
executable = blast_wrapper
queue
```

Another New Feature

- As of 7.6.0, Condor can now move your output file using plugins
- Relies on existing plugin architecture
- Just add `output_destination` to your submit file
- Condor will invoke the plugin to move your output files

Example Submit File

```
universe = vanilla
should_transfer_files = ALWAYS
when_to_transfer_output = ON_EXIT_OR_EVICT
transfer_input_files = input_query,
    ftp://ftp.ncbi.nih.gov/blast/db/FASTA/nr.gz
output_destination =
    gridftp://newbio.cs.wisc.edu/zmiller/output/
executable = blast_wrap.sh
x509userproxy = /tmp/x509up_u24842
queue
```

Notes on Using output_destination

- > The `transfer_output_files` parameter in the submit file can still be used to select a subset of files to transfer
- > If any of the plugins report failure, the job is placed on Hold
- > Nothing is transferred back to the submit node (including stdout and stderr)

Notes on Using output_destination

- > The plugins registered in the condor_config file must be able to handle both input and output:

```
curl_plugin ftp://ftp.wisc.edu/remote/file /tmp/local/file
```

```
curl_plugin /tmp/local/file ftp://ftp.wisc.edu/remote/file
```

globusonline and Condor Integration

- There was significant effort put into integrating the Globus Online service with Condor using this plugin architecture
- Was a driving force for some of the modifications made to Condor, in particular X.509 support
- Was also a driving force for modifications made to the globusonline API

Assumptions about using globusonline

- In order to work, some things must already be in place:
 - User must have a valid X.509 proxy
 - User must already have a globusonline account
 - There must be an endpoint registered and activated that you will transfer to or from

The globusonline Plugin

- Starts a "personal" gridftp server using the user's proxy
- Registers the new server as a new globusonline endpoint
- Activates the endpoint with the user's proxy
- Performs the transfer to/from job sandbox
- Removes the endpoint
- Shuts down the gridftp server

Some Limitations

- By default, when Condor performs credential delegation, it creates "limited" proxies
- However, creating a dynamic endpoint requires delegating the credential again to the globusonline service
- This is not allowed with limited proxies
- You need to instruct Condor not to create limited proxies if you plan to use globusonline

Some Limitations

- Currently, input files can come from several different places, using different plugins
- All output files that are transferred must go to the same URL
- It would be great to be able to send different files using different plugins

Some Limitations

- Condor currently invokes file transfer plugins once for each file
- This creates more overhead than necessary because of standing up, registering, activating, and removing endpoints
- Particularly when using globusonline, Condor could batch all files into one transfer
- This would require modification to Condor's plugin API

Some Limitations

- Error propagation needs improvement
 - You get status codes, but no text
- Condor currently relies on the plugin to do what it's told, and does not verify it did the correct thing
- Plugins should have a heartbeat so we know if they are still working or if they have become stuck

Some Ideas

- > You can use the plugins to fetch a Virtual Machine image for VM universe, allowing the images to be cached at the execute site instead of being transferred for each job
- > You could use plugins to always fetch your files via http, allowing squid caches to save on bandwidth

Some Ideas

- > Plugins can (/should/will) be written for:
 - chirp
 - hdfs
 - bittorrent
 - Anything new under the sun?

Questions?

- > <http://www.cs.wisc.edu/condor/>
- > Email me: zmiller@cs.wisc.edu