



# Beyond the SAGA

From Simple APIs to  
Dynamic Abstractions

- I. A Simple API for Grid Applications
- II. Distributed Programming Abstractions (DPA)
- III. A Good Research Idea ?
- IV. *A Somewhat Concrete Plan*

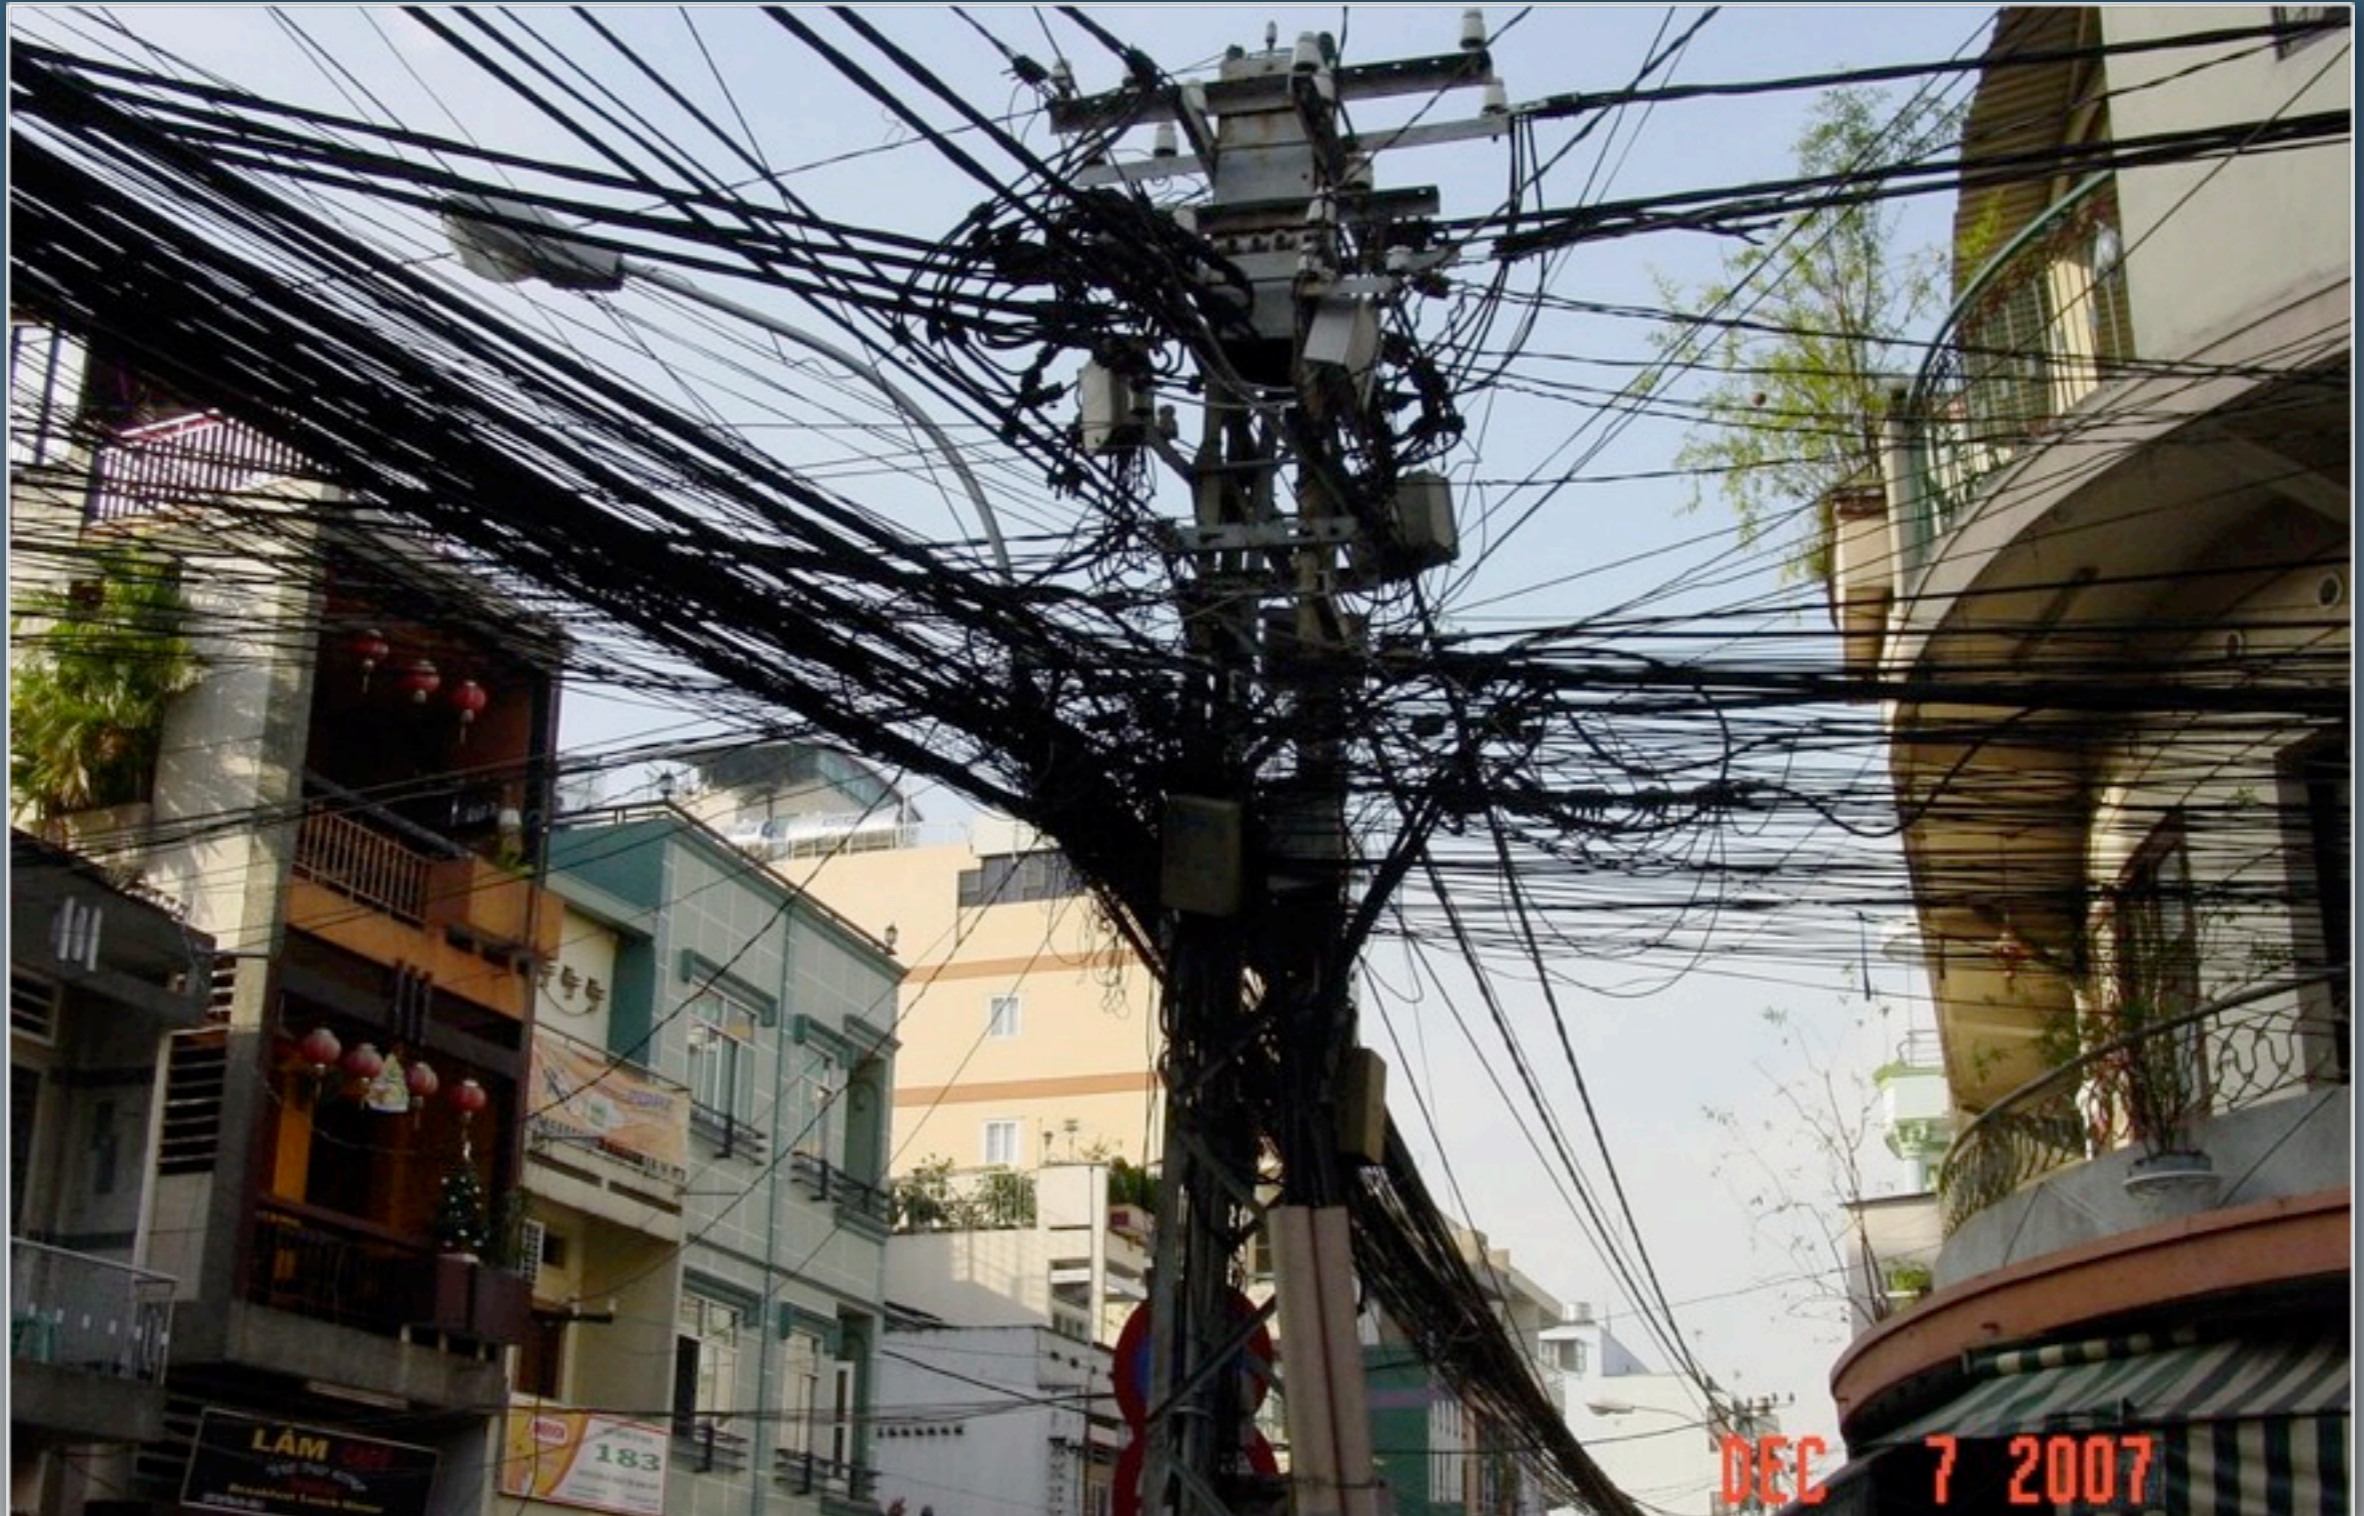


- Simple API for Grid (Distributed) Application
- Open Grid Forum community standard (GFD-R-P.90)
  - Describes a language-independent (SIDL), object oriented API for high-level tasks *considered useful* in distributed applications, like job submission, file transfer, communication, etc...
- Implementations of the GFD-R-P.90 standard:
  - JSAGA (Centre de Calcul IN2P3/CNRS, Lyon, France)
  - JavaSAGA (Vrije Universiteit, Amsterdam, NL)
  - C++ / Python Implementation (LSU, Baton Rouge, USA)

- 1988: Start of the Condor project
- 1998: Globus 1.0 released
- 2002: GridLab's Grid Application Toolkit (GAT)
- 2004: OGF (GGF) SAGA working group formed
- 2006: SAGA C++ development of the reference implementation starts at Louisiana State University
- 2010: GFD-90 1.0 standard (SAGA) released by OGF
- 2011: SAGA is being deployed across US TeraGrid, deployment on EGI, FutureGrid and XD is under review



# A Brief History





# A Brief History



- 1988: Start of the Condor project
- 1998: Globus 1.0 released
- 2002: GridLab's Grid Application Toolkit (GAT)
- 2004: OGF (GGF) SAGA working group formed
- 2006: SAGA C++ development of the reference implementation starts at Louisiana State University
- 2010: GFD-90 1.0 standard (SAGA) released by OGF
- 2011: SAGA is being deployed across US TeraGrid, deployment on EGI, FutureGrid and XD is under review

# Why a Middleware-Independent API ?

- (Lifetime of applications) > (Lifetime of infrastructure and interfaces)
- Portability / adaptability = protection of assets
- Opens new opportunities for large-scale distributed systems research and evaluation
- Newly emerging extreme-scale simulations may have to span (scale-out) across several different infrastructures
- Distributed computing's counterpart of MPI ?



- Designed after the *adaptor pattern*
- A set of C++ libraries and headers grouped into functional packages
- *Adaptors* (plug-ins) that provide access to distributed middleware (Globus, gLite, Condor, etc...)
- A light-weight runtime/dispatcher that the right adaptor for an API call at runtime
- No services, daemons, etc...



# How Does it Work ?

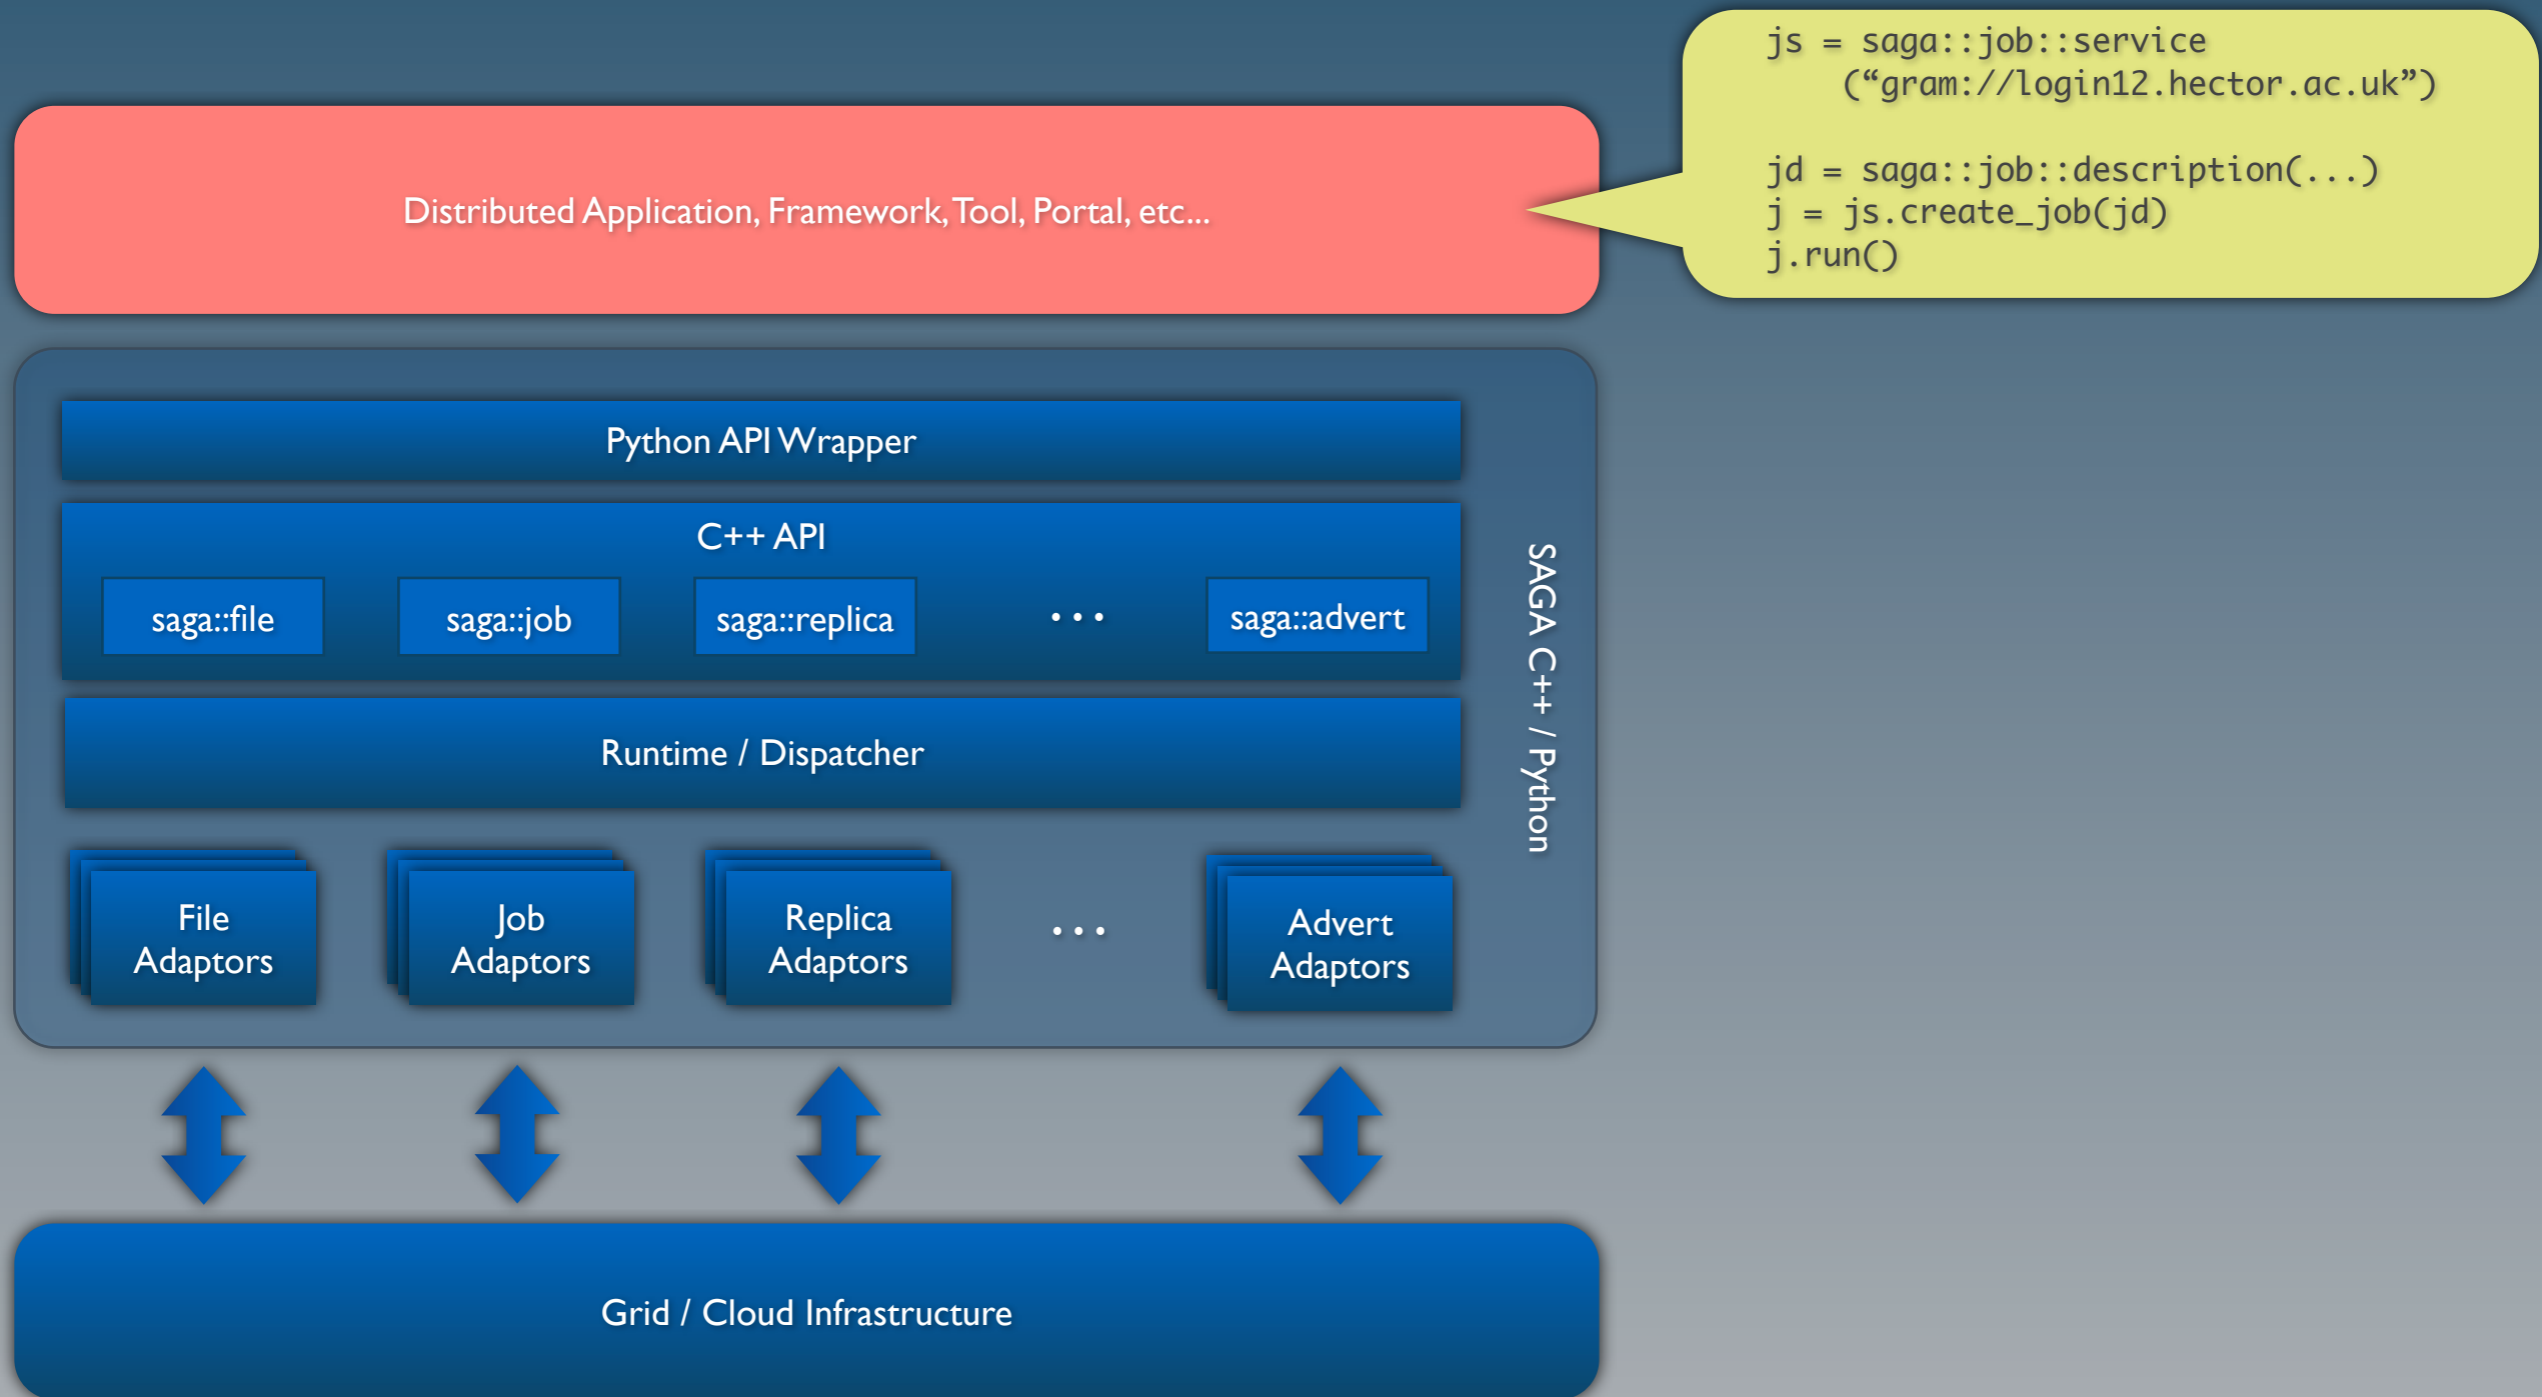




# How Does it Work ?

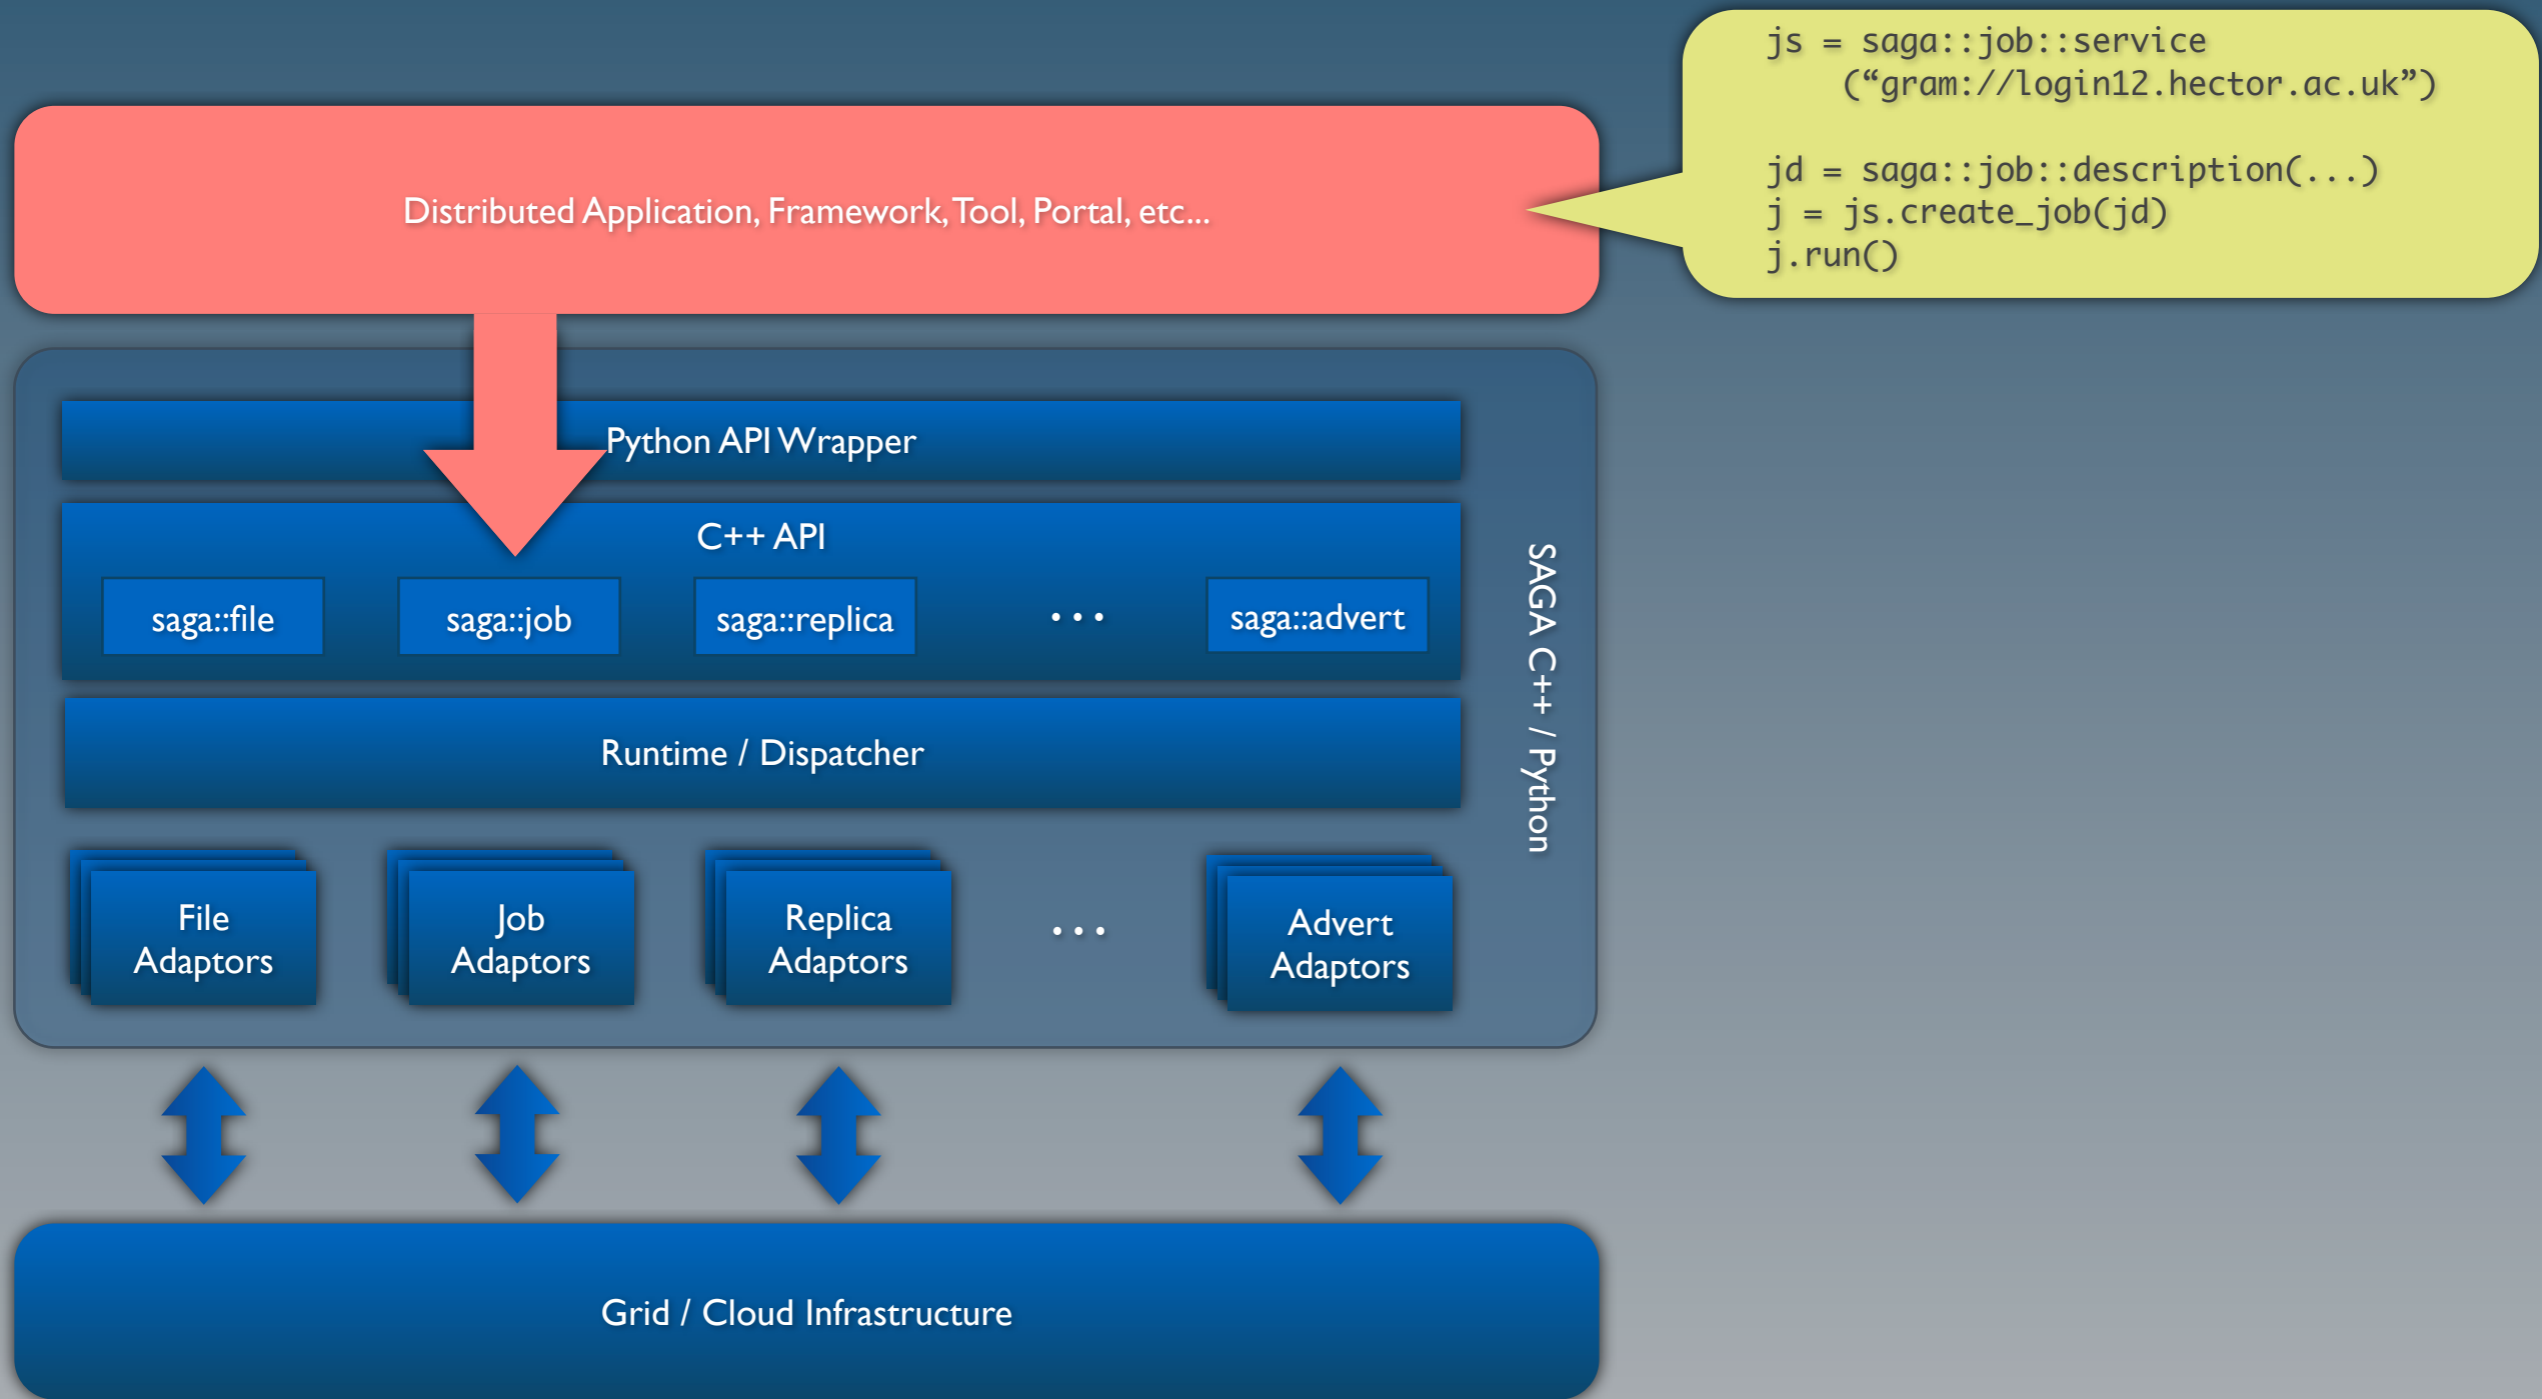


# How Does it Work ?

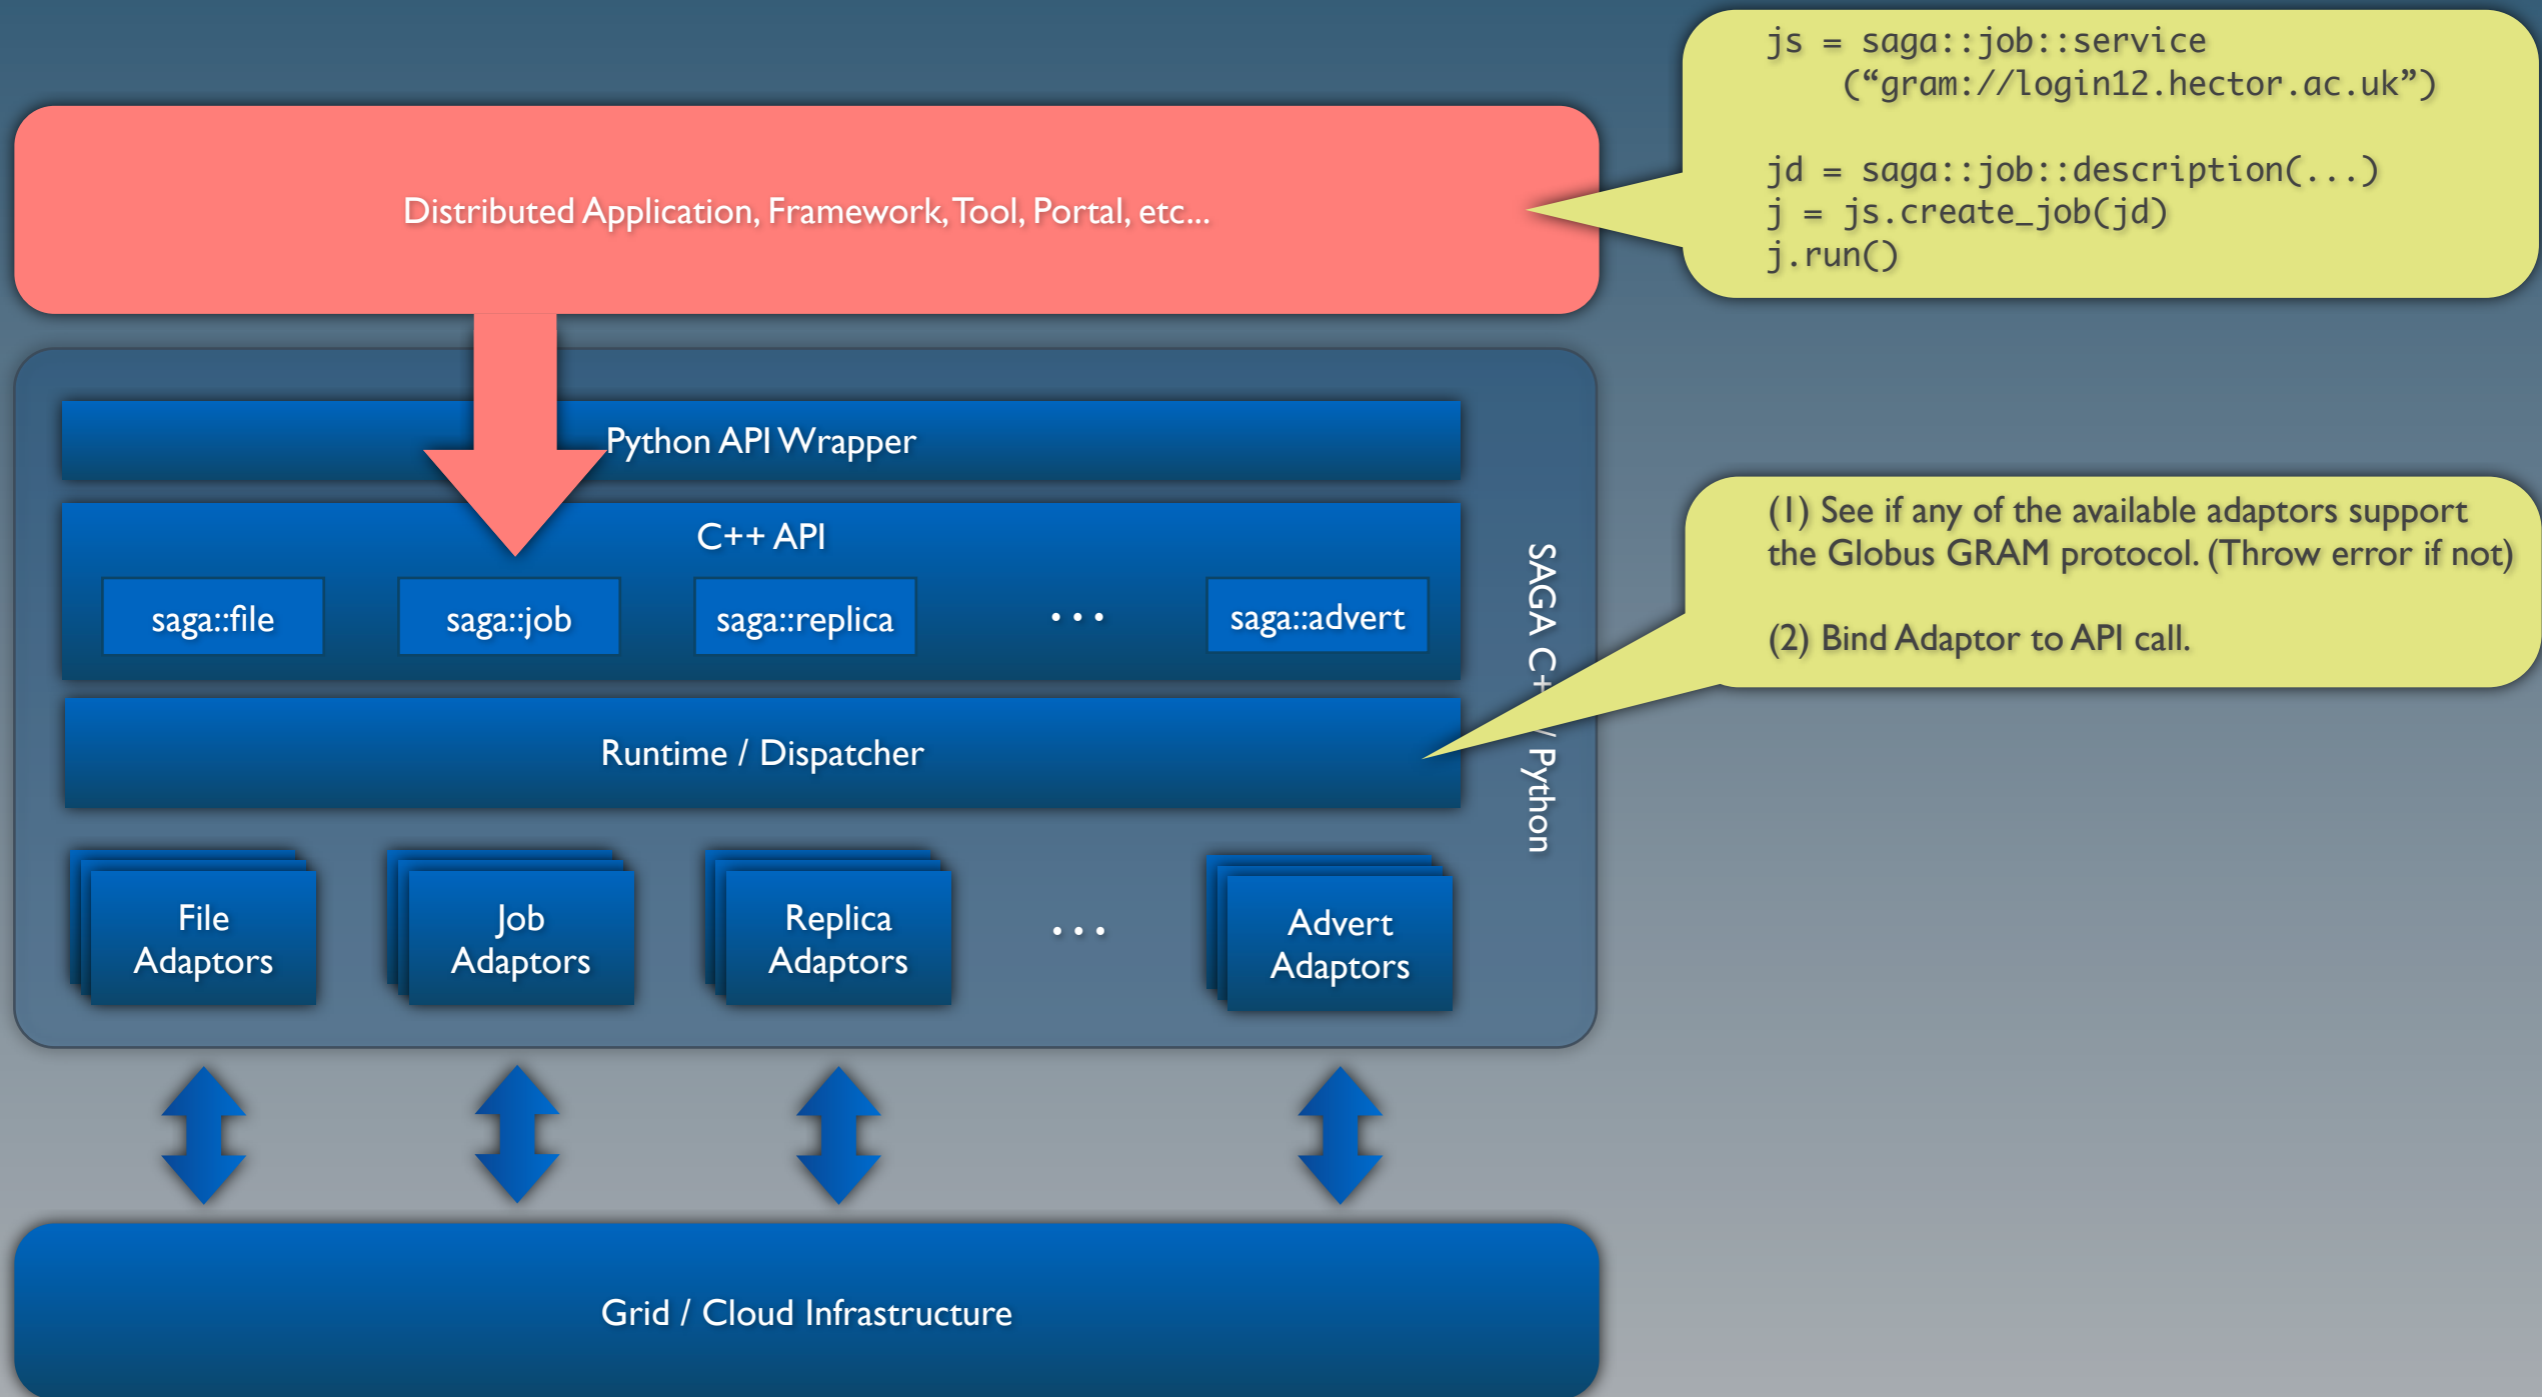




# How Does it Work ?

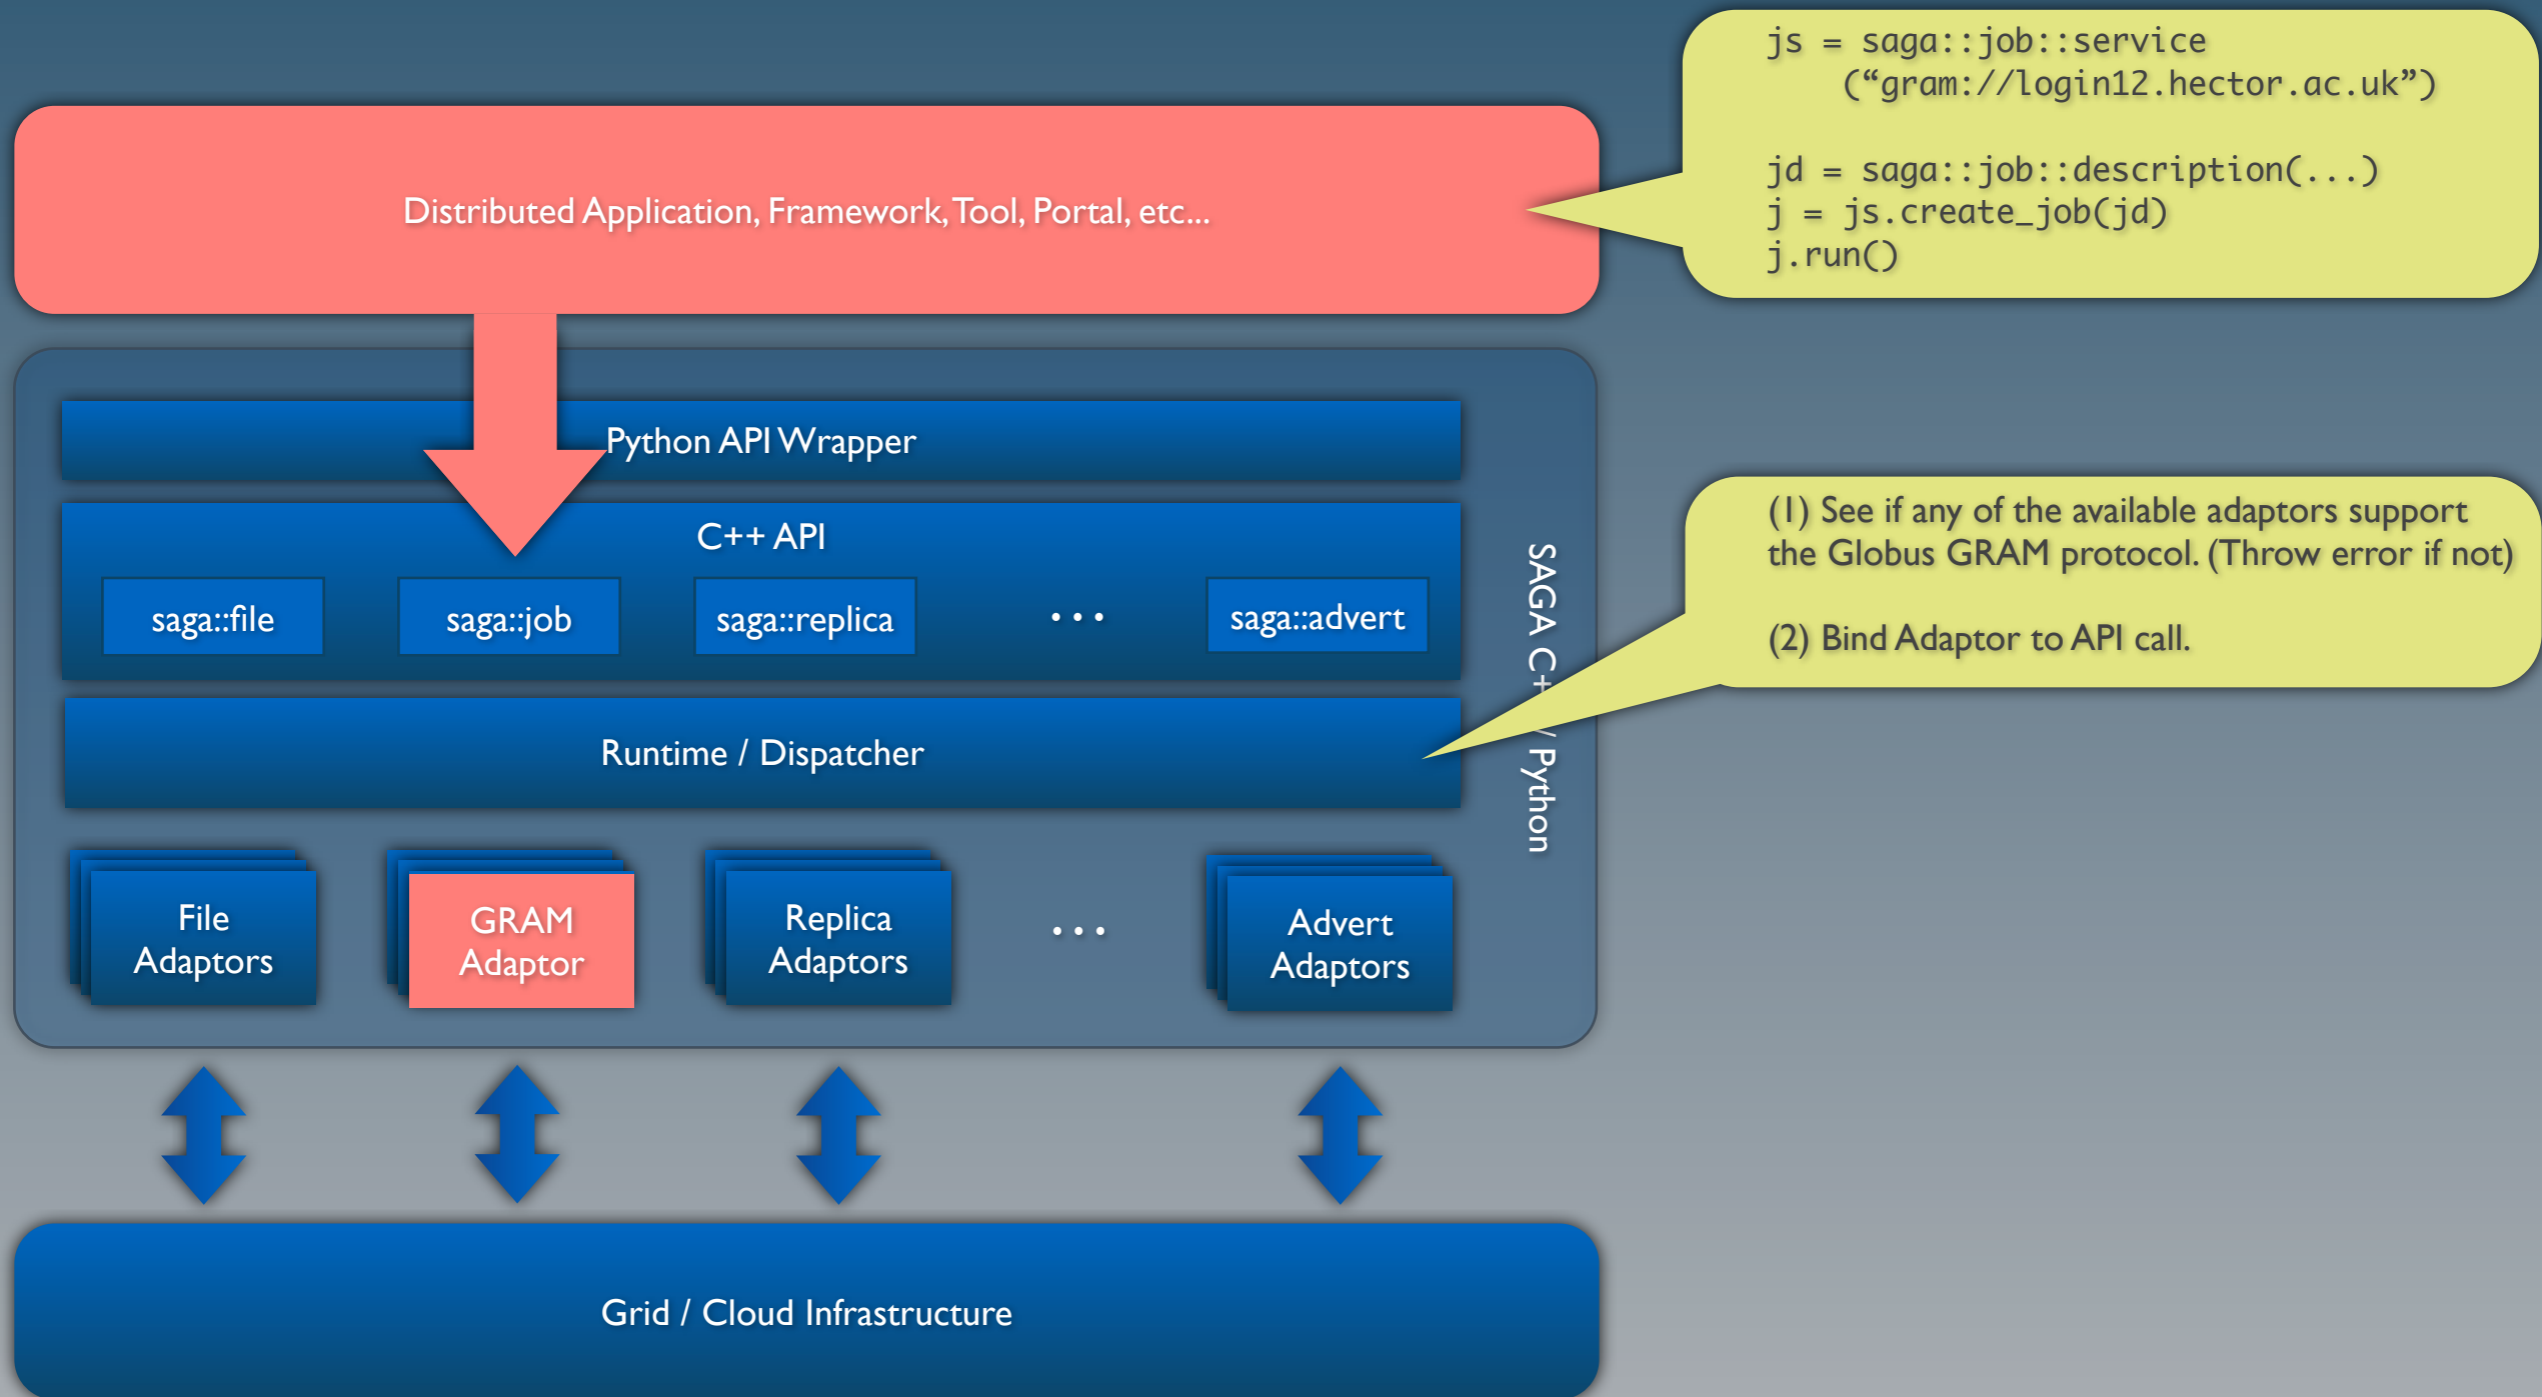


# How Does it Work ?

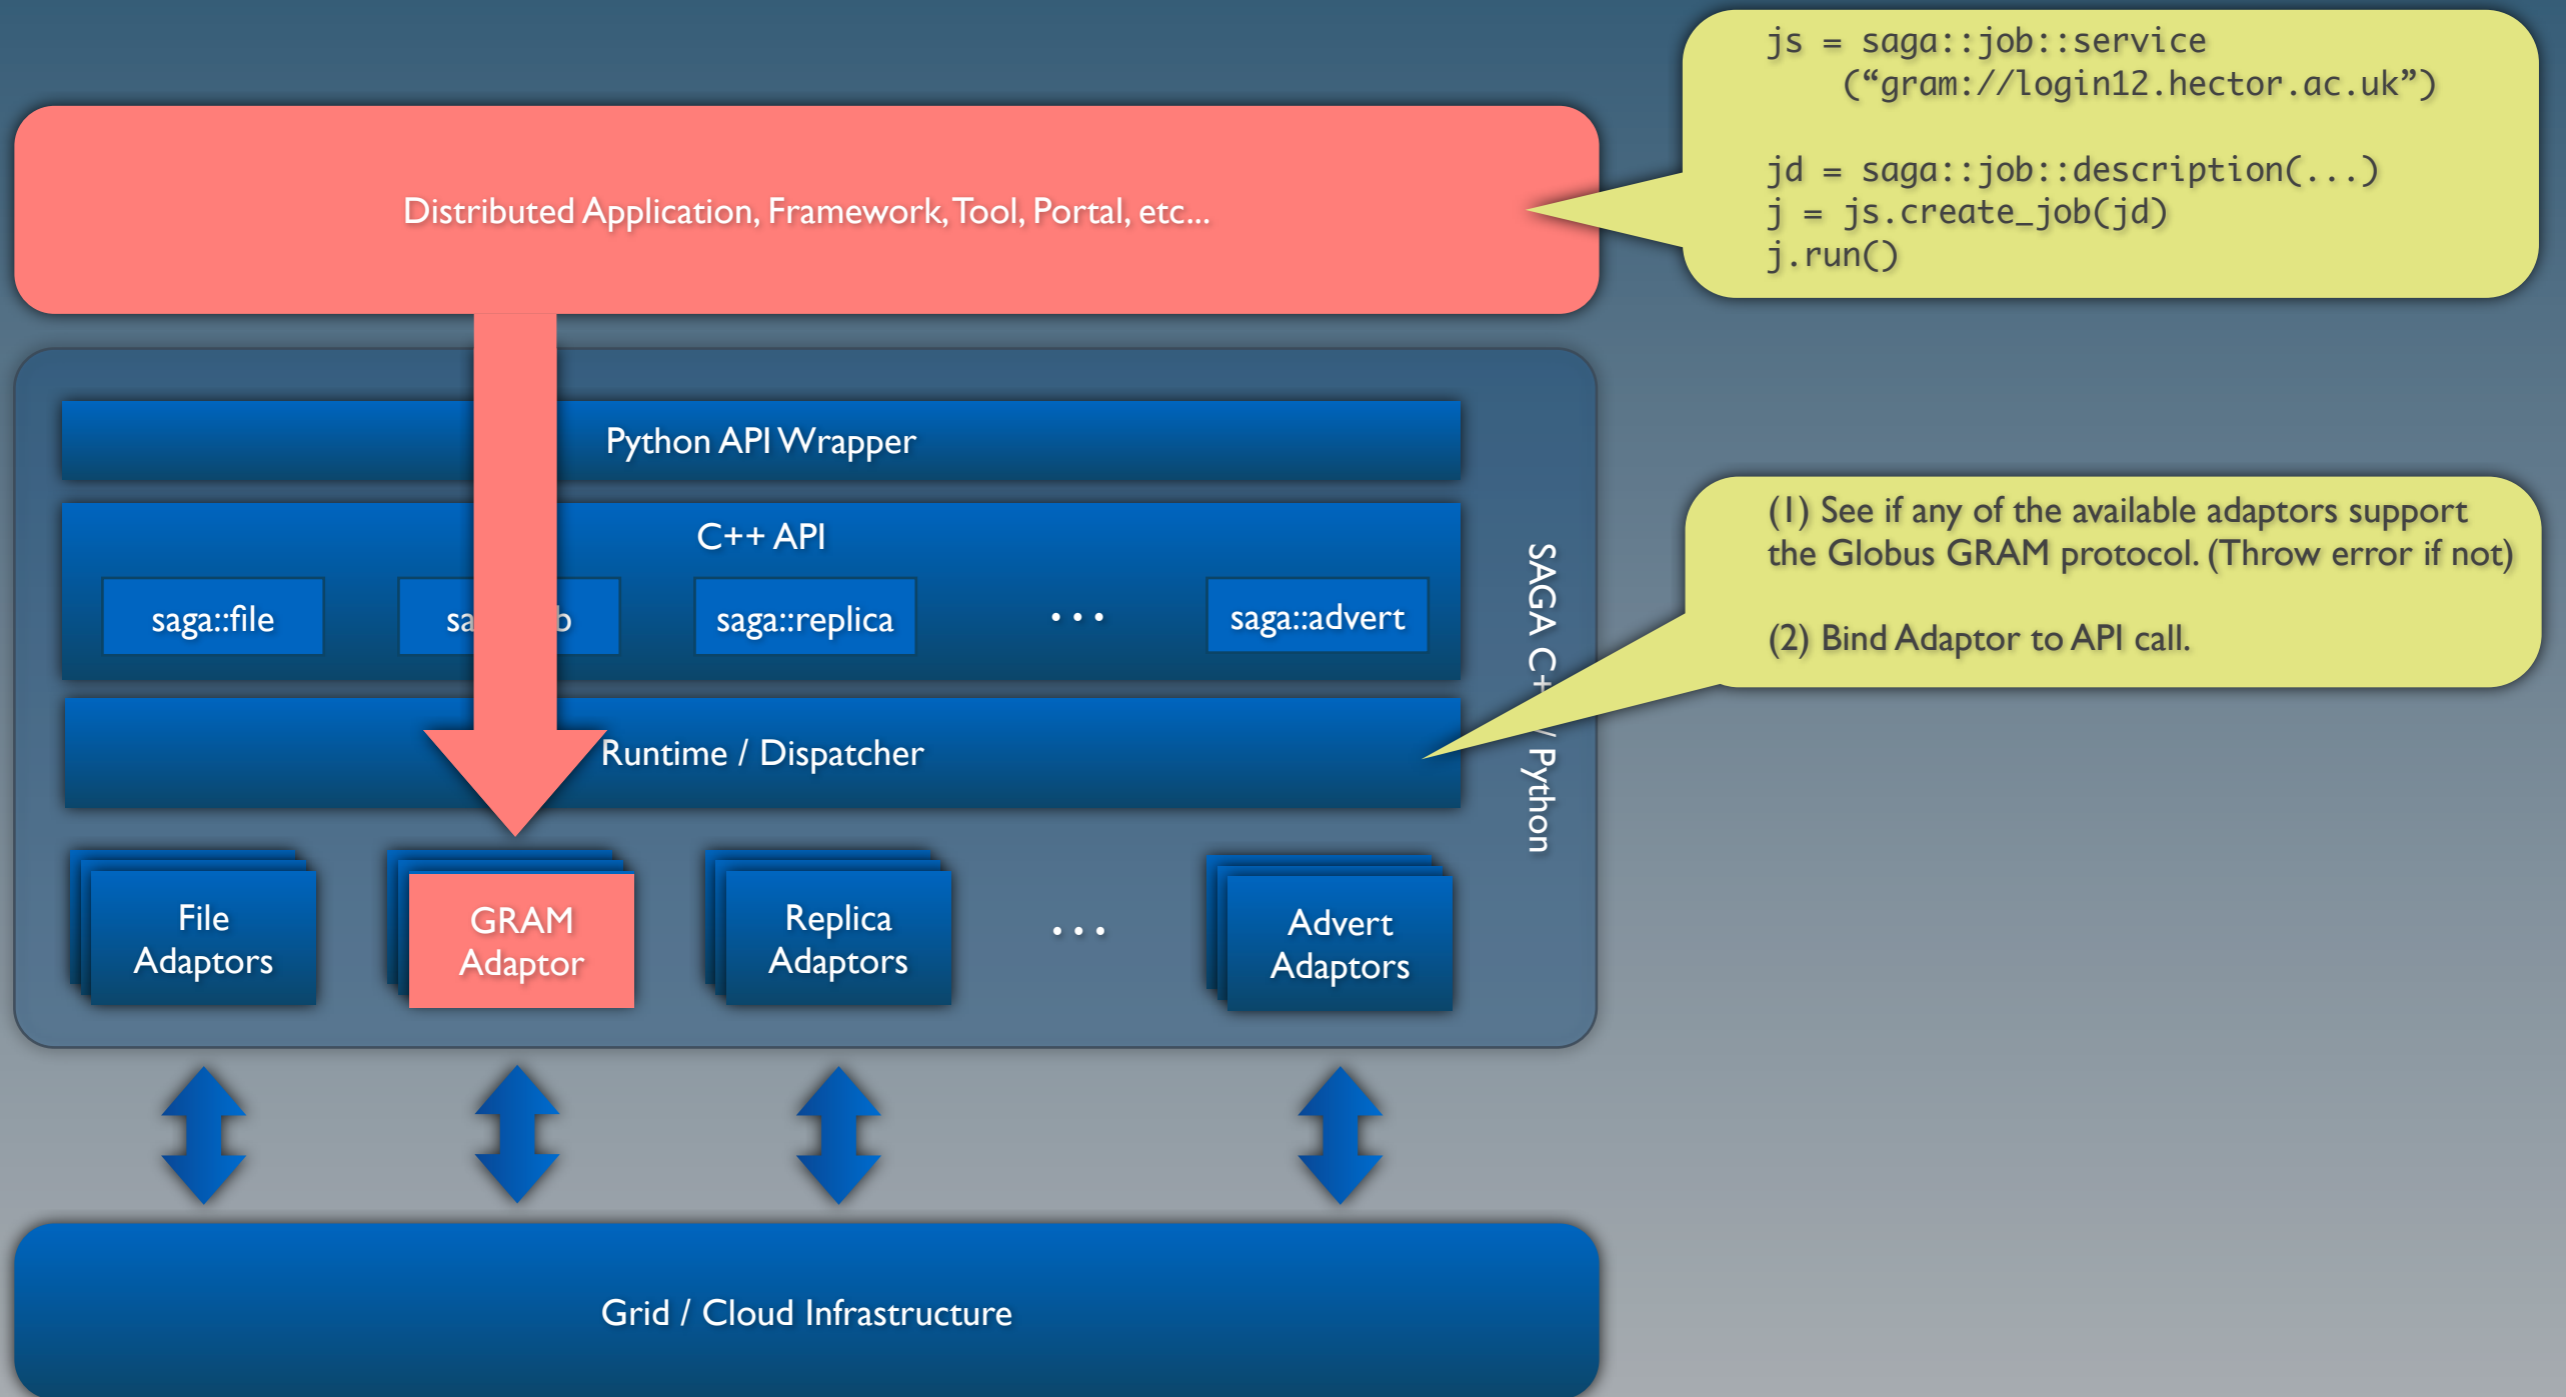




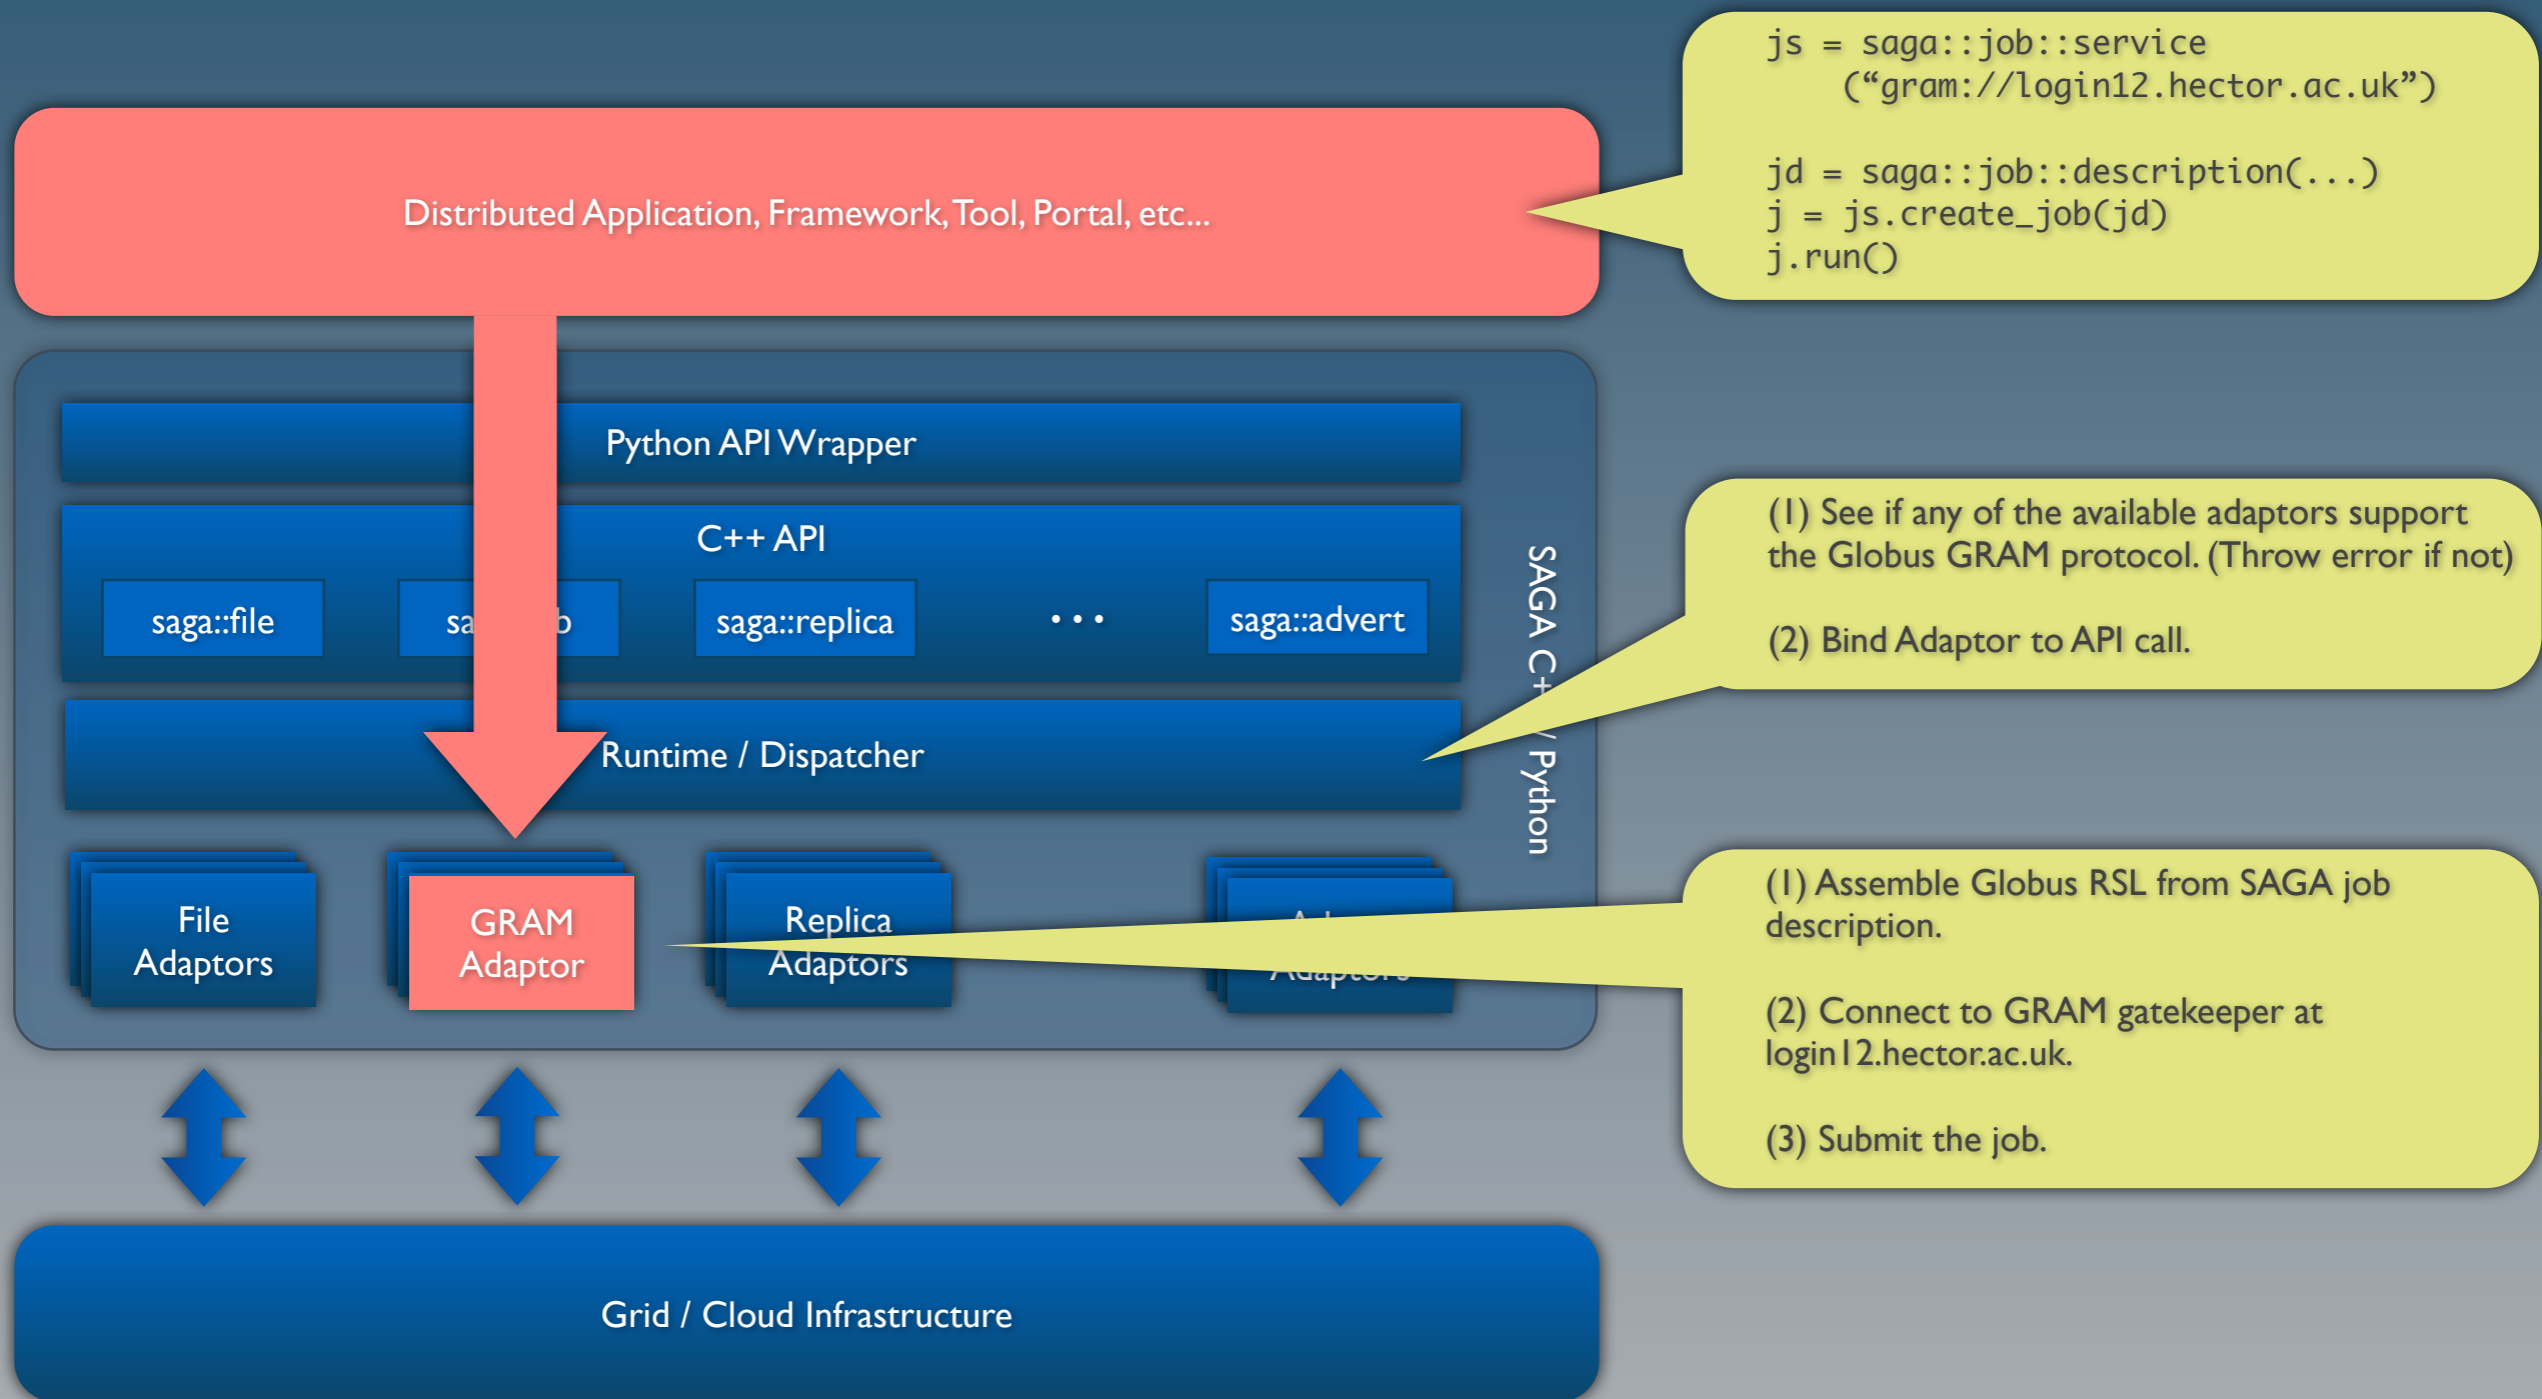
# How Does it Work ?



# How Does it Work ?

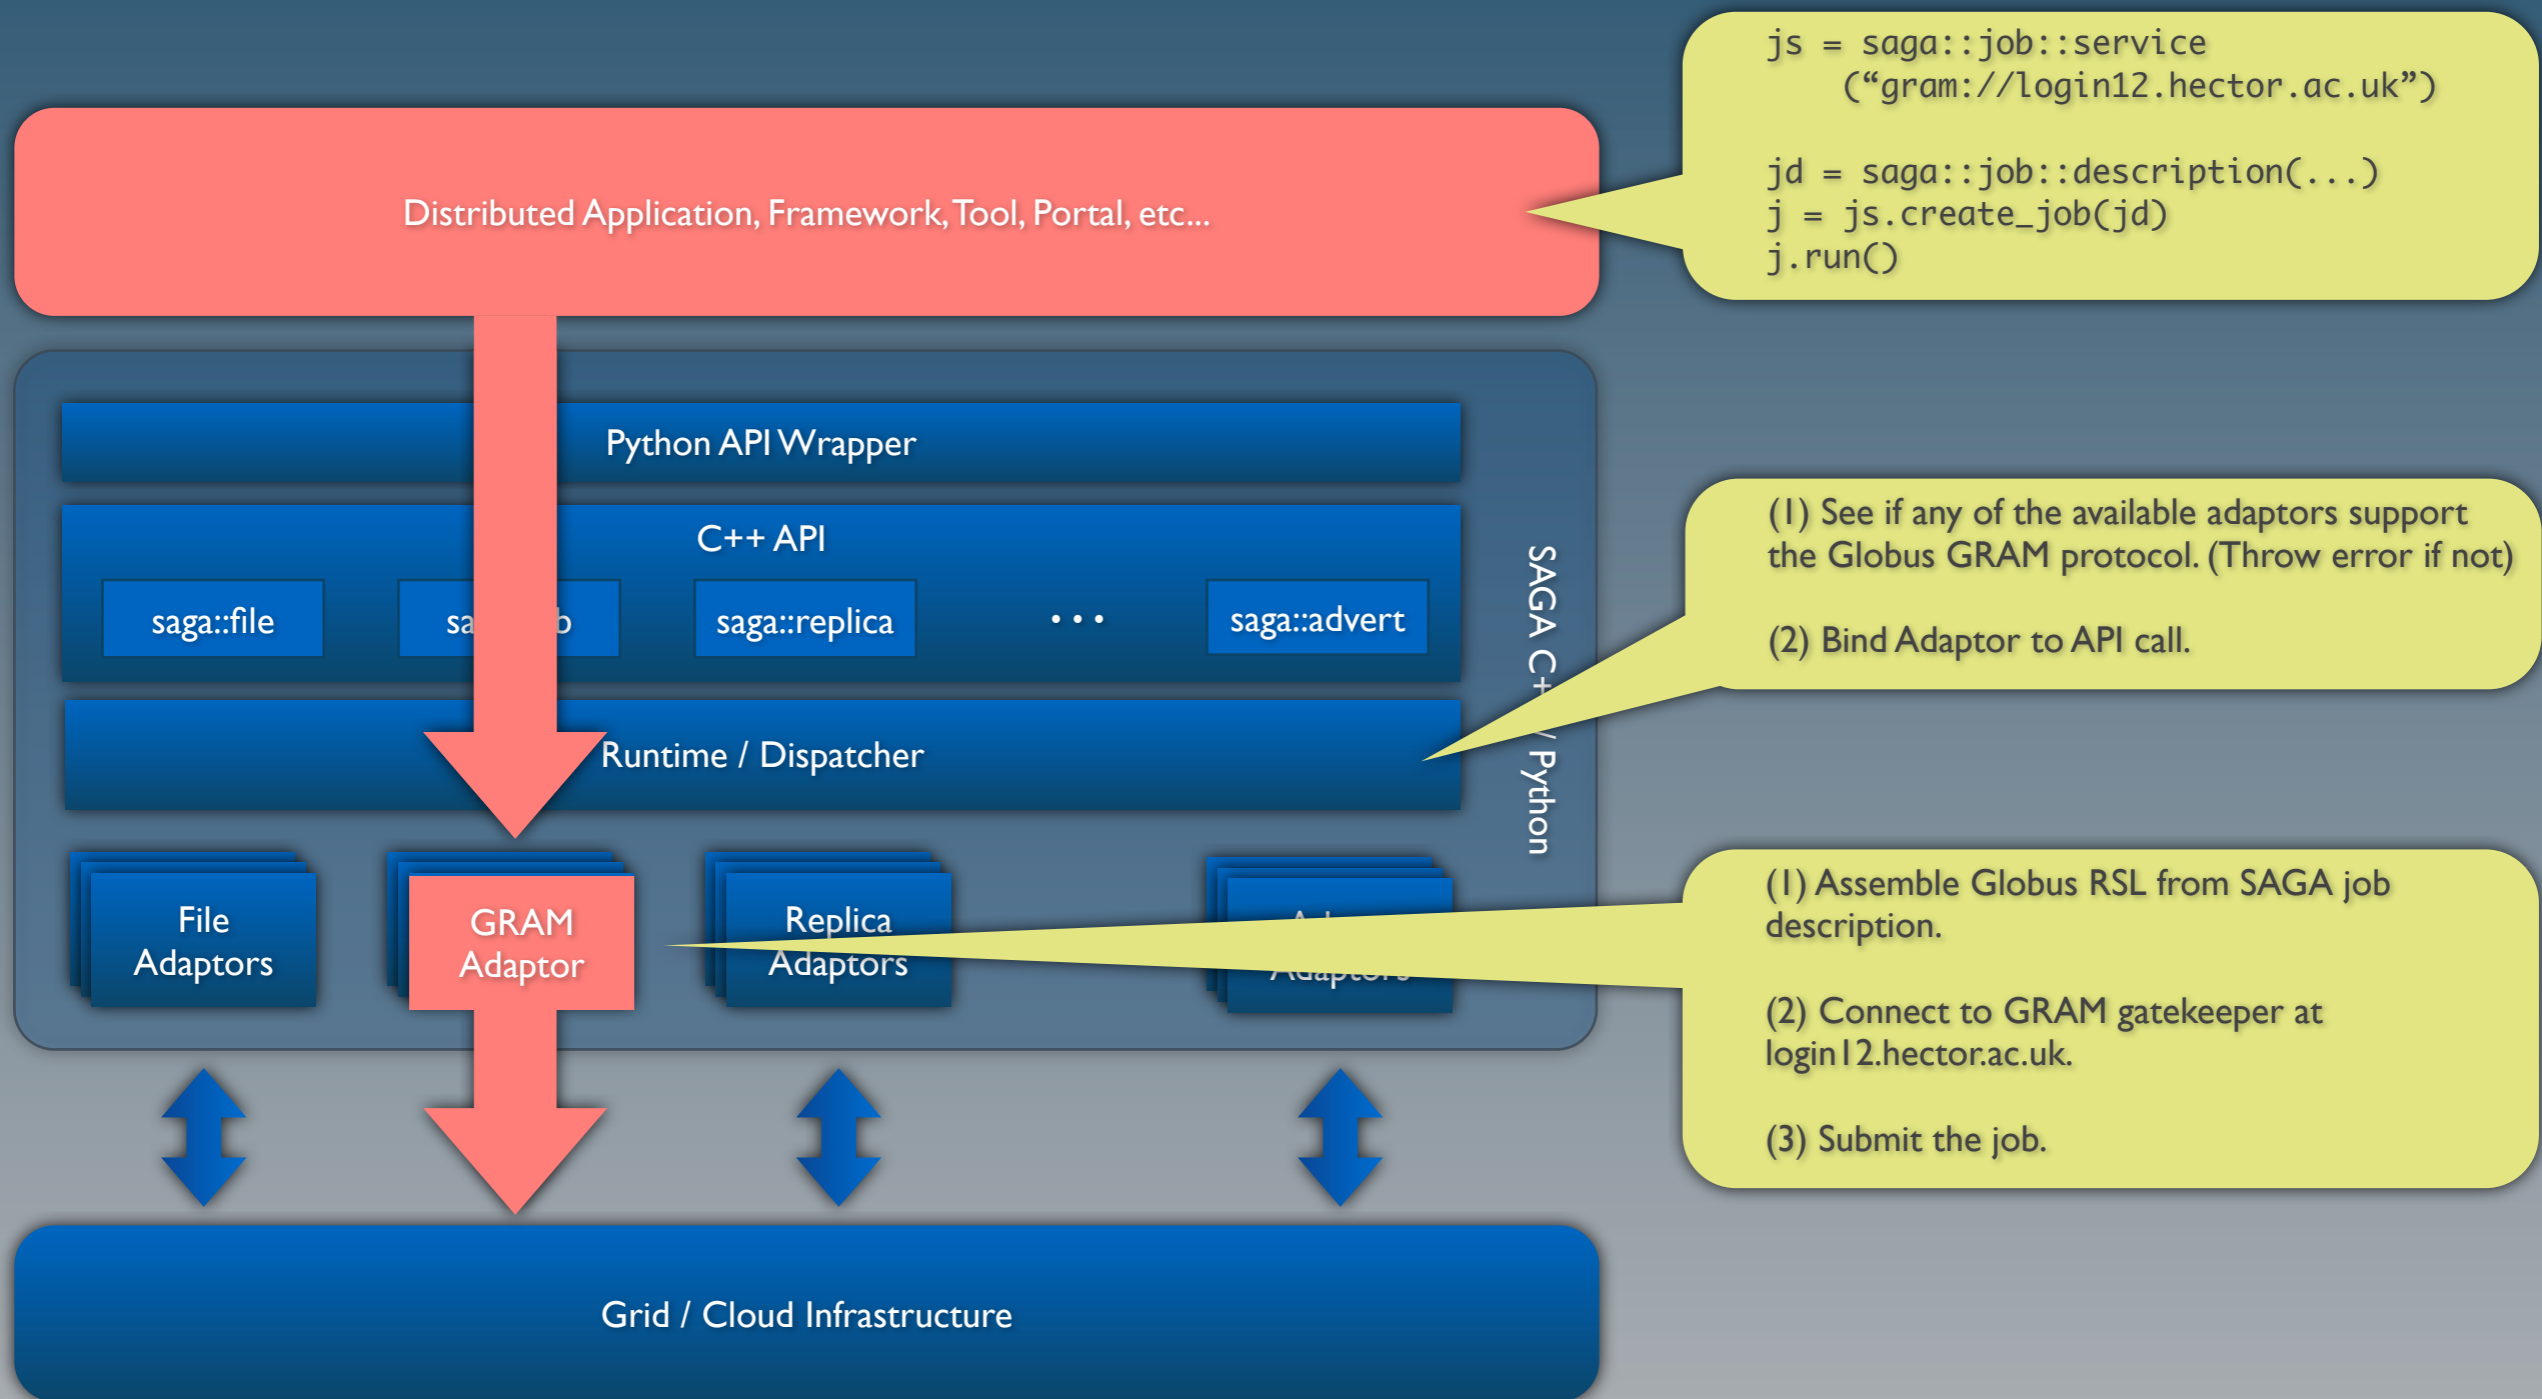


# How Does it Work ?

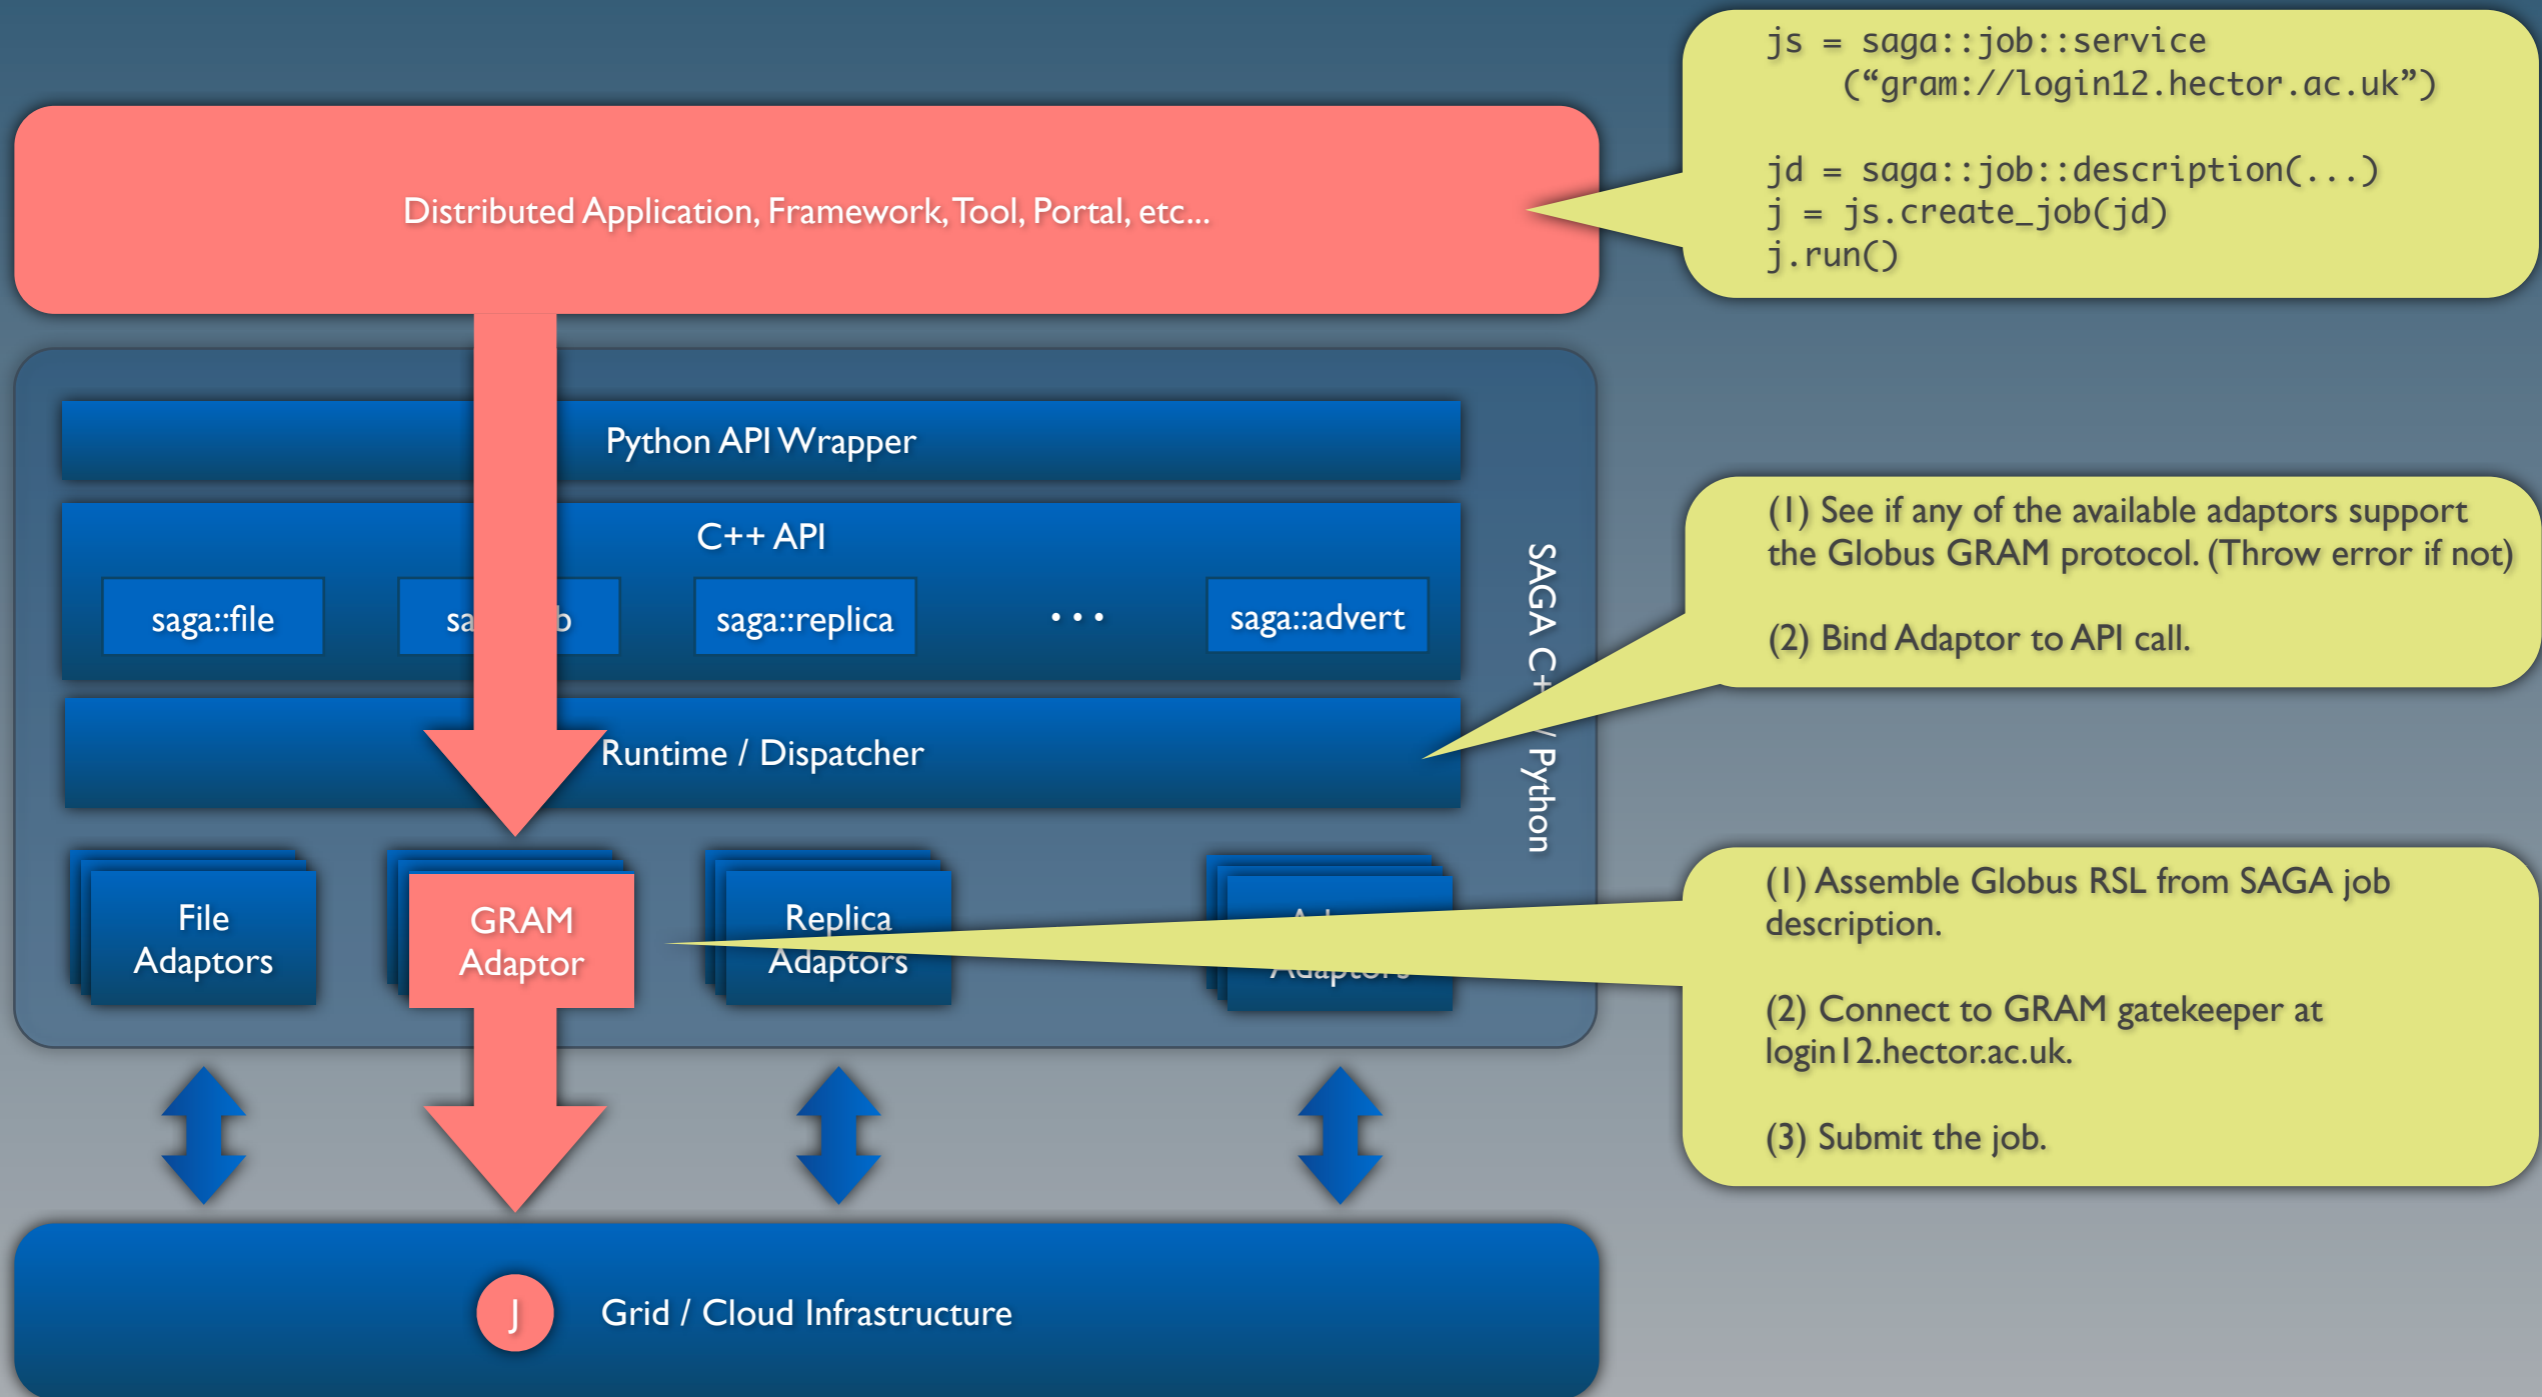




# How Does it Work ?



# How Does it Work ?



- `saga::advert` - *Advert Service Access API*
- `saga::filesystem` - *Filesystem Access API*
- `saga::job` - *Job Submission & Management API*
- `saga::replica` - *Replica Catalog Management API*
- `saga::rpc` - *Remote Procedure Call API*
- `saga::sd` - *Service Discovery API*
- `saga::stream` - *Data Stream Client & Server*
- API Extensions under development: `saga::messaging`



- `saga::job` - Arc, Amazon EC2, Condor, Eucalyptus, Globus GRAM (2&5), Fork, gLite, SSH, Nimubs, OGSA BES, PBS (Pro), Platform LSF, SSH, TORQUE, ...
- `saga::filesystem` - Globus GridFTP, Hadoop HDFS, Local Filesystem, SSHFS, ...
- `saga::replica` - Globus RLS, SQL Replica Service
- `saga::advert` - SQL Advert Service
- `saga::stream` - TCP-based
- `saga::sd` - gLite SD

# A Small Dilemma ?

---

- After several years of prototyping, testing and hardening, we presented SAGA as the holy grail of distributed computing to the user communities
- Lots of advertising, demos, workshops, tutorials
- But: Uptake very slow and not as expected
- Almost by accident, the problem got solved with the “SAGA *Big-Job*” framework

# A Small Dilemma ?

```
/*-----*/
#
#-----*/
# Description : Transfer the data file from client site to each remote site.
# and return the status of the file transfer.
#
# Input      : Source URL string of the source server , and the Destination URL
#              string of the destination server .
#              The format for URL string is
#              Source URL      : full path of the file to transferred
#              Destination URL : gsiftp://server-hostname/fulpathname/ofthefile/tobetranferred
#
# Output     : Displays whether the transfer is SUCCESSFUL/FAILED.
#-----*/
#include <fcntl.h>
#include <iomanip>

#include "globus_ftp_client.h"
#include "globus_common.h"

#include <iostream>
#include <string>
#include <fstream>

#include "common.C"
#include "globusCallback.C"
using namespace std;
//FILE *fd;
#define MAX_BUFFER_SIZE 2048
#define ERROR -1
#define SUCCESS 0

class GridFTP:public InputConfig,public GlobusCallback {

    globus_ftp_client_handleattr_t    hattr;
    globus_ftp_client_operationattr_t oattr;
    globus_ftp_client_handle_t        handle;
    globus_byte_t                      buffer[MAX_BUFFER_SIZE];
    globus_size_t                      buffer_length ;
    globus_result_t                    status;
    char *                              tmpstr;

public :
    GridFTP();
    ~GridFTP();
    int gridftpClientToServer(string ,string);

};

/*-----*/
Function      : done_cb(...)

Used          : gridftpClientToServer(...),gridftpThirdPartyTransfer(...),gridftpServerToClient(...)

Description   : It is a callback function ,it is called when the transfer is
                completely finished, i.e. both the data channel and control channel
                exchange.
                Here it simply sets a global variable (done) to true so the main
                program will exit the while loop.
/*-----*/

namespace gridftp_cb {

    static void done_cb ( void *user_arg , globus_ftp_client_handle_t *handle , globus_object_t *err) {

        GridFTP *monitor=(GridFTP*)user_arg;

        char * tmpstr;
        if ( err ) {
            cout<<"\t\t Status : File Transferred Failed "<<endl;
            cout<<"\t\t ERROR : "<<globus_object_printable_to_string(err)<<endl;
        }
        else
            cout<<"\t\t Status : File Transferred Successfully"<<endl;

        return;
        monitor->setDoneValue();
    };

}

/*-----*/
Function      : data_cb(...)

Used          : gridftpClientToServer(..)

Description   : read or write operation in the FTP Client library is asynchronous.A
                callback of this type is passed to such data operation function calls.
                It is called when the user supplied buffer has been successfully
                transferred to the kernel.
                Note: That does not mean it has been successfully transmitted,instead it
                just reads the next block of data and calls register_write/register_read again.
/*-----*/

static void data_cb (void *user_arg , globus_ftp_client_handle_t *handle,globus_object_t * err, globus_byte_t *buffer,
                    globus_size_t length , globus_off_t offset,globus_bool_t eof)
{
    if ( err ) {
        cout<<"\t\t ERROR : "<<globus_object_printable_to_string(err)<<endl;
    }
    else {
        if ( !eof ) {
            FILE *fd = (FILE *) user_arg;
            int rc;
            rc = fread(buffer, 1, MAX_BUFFER_SIZE, fd);
            if ( ferror(fd) != SUCCESS ) {
                cout<<"\t\t Error : Function data_cb\n"<<"\t\t Error code : "<< errno<<endl;
                return;
            }
            globus_ftp_client_register_write(handle, buffer,rc,offset + length,feof(fd) != SUCCESS,
            data_cb, (void *) fd);
        }
        return;
    }
};
}


```

```
/*-----*/
Class      : GridFTP

Function    : GridFTP ( Constructor of class GridFTP )

Description : Used to
            a) Activate the GridFTP client module
            b) Initialize the mutex lock and condition variables
            c) Initialize the GridFTP client handle
/*-----*/

GridFTP::GridFTP() {

    buffer_length = MAX_BUFFER_SIZE;
    status = (globus_result_t)globus_module_activate(GLOBUS_FTP_CLIENT_MODULE);

    if ( status != GLOBUS_SUCCESS ) {

        tmpstr = globus_object_printable_to_string(globus_error_get(status));
        cout<<"\n\t Error: Failed to load GLOBUS_FTP_CLIENT_MODULE.\n\t Error Code "<<status<<"\n\t"<<tmpstr<<endl;
        exit(1);
    }
};

/*-----*/
Class      : GridFTP

Function    : ~GridFTP (Destructor of class GridFTP )

Description : Used to
            a) Destroy the GridFTP client handle
            b) Deactivate the GridFTP client module
/*-----*/

GridFTP::~GridFTP() {
    globus_module_deactivate_all();
};

/*-----*/
Class      : GridFTP

Function    : gridftpClientToServer(...)

Description : To transfer the file from client site to remote site
Input       : a) Source file path which is to transferred
            b) Destination url to store the file
            c) LogFile
/*-----*/

int GridFTP::gridftpClientToServer(string src,string dst) {

    int rc;
    FILE *fd;

    /* Initialize the handle attribute */
    if (globus_ftp_client_handleattr_init(&hattr) != GLOBUS_SUCCESS) {
        cout<<"\n\t\t ERROR : Failed to activate the ftp client handleattr\n";
        return 1;
    }

    /* Initialize the operation attribute */
    if (globus_ftp_client_operationattr_init(&oattr) != GLOBUS_SUCCESS) {
        cout<<"\n\t\t ERROR : Failed to initialize operationattr\n";
        return 1;
    }

    /* Initialize the handle */
    if (globus_ftp_client_handle_init(&handle,&hattr) != GLOBUS_SUCCESS) {
        cout<<"\n\t\t ERROR : Failed to initialize the handle\n";
        return 1;
    }

    continueOnCond();

    fd = fopen(src.c_str(),"r");
    if ( fd == NULL ) {
        cout<<"Error in opening local file"<<src;
        return 1;
    }

    /* Gridftp API call to start the put operation */
    status = globus_ftp_client_put(&handle,dst.c_str(),GLOBUS_NULL,GLOBUS_NULL,gridftp_cb::done_cb,this);
    if ( status != GLOBUS_SUCCESS ) {

        globus_object_t * err;
        err = globus_error_get(status);
        cout<<endl;
        fprintf(stderr, "\tError : %s", globus_object_printable_to_string(err));
        done = GLOBUS_TRUE;
    }
    else {
        rc = fread(buffer,1,MAX_BUFFER_SIZE,fd);
        globus_ftp_client_register_write(
            &handle,
            buffer,
            rc,
            0,
            feof(fd) != SUCCESS,
            gridftp_cb::data_cb,
            (void *) fd);
    }

    /* lock on condition */
    waitOnCond();
    fclose(fd);

    globus_ftp_client_handle_destroy(&handle);
    return 0;
};

/* start of main */
int main(int argc , char **argv) {

    int          hostCount;
    string        destinationUrl,sourceUrl,tmpUrl;
    string        sourceFile,destinationFile,outputFile,tmpFile;

    /* Creating the object */
    GridFTP      gridftp;

}


```

toty  
e ho  
mm  
, wo  
no  
obl  
<

10





# A Small Dilemma ?

```
/*-----*/
# gridftpClientToServer.C
#
# Description : Transfer the data file from client site to each remote site.
# and return the status of the file transfer.
#
# Input      : Source URL string of the source server , and the Destination URL
#             string of the destination server .
#             The format for URL string is
#             Source URL      : full path of the file to transferred
#             Destination URL : gsiftp://server-hostname/fulpathname/ofthefile/tobetransferred
#
# Output     : Displays whether the transfer is SUCCESSFUL/FAILED.
#-----*/

#include <fcntl.h>
#include <iomanip>

#include "globus_ftp_client.h"
#include "globus_common.h"

#include <iostream>
#include <string>
#include <fstream>

#include "common.C"
#include "globusCallback.C"
using namespace std;
//FILE *fd;
#define MAX_BUFFER_SIZE 2048
#define ERROR -1
#define SUCCESS 0

class GridFTP:public InputConfig,public GlobusCallback {
public:
    globus_ftp_client_handleattr_t hattr;
    globus_ftp_client_operationattr_t oattr;
    globus_ftp_client_handle_t handle;
    globus_byte_t buffer;
    globus_size_t buffer_size;
    globus_result_t status;
    char * tmpstr;

public:
    GridFTP();
    ~GridFTP();
    int gridftpClientToServer(string ,string);
};

/*-----*/
Function : done_cb(...)
Used      : gridftpClientToServer(...),gridftpThirdParty

Description : It is a callback function ,it is called when the data exchange
              completely finished, i.e. both the data of the client and server
              Here it simply sets a global variable (done) to true so that the
              program will exit the while loop.
/*-----*/

namespace gridftp_cb {
    static void done_cb ( void *user_arg , globus_ftp_client_handle_t *handle,
                          globus_size_t length, globus_off_t offset, globus_bool_t eof)
    {
        GridFTP *monitor=(GridFTP*)user_arg;
        char * tmpstr;
        if ( err ){
            cout<<"\t\t Status : File Transfer Failed\n";
            cout<<"\t\t ERROR : "<<globus_object_printable_to_string(err)<<endl;
        }
        else
            cout<<"\t\t Status : File Transfer Successful\n";

        monitor->setDoneValue();
        return;
    }
};

/*-----*/
Function : data_cb(...)
Used      : gridftpClientToServer(...)

Description : read or write operation in the FTP Client library is asynchronous.A
              callback of this type is passed to such data operation function calls.
              It is called when the user supplied buffer has been successfully
              transferred to the kernel.
              Note: That does not mean it has been successfully transmitted,instead it
              just reads the next block of data and calls register_write/register_read again.
/*-----*/

static void data_cb (void *user_arg , globus_ftp_client_handle_t *handle,globus_object_t * err, globus_byte_t *buffer,
                    globus_size_t length , globus_off_t offset,globus_bool_t eof)
{
    if ( err ) {
        cout<<"\t\t ERROR : "<<globus_object_printable_to_string(err)<<endl;
    }
    else {
        if ( !eof ) {
            FILE *fd = (FILE *) user_arg;
            int rc;
            rc = fread(buffer, 1, MAX_BUFFER_SIZE, fd);
            if ( ferror(fd) != SUCCESS ) {
                cout<<"\t\t Error : Function data_cb\n"<<"\t\t Error code : "<< errno<<endl;
                return;
            }
            globus_ftp_client_register_write(handle, buffer,rc,offset + length,feof(fd) != SUCCESS,
                                           data_cb, (void *) fd);
        }
        return;
    }
};
}

/*-----*/
#include <saga/saga.hpp>
#include <iostream>

int main (int argc, char** argv)
{
    try {
        saga::filesystem::file f( std::string(argv[1]) );
        f.copy( std::string(argv[2]) );
    }
    catch(saga::exception const & e)
    {
        std::cerr << "Error: " << e.what() << std::endl;
    }
}
/*-----*/
```

```
/*-----*/
Class      : GridFTP
Function   : GridFTP ( Constructor of class GridFTP )
Description : Used to
            a) Activate the GridFTP client module
            b) Initialize the mutex lock and condition variables
            c) Initialize the GridFTP client handle
/*-----*/

GridFTP::GridFTP() {
    buffer_length = MAX_BUFFER_SIZE;
    status = (globus_result_t)globus_module_activate(GLOBUS_FTP_CLIENT_MODULE);
    if ( status != GLOBUS_SUCCESS ) {
        tmpstr = globus_object_printable_to_string(globus_error_get(status));
        cout<<"\n\t Error: Failed to load GLOBUS_FTP_CLIENT_MODULE.\n\t Error Code : "<<status<<"\n\t"<<tmpstr<<endl;
        exit(1);
    }
};

/*-----*/
Class      : GridFTP
Function   : ~GridFTP (Destructor of class GridFTP )
Description : Used to
            a) Destroy the GridFTP client handle
            b) Deactivate the GridFTP client module
/*-----*/

globus_object_t * err;
err = globus_error_get(status);
cout<<endl;
fprintf(stderr, "\tError : %s", globus_object_printable_to_string(err));
done = GLOBUS_TRUE;
}
else {
    rc = fread(buffer,1,MAX_BUFFER_SIZE,fd);
    globus_ftp_client_register_write(
        shandle,
        buffer,
        rc,
        0,
        feof(fd) != SUCCESS,
        gridftp_cb::data_cb,
        (void *) fd);
}

/* lock on condition */
waitOnCond();
fclose(fd);

globus_ftp_client_handle_destroy(shandle);
return 0;
};

/* start of main */
int main(int argc , char **argv) {
    int hostCount;
    string destinationUrl,sourceUrl,tmpUrl;
    string sourceFile,destinationFile,outputFile,tmpFile;

    /* Creating the object */
    GridFTP gridftp;
}

/*-----*/
```

toty  
e ho  
mm

ad

# A Small Dilemma ?

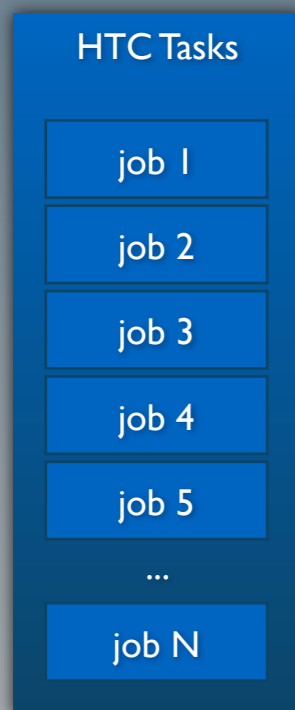
---

- After several years of prototyping, testing and hardening, we presented SAGA as the holy grail of distributed computing to the user communities
- Lots of advertising, demos, workshops, tutorials
- But: Uptake very slow and not as expected
- Almost by accident, the problem got solved with the “SAGA *Big-Job*” framework

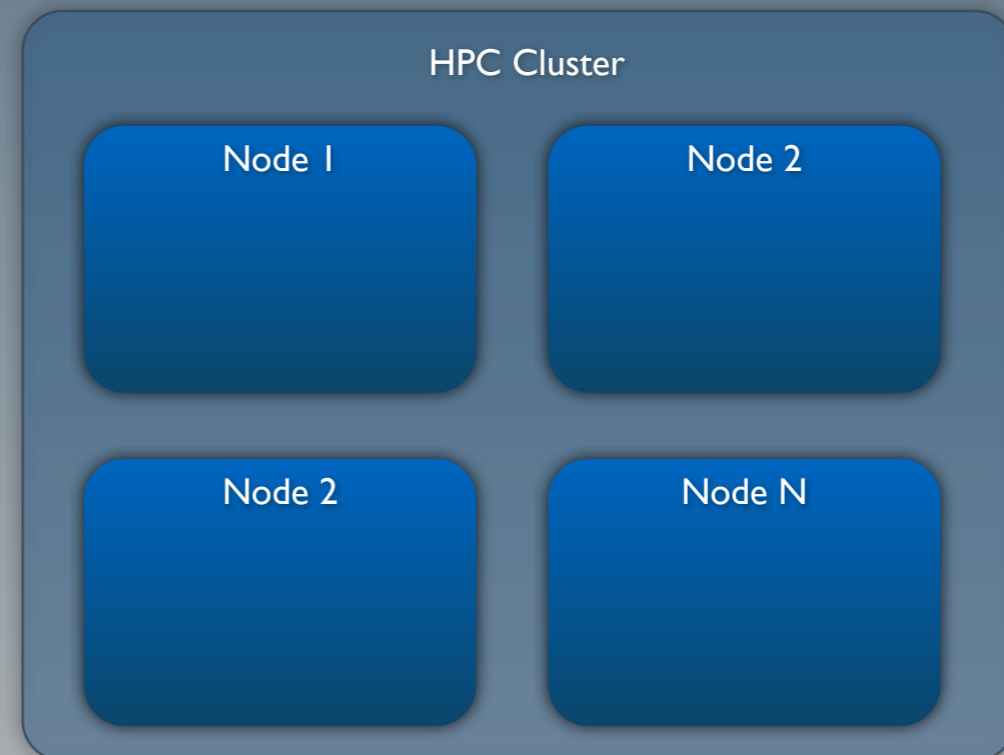
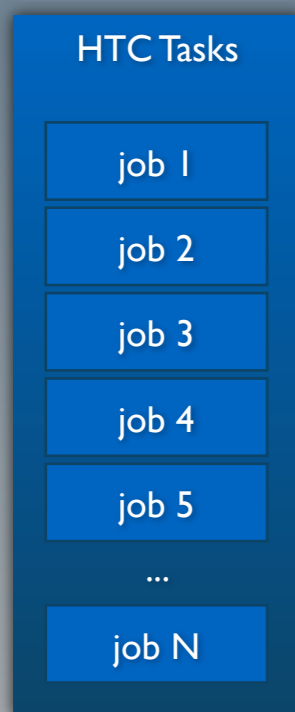
- *Big-Job* (a.k.a. *glide-in*, or *pilot-job*) is a simple distributed programming abstraction written in SAGA (Python)
- It allows to run lots of HTC jobs (often single-core) transparently on HPC machines using overlay scheduling



- *Big-Job* (a.k.a. *glide-in*, or *pilot-job*) is a simple distributed programming abstraction written in SAGA (Python)
- It allows to run lots of HTC jobs (often single-core) transparently on HPC machines using overlay scheduling

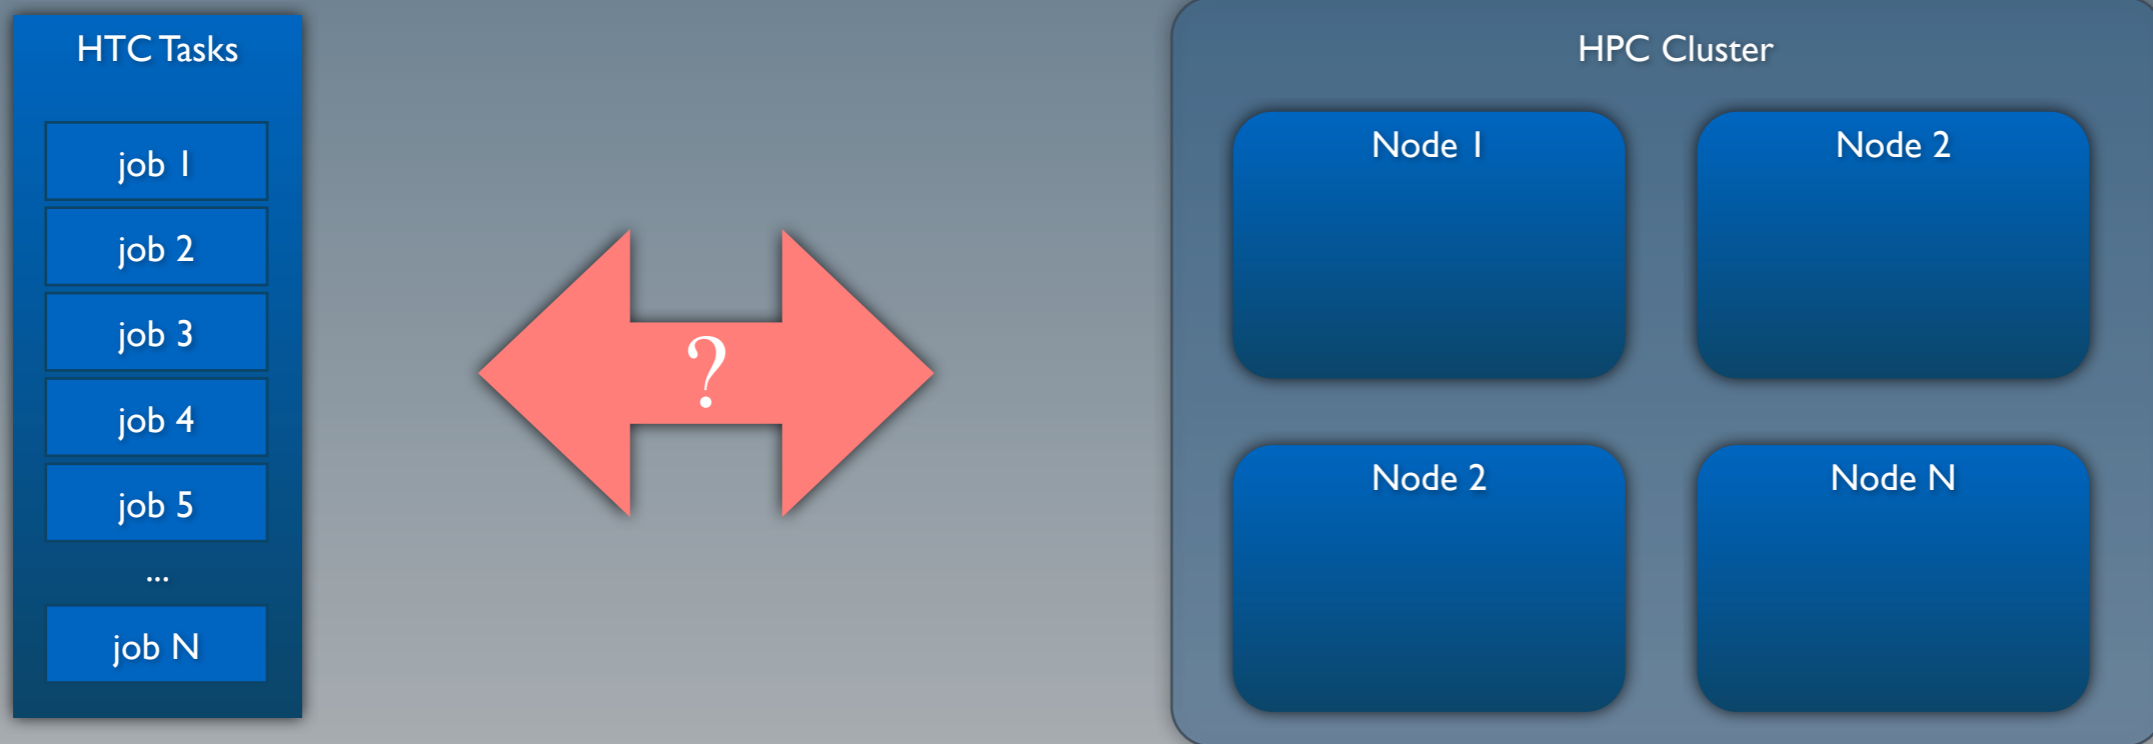


- *Big-Job* (a.k.a. *glide-in*, or *pilot-job*) is a simple distributed programming abstraction written in SAGA (Python)
- It allows to run lots of HTC jobs (often single-core) transparently on HPC machines using overlay scheduling



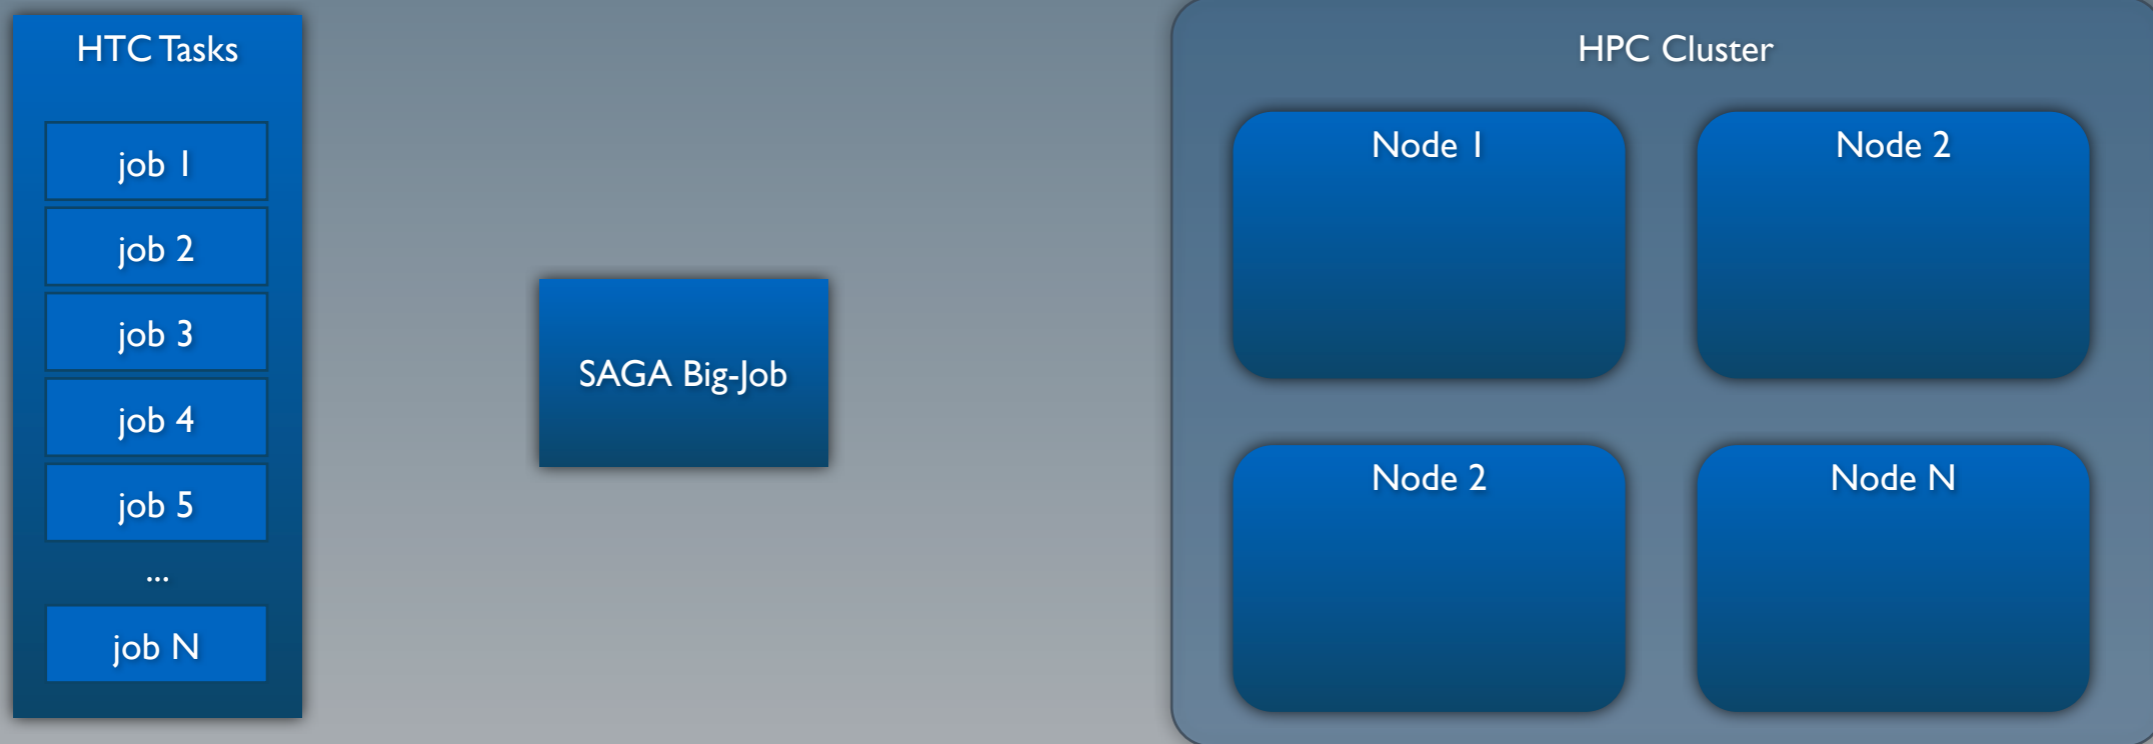
||

- *Big-Job* (a.k.a. *glide-in*, or *pilot-job*) is a simple distributed programming abstraction written in SAGA (Python)
- It allows to run lots of HTC jobs (often single-core) transparently on HPC machines using overlay scheduling



||

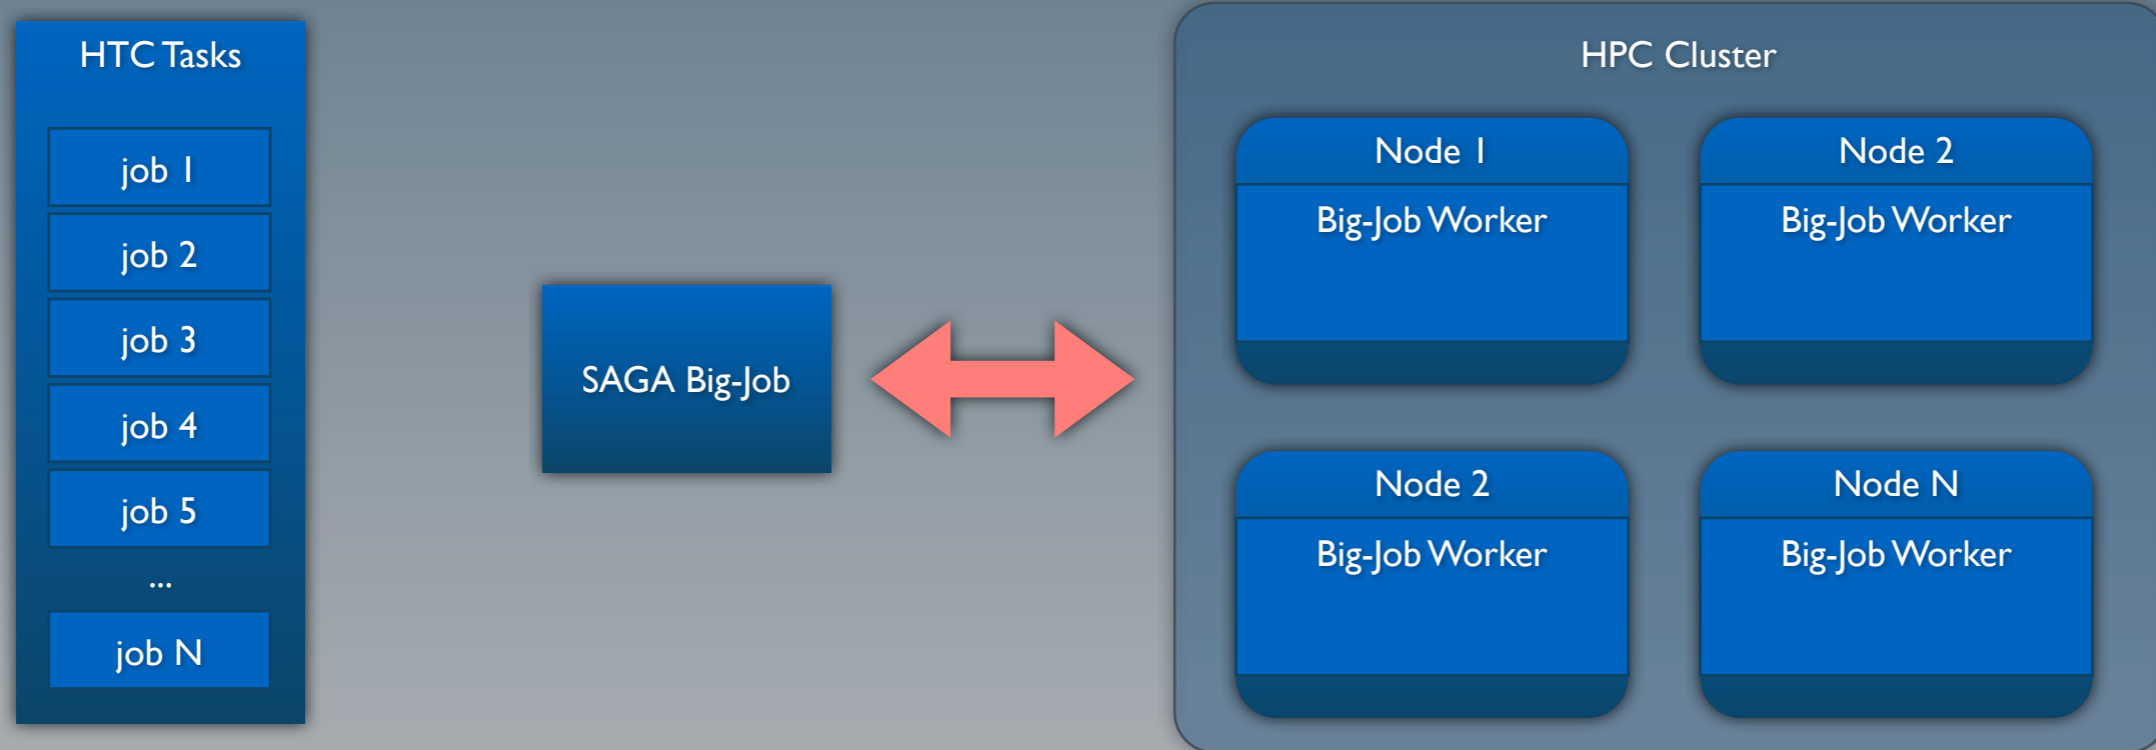
- *Big-Job* (a.k.a. *glide-in*, or *pilot-job*) is a simple distributed programming abstraction written in SAGA (Python)
- It allows to run lots of HTC jobs (often single-core) transparently on HPC machines using overlay scheduling



||

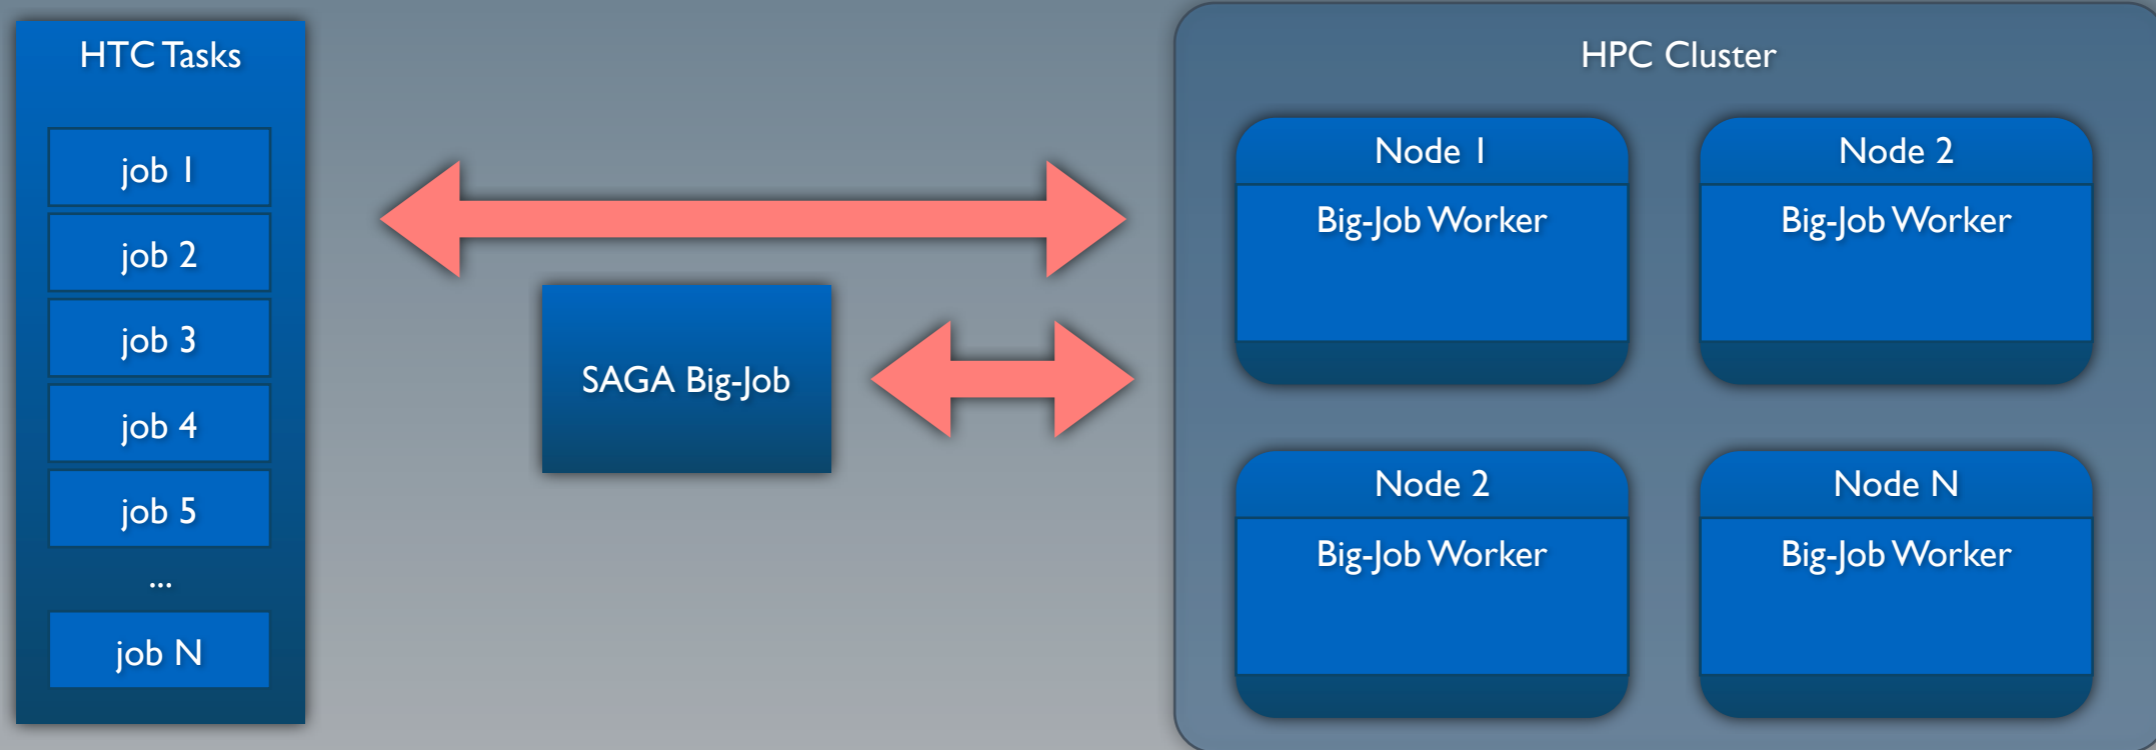


- *Big-Job* (a.k.a. *glide-in*, or *pilot-job*) is a simple distributed programming abstraction written in SAGA (Python)
- It allows to run lots of HTC jobs (often single-core) transparently on HPC machines using overlay scheduling



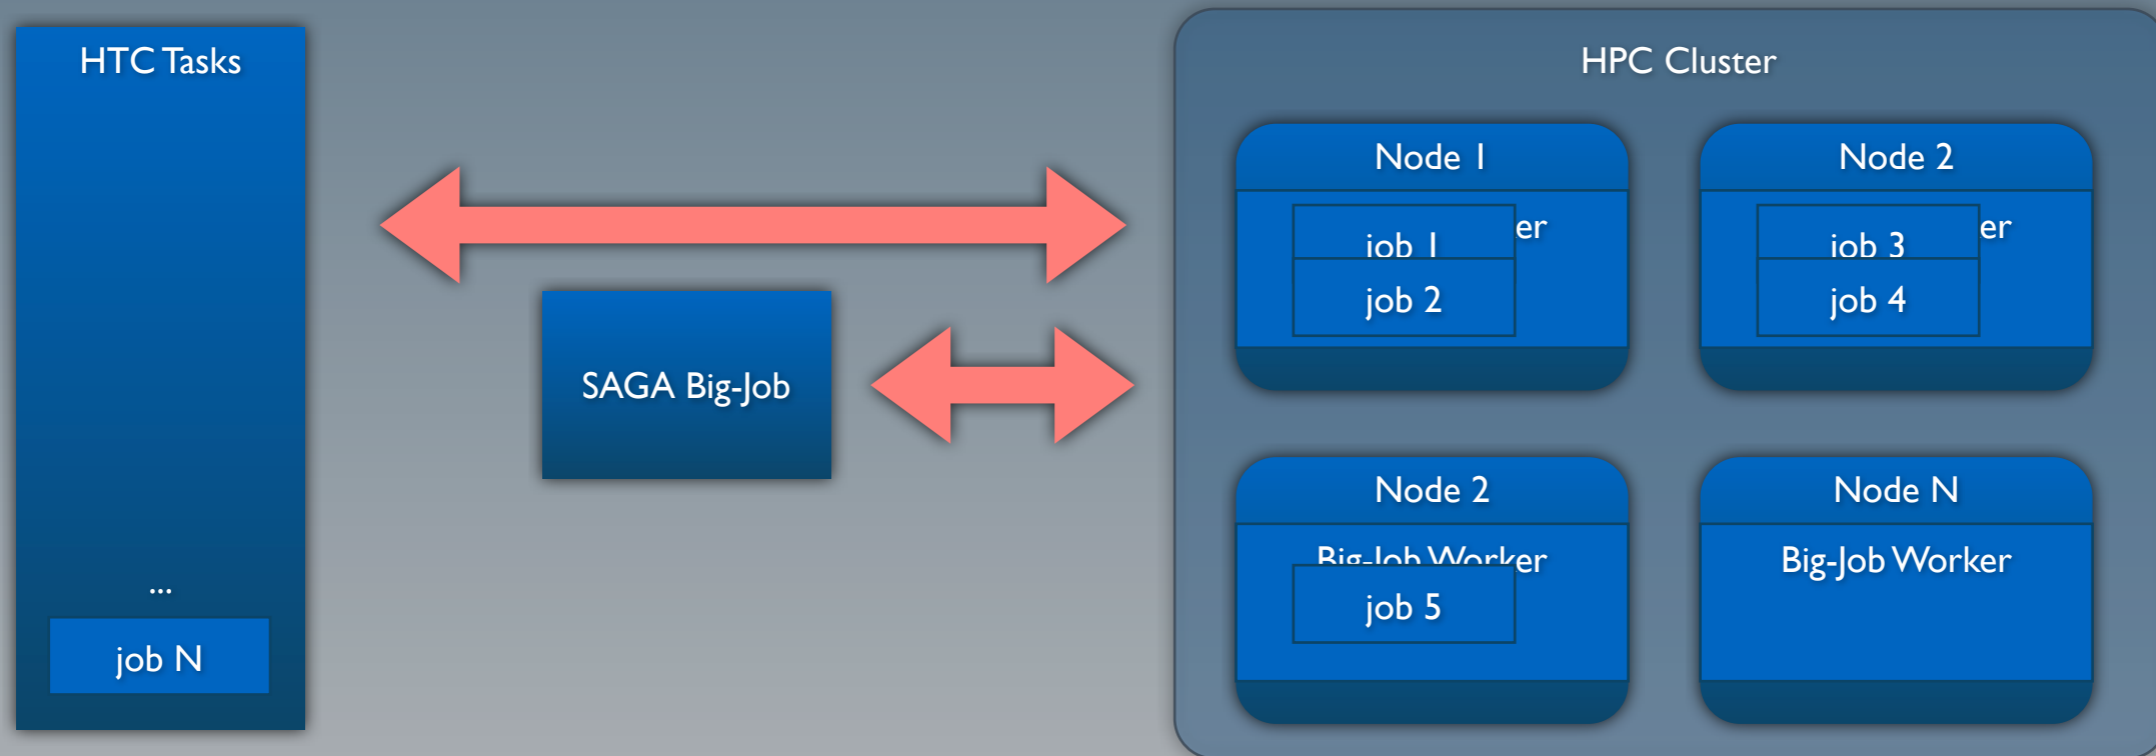
||

- *Big-Job* (a.k.a. *glide-in*, or *pilot-job*) is a simple distributed programming abstraction written in SAGA (Python)
- It allows to run lots of HTC jobs (often single-core) transparently on HPC machines using overlay scheduling



||

- *Big-Job* (a.k.a. *glide-in*, or *pilot-job*) is a simple distributed programming abstraction written in SAGA (Python)
- It allows to run lots of HTC jobs (often single-core) transparently on HPC machines using overlay scheduling



||

- *Big-Job* (a.k.a. *glide-in*, or *pilot-job*) is a simple distributed programming abstraction written in SAGA (Python)
- It allows to run lots of HTC jobs (often single-core) transparently on HPC machines using overlay scheduling
- *Big-Job* suddenly sparked a lot of interest in SAGA
  - Makes a huge problem simply disappear
  - Allows to use legacy code (non-intrusive)
  - Runs everywhere - even on Hector (CRAY) and EC2



- Apparently, we didn't understand the user communities and their requirements properly
- Users (especially the non-technical users) don't want another API. They want simple solutions for their every day problems. Abstractions can help !
- It turns out that SAGA is perfect to develop distributed programming abstractions:
  - Hides specific middleware implementation details
  - Allows concurrent cross-infrastructure resource usage
  - Optimisation and “adaptation” can happen “behind the scenes”

# Abstractions, Abstractions, Abstractions !

- Abstractions: from generic to specialised:
  - Pilot-Job (a.k.a. Big-job, a.k.a. Glide-in)
  - Pilot-Data (data affinity)
  - Master-Worker
  - Workflow (a.k.a. DAG)
  - Peer-to-Peer
  - Replica exchange
  - Map-Reduce
  - All-Pairs
- See eSI 3DPAS theme: <http://www.esi.ac.uk/research-themes/5>

# Abstractions, Abstractions, Abstractions !

- Abstractions: from generic to specialised:
  - Pilot-Job (a.k.a. Big-job, a.k.a. Glide-in)
  - Pilot-Data (data affinity)
  - Master-Worker
  - Workflow (a.k.a. DAG)
    - Peer-to-Peer
  - Replica exchange
    - Map-Reduce
    - All-Pairs
- See eSI 3DPAS theme: <http://www.esi.ac.uk/research-themes/5>

# Is There Potential for (PhD) Research ?

## ANOVA: ANALYSIS OF VALUE

IS YOUR RESEARCH WORTH ANYTHING?

Developed in 1912 by geneticist R.A. Fisher, the Analysis of Value is a powerful statistical tool designed to test the significance of one's work.



am i  
wasting  
my time?

Significance is determined by comparing one's research with the **Dull Hypothesis**:

$$H_0: \mu_1 = \mu_2 ?$$

where,

$H_0$  : the Dull Hypothesis

$\mu_1$  : significance of your research

$\mu_2$  : significance of a monkey typing randomly on a typewriter in a forest where no one hears it.

WWW.PHDCOMICS.COM

JORGE GUAN © 2007

The test involves computation of the  $F'd$  ratio:

$$F'd = \frac{\text{sum(people who care about your research)}}{\text{world population}}$$

This ratio is compared to the F distribution with  $I-1$ ,  $N_T$  degrees of freedom to determine a  $p(\text{in your pants})$  value. A low  $p(\text{in your pants})$  value means you're on to something good (though statistically improbable).

Type I/II Errors

The Analysis of Value must be used carefully to avoid the following two types of errors:

Type I: You incorrectly believe your research is not Dull.

Type II: No conclusions can be made. Good luck graduating.

Of course, this test assumes both Independence and Normality on your part, neither of which is likely true, which means *it's not your problem*.



# Is There Potential for (PhD) Research ?

- Previous work and experience has shown the relevance and importance of abstractions in distributed computing
  - Especially interesting: with abstractions, optimisation can be done “behind the scenes” (hidden from the user)
- Some fundamental questions:
  - What are interesting upcoming challenges in future distributed applications and systems ?
  - Can some of the challenges possibly be address by using distributed programming abstractions ?
  - What is the current state of research in distributed programming abstractions ?

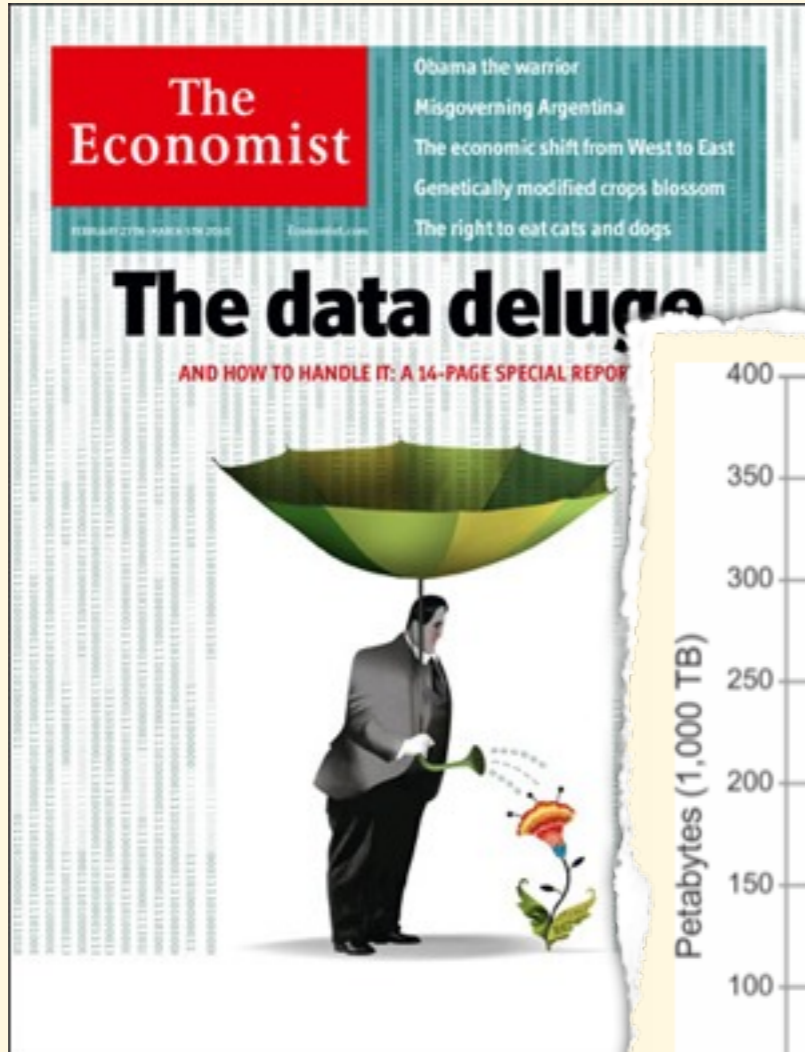
# What are the Upcoming Challenges ?

- The *Data Deluge*

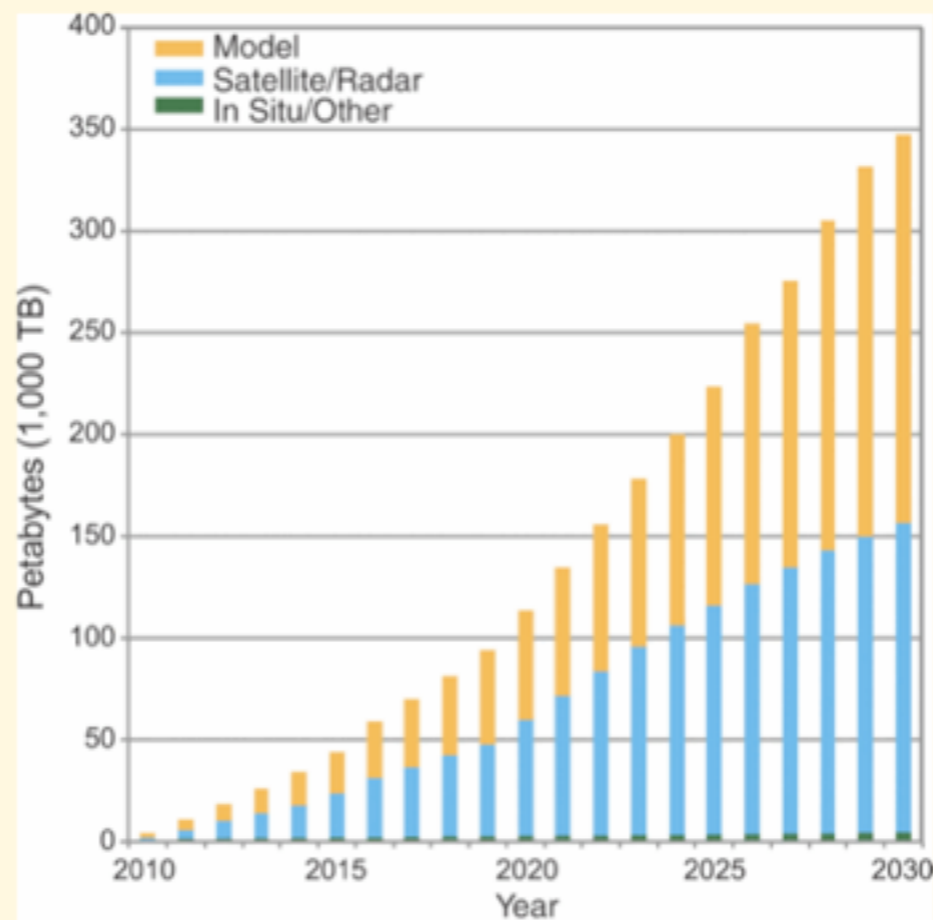
*“The amount of genomic data available for study is increasing at a rate similar to that of Moore’s Law”* <http://www.mcs.anl.gov/uploads/cels/papers/P1238.pdf>

- *“Store now - process later”* might not always be an option
- *Adaptive* algorithms will be required to handle vast amounts of dynamic (streaming) data
- Methods to characterise/predict data and especially dynamic changes in data will play a key role
- Infrastructure Challenge
  - Applications will have to spread and run across different heterogeneous infrastructures (possibly even simultaneously)
  - Probably other objectives besides *“minimise makespan”*

# What are the Upcoming Challenges ?



available for st  
's Law" <http://www.r>



heterogeneous

- Probably other

in simultaneously)

makespan"

# What are the Upcoming Challenges ?

- The *Data Deluge*

*“The amount of genomic data available for study is increasing at a rate similar to that of Moore’s Law”* <http://www.mcs.anl.gov/uploads/cels/papers/P1238.pdf>

- *“Store now - process later”* might not always be an option
- *Adaptive* algorithms will be required to handle vast amounts of dynamic (streaming) data
- Methods to characterise/predict data and especially dynamic changes in data will play a key role
- Infrastructure Challenge
  - Applications will have to spread and run across different heterogeneous infrastructures (possibly even simultaneously)
  - Probably other objectives besides *“minimise makespan”*

# Can Distributed Abstractions Help ?

- There might be lots of new challenges, but from an application perspective things will mostly stay the same (!)
  - Applications will still be using common patterns and abstractions
  - But: input and objectives might change and they might change *DYNAMICALLY*
- Once we have understood the details and dynamics of these new challenges, we can *encapsulate* them inside distributed programming abstractions
- If this hasn't been done, it could be a nice intellectual and practically relevant contribution to the field



# A Case for Dynamic Optimisation

- DPAs supporting **dynamic** optimisation to address dynamic data challenges:
  - Adaptive execution strategy based on input, output and system characteristics
  - Ability to “understand” and “predict” data (characteristics)
- DPAs supporting **autonomous** adaption to address rising complexity in infrastructure

- There doesn't seem to be a lot of work that revisits DPAs in the context of:
  - Streaming input / output data
  - Data with dynamically changing characteristics
  - Autonomic | adaptive | dynamic | optimisation
- Lots of work in static optimisation (“*fine-tuning*”)
  - Optimise throughput for workload X on (idealised) platform Y
  - Usually confined to a specific middleware
- There seems to be a lot of uncharted territory, but that remains to be proved (more reading!)

IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, VOL. 21, NO. 1, JANUARY 2010

33

## All-Pairs: An Abstraction for Data-Intensive Computing on Campus Grids

Christopher Moretti, *Student Member, IEEE*, Hoang Bui, *Student Member, IEEE*, Karen Hollingsworth, Brandon Rich, Patrick Flynn, *Senior Member, IEEE*, and Douglas Thain, *Member, IEEE*

**Abstract**—Today, campus grids provide users with easy access to thousands of CPUs. However, it is not always easy for nonexpert users to harness these systems effectively. A large workload composed in what seems to be the obvious way by a naive user may accidentally abuse shared resources and achieve very poor performance. To address this problem, we argue that campus grids should provide end users with high-level abstractions that allow for the easy expression and efficient execution of data-intensive workloads. We present one example of an abstraction—All-Pairs—that fits the needs of several applications in biometrics, bioinformatics, and data mining. We demonstrate that an optimized All-Pairs abstraction is both easier to use than the underlying system, achieve performance orders of magnitude better than the obvious but naive approach, and is both faster and more efficient than a tuned conventional approach. This abstraction has been in production use for one year on a 500 CPU campus grid at the University of Notre Dame and has been used to carry out a groundbreaking analysis of biometric data.

**Index Terms**—All-pairs, biometrics, cloud computing, data intensive computing, grid computing.

### 1 INTRODUCTION

MANY fields of science and engineering have the potential to use large numbers of CPUs to attack problems of enormous scale. Campus-scale computing grids are now a standard tool employed by many academic institutions to provide large-scale computing power. Using middleware such as Condor [40] or Globus [21], many disparate clusters and stand-alone machines may be joined into a single computing system with many providers and consumers. Today, campus grids of about one thousand machines are commonplace [41], and are being grouped into larger structures, such as the 20,000-CPU Indiana Diagrid and the 40,000-CPU Open Science Grid [36].

Campus grids have the unique property that consumers of the system must always defer to the needs of the resource providers. For example, if a desktop computer is donated to the campus grid, then a visiting job may use it during idle times, but will be preempted when the owner is busy at the keyboard. If a research cluster is donated to the campus grid, visiting jobs may use it, but might be preempted by higher priority batch jobs submitted by the owner of the cluster. In short, the user of the system has access to an enormous number of CPUs, but must expect to be preempted from many of them as a normal condition.

Because of this property, scaling up an application to a campus grid is a nontrivial undertaking. Parallel libraries and languages such as MPI [18], OpenMP [14], and Cilk [8] are not usable in this context because they do not explicitly address preemption and failure as a normal case. Instead,

• The authors are with the Department of Computer Science and Engineering, University of Notre Dame, 384 Fitzpatrick Hall, Notre Dame, IN 46556. E-mail: {cmoretti, hbui, kholling, brich, flynn, dthain}@nd.edu.

Manuscript received 22 July 2008; revised 11 Feb. 2009; accepted 9 Mar. 2009; published online 13 Mar. 2009.

Recommended for acceptance by D. Bader.

For information on obtaining reprints of this article, please send e-mail to: tpd@scomputer.org, and reference IEEECS Log Number TPDS-2008-07-0277. Digital Object Identifier no. 10.1109/TPDS.2009.49.

1045-9219/10/\$26.00 © 2010 IEEE

large workloads must be specified as a set of sequential processes connected by files. End users must carefully arrange the I/O behavior of their workloads. Bad configurations can result in poor performance, outright failure of the application, and abuse of physical resources shared by others. All too often, an end user composes a workload that runs correctly on one machine, then on 10 machines, but fails disastrously on 1,000 machines.

Providing an *abstraction* is one approach to avoiding these problems. An abstraction allows a user to declare a workload composed of multiple sequential programs and the data that they process, while hiding the details of how the workload will be realized in the system. Abstracting away details hides complications that are not apparent or important to a novice, limiting the opportunity for disasters. Because an abstraction states a workload in a declarative way, it can be realized within the grid in whatever way satisfies cost, policy, and performance constraints. Abstractions could also be implemented in other kinds of systems, such as dedicated clusters or multicore CPUs, but we do not address those here.

We have implemented one such abstraction—All-Pairs—for a class of problems found in many fields. All-Pairs is the Cartesian product of a large number of objects with a custom comparison function. While simple to state, it is nontrivial to carry out on large problems that require hundreds of nodes running for several days. All-Pairs is similar in spirit to other abstractions such as Dryad [28], Map-Reduce [16], Pegasus [17], and Swift [42], but it addresses a different category of applications.

Our implementation of All-Pairs is currently in production use on a 500 CPU campus grid at the University of Notre Dame, using Condor [40] to manage the CPUs and Chirp [39] to manage the storage. To demonstrate the performance benefits of using an abstraction, we compare two different implementations. The *conventional* implementation executes the specification by simply submitting a series of batch jobs that use a central file server to read data on demand and write

Published by the IEEE Computer Society

be a lot of work that revisits

ut data

changing characteristics

dynamic | optimisation

optimisation (“fine-tuning”)

or workload X on (idealised) platform Y

pecific middleware

t of uncharted territory, but that  
more reading!)

IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, VOL. 21, NO. 1, JANUARY 2010

## All-Pairs: An Abstraction for Data-Intensive Computing on Campus Clusters

Christopher Moretti, *Student Member, IEEE*, Hoang Bui, *Student Member, IEEE*, Brandon Rich, Patrick Flynn, *Senior Member, IEEE*, and Douglas Thain

**Abstract**—Today, campus grids provide users with easy access to thousands of CPUs. However, users do not always harness these systems effectively. A large workload composed in what seems to be an ad hoc manner can accidentally abuse shared resources and achieve very poor performance. To address this problem, we present a framework that provides end users with high-level abstractions that allow for the easy expression and efficient execution of data-intensive applications. We present one example of an abstraction—All-Pairs—that fits the needs of several applications in the field of data mining. We demonstrate that an optimized All-Pairs abstraction is both easier to use than a naive approach and is both faster and more accurate than a conventional approach. This abstraction has been in production use for one year on a 500 CPU campus grid and has been used to carry out a groundbreaking analysis of biometric data.

**Index Terms**—All-pairs, biometrics, cloud computing, data intensive computing, grid computing

### 1 INTRODUCTION

MANY fields of science and engineering have the potential to use large numbers of CPUs to attack problems of enormous scale. Campus-scale computing grids are now a standard tool employed by many academic institutions to provide large-scale computing power. Using middleware such as Condor [40] or Globus [21], many disparate clusters and stand-alone machines may be joined into a single computing system with many providers and consumers. Today, campus grids of about one thousand machines are commonplace [41], and are being grouped into larger structures, such as the 20,000-CPU Indiana Diagrid and the 40,000-CPU Open Science Grid [36].

Campus grids have the unique property that consumers of the system must always defer to the needs of the resource providers. For example, if a desktop computer is donated to the campus grid, then a visiting job may use it during idle times, but will be preempted when the owner is busy at the keyboard. If a research cluster is donated to the campus grid, visiting jobs may use it, but might be preempted by higher priority batch jobs submitted by the owner of the cluster. In short, the user of the system has access to an enormous number of CPUs, but must expect to be preempted from many of them as a normal condition.

Because of this property, scaling up an application to a campus grid is a nontrivial undertaking. Parallel libraries and languages such as MPI [18], OpenMP [14], and Cilk [8] are not usable in this context because they do not explicitly address preemption and failure as a normal case. Instead,

• The authors are with the Department of Computer Science and Engineering, University of Notre Dame, 384 Fitzpatrick Hall, Notre Dame, IN 46556. E-mail: {cmoretti, hbui, kholling, brich, flynn, dthain}@nd.edu.

Manuscript received 22 July 2008; revised 11 Feb. 2009; accepted 9 Mar. 2009; published online 13 Mar. 2009.

Recommended for acceptance by D. Bader. For information on obtaining reprints of this article, please send e-mail to: [tpds@computer.org](mailto:tpds@computer.org), and reference IEEECS Log Number TPDS-2008-07-0277. Digital Object Identifier no. 10.1109/TPDS.2009.49.

1045-9219/10/\$26.00 © 2010 IEEE

large workloads must be broken into smaller processes connected to the grid. The way to arrange the I/O behavior of the application, and the way to manage the data, are important to a novice, like a biologist. Because an abstraction way, it can be realized in a way that satisfies cost, policy, and other concerns. Implementations could also be implemented in a way such as dedicated clusters to address those here.

Providing an abstraction to address these problems. An abstraction of a workload composed of the data that they process. The way they process the workload will be different. Away details hides complexity that is important to a novice, like a biologist. Because an abstraction way, it can be realized in a way that satisfies cost, policy, and other concerns. Implementations could also be implemented in a way such as dedicated clusters to address those here.

We have implemented an abstraction for a class of problems. The Cartesian product of a custom comparison function. It is nontrivial to carry out hundreds of nodes in a similar in spirit to Map-Reduce [16]. It addresses a different class of problems.

Our implementation uses a 500 CPU campus grid, using Condor to manage the storage and benefits of using an abstraction. The implementation uses the specification by the user that use a central file system.

Cluster Comput (2010) 13: 243–256  
DOI 10.1007/s10586-010-0134-7

## Harnessing parallelism in multicore clusters with the All-Pairs, Wavefront, and Makeflow abstractions

Li Yu · Christopher Moretti · Andrew Thrasher · Scott Emrich · Kenneth Judd · Douglas Thain

Received: 9 November 2009 / Accepted: 16 March 2010 / Published online: 23 April 2010  
© Springer Science+Business Media, LLC 2010

**Abstract** Both distributed systems and multicore systems are difficult programming environments. Although the expert programmer may be able to carefully tune these systems to achieve high performance, the non-expert may struggle. We argue that high level abstractions are an effective way of making parallel computing accessible to the non-expert. An abstraction is a regularly structured framework into which a user may plug in simple sequential programs to create very large parallel programs. By virtue of a regular structure and declarative specification, abstractions may be materialized on distributed, multicore, and distributed multicore systems with robust performance across a wide range of problem sizes. In previous work, we presented the All-Pairs abstraction for computing on distributed systems of single CPUs. In this paper, we extend All-Pairs to multicore systems, and introduce the Wavefront and Makeflow abstractions, which represent a number of problems in economics and bioinformatics. We demonstrate good scaling of both abstractions up to 32 cores on one machine and hundreds of cores in a distributed system.

**Keywords** Abstractions · Multicore · Distributed systems · Bioinformatics · Economics

L. Yu (✉) · C. Moretti · A. Thrasher · S. Emrich · D. Thain  
Department of Computer Science and Engineering, University of Notre Dame, South Bend, USA  
e-mail: [lyu2@nd.edu](mailto:lyu2@nd.edu)

K. Judd  
Hoover Institution, Stanford University, Stanford, USA

that revisits

ics

on

uning”)

realised) platform Y

territory, but that



## All-Pairs: An Abstraction for Data-Intensive Computing on Campus Grids

Christopher Moretti, *Student Member, IEEE*, Hoang Bui, *Student Member, IEEE*, Brandon Rich, Patrick Flynn, *Senior Member, IEEE*, and Douglas Thain

**Abstract**—Today, campus grids provide users with easy access to thousands of CPUs. However, users do not always harness these systems effectively. A large workload composed in what seems to be an efficient way can accidentally abuse shared resources and achieve very poor performance. To address this problem, we propose high-level abstractions that allow for the easy expression and efficient execution of data-intensive workloads. We present one example of an abstraction—All-Pairs—that fits the needs of several applications, such as data mining. We demonstrate that an optimized All-Pairs abstraction is both easier to use than the obvious but naive approach, and is both faster and more efficient than a conventional approach. This abstraction has been in production use for one year on a 500-CPU campus grid and has been used to carry out a groundbreaking analysis of biometric data.

**Index Terms**—All-pairs, biometrics, cloud computing, data intensive computing, grid computing

### 1 INTRODUCTION

MANY fields of science and engineering have the potential to use large numbers of CPUs to attack problems of enormous scale. Campus-scale computing grids are now a standard tool employed by many academic institutions to provide large-scale computing power. Using middleware such as Condor [40] or Globus [21], many disparate clusters and stand-alone machines may be joined into a single computing system with many providers and consumers. Today, campus grids of about one thousand machines are commonplace [41], and are being grouped into larger structures, such as the 20,000-CPU Indiana Diagrid and the 40,000-CPU Open Science Grid [36].

Campus grids have the unique property that consumers of the system must always defer to the needs of the resource providers. For example, if a desktop computer is donated to the campus grid, then a visiting job may use it during idle times, but will be preempted when the owner is busy at the keyboard. If a research cluster is donated to the campus grid, visiting jobs may use it, but might be preempted by higher priority batch jobs submitted by the owner of the cluster. In short, the user of the system has access to an enormous number of CPUs, but must expect to be preempted from many of them as a normal condition.

Because of this property, scaling up an application to a campus grid is a nontrivial undertaking. Parallel libraries and languages such as MPI [18], OpenMP [14], and Cilk [8] are not usable in this context because they do not explicitly address preemption and failure as a normal case. Instead,

• The authors are with the Department of Computer Science and Engineering, University of Notre Dame, 384 Fitzpatrick Hall, Notre Dame, IN 46556. E-mail: {cmoretti, hbui, kholling, brich, flynn, dthain}@nd.edu.

Manuscript received 22 July 2008; revised 11 Feb. 2009; accepted 9 Mar. 2009; published online 13 Mar. 2009.

Recommended for acceptance by D. Bader. For information on obtaining reprints of this article, please send e-mail to: tpd@computer.org, and reference IEEECS Log Number TPDS-2008-07-0277. Digital Object Identifier no. 10.1109/TPDS.2009.49.

large workloads must be broken into smaller processes connected by a network. The I/O behavior of these processes can result in bottlenecks in the application, and in some cases, the application may not run correctly on one machine and fails disastrously on another.

Providing an abstraction that hides these problems. An abstraction of a workload composed of many small processes of the data that they process. The abstraction will be important to a novice, like a user of a desktop computer. Because an abstraction hides away details, it can be realized in a way that satisfies cost, policy, and other concerns. Abstractions could also be implemented such as dedicated clusters to address those here.

We have implemented an abstraction for a class of problems in the Cartesian product of a custom comparison. It is nontrivial to carry out hundreds of nodes in a distributed system similar in spirit to Map-Reduce [16]. It addresses a different set of concerns.

Our implementation uses a 500-CPU campus grid, using Condor [40] to manage the storage and retrieval of data. The benefits of using an abstraction are that the specification by which the application is run that use a central file system.

## Harnessing parallelism in multicore clusters with Wavefront, and Makeflow abstractions

Li Yu · Christopher Moretti · Andrew Thrasher · Scott Emrich · Kenneth Judd · Douglas Thain

Received: 9 November 2009 / Accepted: 16 March 2010 / Published online: 23 April 2010  
© Springer Science+Business Media, LLC 2010

**Abstract** Both distributed systems and multicore systems are difficult programming environments. Although the expert programmer may be able to carefully tune these systems to achieve high performance, the non-expert may struggle. We argue that high level abstractions are an effective way of making parallel computing accessible to the non-expert. An abstraction is a regularly structured framework into which a user may plug in simple sequential programs to create very large parallel programs. By virtue of a regular structure and declarative specification, abstractions may be materialized on distributed, multicore, and distributed multicore systems with robust performance across a wide range of problem sizes. In previous work, we presented the All-Pairs abstraction for computing on distributed systems of single CPUs. In this paper, we extend All-Pairs to multicore systems, and introduce the Wavefront and Makeflow abstractions, which represent a number of problems in economics and bioinformatics. We demonstrate good scaling of both abstractions up to 32 cores on one machine and hundreds of cores in a distributed system.

**Keywords** Abstractions · Multicore · Distributed systems · Bioinformatics · Economics

L. Yu (✉) · C. Moretti · A. Thrasher · S. Emrich · D. Thain  
Department of Computer Science and Engineering, University of Notre Dame, South Bend, USA  
e-mail: lyu2@nd.edu

K. Judd  
Hoover Institution, Stanford University, Stanford, USA

### 1 Introduction

Distributed systems are very challenging to program. We refer to all of the problems we wish to execute a parallelism is confronted. The workload should be distributed in a way that is most efficient to be used? Will the number of processors in the system and workload be such as the size of a problem? a completely different set of concerns.

Multicore systems are also very challenging. The orders of magnitude of the problem are similar. How should work be distributed? How should message passing be used? Will the software that harnesses these resources. Inadvertent poor choices can result in poor performance or even outright failures of the application. Poor choices can also lead to inefficient use of shared resources and abuse of the distributed system's infrastructure such as job queues and matchmaking software.

We argue that an abstraction is a declarative structure that allows a user to scale to very large problems. Such an abstraction gives the user an interface to define their problem in terms of data and computation requirements, while hiding the details of how the problem will be realized in the system. Abstracting away details also hides away the complexity of the underlying hardware. Abstractions allow non-expert users to solve data-intensive problems that run on distributed batch systems.

## All-Pairs: An Abstraction for Data-Intensive Cloud Computing

Christopher Moretti, Jared Bulosan, Douglas Thain, and Patrick J. Flynn  
Department of Computer Science and Engineering, University of Notre Dame \*

### Abstract

Although modern parallel and distributed computing systems provide easy access to large amounts of computing power, it is not always easy for non-expert users to harness these large systems effectively. A large workload composed in what seems to be the obvious way by a naive user may accidentally abuse shared resources and achieve very poor performance. To address this problem, we propose that production systems should provide end users with high-level abstractions that allow for the easy expression and efficient execution of data intensive workloads. We present one example of an abstraction – All-Pairs – that fits the needs of several data-intensive scientific applications. We demonstrate that an optimized All-Pairs abstraction is both easier to use than the underlying system, and achieves performance orders of magnitude better than the obvious but naive approach, and twice as fast as a hand-optimized conventional approach.

### 1 Introduction

Many scientists have large problems that can make use of distributed computing; however, most also are not distributed computing experts. Without distributed computing experience and expertise, it is difficult to navigate the large number of factors involved in large distributed systems and the software that harnesses these resources. Inadvertent poor choices can result in poor performance or even outright failures of the application. Poor choices can also lead to inefficient use of shared resources and abuse of the distributed system's infrastructure such as job queues and matchmaking software.

Providing an abstraction useful for a class of problems is one approach to avoiding the pitfalls of distributed computing. Such an abstraction gives the user an interface to define their problem in terms of data and computation requirements, while hiding the details of how the problem will be realized in the system. Abstracting away details also hides

\*This work was supported in part by National Science Foundation grants CCF-06-21434 and CNS-06-43229.

978-1-4244-1694-3/08/\$25.00 ©2008 IEEE

complications that are not apparent to a novice, limiting the opportunity for disastrous decisions that result in pathological cases. The goal is not to strip power from smart users, but rather to make distributed computing accessible to non-experts.

We have implemented one such abstraction – All-Pairs – for a class of problems found in several scientific fields. This implementation has several broad steps. First, we model the workflow so that we may predict execution based on grid and workload parameters, such as the number of hosts. We distribute the data to the compute nodes via a spanning tree, choosing sources and targets in a flexible manner. We dispatch batch jobs that are structured to provide good results based on the model. Once the batch jobs have completed, we collect the results into a canonical form for the end-user, and delete the scratch data left on the compute nodes.

We also examine two algorithms for serving the workload's data requirement: demand paging similar to a traditional cluster and active storage. Active storage delivers higher throughput and efficiency for several large workloads on a shared distributed system, and can result in total workload turnaround times that are up to twice as fast.

We evaluate the abstraction's model, execution, and data delivery on All-Pairs problems in biometrics and data mining on a 500-CPU shared computing system. We have found turnaround time with the abstraction is orders of magnitude faster than for workloads configured using non-experts' choices.

### 2 The All-Pairs Problem

The All-Pairs problem is easily stated:

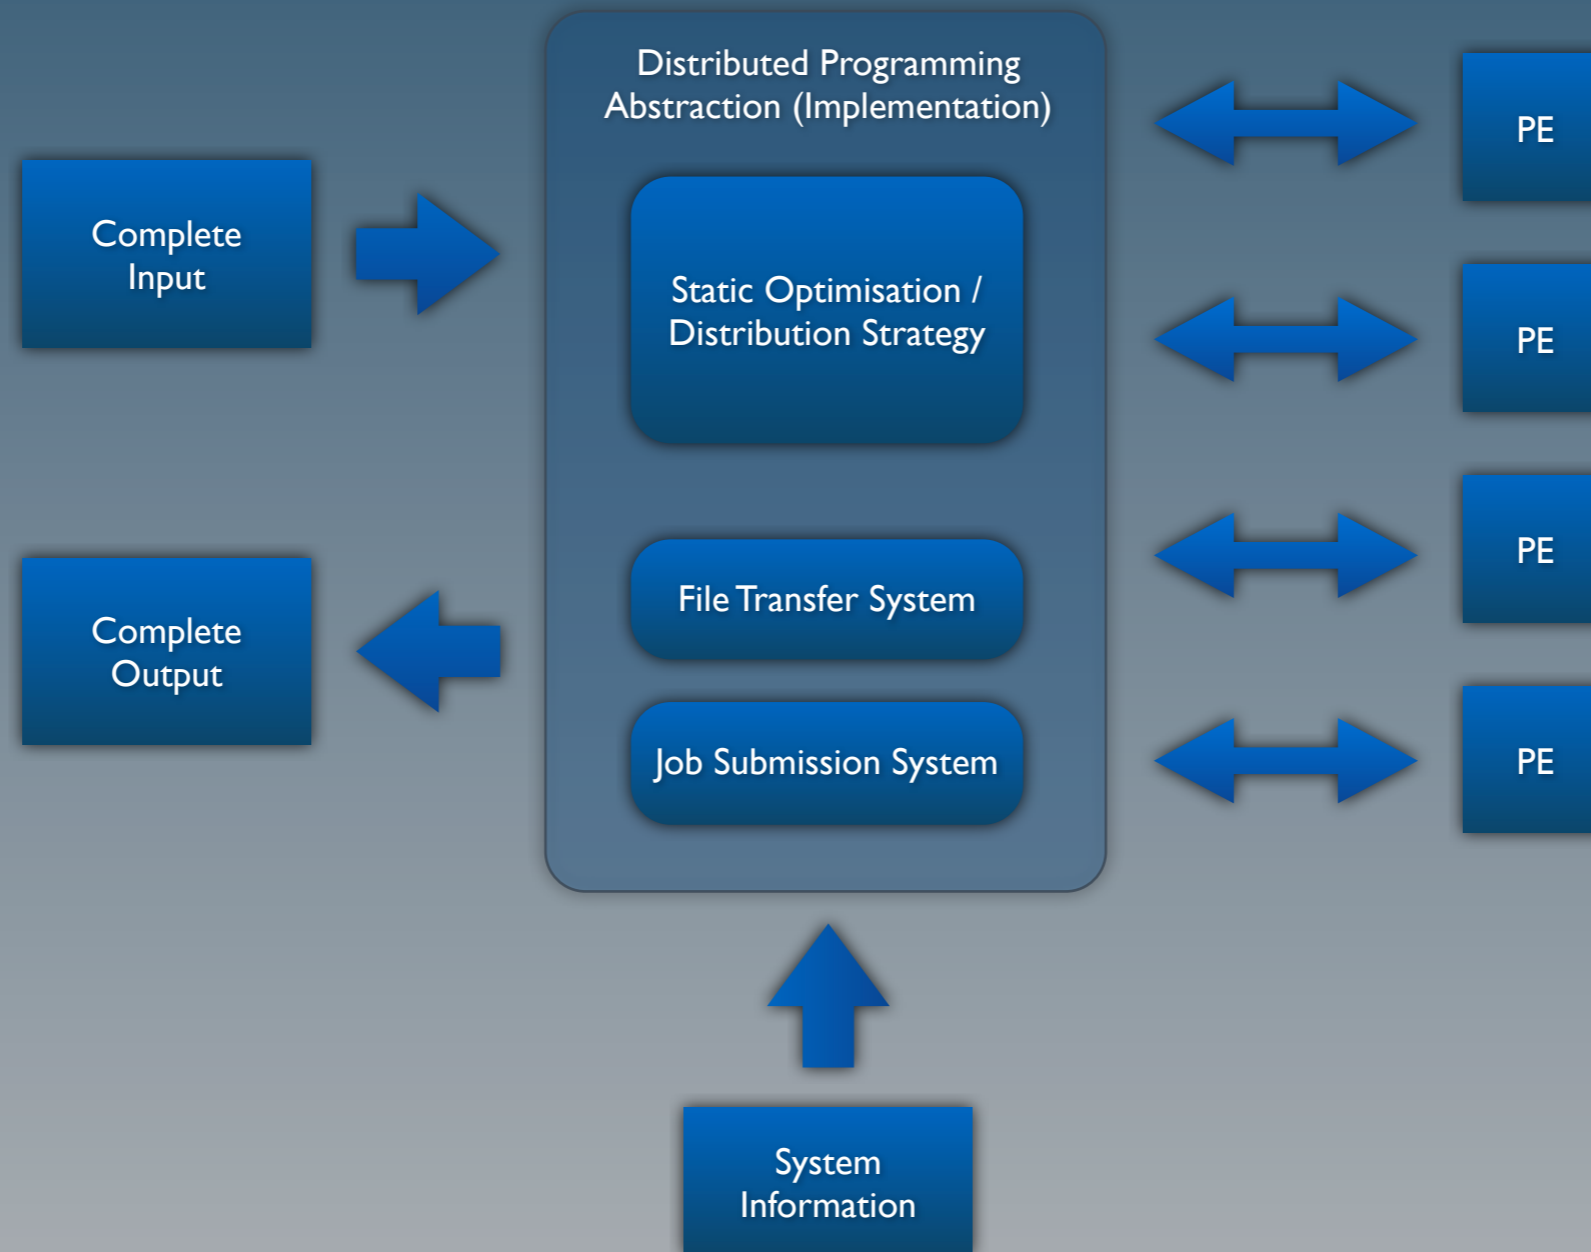
**All-Pairs( set A, set B, function F ) returns matrix M:**  
Compare all elements of set A to all elements of set B via function F, yielding matrix M, such that  $M[i,j] = F(A[i],B[j])$ .

Variations of the All-Pairs problem occur in many branches of science and engineering, where the goal is either to understand the behavior of a newly created function

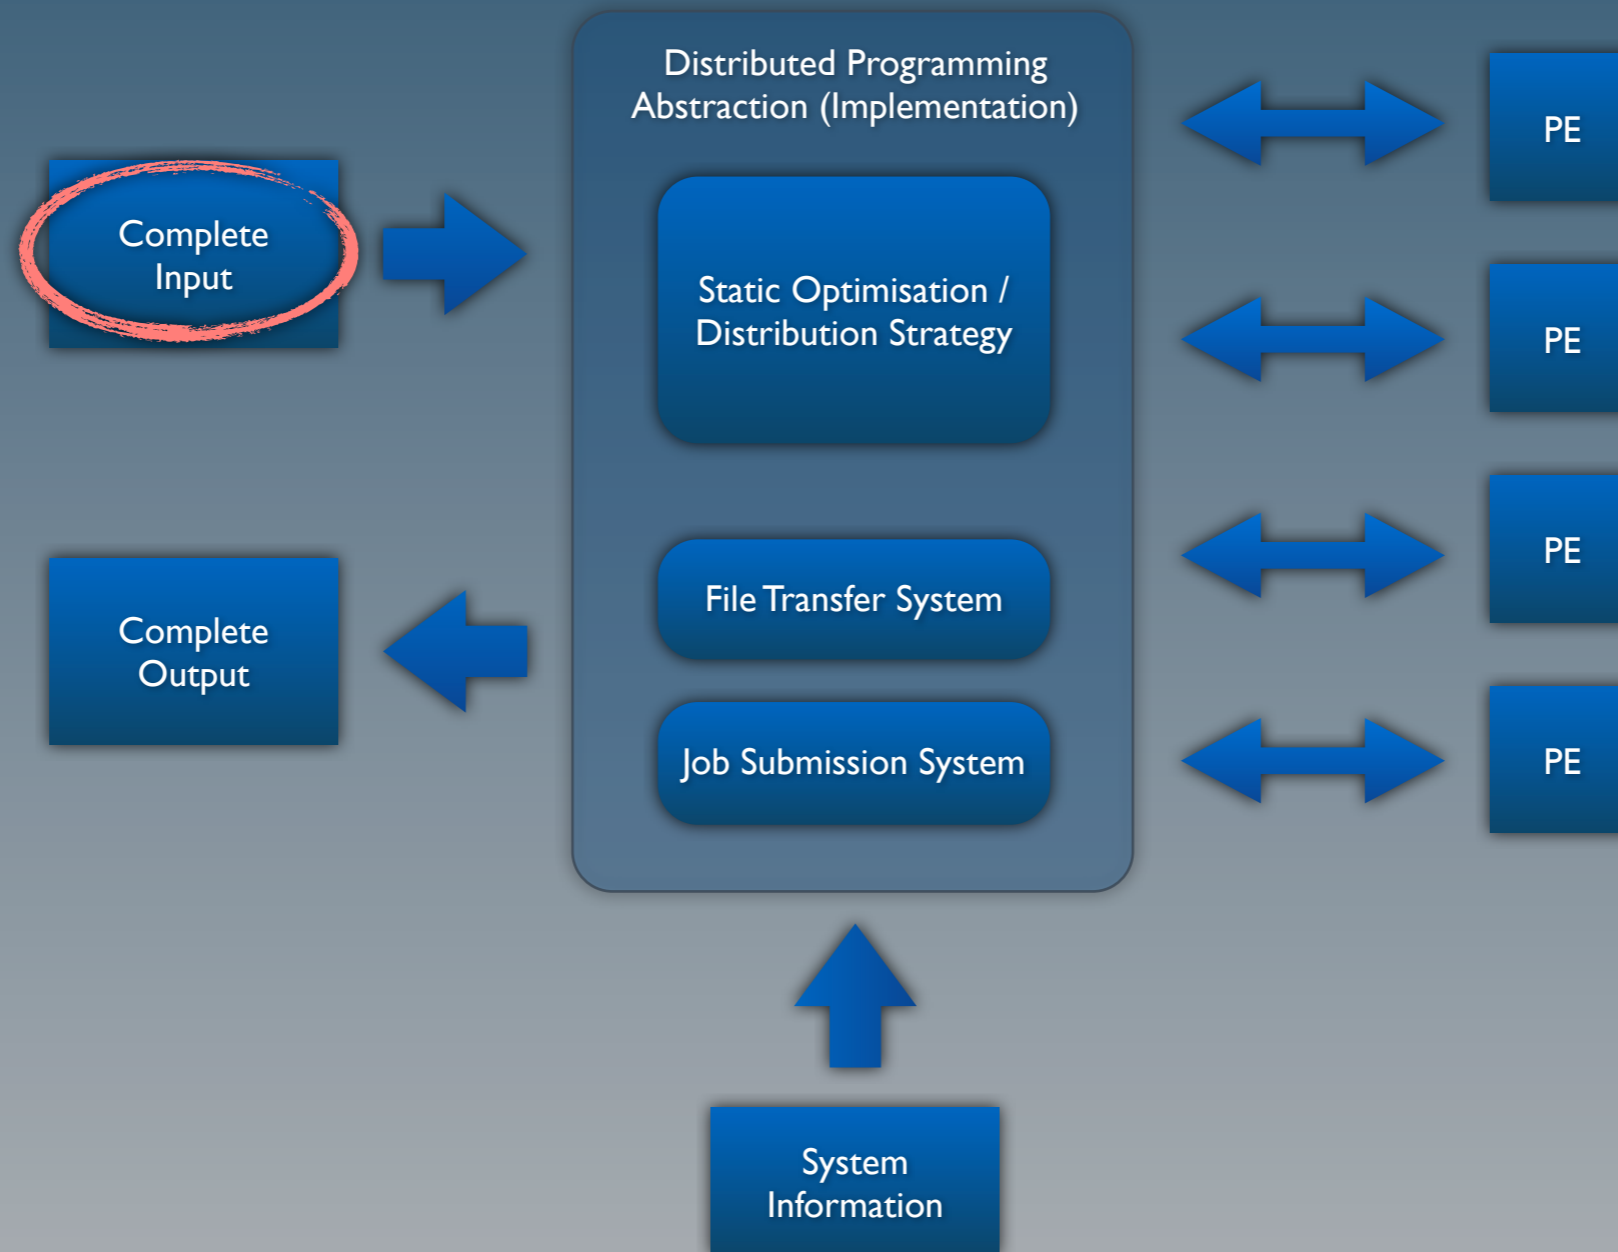


- There doesn't seem to be a lot of work that revisits DPAs in the context of:
  - Streaming input / output data
  - Data with dynamically changing characteristics
  - Autonomic | adaptive | dynamic | optimisation
- Lots of work in static optimisation (“*fine-tuning*”)
  - Optimise throughput for workload X on (idealised) platform Y
  - Usually confined to a specific middleware
- There seems to be a lot of uncharted territory, but that remains to be proved (more reading!)

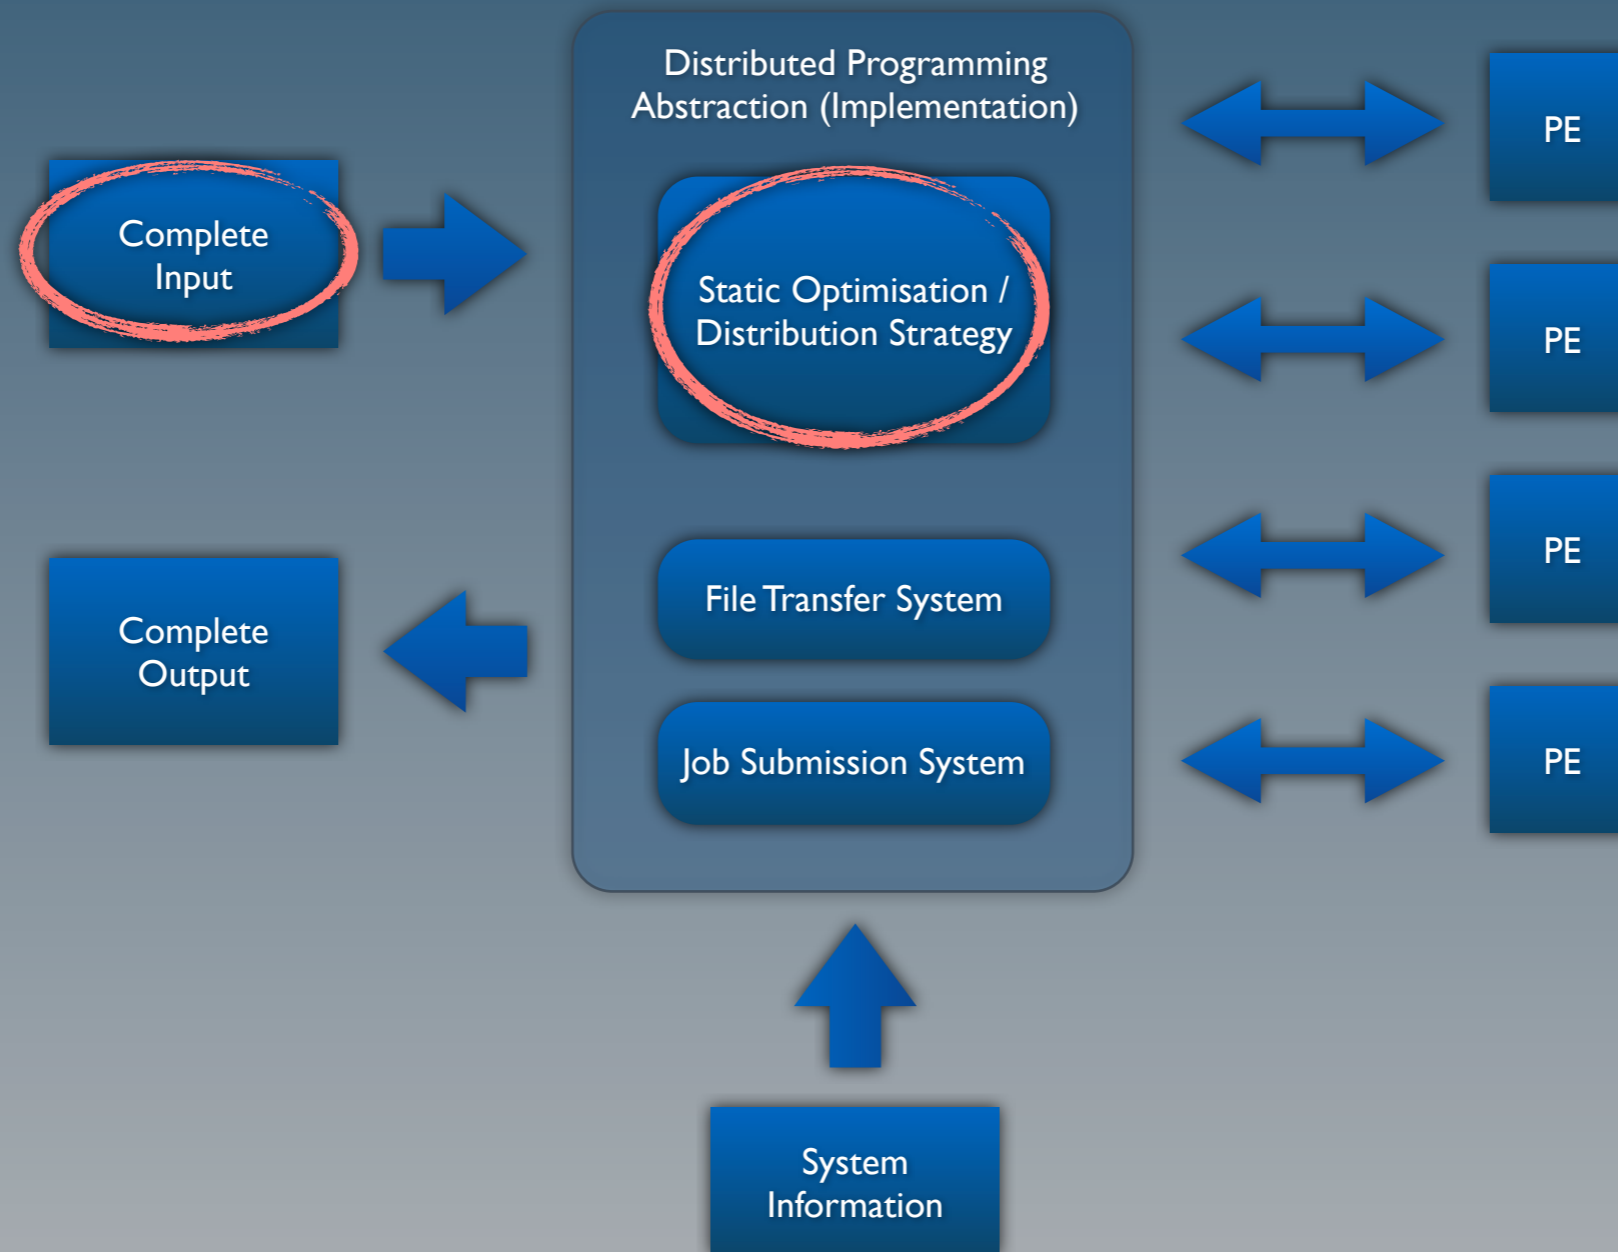
# Everything is Static



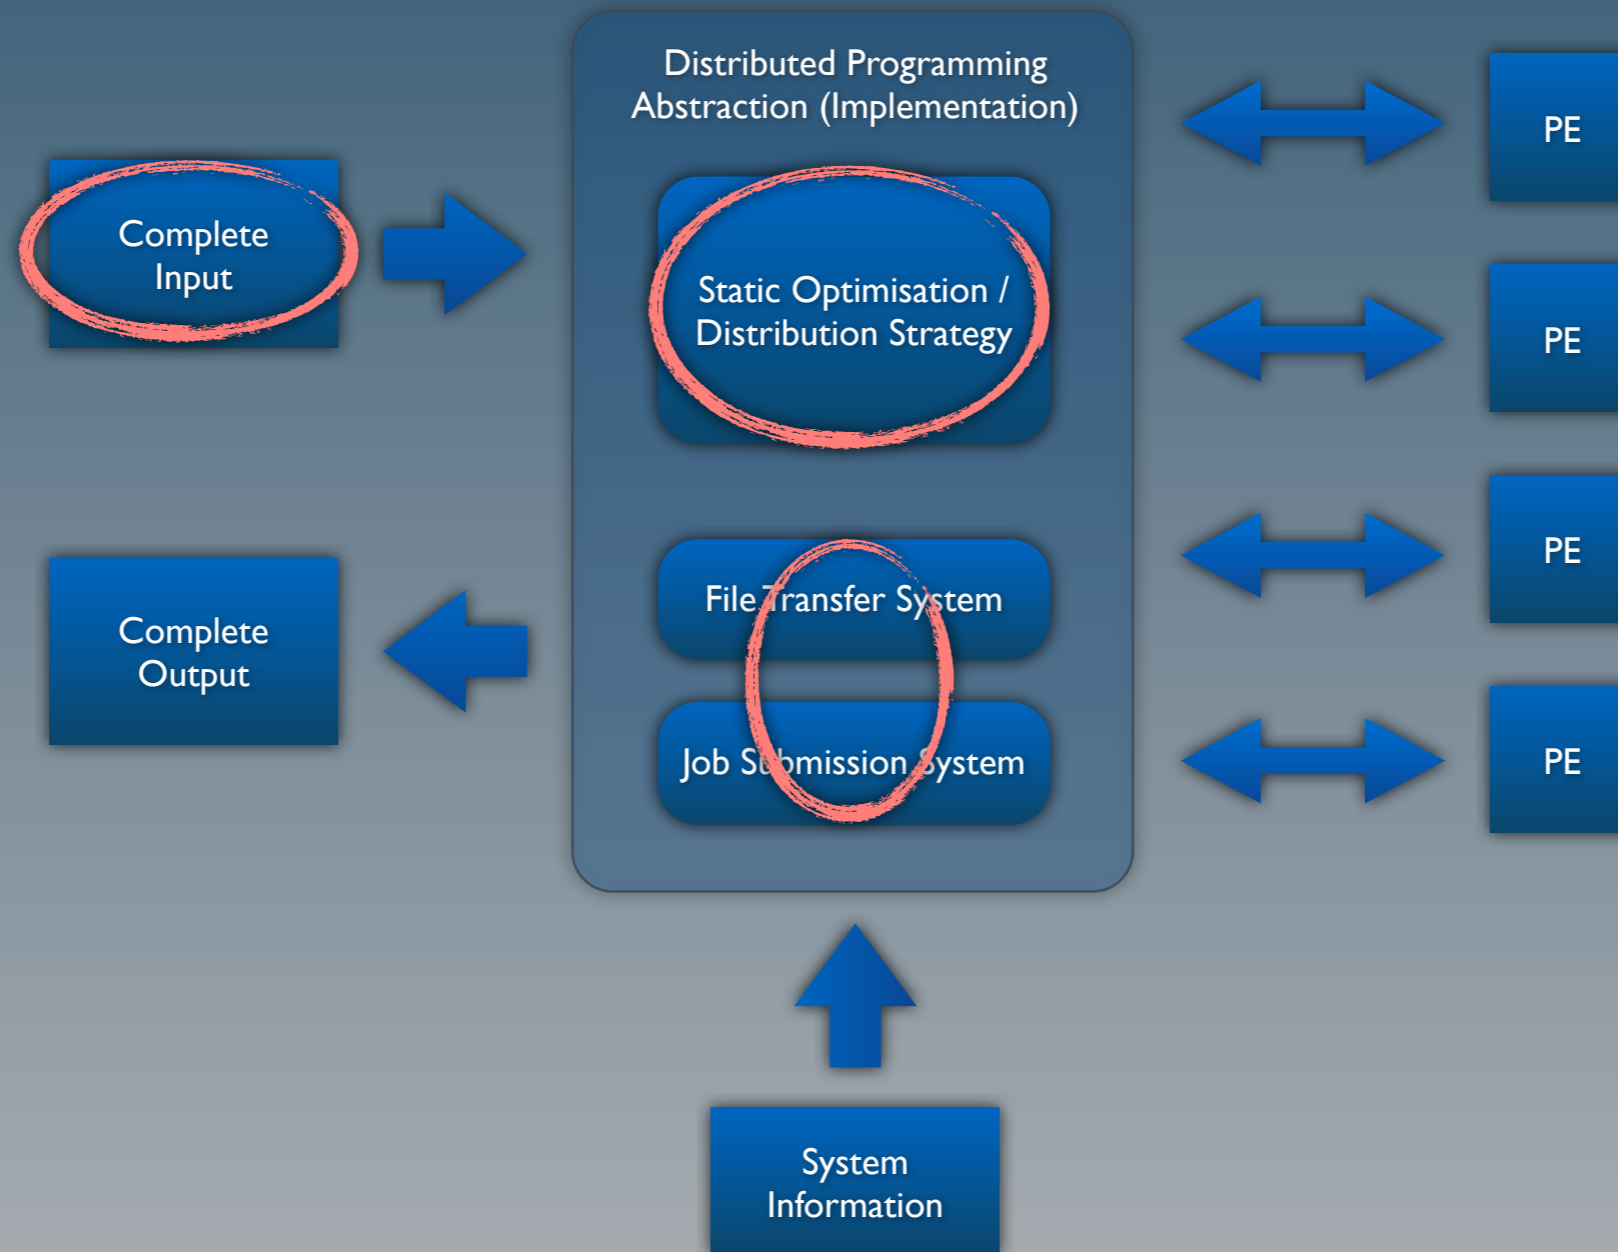
# Everything is Static



# Everything is Static

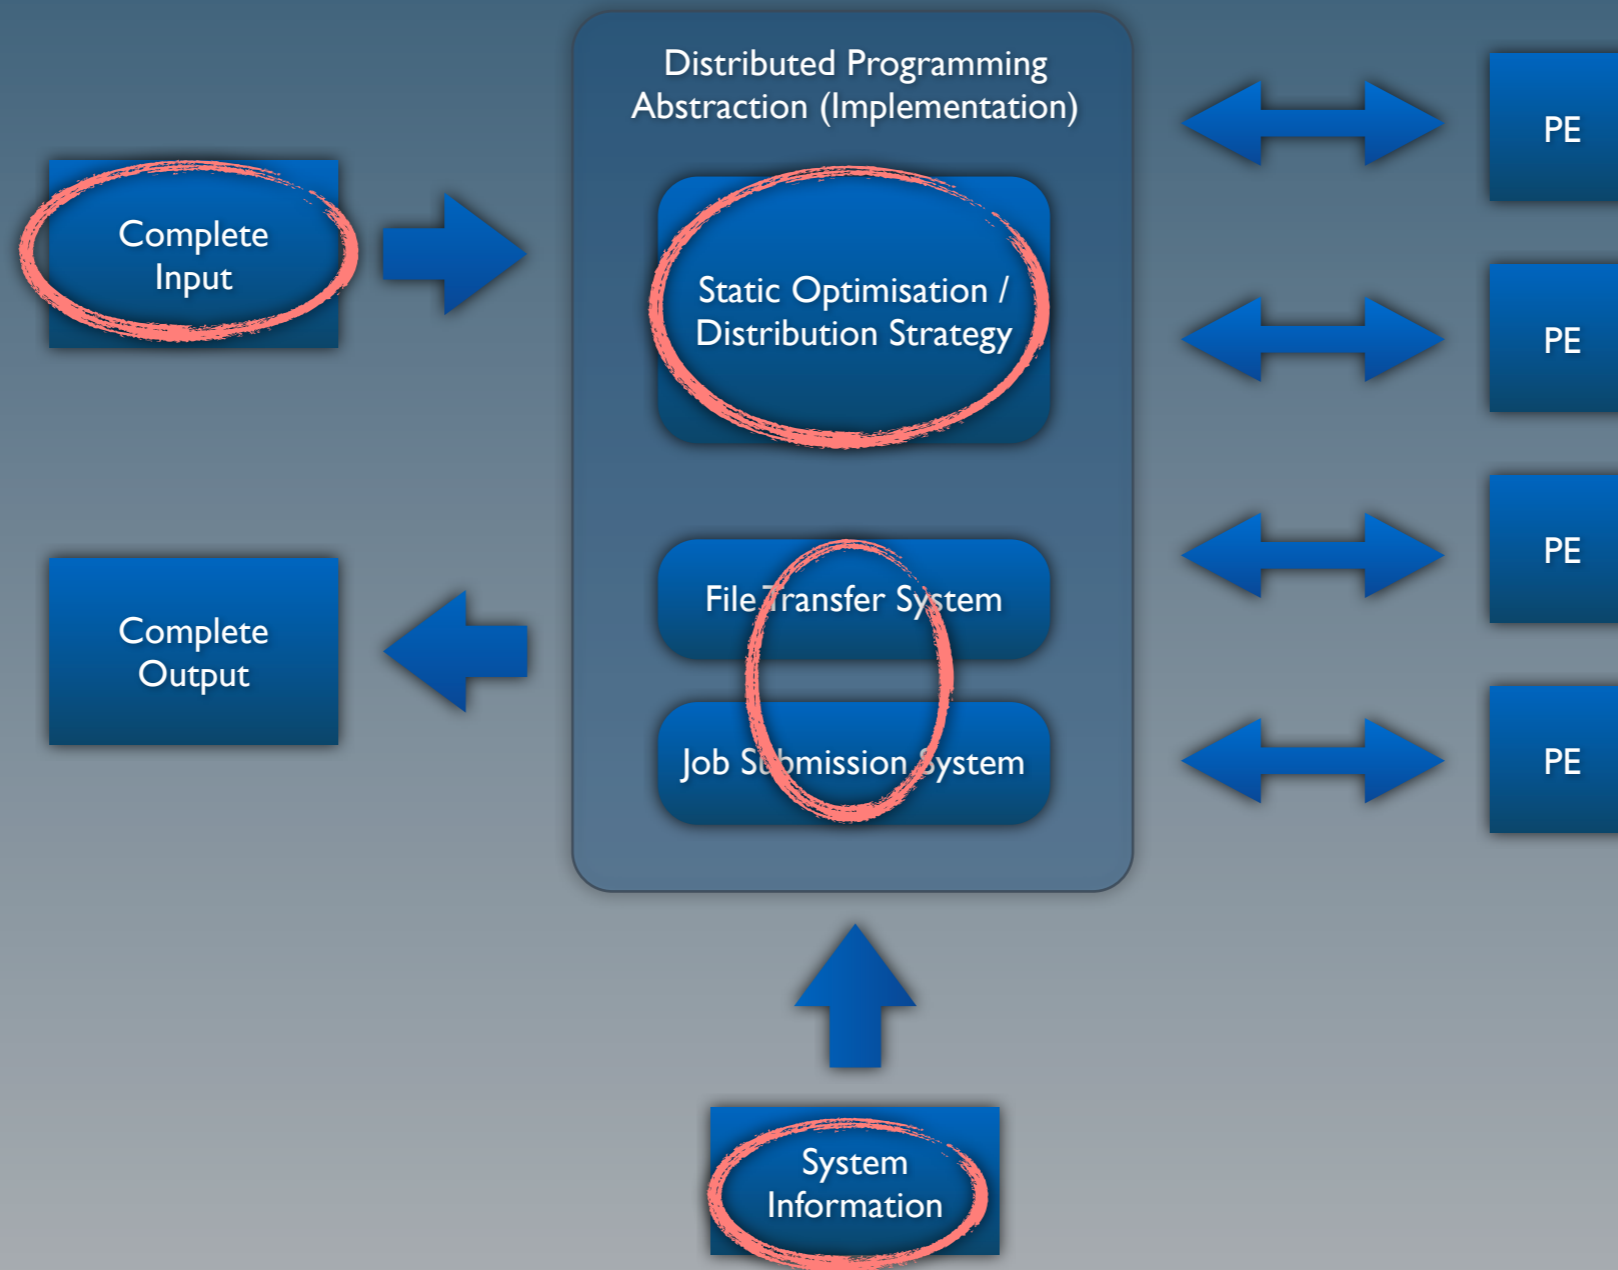


# Everything is Static

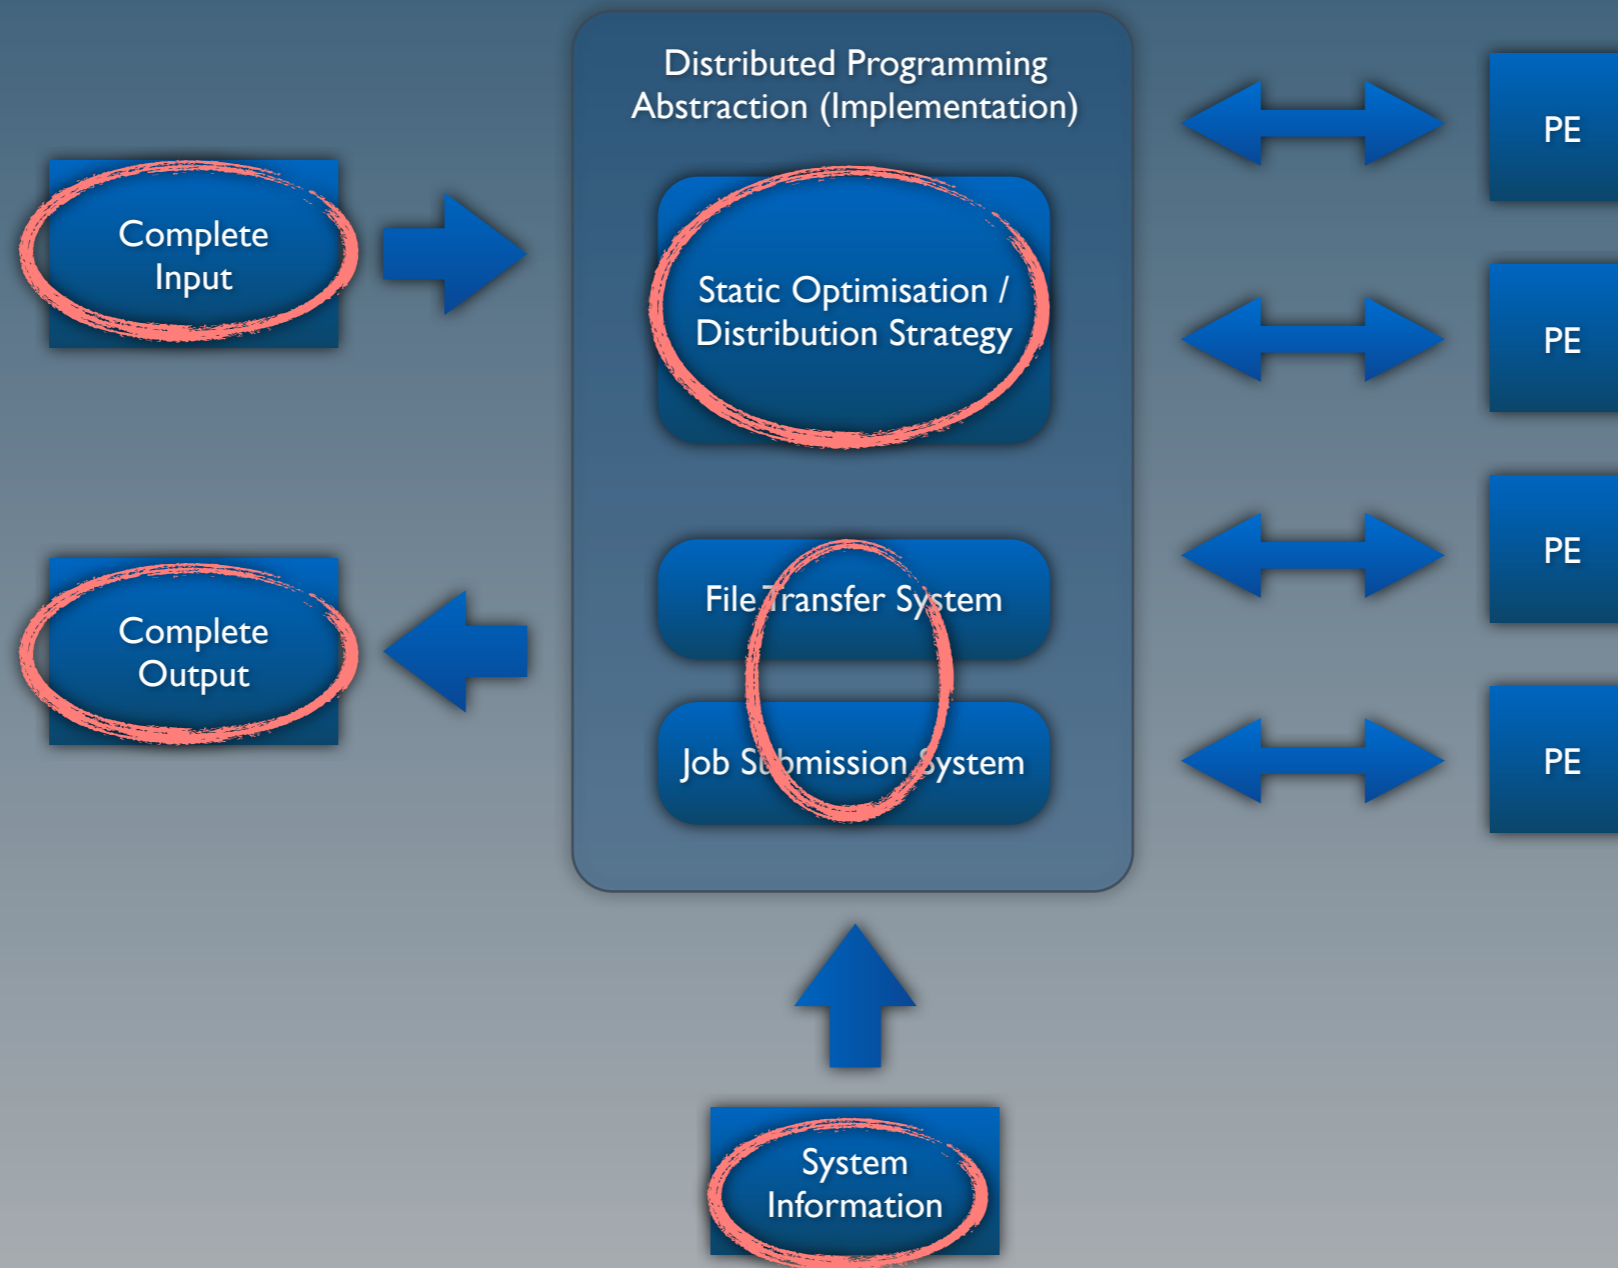




# Everything is Static



# Everything is Static



- Pick a common distributed abstraction and make it a “*Dynamically Optimising and Adaptive Abstraction*” (working title)
- A really good candidate so far seems to be the *All-Pairs* abstraction:
  - Simple and generic
  - Well defined input and output
  - Lots of real-world applications and data available
  - Many potential dynamic and big-data use-cases
  - Not confined to a specific type of infrastructure



- Pick a common distributed abstraction and make it a “*Dynamically Optimising and Adaptive Abstraction*” (working title)
- A really good candidate so far seems to be the *All-Pairs* abstraction:
  - Simple and generic
  - Well defined input and output
  - Lots of real-world applications and data available
  - Many potential dynamic and big-data use-cases
  - Not confined to a specific type of infrastructure



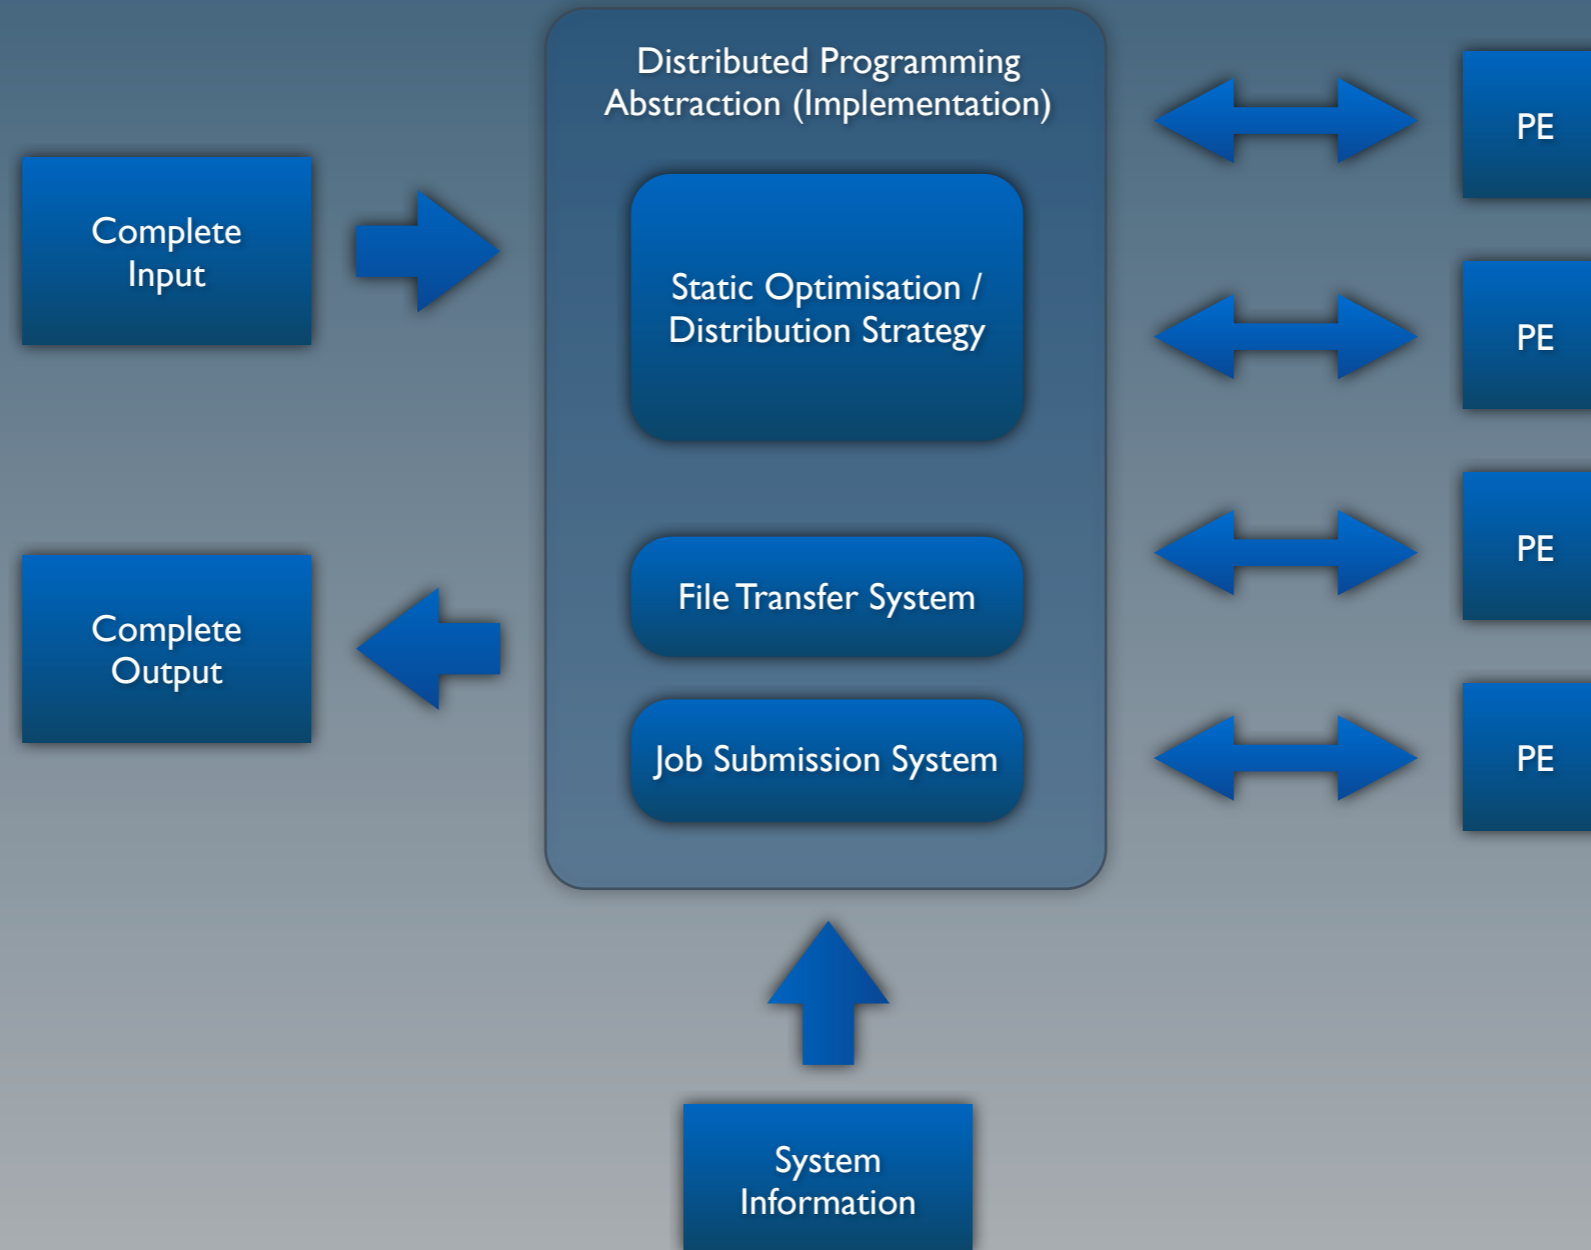
- **Bioinformatics**
  - Phylogenetic tree generation
  - Sequence Alignment
- **Biometrics**
  - Feature (e.g., face) detection
- **Machine Learning**
  - Unsupervised learning (e.g., clustering, density estimation)
  - Evaluation of new learning functions
- *N*-body simulation (Cosmology)

- Web / Data-mining
  - (Near) duplicate document detection
    - “*more-like-this queries*”, collaborative filtering
    - “*data-shedding*”, duplicate deletion
  - Query refinement
  - Coalition detection
- Interesting: some of the matching algorithms in data-mining already support speed v.s. accuracy tuning

- Develop a model that helps us to understand the relevant aspects of data in an *All-Pairs* context
- Define the characteristics we need to capture / extract in order to support dynamic decision-making
  - Static / dynamic ?
  - Rate ?
  - Dependency ?
  - Affinity / Relationship ?

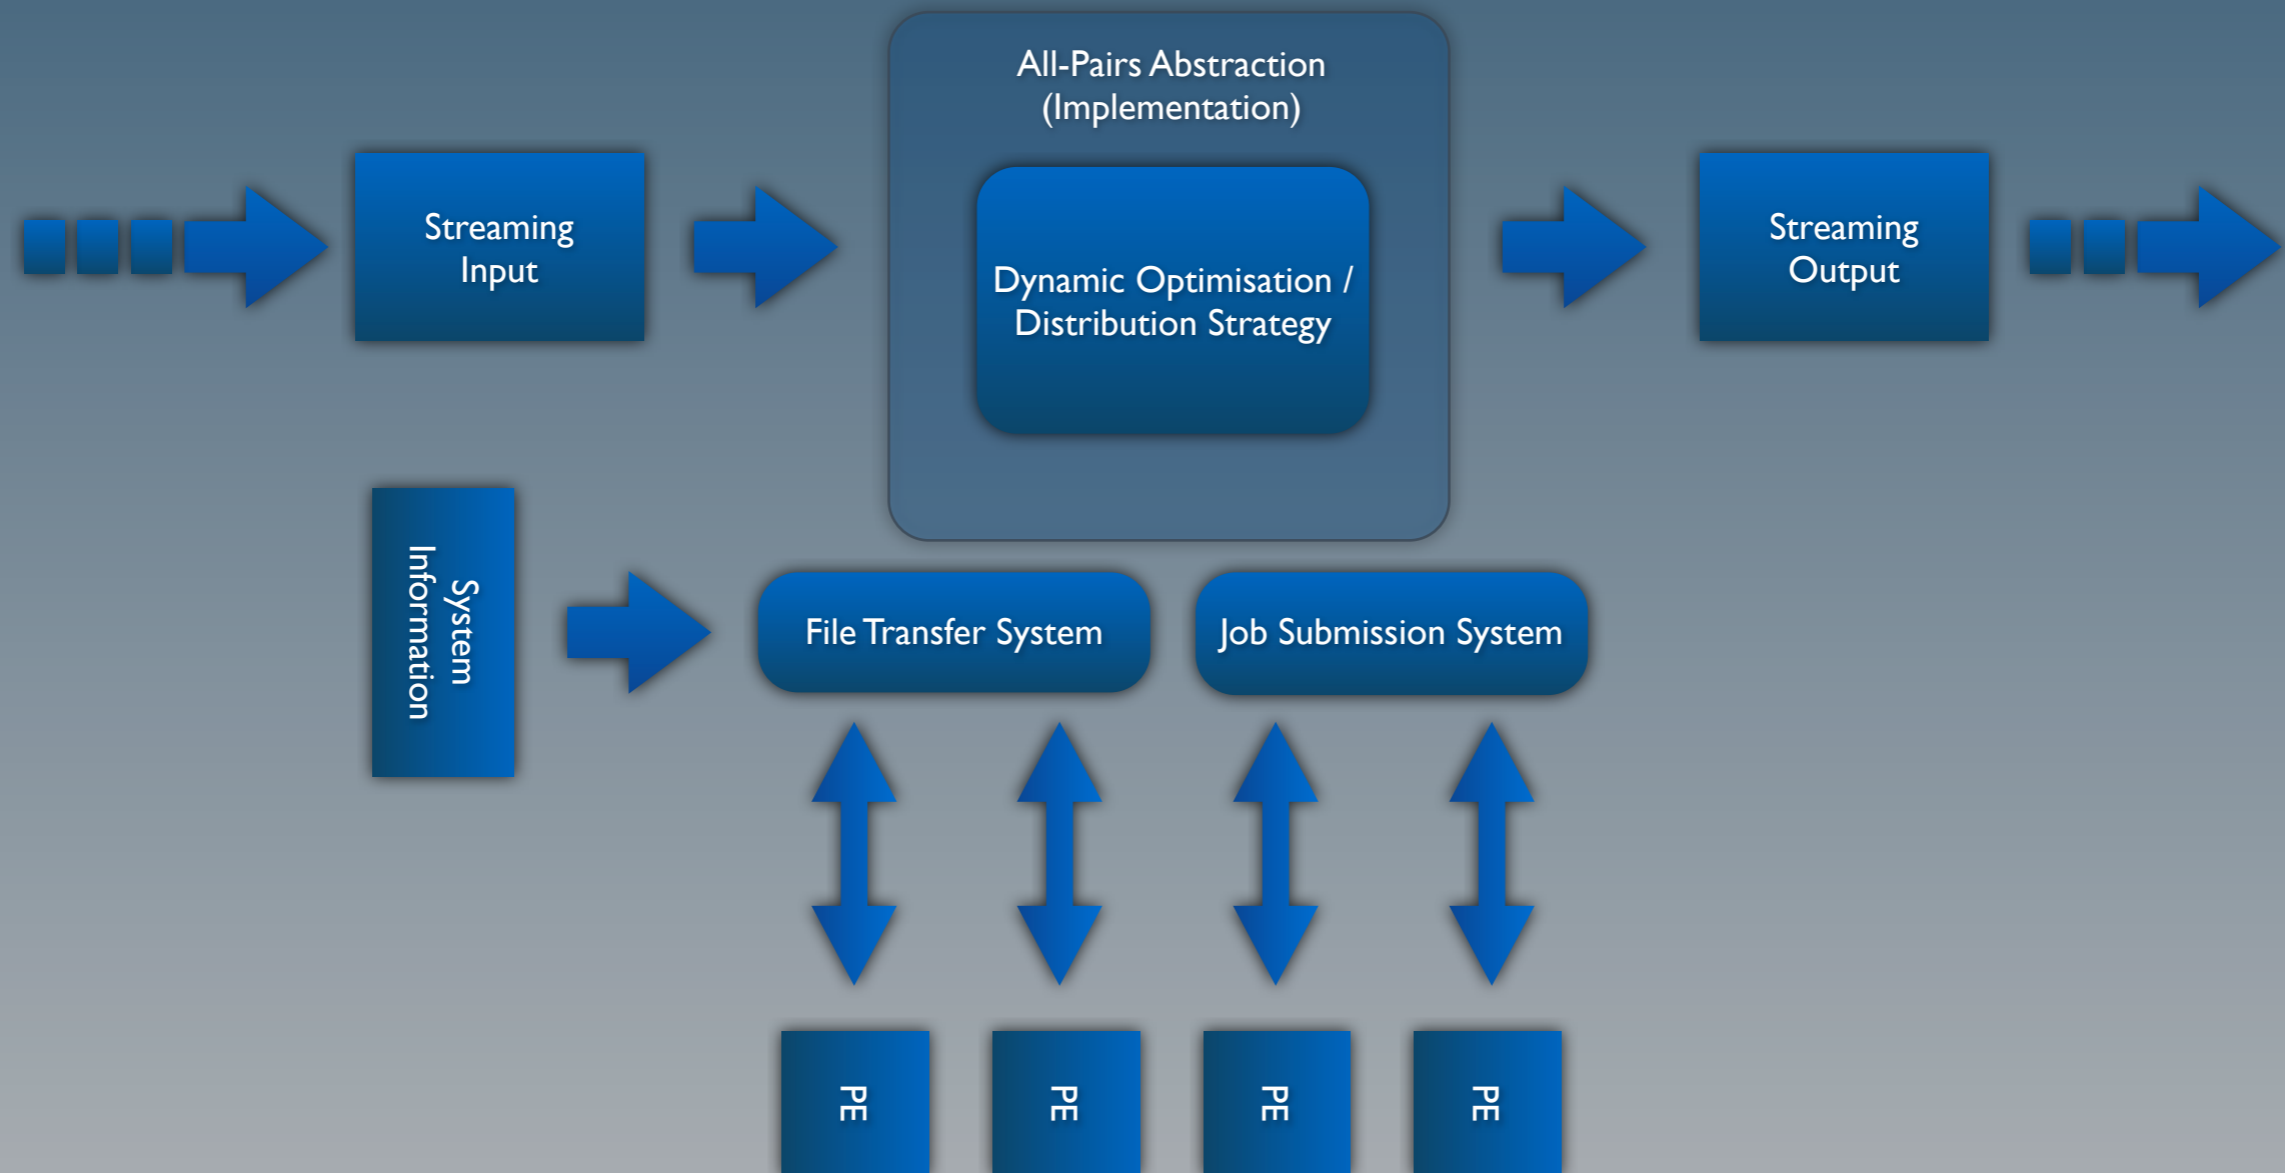
- Revisit the issue of optimisation objectives
  - Minimise energy consumption ?
  - Get me the best possible solution within  $X$  hours ?
  - Don't exceed bandwidth  $X$  or storage  $Y$  ?
- Develop an optimisation framework
  - Possibly based on control theory concepts
  - “Control plant” being the All-Pairs execution framework
  - “Sensors” measuring input and output data
  - “Controller” objective-specific optimisation rules and functions

# A Somewhat Concrete Plan (cont.)

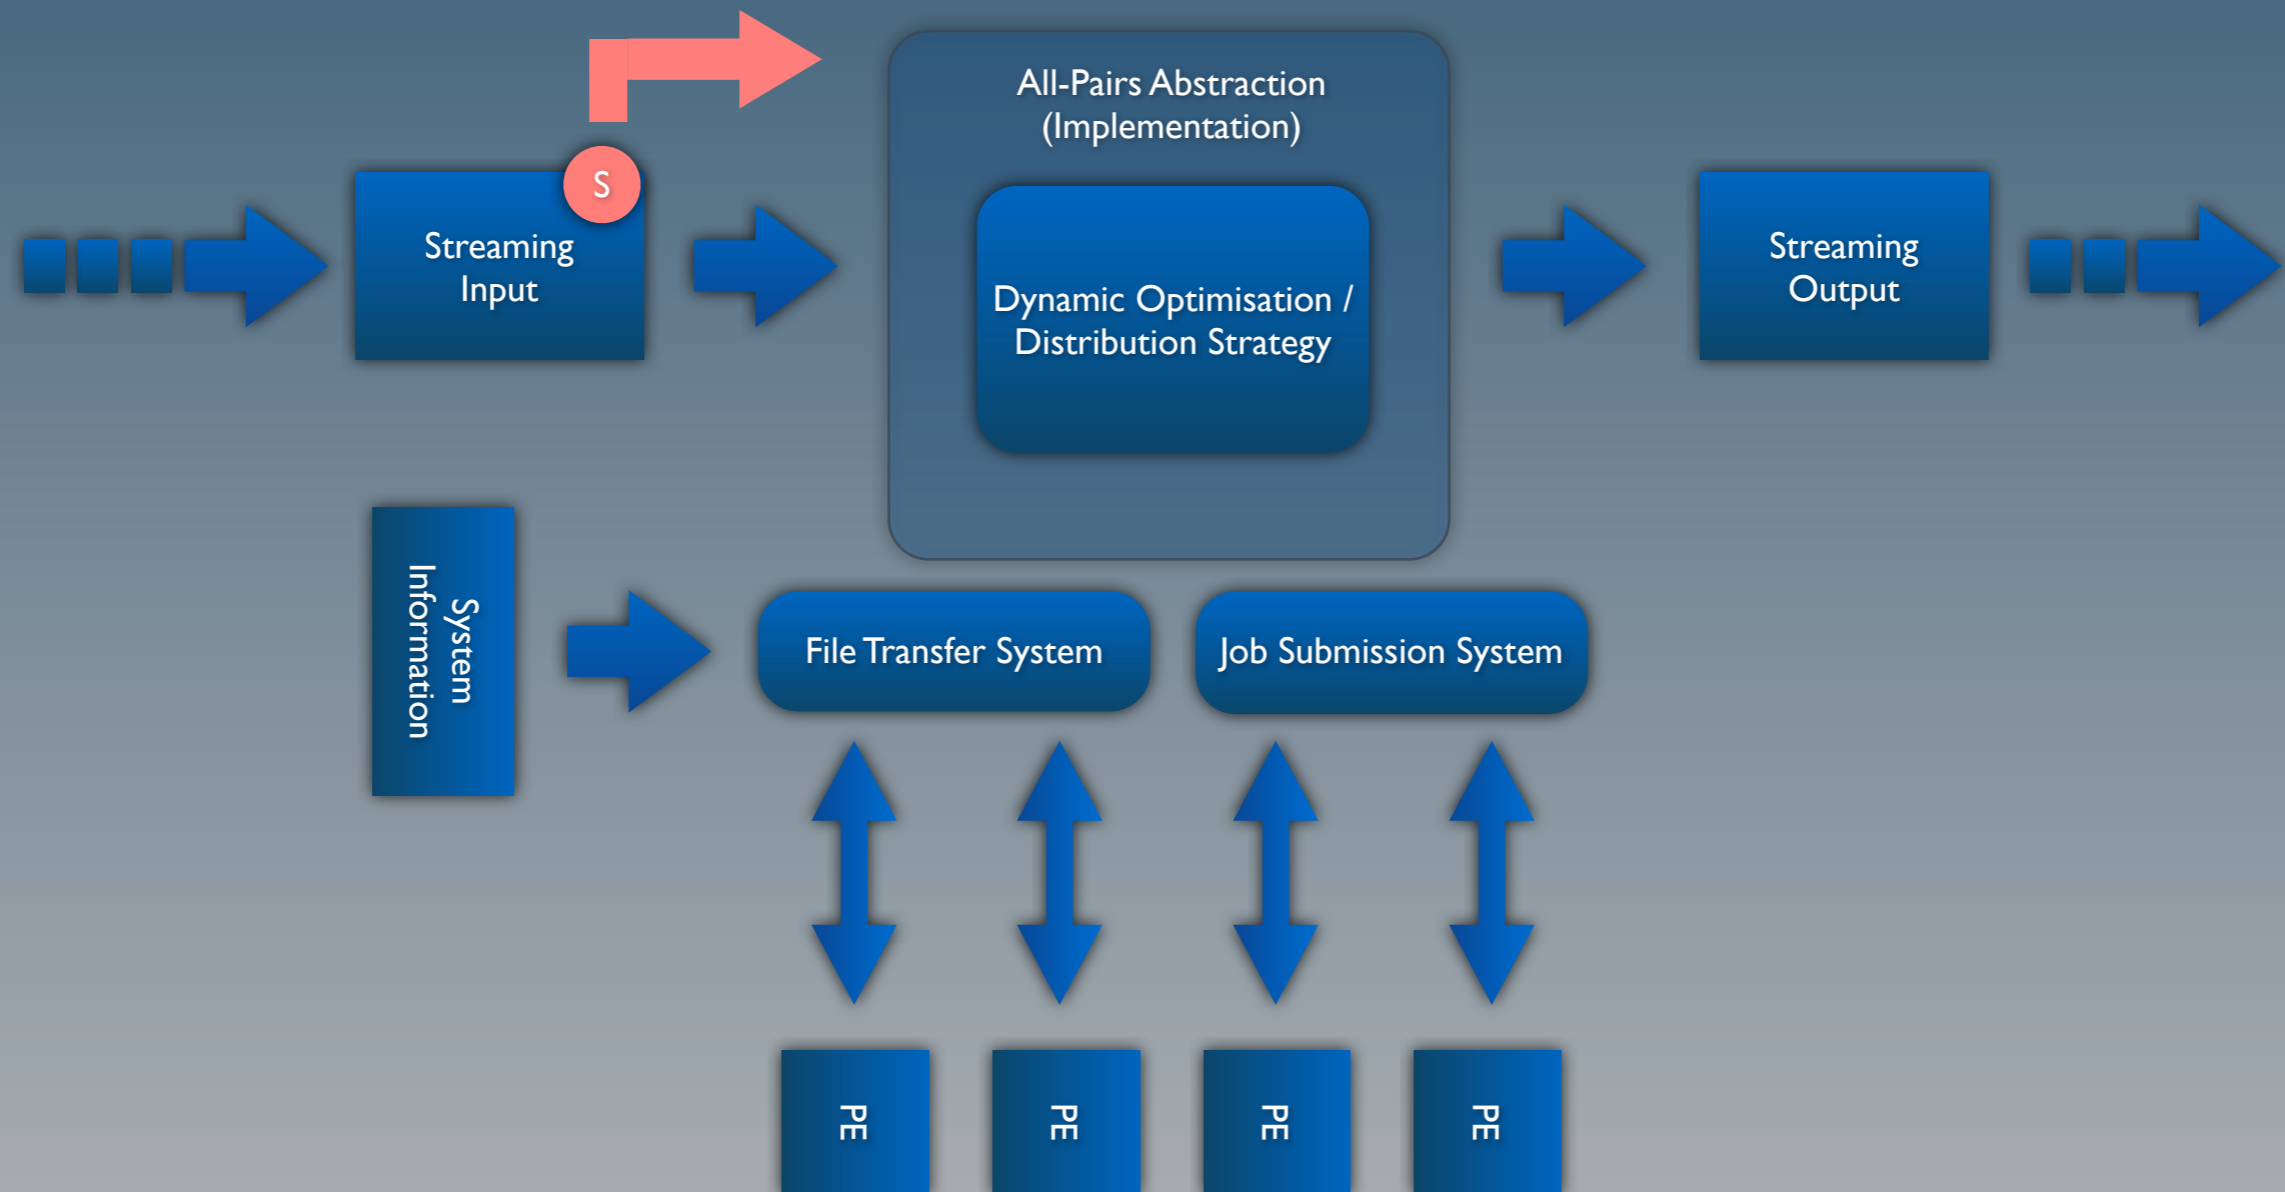




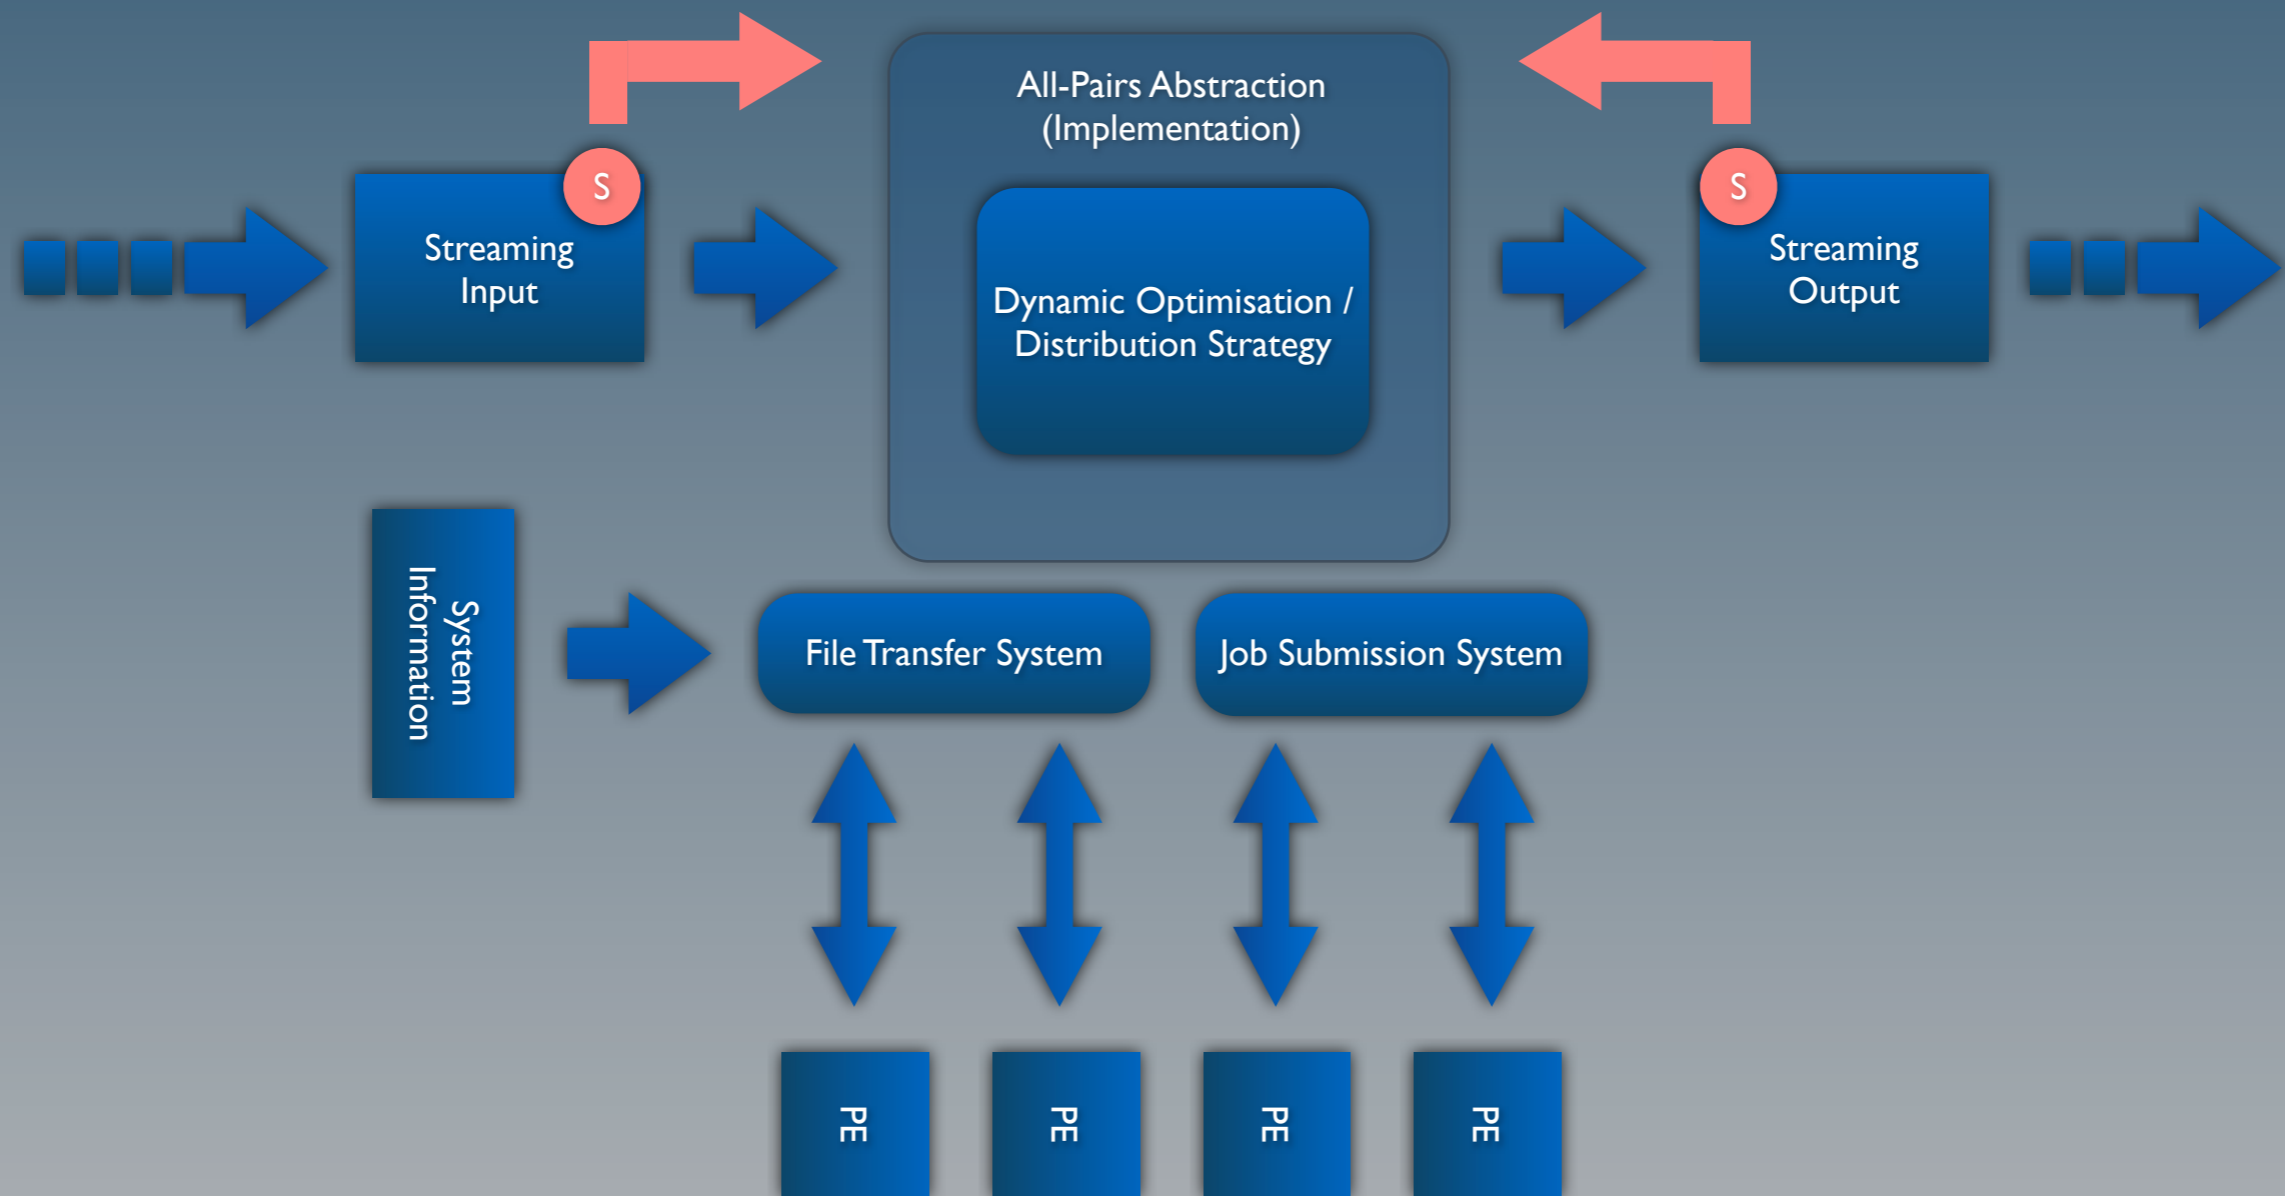
# A Somewhat Concrete Plan (cont.)



# A Somewhat Concrete Plan (cont.)



# A Somewhat Concrete Plan (cont.)



- Can we predict data ?
- How do the optimisation rules and functions look like ?
  - Finite set of rules ?
  - Machine Learning ?
  - Bayesian Networks ?
  - ...
  - A combination of the above ?
- Is dynamic optimisation cheap enough to use it with static problems as well ?
- ...

- What is the right level of abstraction ?
- How do we implement sensors ?
- How do we implement actuators ?
- How would an unobtrusive *'user-interface'* look like ?
- Can we be completely infrastructure independent ?
- What about fault-tolerance ?
- Can we develop a *generic* dynamic optimisation framework that can be used not just for All-Pairs ?



- How can we find real-world dynamic applications ?
  - “*Social engineering*”
- Can we generate synthetic workload ?
- What are the testbeds we should use ?
  - TeraGrid (HPC Grid)
  - EGI (HTC Grid)
  - Clouds ?

# More Information ?

---

## SAGA:

C++/Python: <http://saga.cct.lsu.edu>

OGF GFD.90: <http://www.ogf.org/documents/GFD.90.pdf>

Mailing-Lists: [saga-users@cct.lsu.edu](mailto:saga-users@cct.lsu.edu)

## My Ph.D. Research:

Homepage: <http://www.oleweidner.com>

Code Repo.: <https://github.com/oweidner>

Email: [ole.weidner@ed.ac.uk](mailto:ole.weidner@ed.ac.uk)