



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

FutureGrid 101

Part 2: Getting Started

Craig Stewart

Indianapolis, USA,
Nov 30 – Dec 3, 2010

Paper Submission Deadline:
August 20th, 2010



2nd IEEE International Conference on Cloud Computing Technology and Science

[Home](#) [Program](#) [Resources](#) [Get Involved](#) [Submission](#) [Committee Info](#) [Attending](#) [About](#)

Welcome to CloudCom2010!

Downloadable CFP version   

IEEE CloudCom 2010 - New paper submission deadline, workshops and keynotes. Registration now open!

Submit a paper, poster, demo or exhibit and join top researchers in the field of grid and cloud computing November 30-December 3, 2010 for the 2nd IEEE International Conference on Cloud Computing Technology and Science. The conference is hosted by Pervasive Technology Institute at Indiana University and will be held at the University Place Conference Center and Hotel on the campus of IUPUI in downtown Indianapolis, Indiana, USA.

CloudCom will consist of **research presentations** and **instructional hands-on** activities. The **research** section includes our main conferences, workshops, and a "Cloud Research Issues" panel; the **instructional** section has tutorials and an "Adopting Clouds (Are Clouds for Me?)" Panel to find out if Cloud Computing is right for your situation. All Keynotes, "Cloud Standards", Panel materials, and Posters/Demos are shared to attendees. All these events are sure to be informative for you and your projects. **Be sure to check out some of the planned exhibitions [here](#).**

Registration is now open! Register before October 1 to receive early registration discount.

Important Dates

| | |
|---|--------------------|
| Submission Deadline | August 20, 2010 |
| Author Notification | September 15, 2010 |
| Camera-ready Manuscript | October 1, 2010 |
| Author Registration | October 1, 2010 |
| Poster/Demo/Exhibition Proposals Due | September 15, 2010 |
| Poster/Demo/Exhibition Notification of Acceptance | September 22, 2010 |

Keynote Speakers

Ian T. Foster, Director, CI,
Distinguished Fellow, Argonne National
Laboratory, IL, USA

Dennis Gannon, Director of
Applications for the Cloud Computing,
Microsoft Research, WA, USA

Kai Hwang, Professor, IEEE Fellow,
Director, Internet and Grid Comp. Lab.,
USC, USA

Tutorials

[MapReduce/Hadoop](#)
Sponsored by Yahoo

[Cloud Computing with Windows Azure](#)
Sponsored by Microsoft

Academic Tutorials

Workshops

[CPSRT 2010](#)
[Cloud Computing, HCI, & Design:
Sustainability and Social Impacts](#)
[MAPRED'2010](#)
[CLOUD/'10](#)

Plenary Panels


Cloud Standards
Adopting Cloud
Cloud Research Issues




Futuregrid.org



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

[My Account](#) [Logout](#) [@ Contact FG](#) [Forums](#) [? Help & Support](#)

[Home](#) | [Prospective Users](#) | [Current Users](#) | [About](#)



Is FutureGrid for You?

1 2 3 4

WHAT IS FUTUREGRID?

The FutureGrid is an NSF-funded project which will provide an experimental platform that accommodates batch, grid and cloud computing, allowing researchers to attack a range of research questions associated with optimizing, integrating and scheduling the different service models.

VIRTUAL MACHINE SUPPORT

The FutureGrid(FG) test-bed will support virtual machine-based environments, as well as native operating systems for experiments aimed at minimizing overhead and maximizing performance. The project partners will integrate existing open-source software packages to create an easy-to-use software environment that supports the instantiation, execution and recording of grid and cloud computing experiments.


EXPERIMENT MANAGEMENT



The FutureGrid (FG) test-bed will make it possible for researchers to conduct experiments by submitting an experiment plan that is then executed via a sophisticated workflow engine, preserving the provenance and state information necessary to allow reproducibility.

CLOUD TEST-BED


FutureGrid(FG) will provide a significant new experimental computing grid and cloud computing test-bed to the research community, together with user support for third-party researchers conducting experiments on FutureGrid. The FutureGrid (FG) test-bed includes a geographically distributed set of heterogeneous computing systems, a data management system that will hold both metadata and a growing library of software images, and a dedicated network allowing isolatable, secure experiments.



NEWS / ANNOUNCEMENTS



-  **IEEE CloudCom 2010 - New paper submission deadline, workshops and keynotes** Submit a paper, poster, demo or exhibit and join top researchers in the cloud. [Read More ...](#)
-  **TACC to participate in 'FutureGrid' project** [TACC](#) [Read More ...](#)

EVENTS



-  **Virtual Summer School in "Big Data for Science" is** NCSA is organizing virtual summer schools this year <http://www.vscse.org/> including Big [Read More ...](#)
-  **Apply for "Broadening Participation" Travel Grants** The TeraGrid '10 (TG'10) Conference invites faculty among under-represented communities to apply for [Read More ...](#)



We're still in 'early adopter' mode



- But we are very much interested in applications experiments, computational science experiments, and computer science experiments
- Some early examples:
 - Genome assembly
 - Industrial workflows
 - New middleware



If you remember one thing from
this talk



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- For help with FutureGrid:
 1. Go to kb.indiana.edu and search for futuregrid + some other topic
 2. If that does not give you the information you need to solve your problem send email to help@futuregrid.org
 3. Please note that KB is still a work in progress as of this time!



Applying for an account



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

FutureGrid has a \$10.1 M budget from the NSF; an additional match of \$5M from IU; and more match beyond that from participating partners. Applying for use of FutureGrid is a bit like applying for a grant from the NSF. Less complicated, but it's a nontrivial resource and we as project organizers are doing our best to be responsible stewards of a valuable national resource.



Authentication



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- At present authentication is handled via ssh keys
- If you already have a public ssh key you can upload it as part of application process



Making ssh keys



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- `mkdir ~/.ssh`
- `chmod 700 ~/.ssh`
- `ssh-keygen -q -f ~/.ssh/id_rsa -t rsa`
 - Enter passphrase (empty for no passphrase):
 - Enter same passphrase again:
- `chmod go-w ~/`
- `chmod 700 ~/.ssh`
- `chmod go-rwx ~/.ssh/*`
- Your public ssh keys will be in `id_rsa.pub`



Application process



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

The screenshot shows a web browser window with the URL <http://dev2.futuregrid.org/node/821>. The page title is "FutureGrid Early Adopter Account and Project Registration".

FOR ALL USERS

- Forums

TUTORIALS

- How to use Nimbus

FOR CURRENT USERS

- Hardware/Software Status

FutureGrid Early Adopter Account and Project Registration

This six-page online form constitutes your application to use FutureGrid resources, provided under the auspices of and funding from the National Science Foundation.

As of this date, FutureGrid is accepting applications for 'early adopters' - meaning that the FutureGrid will still be changing rapidly and stability of resources may not yet be at the level anticipated when FutureGrid as a whole goes into full production mode.

These initial questions ask basic background information about you that we need to have for purposes of communicating with you, managing security, and reporting about usage to the National Science Foundation.

Name: *

Please provide your name (lastname, firstname)

Email Address: *

Please provide your primary email address (the one we should use for communication with you).

Institutional Affiliation: *

Please provide your institutional affiliation and current physical (postal) mailing address.

Web Page:

If you have an official web page for your research activities, please provide the URL for it.

Citizenship: *

Of what nation are you a citizen?

[NEXT PAGE >](#)

Vertical navigation links on the left: feedback; on the right: sitemap.



Application process



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

The screenshot shows a web browser window displaying the FutureGrid Early Adopter Account and Project Registration page. The browser's address bar shows the URL `http://dev2.futuregrid.org/node/821`. The page features a navigation menu with links for Home, Prospective Users, Current Users, About, and Training, Education & Outreach. A sidebar on the left contains sections for 'FOR ALL USERS' (Forums), 'TUTORIALS' (How to use Nimbus), and 'FOR CURRENT USERS' (Hardware/Software Status). The main content area is titled 'FutureGrid Early Adopter Account and Project Registration' and indicates it is page 2 of 6. It includes a 'Project Orientation' section with checkboxes for 'Research' (checked) and 'Training, Education, and Outreach'. Below this, there are input fields for 'Primary Discipline' (filled with 'Computer Science') and 'Primary Subdiscipline' (filled with 'Distributed computing'). The page also includes a 'feedback' button on the left and a 'sitemap' button on the right.



Application process



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

Getting Started | post to del.icio.us | my del.icio.us | CITeam Bookmark : c | Other Bookmarks

FutureGrid Early Adopter Account and Project Registration

This is page 3 of 6

FOR ALL USERS
Forums

TUTORIALS
How to use Nimbus

FOR CURRENT USERS
Hardware/Software Status

Systems Resources You Need

Currently Available Systems:

- hotel.futuregrid.org (IBM iDataPlex at U Chicago)
- india.futuregrid.org (IBM iDataPlex at IU)
- sierra.futuregrid.org (IBM iDataPlex at SDSC)
- xray.futuregrid.org (Cray XM5 at IU)
- Not sure / don't care (what I really need is a software environment and I don't care where it runs)

Which currently available FutureGrid hardware systems do you wish to use in your research (check all that apply)

Other:

Other hardware system not specified above

Software Resources You Need

Currently Available Software Environments:

- Common TeraGrid Software Stack
- Eucalyptus VM
- Genesis II
- gLite
- Nimbus VM
- Twister
- Unicore 6

Which currently available software environment do you need to use (check all that apply)

Other:

feedback

stiemann



Application process



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

Google | NMI Build & Test System : R... | FutureGrid Early Adopter Acc... |

http://dev2.futuregrid.org/node/821

Getting Started | post to del.icio.us | my del.icio.us | CITeam Bookmark : c | Other Bookmarks

Software Resources You Need

Currently Available Software Environments:

- Common TeraGrid Software Stack
- Eucalyptus VM
- Genesis II
- gLite
- Nimbus VM
- Twister
- Unicore 6

Which currently available software environment do you need to use (check all that apply)

Other:

Other software environment not specified above

Description of Scale: *

I want to try all possible configurations of FutureGrid.

Briefly describe the scale of resources you expect to need (e.g. "every system you have for a week for a class"; "a few VMs for an experiment"; "I want to run a set of comparisons on entire systems and for each I'll need about ____ days to do that")

< PREVIOUS PAGE | NEXT PAGE >

FutureGrid Project is funded by NSF Grant # OCI-0910812



Application process



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

Getting Started | post to del.icio.us | my del.icio.us | CITeam Bookmark : c | Other Bookmarks

FutureGrid Early Adopter Account and Project Registration

This is page 4 of 6

In order to fulfill the reporting requirements with the National Science Foundation under the Cooperative Service Agreement through which we operate FutureGrid, it is necessary that we have your help in preparing reports and sharing information with the NSF and as appropriate other scientists. Please indicate your agreement to provide each of the following promptly (within 21 days of request).

This seems like a lot of text but it's very straightforward; we simply want to make the reporting needs fully clear from the start.

Reporting Participation Agreement: *

| | No | Yes |
|--|--------------------------|-------------------------------------|
| Citations to your peer reviewed published papers that were developed at least in part based on use of FutureGrid or which make reference to FutureGrid | <input type="checkbox"/> | <input checked="" type="checkbox"/> |
| Citations to your technical reports and other non-peer reviewed published papers that were developed at least in part based on use of FutureGrid or which make reference to FutureGrid | <input type="checkbox"/> | <input checked="" type="checkbox"/> |
| Citations to web pages that include information technical reports and other non-peer reviewed published papers that were developed base on use of FutureGrid or which make reference to FutureGrid | <input type="checkbox"/> | <input checked="" type="checkbox"/> |
| Electronic copies of slides from talks that reference your work done with FutureGrid or which mention FutureGrid (.pdfs or other 'not easily reusable' format o.k.; we will ask you for your permission to post slides publicly and will *not* post them publicly without your permission) | <input type="checkbox"/> | <input checked="" type="checkbox"/> |
| A very brief result of key performance results at the end of your work with FutureGrid (or once per year summarizing the year's work for multi-year projects) | <input type="checkbox"/> | <input checked="" type="checkbox"/> |

Do you agree to participate in preparing reports and sharing information with the NSF and as appropriate other scientists as described above? Please indicate your agreement to each of the above by selecting "Yes"

NSF Acknowledgements: *

| | No | Yes |
|--|--------------------------|-------------------------------------|
| This material is based upon work supported in part by the National Science Foundation under Grant No. 0910812 to Indiana University for "FutureGrid: An Experimental, High-Performance Grid Test-bed." Partners in the FutureGrid project include U. Chicago, U. Florida, San Diego Supercomputer Center - UC San Diego, U. Southern California, U. Texas at Austin, U. Tennessee at Knoxville, U. of Virginia, Purdue I., and T-U. Dresden. -Or the shorter- This material is based upon work supported in part by the National Science Foundation under Grant No. 0910812. | <input type="checkbox"/> | <input checked="" type="checkbox"/> |
| Any nontechnical publications, presentations, web pages, etc. should also include Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF. | <input type="checkbox"/> | <input checked="" type="checkbox"/> |

Do you agree to include in your technical referred papers and other papers, web pages, and presentations, acknowledgments as above (these are derived from the NSF-specified language)? Please indicate your agreement to the above by selecting "Yes"

< PREVIOUS PAGE | NEXT PAGE >



Application process



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

Getting Started | post to del.icio.us | my del.icio.us | CITeam Bookmark : c | Other Bookmarks

Future Grid | Apply for Account | Access Portal | Search | @ Contact FG | Forums | Help & Support

Home | Prospective Users | Current Users | About | Training, Education & Outreach

FutureGrid Early Adopter Account and Project Registration

FOR ALL USERS
Forums

TUTORIALS
How to use Nimbus

FOR CURRENT USERS
Hardware/Software Status

This is page 5 of 6

After you complete your application, you will receive an email message asking you to sign and return the TeraGrid standard User Responsibility Form. A copy of the text of this form is included below:

FutureGrid User Responsibility Form

Based on TeraGrid User Responsibility Form Last updated 8/06/09 (v2.2.2), modified specifically for FutureGrid by changing contact information as appropriate for interacting with FutureGrid.

Introduction

All TeraGrid resource provider (RP) sites have legal and other obligations to protect shared resources as well as the intellectual property of users, including the FutureGrid project. Users share this responsibility by observing the rules of acceptable use that are outlined in this document. TeraGrid resources include hardware, software, network connections, and storage. Each resource is finite and shared by the entire research community. Responsible conduct on the part of each user is essential to ensure equitable and secure access for all. Failure to use TeraGrid resources properly may result in the penalties outlined in section 5, including those imposed by TeraGrid, civil, and/or criminal penalties.

Each time an application for TeraGrid resources is submitted, the Acceptance Statement, or last page of this form, must be received by the TeraGrid Allocations Department within 30 days. If it is received in time, the account will be enabled or continued—if not, it will be deactivated. Your signature serves as acknowledgment that you have read and understand your responsibilities as a user. The Acceptance Statement should be sent via email, fax, or US mail to the Pervasive Technology Institute, 2719 E 10th Street, Bloomington IN 47408 ATTN: FutureGrid project. If you have questions, please write to help@futuregrid.org.

1.0: Account, Password, and Certificate Management

FutureGrid will provide you with the accounts necessary to access allocated systems. An account is assigned for one user only and must not be shared with others—including students and/or collaborators. For community accounts, see section 1.2 below.

Passwords and certificates are the keys to your account. Never share a password out loud, or write it down where it could be found and/or associated with your account. Never use tools which openly expose them on the network, such as telnet. Make sure that file and directory permissions prevent others from reading or copying the private key portion of certificates, which is the equivalent of a password. Do not store your password(s) in unencrypted files or even in encrypted files if possible.

TeraGrid and FutureGrid support staff will never ask for your password, and will never send a password via e-mail, set them to a requested string, or perform any other activity which could reveal it to others. If a support person insists that you share your password, report it to the TeraGrid helpdesk: help@futuregrid.org

feedback | sitemap



Application process



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

Getting Started | post on del.icio.us | my del.icio.us | CITeam Bookmark : c | Other Bookmarks

TUTORIALS
How to use Nimbus

FOR CURRENT USERS
Hardware/Software Status

About your research

Abstract: *
This is the abstract.
Please provide a short abstract of your planned research or educational project using FutureGrid

Intellectual Merit: *
It will slice bread.
In reference to NSF merit review criteria, please briefly describe the intellectual merit of your proposed research or educational activity.

Broader Impacts: *
Bread will feed the hungry.
In reference to NSF merit review criteria, please briefly describe the broader impacts of your proposed research or educational activity.

Software Contributions:

| | No | Yes |
|---|--------------------------|-------------------------------------|
| Will your project generate any software that can be used by other Futuregrid Users? If yes, please also answer a) and b) below. | <input type="checkbox"/> | <input checked="" type="checkbox"/> |
| a) Will you be able to generate documentation for your software? | <input type="checkbox"/> | <input checked="" type="checkbox"/> |
| b) Will you be able to provide support for your software? | <input type="checkbox"/> | <input checked="" type="checkbox"/> |

Biographical Sketch:
 No file chosen
Please attach your NSF or NIH format short (2 page) biographical sketch.

Account Information

Username:
m0j0c0c0
Do you have a preferred username? If so, please provide it here.

Your SSH Public Key:
dfjldjdfkjdekldlehj3oihedhjehrelslnkehrlehrleh|
Please provide the SSH Public Key you will use for authenticating to FutureGrid resources

feedback | stiemma



Logging in



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- Different sites are using different systems so login directions are included in your site's handouts!



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

FutureGrid 101

Part 3: Eucalyptus

Craig Stewart



Eucalyptus

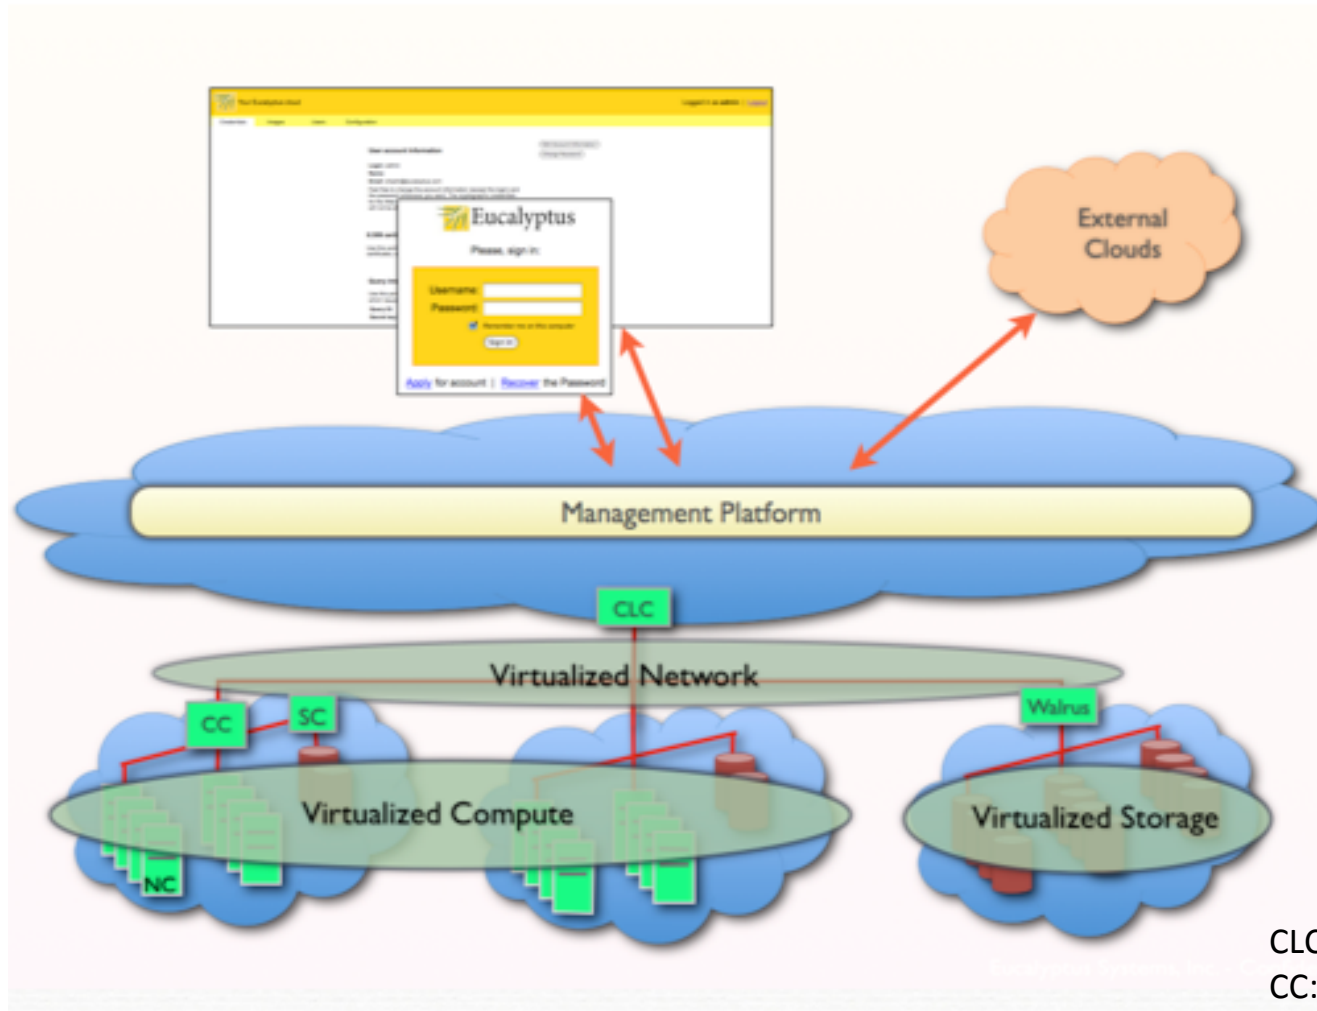


PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

- Elastic Utility Computing Architecture Linking Your Programs To Useful Systems
 - Eucalyptus is an open-source software platform that implements IaaS-style cloud computing using the existing Linux-based infrastructure
 - IaaS Cloud Services providing atomic allocation for
 - Set of VMs
 - Set of Storage resources
 - Networking



Eucalyptus Architecture



Legend
 CLC: Cloud Controller
 CC: Cluster Controller
 NC: Node Controller



Open Source Eucalyptus



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- **Eucalyptus Features**
 - Amazon AWS Interface Compatibility
 - Web-based interface for cloud configuration and credential management.
 - Flexible Clustering and Availability Zones.
 - Network Management, Security Groups, Traffic Isolation
 - Elastic IPs, Group based firewalls etc.
 - Cloud Semantics and Self-Service Capability
 - Image registration and image attribute manipulation
 - Bucket-Based Storage Abstraction (S3-Compatible)
 - Block-Based Storage Abstraction (EBS-Compatible)
 - Xen and KVM Hypervisor Support

Source: <http://www.eucalyptus.com>



Eucalyptus Testbed



- Eucalyptus is available to FutureGrid users on the India and Sierra clusters
- Users can use a maximum of 50 nodes on India and 21 on Sierra
 - Each node supports up to 8 small VMs
 - Different availability zones provide VMs with different compute and memory capacities

```
AVAILABILITYZONE india 149.165.146.135
AVAILABILITYZONE |- vm types free / max cpu ram disk
AVAILABILITYZONE |- m1.small 0400 / 0400 1 512 5
AVAILABILITYZONE |- c1.medium 0400 / 0400 1 1024 7
AVAILABILITYZONE |- m1.large 0200 / 0200 2 6000 10
AVAILABILITYZONE |- m1.xlarge 0100 / 0100 2 12000 10
AVAILABILITYZONE |- c1.xlarge 0050 / 0050 8 20000 10

AVAILABILITYZONE sierra 198.202.120.90
AVAILABILITYZONE |- vm types free / max cpu ram disk
AVAILABILITYZONE |- m1.small 0160 / 0160 1 512 5
AVAILABILITYZONE |- c1.medium 0160 / 0160 1 1024 7
AVAILABILITYZONE |- m1.large 0080 / 0080 2 6000 10
AVAILABILITYZONE |- m1.xlarge 0040 / 0040 2 12000 10
AVAILABILITYZONE |- c1.xlarge 0020 / 0020 8 30000 10
```



Account Creation



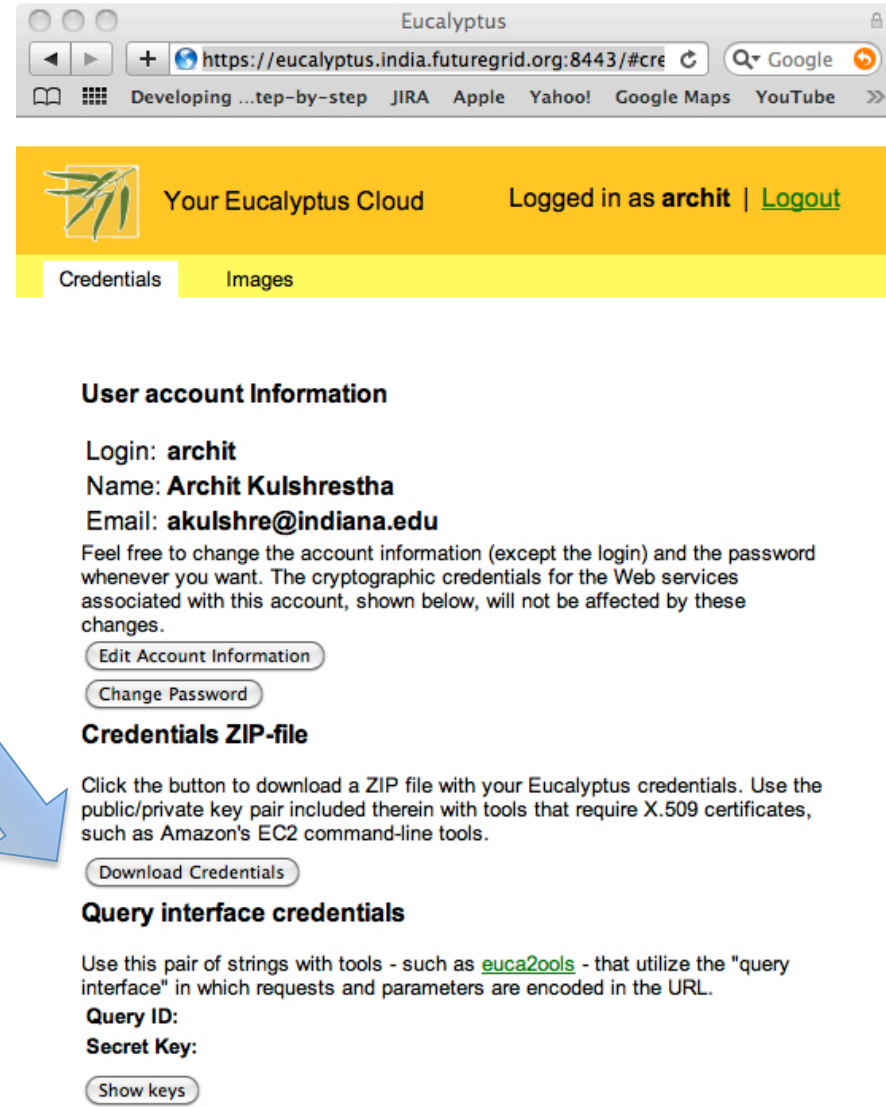
PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- In order to be able to use Eucalyptus and obtain keys, users will need to request accounts at the Eucalyptus Web Interfaces:
 - <https://eucalyptus.india.futuregrid.org:8443/>
 - <https://eucalyptus.sierra.futuregrid.org:8443/>
- On the Login page click *apply for account*
- On the next page that pops up fill out the mandatory and optional sections of the form
- Once complete, click *signup* and the Eucalyptus administrator will be notified of the account request
- You will get an email once the account has been approved
- Click on the link provided in the email to confirm and complete the account creation process

Obtaining Credentials

- Download your credentials as a zip file from the web interface for use with euca2ools
- Save this file and extract it
- On the command prompt change to the euca2-`{username}`-x509 folder which was just created
 - `cd euca2-username-x509`
- Source the eucarc file using the command `source eucarc`
 - `source ./eucarc`



The screenshot shows a web browser window titled "Eucalyptus" with the URL `https://eucalyptus.india.futuregrid.org:8443/#cre`. The page header includes the Future Grid logo and the text "Your Eucalyptus Cloud" and "Logged in as archit | Logout". Below the header, there are two tabs: "Credentials" (selected) and "Images". The main content area is titled "User account Information" and displays the following details:

- Login: **archit**
- Name: **Archit Kulshrestha**
- Email: **akulshre@indiana.edu**

Below the details, there is a note: "Feel free to change the account information (except the login) and the password whenever you want. The cryptographic credentials for the Web services associated with this account, shown below, will not be affected by these changes." There are two buttons: "Edit Account Information" and "Change Password".

The next section is titled "Credentials ZIP-file" and contains the text: "Click the button to download a ZIP file with your Eucalyptus credentials. Use the public/private key pair included therein with tools that require X.509 certificates, such as Amazon's EC2 command-line tools." There is a button labeled "Download Credentials".

The final section is titled "Query interface credentials" and contains the text: "Use this pair of strings with tools - such as [euca2ools](#) - that utilize the "query interface" in which requests and parameters are encoded in the URL." There are two labels: "Query ID:" and "Secret Key:". There is a button labeled "Show keys".



Install Euca2ools



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- Euca2ools are the command line clients used to interact with Eucalyptus
- Install euca2ools bundle from <http://open.eucalyptus.com/downloads>
- Instructions for various Linux platforms are available on the download page



Euca2ools

- Testing your setup
 - Use euca-describe-availability-zones to test the setup

```
euca-describe-availability-zones  
AVAILABILITYZONE india 149.165.146.135
```

- List the existing images using euca-describe-images

```
$ euca-describe-images  
IMAGE emi-0B951139 centos53/centos.5-3.x86-64.img.manifest.xml admin  
available public x86_64 machine  
IMAGE emi-409D0D73 rhel55/rhel55.img.manifest.xml admin available  
public x86_64 machine  
...
```



Key management



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

- Create a keypair and add the public key to eucalyptus

```
euca-add-keypair userkey > userkey.pem
```

- Fix the permissions on the generated private key

```
chmod 0600 userkey.pem
```



Image Deployment



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

- Now we are ready to start a VM using one of the pre-existing images
- We need the emi-id of the image that we wish to start (listed in the output of euca-describe-images command that we saw earlier)
 - We use the euca-run-instances command to start the VM

```
euca-run-instances -k userkey -n 1 emi-0B951139 -t c1.medium
RESERVATION    r-4E730969    archit archit-default
INSTANCE       i-4FC40839    emi-0B951139  0.0.0.0 0.0.0.0 pending userkey
2010-07-20T20:35:47.015Z    eki-78EF12D2    eri-5BB61255
```



Monitoring

- euca-describe-instances shows the status of the VMs

```
$ euca-describe-instances
RESERVATION  r-4E730969  archit default
INSTANCE     i-4FC40839  emi-0B951139  149.165.146.153 10.0.2.194
pending     userkey     0             m1.small      2010-07-20T20:35:47.015Z
india      eki-78EF12D2  eri-5BB61255
```

- Shortly after...

```
$ euca-describe-instances
RESERVATION  r-4E730969  archit default
INSTANCE     i-4FC40839  emi-0B951139  149.165.146.153 10.0.2.194
running     userkey     0             m1.small      2010-07-20T20:35:47.015Z
india      eki-78EF12D2  eri-5BB61255
```



VM Access

- First we must create rules to allow access to the VM over ssh

```
euca-authorize -P tcp -p 22 -s 0.0.0.0/0 default
```

- The ssh private key that was generated earlier can now be used to log in to the VM

```
ssh -i userkey.pem root@149.165.146.153
```



Image Deployment (1/3)



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

- We will use the example Fedora 10 image to test uploading images
 - Download the gzipped tar ball

```
wget http://open.eucalyptus.com/sites/all/modules/pubdlcnt/pubdlcnt.php?file=http://www.eucalyptussoftware.com/downloads/eucalyptus-images/euca-fedora-10-x86_64.tar.gz&nid=1210
```

- Uncompress and Untar the archive

```
tar xzf euca-fedora-10-x86_64.tar.gz
```



Image Deployment (2/3)



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

- Next we bundle the image with a kernel and a ramdisk using the `euca-bundle-image` command
 - We will use the xen kernel already registered
 - `euca-describe-images` returns the kernel and ramdisk IDs that we need

```
$ euca-bundle-image -i euca-fedora-10-x86_64/fedora.  
10.x86-64.img --kernel eki-78EF12D2 --ramdisk eri-5BB61255
```

- Use the generated manifest file to upload the image to Walrus

```
$ euca-upload-bundle -b fedora-image-bucket -m /tmp/fedora.  
10.x86-64.img.manifest.xml
```



Image Deployment (3/3)



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

- Register the image with Eucalyptus

```
euca-register fedora-image-bucket/fedora.10.x86-64.img.manifest.xml
```

- This returns the image ID which can also be seen using `euca-describe-images`

```
$ euca-describe-images
IMAGE emi-FFC3154F fedora-image-bucket/fedora.
10.x86-64.img.manifest.xml archit available public x86_64 machine
eri-5BB61255 eki-78EF12D2
IMAGE emi-0B951139 centos53/centos.5-3.x86-64.img.manifest.xml admin
available public x86_64 machine ...
```




FutureGrid 101

Part 4: More about software
architecture

Gregor von Laszewski



RAIN: Dynamic Provisioning



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

- Change underlying system to support current user demands at different levels
 - Linux, Windows, Xen, KVM, Nimbus, Eucalyptus, Hadoop, Dryad
 - Switching between Linux and Windows possible!
- **Stateless** (means no “controversial” state) images: Defined as any node that that does not store permanent state, or configuration changes, software updates, etc.
 - Shorter boot times
 - Pre-certified; easier to maintain
- **Statefull** installs: Defined as any node that has a mechanism to preserve its state, typically by means of a non-volatile disk drive
 - Windows
 - Linux with custom features
- Encourage use of services: e.g. MyCustomSQL as a service and not MyCustomSQL as part of installed image?
- Runs OUTSIDE virtualization so cloud neutral
- Use **Moab** to trigger changes and **xCAT** to manage installs



xCAT and Moab in detail



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- xCAT
 - Uses Preboot eXecution Environment (PXE) to perform remote network installation from RAM or disk file systems
 - Creates stateless Linux images (today)
 - Changes the boot configuration of the nodes
 - We are intending in future to use remote power control and console to switch on or of the servers (IPMI)
- Moab
 - Meta-schedules over resource managers
 - Such as TORQUE(today) and Windows HPCS
 - control nodes through xCAT
 - Changing the OS
 - Remote power control in future



Command line RAIN Interface

- fg-deploy-image
 - host name
 - image name
 - start time
 - end time
 - label name
- fg-add
 - label name
 - framework hadoop
 - version 1.0
- Deploys an image on a host
- Adds a feature to a deployed image



Draft GUI for FutureGrid Dynamic Provisioning



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

The main GUI window is titled "Firefox" and "Select Image". It features a vertical sidebar on the left with six buttons: "Select Nodes" (yellow), "Select Image" (orange), "Select Application" (blue), "Deploy Application" (cyan), "Run Application" (green), and "Inspect Results" (orange). The main content area is divided into four sections: "System" with checkboxes for iDataPlex IU, iDataPlex SDSC, Dell Cluster, and Tempest; "Operating System" with checkboxes for Red Hat Linux 5.4.1, Red Hat Linux 5.4.2, Red Hat Linux 5.4.3, and Windows 7; "Framework" with checkboxes for Hadoop, Dryad, MPI, and "..."; and a "Loading..." section at the bottom with a progress bar and "Continue" and "Cancel" buttons.

Futuregrid Knowledge Base (KB)

The Google Gadget window is titled "Futuregrid Knowledge Base (KB)". It has a search bar with the text "hadoop" and a "Go" button. Below the search bar is a table with the following content:

| Title |
|----------------------------|
| Where can I find tutorials |
| Help & Support: |
| Help & Support: |

Google Gadget for FutureGrid Support

We are building a FutureGrid Portal!



Experiment Manager

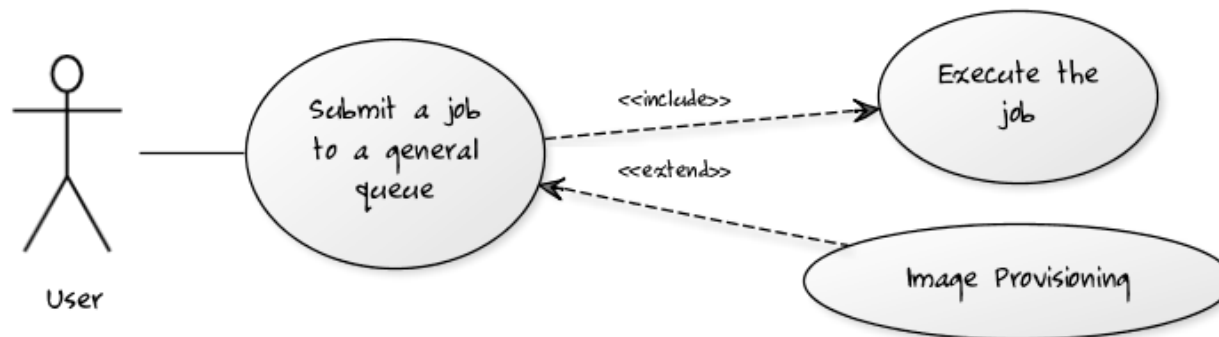
FutureGrid Software Component



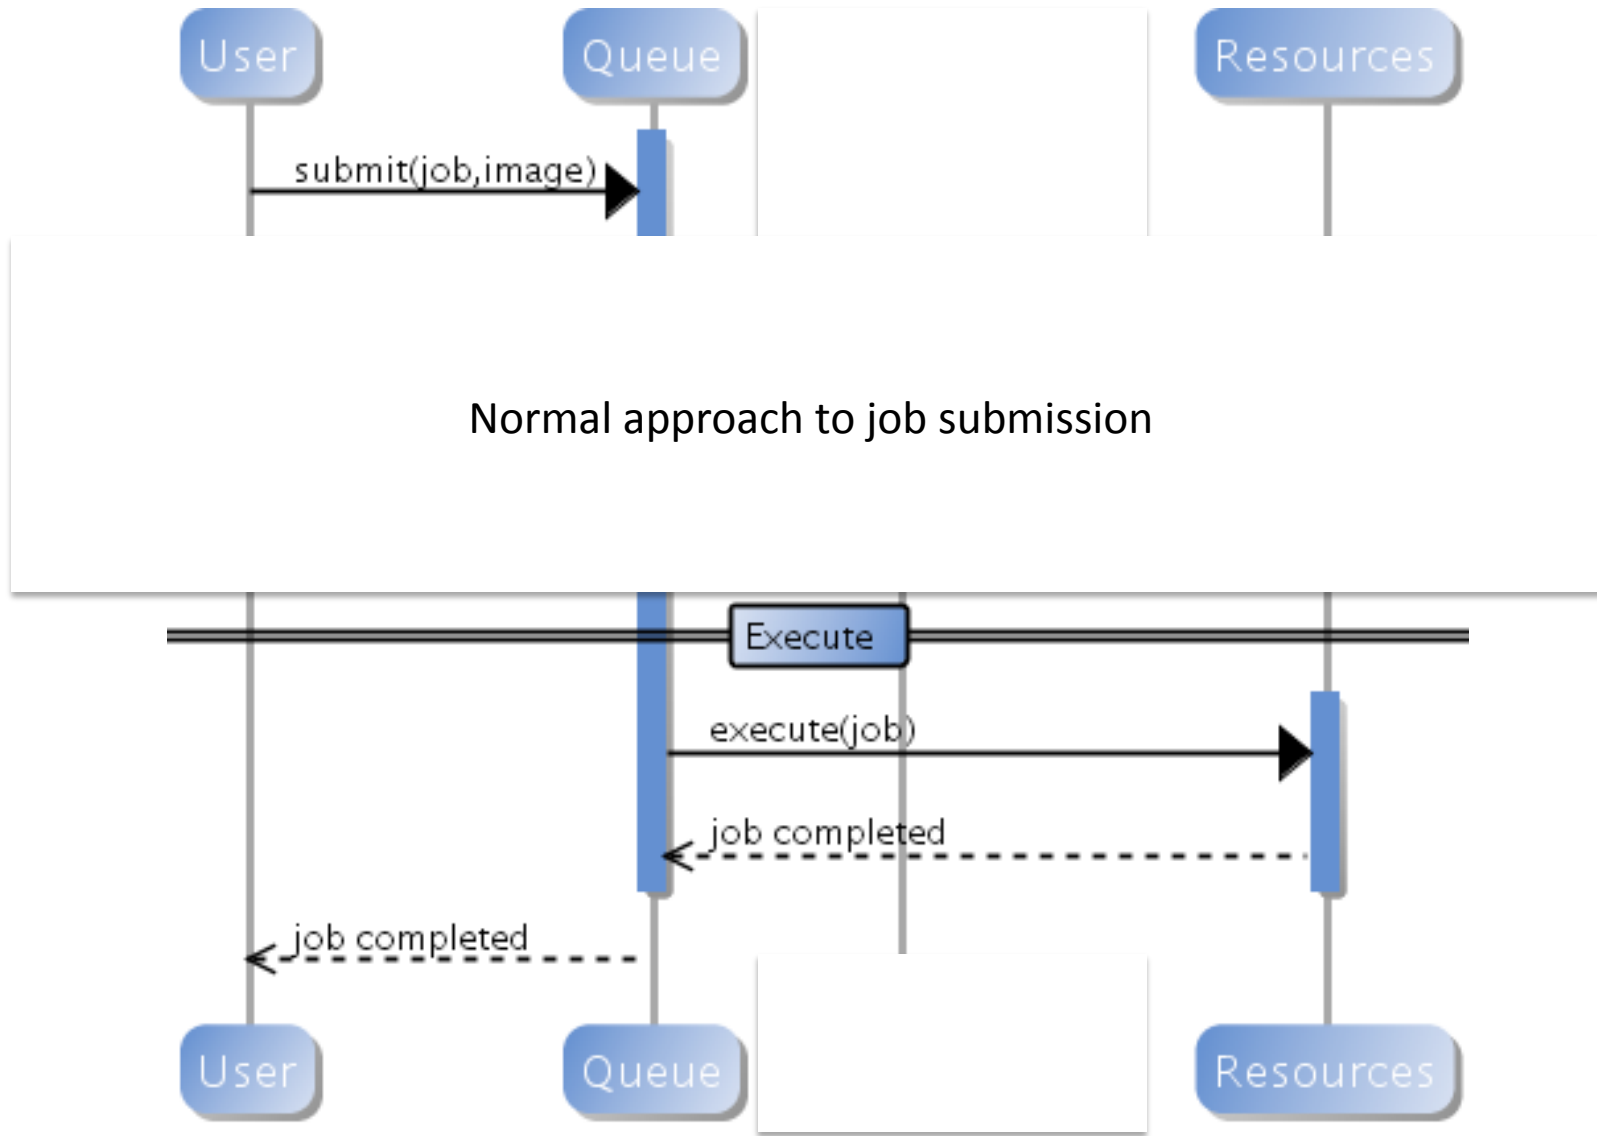
- Objective
 - Manage the provisioning for reproducible experiments
 - Coordinate workflow of experiments
 - Share workflow and experiment images
 - Minimize space through reuse
- Risk
 - Images are large
 - Users have different requirements and need different images

Per Job Reprovisioning

- The user submits a job to a general queue
 - This job specifies a custom Image type attached to it
- The Image gets reprovisioned on the resources
- The job gets executed within that image
- After job is done the Image is no longer needed
- Use case: Many different users with many different images



Custom Re provisioning





Reprovisioning based on prior state

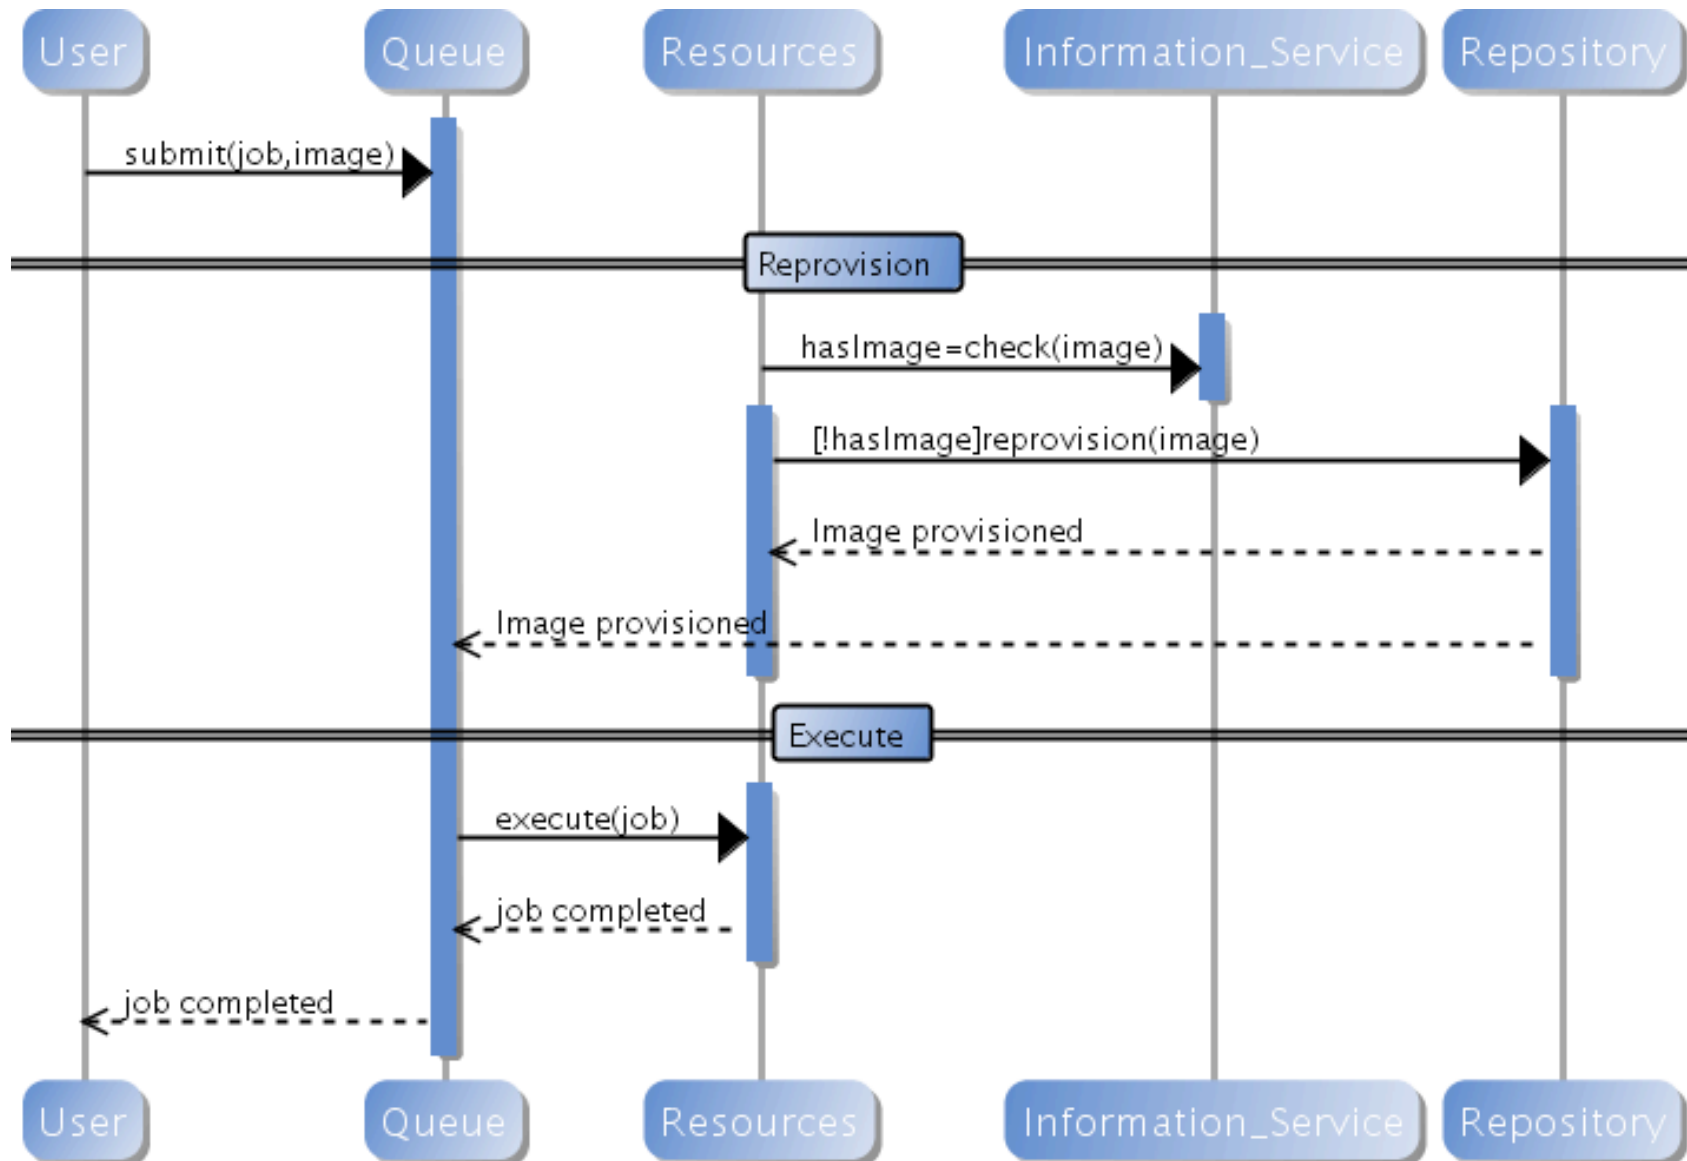


PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- The user submits a job to a general queue
 - This job specifies an OS (re-used stateless image) type attached to it
- The queue evaluates the OS requirement
 - If an available node has OS already running, run the job there
 - If there are no OS types available, reprovision an available node and submit the job to the new node
- Repeat the provisioning steps if the job requires multiple processors (such as a large MPI job)
- Use case: reusing the same stateless image between usages

Generic Reprovisioning





Manage your own VO queue

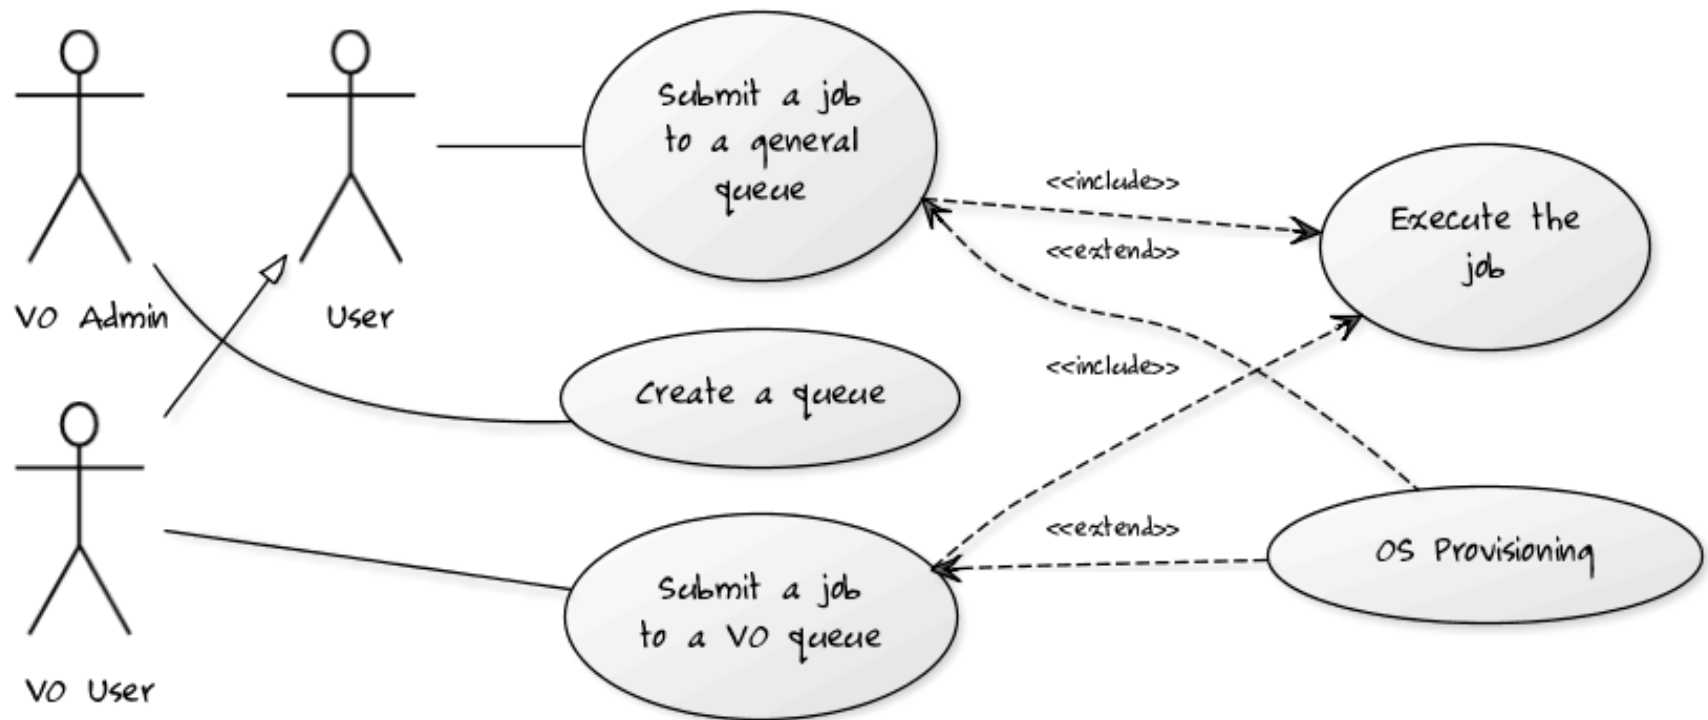


PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- This use case illustrates how a group of users or a Virtual Organization (VO) can handle their own queue to specifically tune their application environment to their specification
- A VO sets up a new queue, and provides an Operating System image that is associated to this image
 - Can aid in image creation through the use of advanced scripts and a configuration management tool
- A user within the VO submits a job to the VO queue
- The queue is evaluated, and determines if there are free resource nodes available
 - If there is an available node and the VO OS is running on it, then the job is scheduled there
 - If an un-provisioned node is available, the VO OS is provisioned and the job is then submitted to that node
 - If there are other idle nodes without jobs running, a node can be re-provisioned to the VO OS and the job is then submitted to that node
- Repeat the provisioning steps if multiple processors are required (such as an MPI job)
- Use case: Provide a service to the users of a VO
 - For example: submit a job that uses particular software
 - For example: provide a queue called Genesis or Hadoop for the associated user community
 - Provisioning is hidden from the users

VO Queue





Current Status of Dynamic Provisioning @ FutureGrid



PERVASIVE TECHNOLOGY
INSTITUTE
INDIANA UNIVERSITY

- FutureGrid now supports the Dynamic Provisioning feature through MOAB
 - Submit a job with 'os1' requested, if there is a node running 'os1' and in idle status, the job will be scheduled
 - If there is no node running 'os1', a provisioning job will be started automatically and change an idle node's OS to the requested one
 - When it's done, the submitted job will be scheduled there
 - In our experiment we used 2 rhel5 OS and dynamically switched between, one stateless and one statefull
 - In our experiment
 - Reprovisioning costs were approximately 4-5 minutes for statefull and stateless
 - Used sierra.futuregrid.org iDataPlex at SDSC



Difficult Issues



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- Performance of VMs poor with Infiniband: FutureGrid does not have resources to address such core VM issues – we can identify issues and report
- What about root access?
 - Typically administrators involved in preparation of images require root access
 - This is part of certification process
 - We will offer certified tools to prepare images
- What about attacks on Infiniband switch? We need to study this
- How does one certify statefull images?
 - All statefull images must have a certain default software included that auto-update the image, which is tested against a security service prior to staging
 - If an image is identified as having a security risk it is no longer allowed to be booted



FutureGrid CloudBench



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY




- Maintain a set of comparative benchmarks for comparable operations on FutureGrid, Azure, Amazon (e.g. <http://azurescope.cloudapp.net/> with Amazon and FutureGrid analogues)
- Need MapReduce as well

Indianapolis, USA,
Nov 30 – Dec 3, 2010



2nd IEEE International Conference on Cloud Computing Technology and Science

- About
- Committee
- Dates
- Program
- Workshops
- Call For Paper
- Submission
- Registration, Travel & Hotel

Downloadable CFP version   

“Cloud” is a common metaphor for an Internet accessible infrastructure (e.g. data storage and computing hardware) which is hidden from users. Cloud Computing makes data truly mobile and a user can simply access a chosen cloud with any internet accessible device. In Cloud Computing, IT-related capabilities are provided as services, accessible without requiring detailed knowledge of the underlying technology. Thus, many mature technologies are used as components in Cloud Computing, but still there are many unresolved and open problems. This conference and workshop series, steered by the [Cloud Computing Association](#), aims to bring together researchers who work on cloud computing and related technologies.

Submission Guidelines

Please submit your paper to the [CloudCom 2010 submission server](#) via an EasyChair account.

(* You are leaving the CloudCom and IU web sites and going to an independent company that is not affiliated with CloudCom or Indiana University.)

Keynote Speakers

- Ian T. Foster, Director, CI, Distinguished Fellow, Argonne National Laboratory, IL, USA
- Dennis Gannon, Director of Applications for the Cloud Computing Futures Group, Microsoft Research, WA, USA
- Kai Hwang, Professor, IEEE Fellow, Director, Internet and Grid Comp. Lab., USC, USA ([Abstract](#))

Tutorial

- [MapReduce/Hadoop Tutorial sponsored by Yahoo](#)
- Microsoft Azure Tutorial

Workshops

- Call for Workshops: Proposals for workshops are welcome.
- [International Workshop on Cloud Privacy, Security, Risk & Trust \(CPSRT 2010\)](#)
 - [Cloud Computing , HCI, & Design: Sustainability and Social Impacts](#)
 - [First International Workshop on Theory and Practice of MapReduce \(MAPRED'2010\)](#)
 - [The 1st International Workshop on Cloud Computing and Quality Assurance \(CLOUDit'10\)](#)



Acknowledgments



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

- This material is based upon work supported in part by the National Science Foundation under Grant No. 0910812 to Indiana University for "FutureGrid: An Experimental, High-Performance Grid Test-bed." Partners in the FutureGrid project include U. Chicago, U. Florida, San Diego Supercomputer Center - UC San Diego, U. Southern California, U. Texas at Austin, U. Tennessee at Knoxville, U. of Virginia, Purdue I., and T-U. Dresden.
- Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.