

**NRC-CNRC**

From *Discovery*  
to *Innovation...*

# Repoman: A Simple RESTful X.509 Virtual Machine Image Repository

Roger Impey




National Research  
Council Canada

Conseil national  
de recherches Canada

Canada

## Project Term



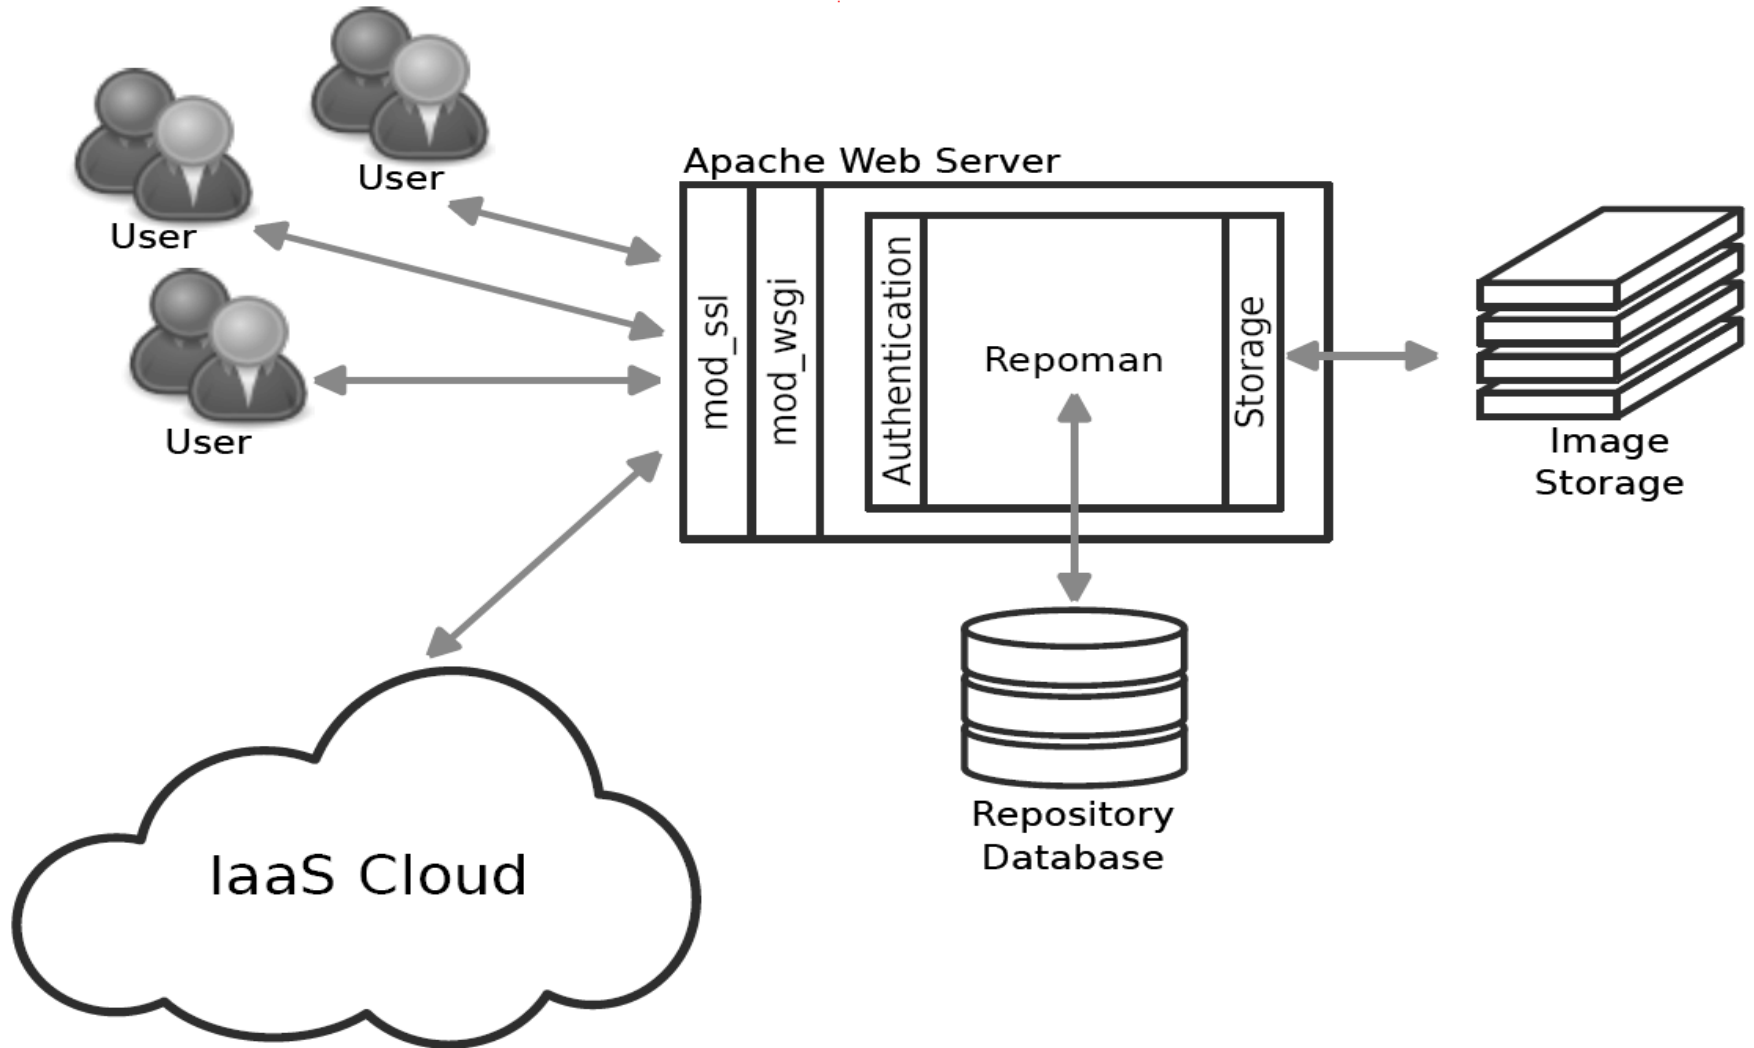
**University  
of Victoria**

- University of Victoria
- R.J. Sobie, M. Anderson, P. Armstrong, A. Agarwal, **Kyle Fransham**, D. Harris, I. Gable, C. Leavett-Brown, M. Paterson, D. Penfold-Brown, and
  - **Matthew Vliet**
- At National Research Council **NRC-CNRC**
- A. Charbonneau, R. Impey, and W. Podaima

# Why do we need Repoman?

- We need Repoman to facilitate the easy sharing, manipulation, storage and management of virtual machines image for use in the work we are doing in with Nimbus, and cloud scheduler.
- Current software was either too restrictive with one or more the following:
  - Not stand alone.
  - Authentication Scheme could not accommodate x509 grid certificates.
  - Sharing files was difficult without prior knowledge of what files are available.
  - Creating a consistent and enforced set of metadata was not easy.

# Basic Architecture of Repoman



# Features of Repoman

- Repoman is an X509 authenticated Web server that stores Virtual machines Images.
  - Nimbus uses X509
- It provides a simple API to store, share and manage VM images.
- Completely “stand alone”
- Installed using pypi or pip.
  
- Amazon S3 clones come the closest to providing similar functionality.
- The other one that comes close is Glance from OpenStack
  - When Repoman was started, this did not exist
  - no X509 authentication

# Implementation Details

- **Written in Python**
  - Pylons web framework**
    - One of the big Python web frameworks (lots of users, been around a while)
    - Easy to customize to your needs. If you don't like something in, change it.
    - Pylons apps are WSGI apps
- **SQLAlchemy**
  - Interface to many different databases without worrying about different implementation details
  - **Apache**
    - mod\_wsgi**
      - Allows for running Python WSGI applications within the Apache Web Server
    - mod\_ssl**
      - Provides SSL encryption for the webserver
      - Allows for checking of client certificates
      - This is KEY for authenticating users with x509 certificates

# Repoman Client

```
impeyr@vm019:~> repoman --help
```

```
usage: repoman [-h] [--help-all] [--version] [-H HOST] [-P PORT]
```

optional arguments:

-h, --help

--help-all

--version

-H HOST, --host HOST Override host setting

-P PORT, --port PORT Override port setting for host

--proxy PROXY Override default proxy certificate

SUBCOMMANDS:

get Download the specified image file

save snapshot and upload current system

rename rename an existing image from 'old' to 'new'

delete Delete an image from the repository

list List a users images stored in the repository

# Repoman Client

```
impeyr@vm019:~> grid-proxy-init
```

```
Your identity: /C=CA/O=Grid/OU=imsb.nrc.ca/CN=Roger Impey
```

```
Enter GRID pass phrase for this identity:
```

```
Creating proxy ..... Done
```

```
Your proxy is valid until: Tue Mar 15 03:40:39 2011
```

```
impeyr@vm019:~> repoman list
```

```
impey/cole.Feb072011.img.gz
```

```
impey/rhee.Feb072011.img.gz
```

```
impey/thames.Jan142011.img.gz
```

```
impey/rideau.Jan172011.img.gz
```

```
impey/cole.Jan142011.img.gz
```

```
impey/thames.Dec172010.img.gz
```

```
impey/vm019.sao.nrc.ca
```

```
impeyr@vm019:~>
```

Note; No GUI yet!



# Repoman Metadata

```
impeyr@vm019:~> repoman describe-image vm019.sao.nrc.ca
{'checksum': {'type': None, 'value': None},
 'description': None,
 'expires': None,
 'file_url': 'https://vmrepo.cloud.nrc.ca/api/images/raw/impeyr/vm019.sao.nrc.ca',
 'http_file_url': None,
 'hypervisor': None,
 'modified': 'Tue Mar 8 20:44:17 2011',
 'name': 'vm019.sao.nrc.ca',
 'os_arch': None,
 'os_type': None,
 'os_variant': None,
 'owner': 'https://vmrepo.cloud.nrc.ca/api/users/impeyr',
 'owner_user_name': 'impeyr',
 'raw_file_uploaded': False,
 'read_only': False,
 'shared_with': {'groups': [], 'users': []},
 'size': None,
 'unauthenticated_access': False,
 'uploaded': 'Tue Mar 8 20:44:17 2011',
 'uuid': 'd5c71b1749c411e087980025648d3a02',
 'version': 0}
```

## JSON (JavaScript Object Notation)

# Topical Use Case End-user

- SysAdmin created a base image(s) that was preconfigured with the required software and libraries.
- Base Image is shared with user or group on Repoman.
- User boots an interactive version of that image, and customizes it to his/her needs.
- Once customized, User saves the image back to the Repoman repository.
- Save option uses rsync on a mounted file system

# Distributed Clouds

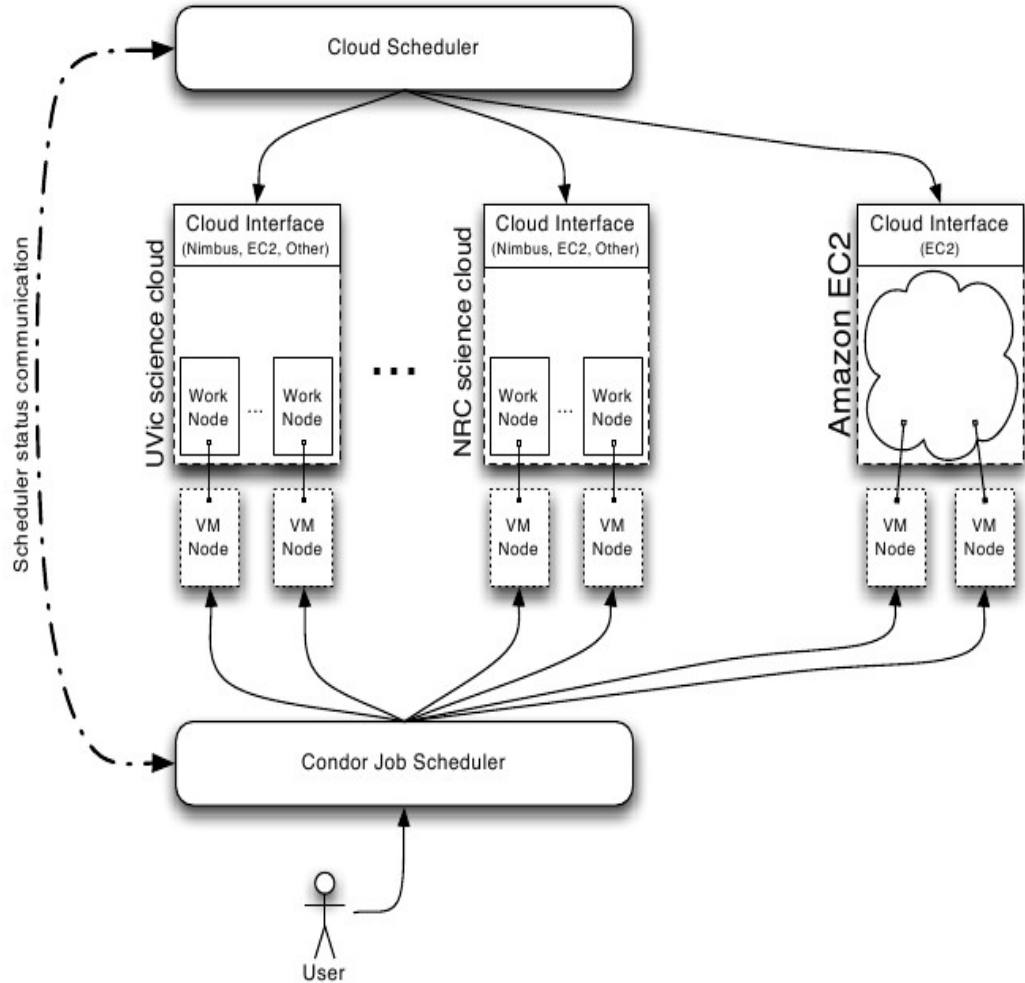
We encapsulate the software in a Virtual Machine, and we can run many types of new and old applications on a single cluster or multiple clusters.

This technology (Infrastructure as a Service) already exists:

- - Amazon EC2 / Rackspace (Commercial providers)
- - Nimbus / OpenNebula / Eucalyptus (Open Source)
- We use Nimbus, together with Job Scheduler and Condor.

# Cloud Scheduler

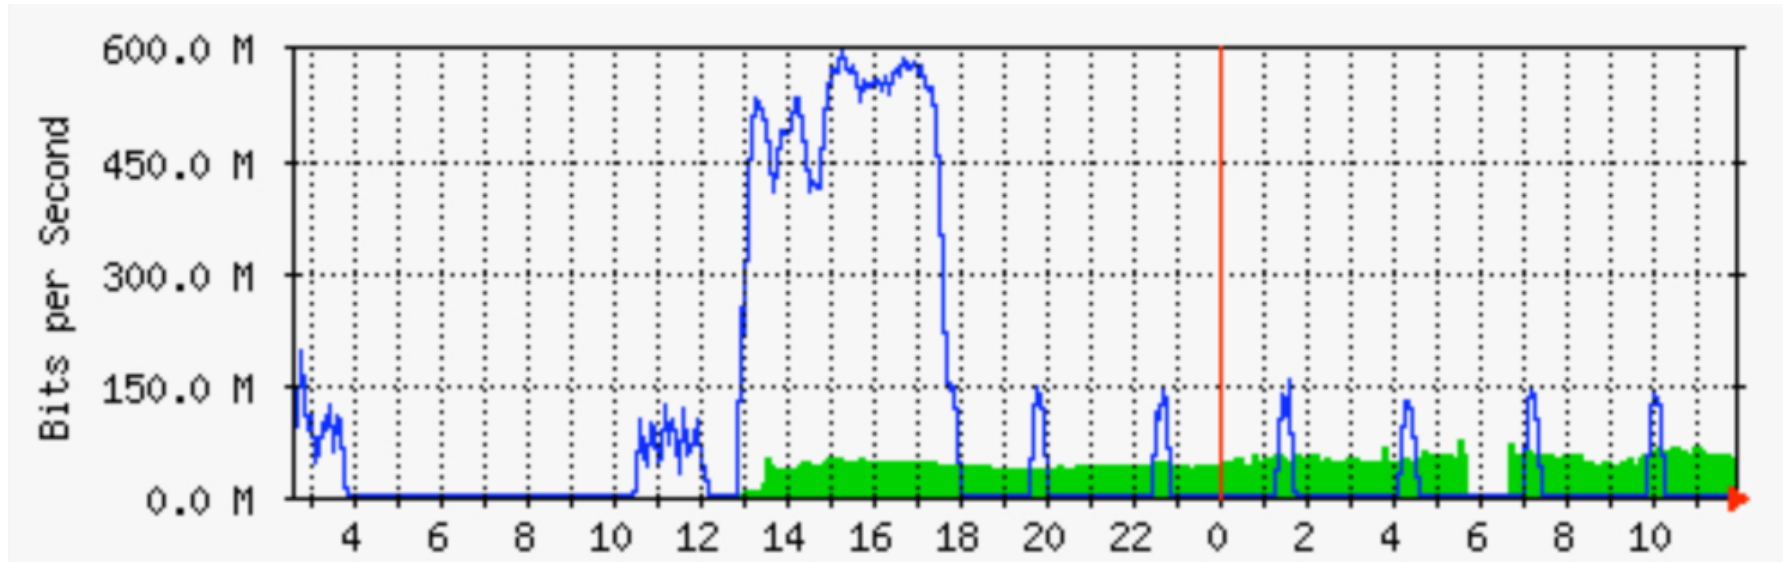
- The Cloud Scheduler knows about the IaaS resources available, and it polls the job scheduler's queue.
- When it sees a new job in the queue, the Cloud Scheduler selects a resource and sends the instruction to boot a VM.
- The VM boots, registers itself with the job scheduler, and runs the job.



# Topical Use Case Production

- User then submits a set of batch jobs to the Condor queue.
- Cloud Scheduler tells Nimbus to boot virtual machines for the batch jobs in the queue
- Nimbus retrieves the needed VM from the Repoman server and boots the VM.
- The jobs run and drain the Queue.

# Network Traffic



The inbound and outbound network traffic at the NRC cluster in Ottawa. The VMs were stored at NRC and the data at UVIC. The outbound traffic (shown by the blue line) shows the transfer of the 60 VMs to UVIC.

## Summary

- Running a distributed Cloud system in production.
- Currently used heavily for HEP (BarBar) and Astrophysics (CANFAR).
- Cloud Scheduler can be used to take advantage of many clouds, including locally owned resources, FutureGrid and Amazon EC2.
- RepoMan is used by end-users to easily share, manipulate, store and manage virtual machines images

# Open Source Software

Systems employs open source software which are readily available.

Missing pieces developed in-house as open source projects.

- **Cloud Scheduler:**
  - <https://github.com/hep-gc/cloud-scheduler>
- **Repoman:**
  - <https://github.com/hep-gc/repoman>



**NRC-CNRC**

From *Discovery*  
to *Innovation...*

## Acknowledgments



**canarie**



Canada's Advanced Research and Innovation Network  
Le réseau évolué de recherche et d'innovation du Canada



National Research  
Council Canada

Conseil national  
de recherches Canada

Canada