

# Virtualizing the Philippine e-Science Grid

International Symposium on Grids and Clouds 2011  
25 March 2011

Rene Mendoza  
rene@asti.dost.gov.ph  
Advanced Science and Technology Institute



# Outline

- Who we are
- ASTI's initiatives on Cluster and Grid Computing – the PsciGrid
- Problems and Headaches
- Going the virtual way
- Summary
- Future plans

# The Advanced Science and Technology Institute (ASTI)

- A research and dev't institute under the Philippine Government's Department of Science and Technology (DOST)
- Our mandate -  
Conduct R&D in the advanced fields of ICT and Electronics



# The Philippine e-Science Grid (PsciGrid)

Funding: Department of Science and Technology Grants-In-Aid (DOST-GIA) Program ~ (1M USD)

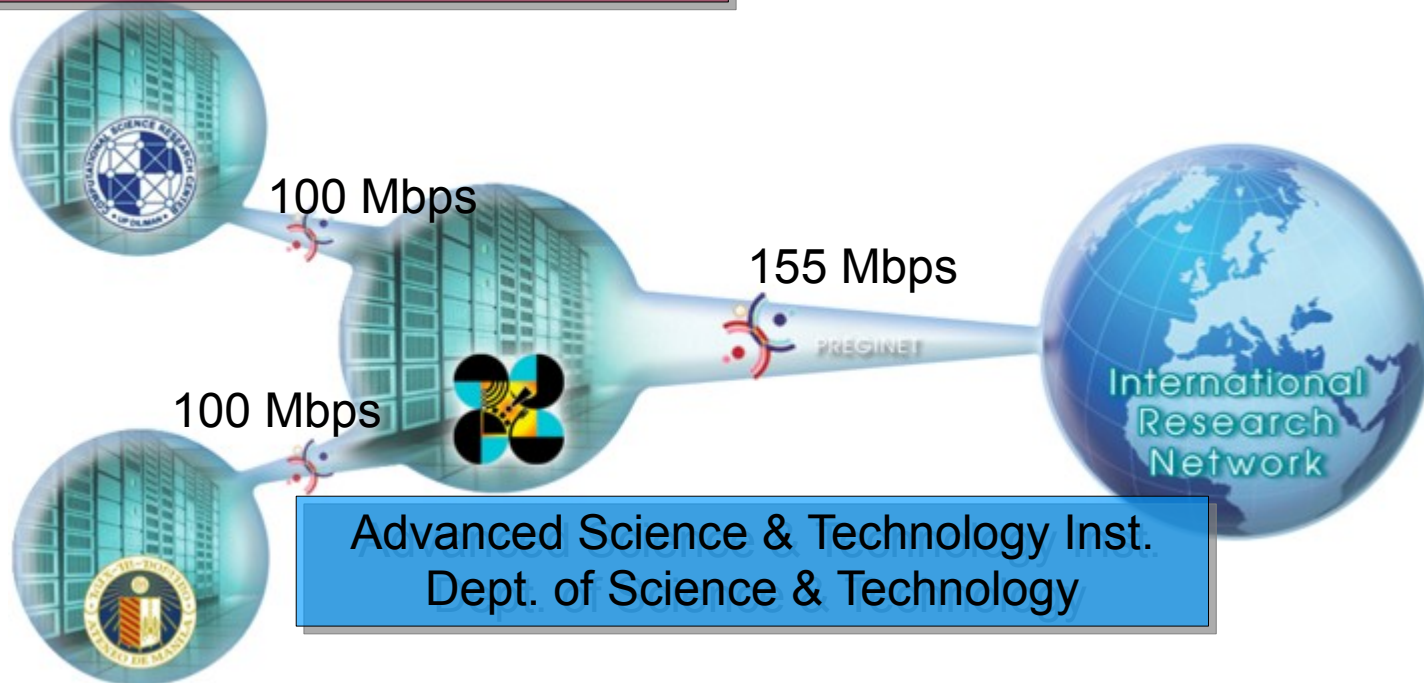
Duration: 3 years (January 2008 – December 2010); Extended until  
**June 2011**

## Objectives:

- Establish a national e-Science grid infrastructure in the Philippines that will enable collaborative research among local educational and research institutions.
- Provide seamless access to high-performance computing resources and applications for Life and Physical Sciences

# The PsciGrid Infrastructure

University of the Philippines  
Computational Science Research Center



Advanced Science & Technology Inst.  
Dept. of Science & Technology

Ateneo De Manila University  
School of Science and Engineering

# ASTI's HPC

## Computing

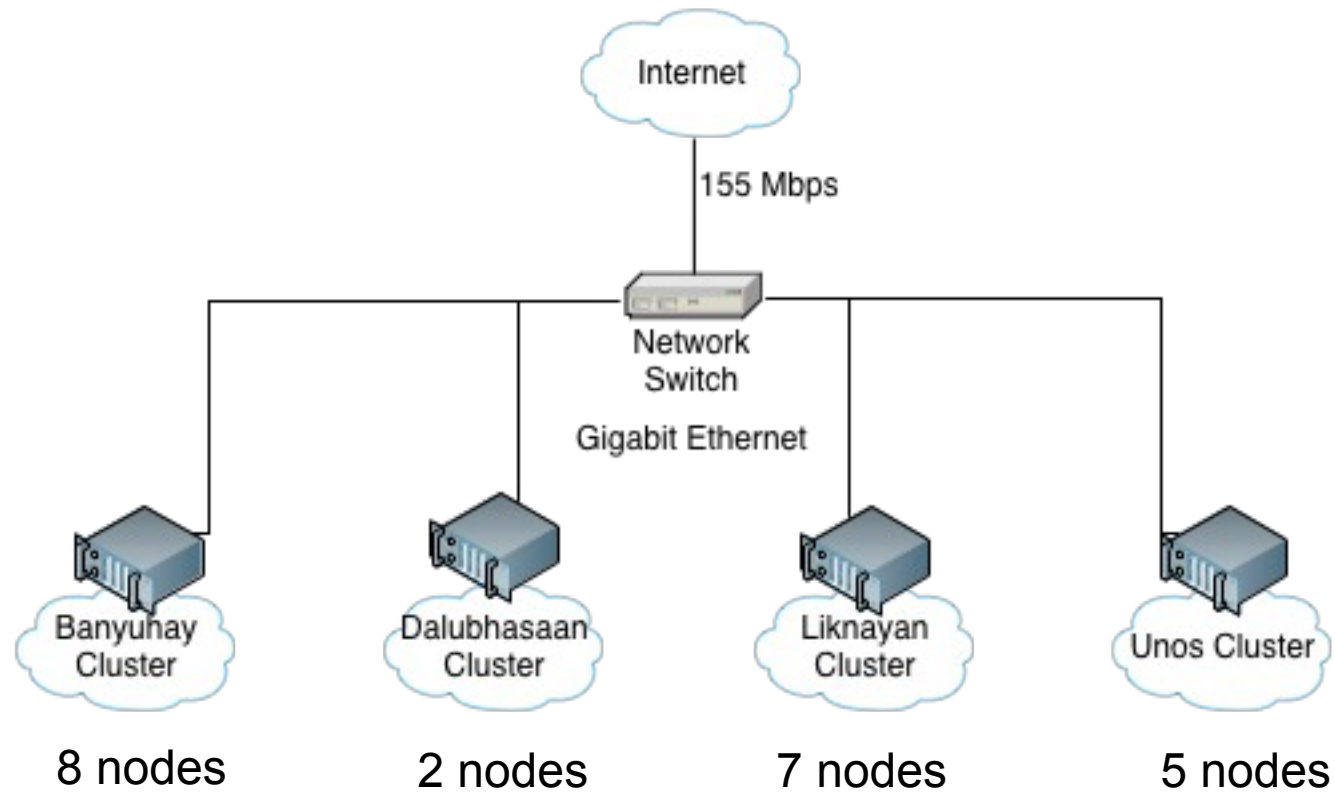
- 51 computing nodes (2 x 2.0 GHz Intel Xeon), 408 cores
- 300GB/500 GB of disk space and 16GB/24GB of RAM per node
- 8 FPGA – based hardware accelerators

## Storage

- 6TB for raw data
- 4TB for DNA and protein sequences (Bio-Mirror)
- 4TB for software mirror
- Additional 16TB (usable 12TB)




# ASTI's HPC



# Users and Applications

- **Users**



- **Applications**

- RegCM, WRF, MM5
- Bioinformatics apps
- SPECFEM3D



# Problems and Headaches

- Job scheduling inefficiency
- Unbalanced utilization of compute nodes between clusters some are unused
- Heterogeneous cluster sets
- Disk space shortage
- Low uptake for targeted users of the HPC



# To charge or not to charge?

- Government projects have very limited funding
- Must come up with a sustainability plan
- Charge users based on some criteria?



# To virtualize or not to virtualize?

Two Options:

1. Setup one big physical cluster and provide access to all local users OR
2. Setup individual virtual clusters running on one common physical cluster

OUR CHOICE? 2



# Why Virtualization?

- Dynamic or on-demand reallocation of unused compute nodes to busy clusters.
- Better user isolation and fine-grained control
- Easier backup
- Flexibility

# Cluster Usage Policy (aka "The Ten Commandments in Using the HPC Cluster")

1. A user needs to apply for a valid login-account.
2. SSH access is allowed through key-based authentication only.
3. No sharing of login accounts.
4. User should report any suspicious activity regarding his/her account.
5. No running of jobs on the head node.

# The Ten Commandments in Using the HPC Cluster

6. Jobs should be run using Torque.

7. Paying users will be prioritized over non-paying users.

8. The maximum number of running jobs per user is 4.

9. No logging in on the compute nodes.

10. Repeated violations of these rules may lead to suspension of the user login account

# Storage Policy

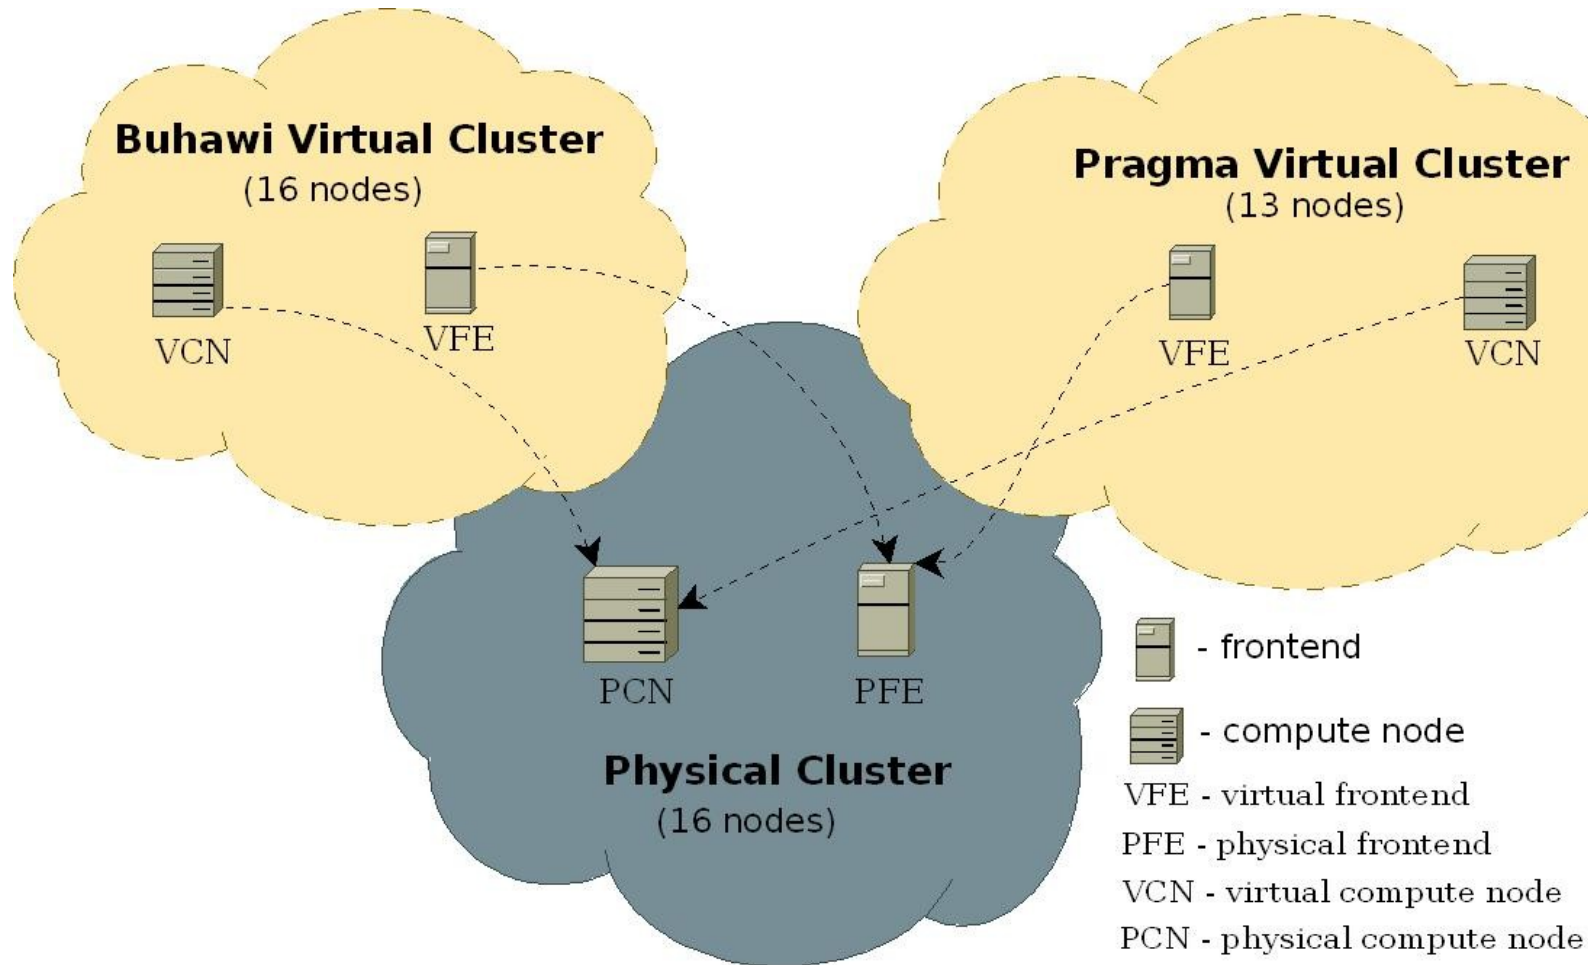
1. Each user is given 3 GB of HD space in his home directory.
2. Files in the home directory will remain as long as the account is valid.
3. A work space will be provided under /work.
4. The work space has no quota, but when the filesystem becomes full, the work space is reduced to 50% (delete oldest first).
5. Larger disk storage may be requested but is subject to disk space availability.

# Experimental Virtual Clusters Setup

- Two clusters
  - Buhawi (meteorology cluster)
    - 16 nodes, 4 cores per node, 8 GB memory, 36 GB HD
  - PRAGMA
    - 13 nodes, 2 cores per node, 4 GB memory, 36 GB HD



# Virtual Clusters Setup



Physical Cluster: Rocks 5.3 with Xen Roll

# Summary

- ASTI's HPC is gradually moving to a full virtual infrastructure (uphill battle)
- Usage and storage policies are finally being implemented.
- Exploring the possibility of setting up a sustainability plan



# Future Plans



- Continue virtualization efforts (3-6 months)
- Procurement of blade servers (next 2-3 years)
- Diskless clusters (within 1 year)
- IPv6 clusters (within 1 year)
- Storage clusters (within 1 year)
- Use of the HPC for the government's disaster mitigation plans (within 1 year)
- Integrate PsciGrid with government-wide data center (within 1 year)



Thank you!