Virtualization Technologies in Support to Scientific Applications

Bruno Schulze

Third Brazil-France workshop
02 – 05 Sept 2013

Computing on Clouds

- Clouds have emerged as a solution to provide resources on demand
- Scientific communities viewing or embracing Clouds as an on-demand-infrastructure for their needs
- Elasticity and almost “unlimited scalability” of Clouds
- Trade-off between Cost and Performance
- Performance dictated by Scientific Applications' behavior

Computing on Clouds

But still some issues remain ...

- Variability of Cloud environments
- Security Concerns / Access Concerns
- Networking and Intra-connectivity
- SLA and the infrastructure served by Cloud Providers

Exploring Example I

Helix Nebula - ATLAS use cases

- What are the financial costs of networking data transfers into and out of cloud resources, short and long-term data storage in the cloud, and CPU resources for running the various ATLAS use cases?
- What are the appropriate Service Level Agreements and how can they be defined for broader usage?
- What are the policy and legal constraints in moving scientific data across academic networks into commercial resources and back again?

Exploring Example II

Uber-Cloud - HPC Experiment: Actual Case Studies Focused on Solving Real Computing Problems

- Benefiting from remote access to HPC
  - Challenges still remain ...
    - trusting in the resource provider
    - giving away some control over applications, data & resources
    - security
    - provider lock-in
    - software licensing
    - unfamiliar pay-per-use computing model, and
    - a general lack of clarity in distinguishing between hype and reality

http://www.hpcwire.com/hpcwire/2012-09-20/half-time_in_the_uber-cloud.html

Virtualization Technologies

- Hypervisor development
- Hardware-assisted virtualization development
- Scientific Application Layer **X** Hypervisor Layer **X** Hardware Layer

Clouds Infrastructure

- Multitude of Cloud Environments
- High-Performance Parallel and Distributed Computing
- How to use clouds to benefit Scientific Applications

Cloud Approach

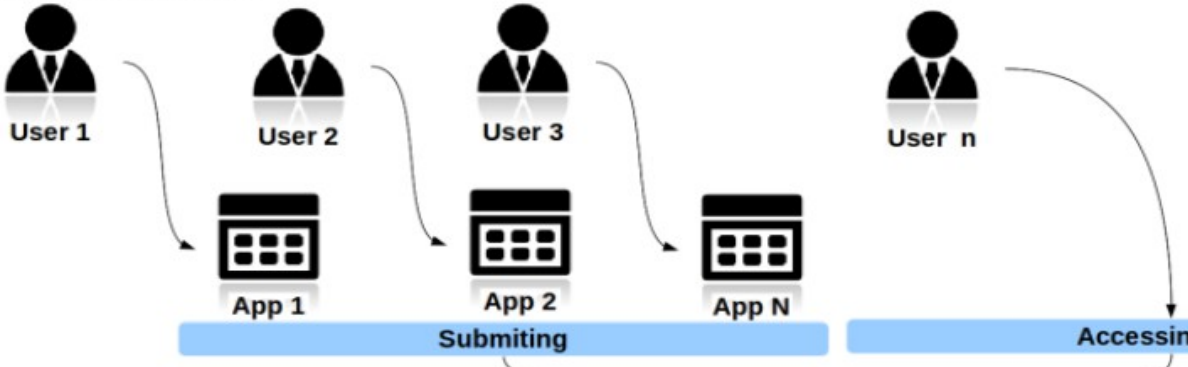

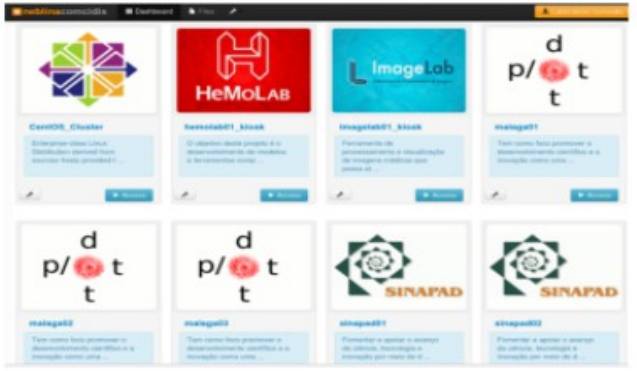

Cloud Approach

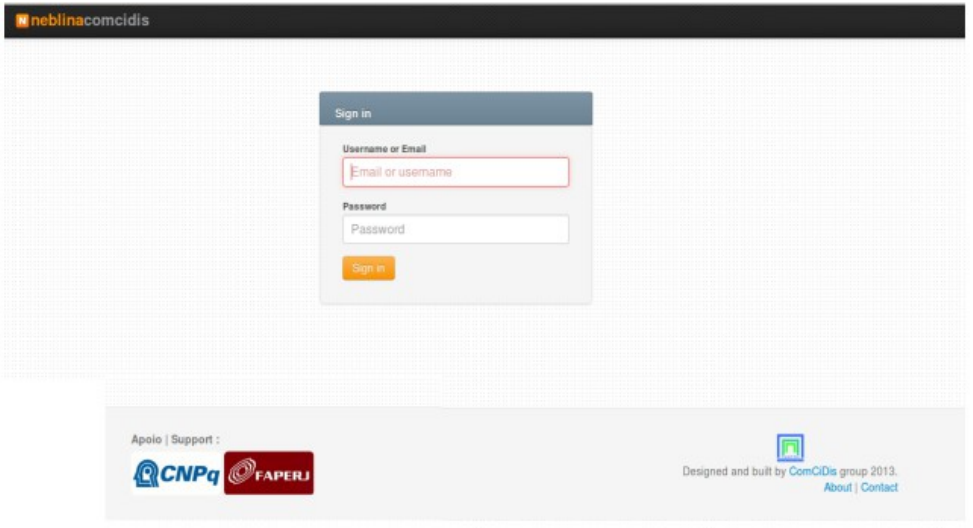

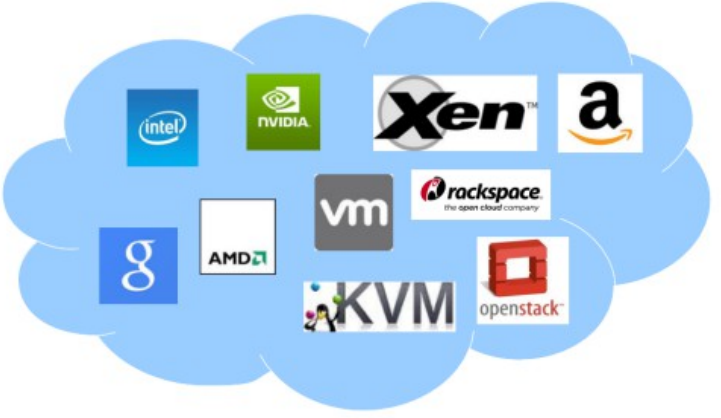

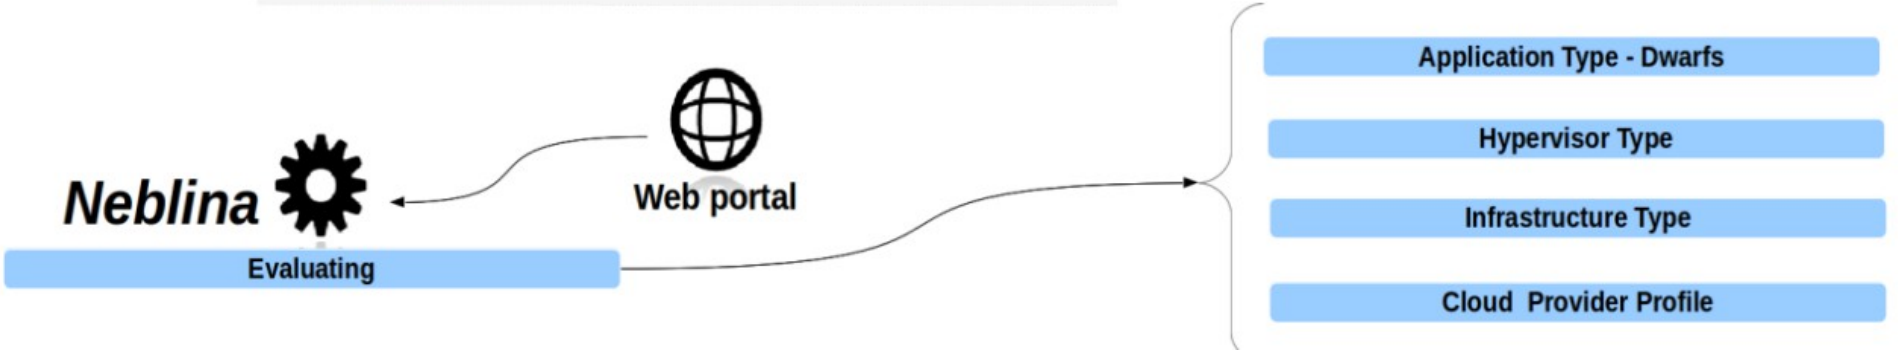

Cloud Approach

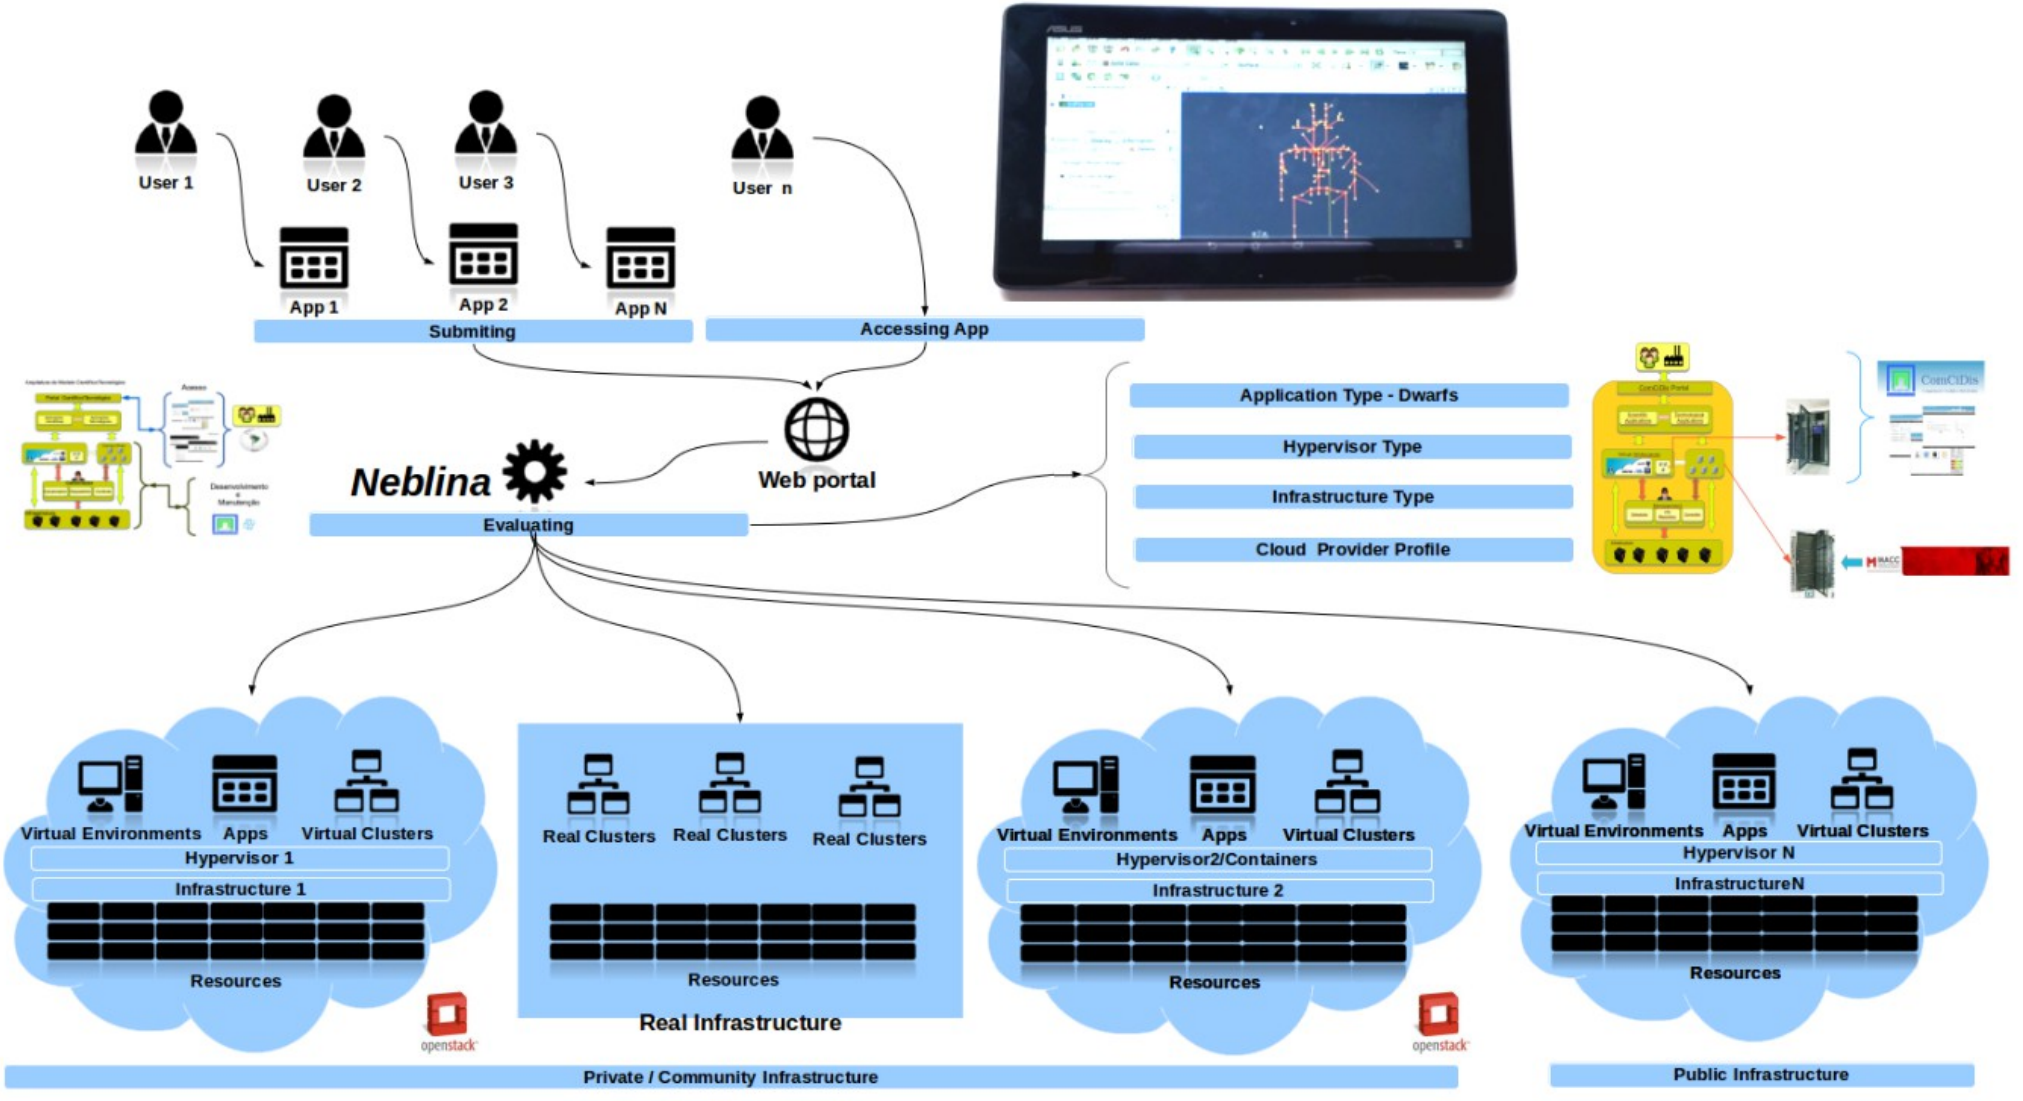

Neblina

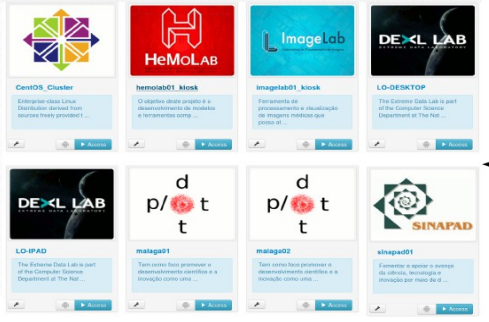

Accessing App

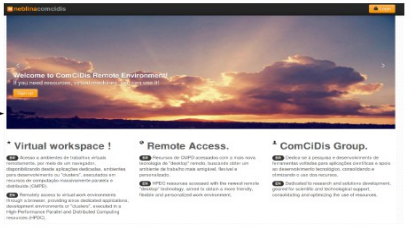

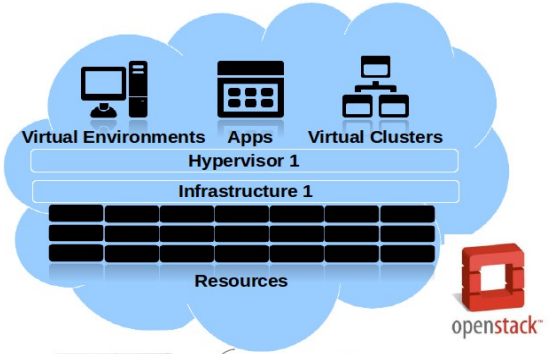

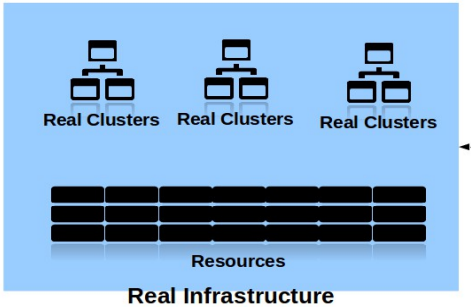

Submiting

O cluster da SGI

94 nós de trabalho perfazendo um total de 1128 núcleos; com rede Ethernet, sendo prevista a sua atualização para InfiniBand.

12 servidores GPGPU cada uma com dois processadores/seis núcleos perfazendo um total de 144 núcleos de CPU.
Cada servidor possui 2 placas Tesla M2040 cada uma com 448 núcleos de GPU, perfazendo um total de 10752 núcleos de GPU.

1 servidor com processador AMD com 64 núcleos de CPU, e

Dois servidores Intel com 12 núcleos cada e com uma placa Xeon Phi cada um com 61 núcleos cada placa.

Cluster Bull HPC

100 nós de trabalho com dois processadores de seis núcleos cada, com um total de 1200 núcleos

Possui ainda 4 nós de GPGPU com um total de 48 núcleos de CPU e 3584 núcleos de GPU

A última expansão (rack novo) veio com:

4 nós de trabalho duplos, cada um com com dois processadoes em cada placa, perfazendo um total de 128 núcleos, e

Um nó de login com 16 núcleos de processamento.

Private Infrastructure

Neblina

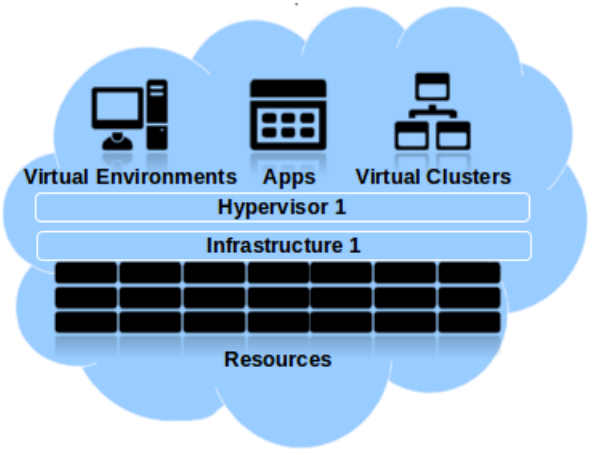

SGI Cluster

94 work nodes, 1128 cores; Ethernet (InfiniBand Installation schedule);

12 GPGPU servers, each server with 2 Tesla M2040, with 448 GPU cores - 10752 GPU cores total.

1 AMD server with 64 CPU cores, and

2 Intel server with Xeon Phi co-processor

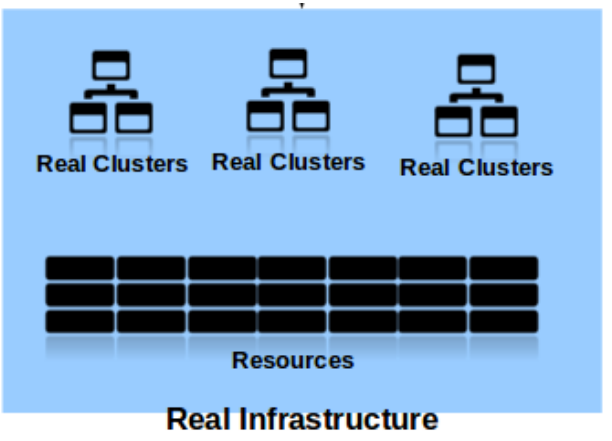

Bull HPC Cluster

104 work nodes, 1328 cores; InfiniBand;

4 GPGPU servers, each server with 2 Tesla M2040, with 448 GPU cores - 3584 GPU cores total.

- Conducting research and development in computer science and its medical applications, especially the computer simulation and modeling of the physiological systems that integrate the human body
- Promoting the development of medical image processing, scientific visualization and virtual reality in the development of medical applications directed to computer-aided diagnosis, treatment, surgical planning, medical training and accreditation
- Training human resources and promoting technology transfer and innovation to the area of health assisted by scientific computing.

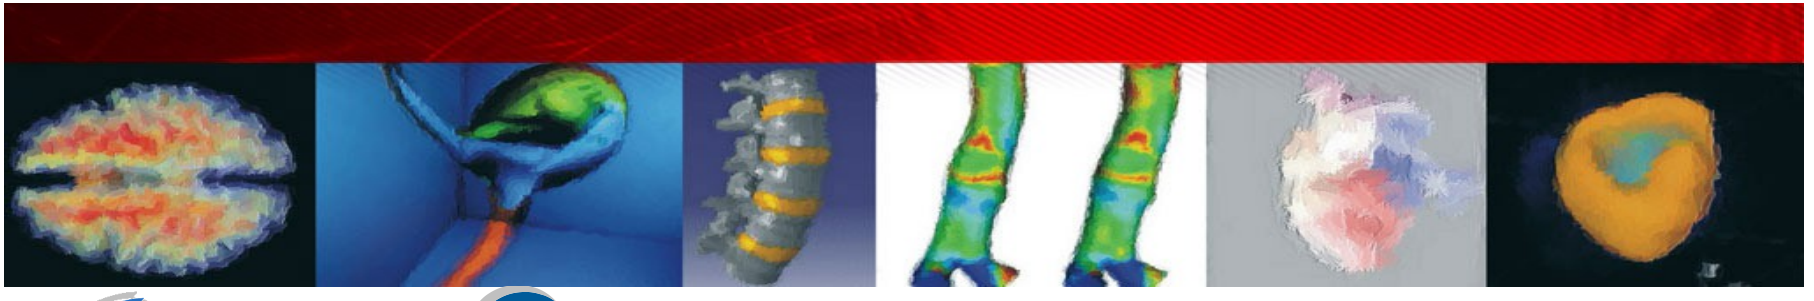

FAPERJ
Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro
CNPq
Conselho Nacional de Desenvolvimento Científico e Tecnológico

Neblina

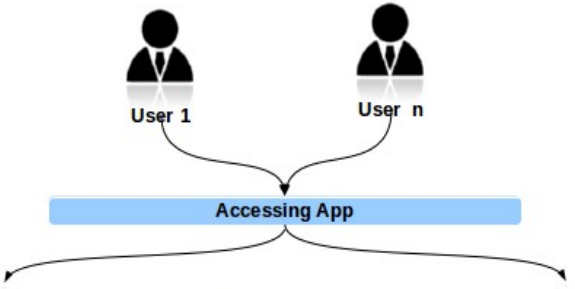
User 1
User n
Accessing App

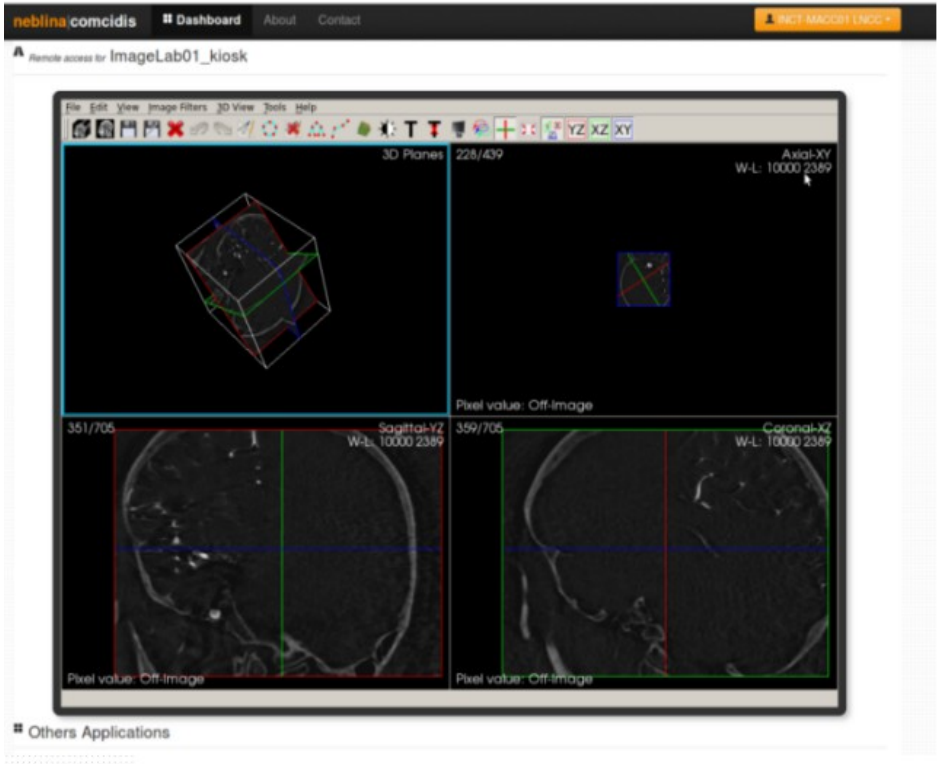
neblina comcidis
Dashboard
About
Contact
ImageLab01_kiosk
3D Planes
Axial-XY
Sagittal-YZ
Coronal-XZ
Pixel value: Off-image
Others Applications

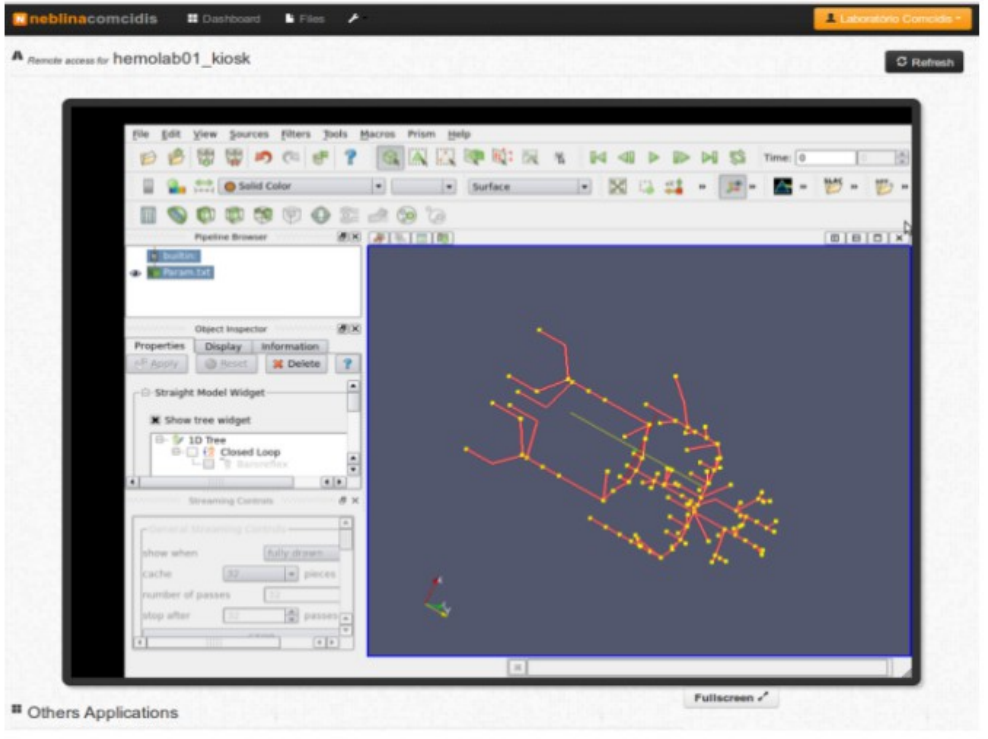
neblinacomcidis
Dashboard
Files
hemolab01_kiosk
Refresh
Others Applications

INCT-MACC/ComCiDis

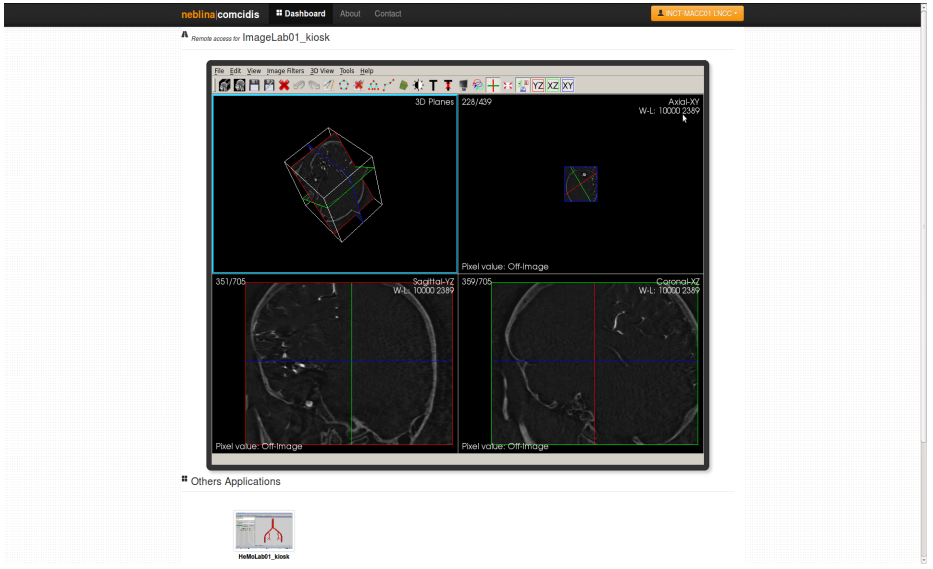

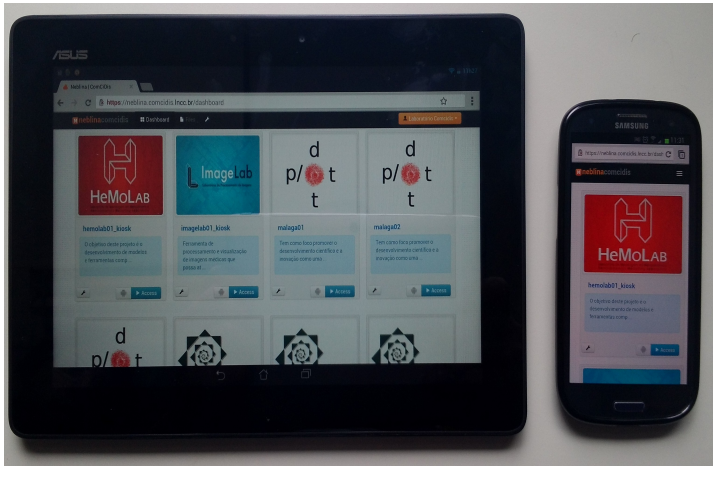

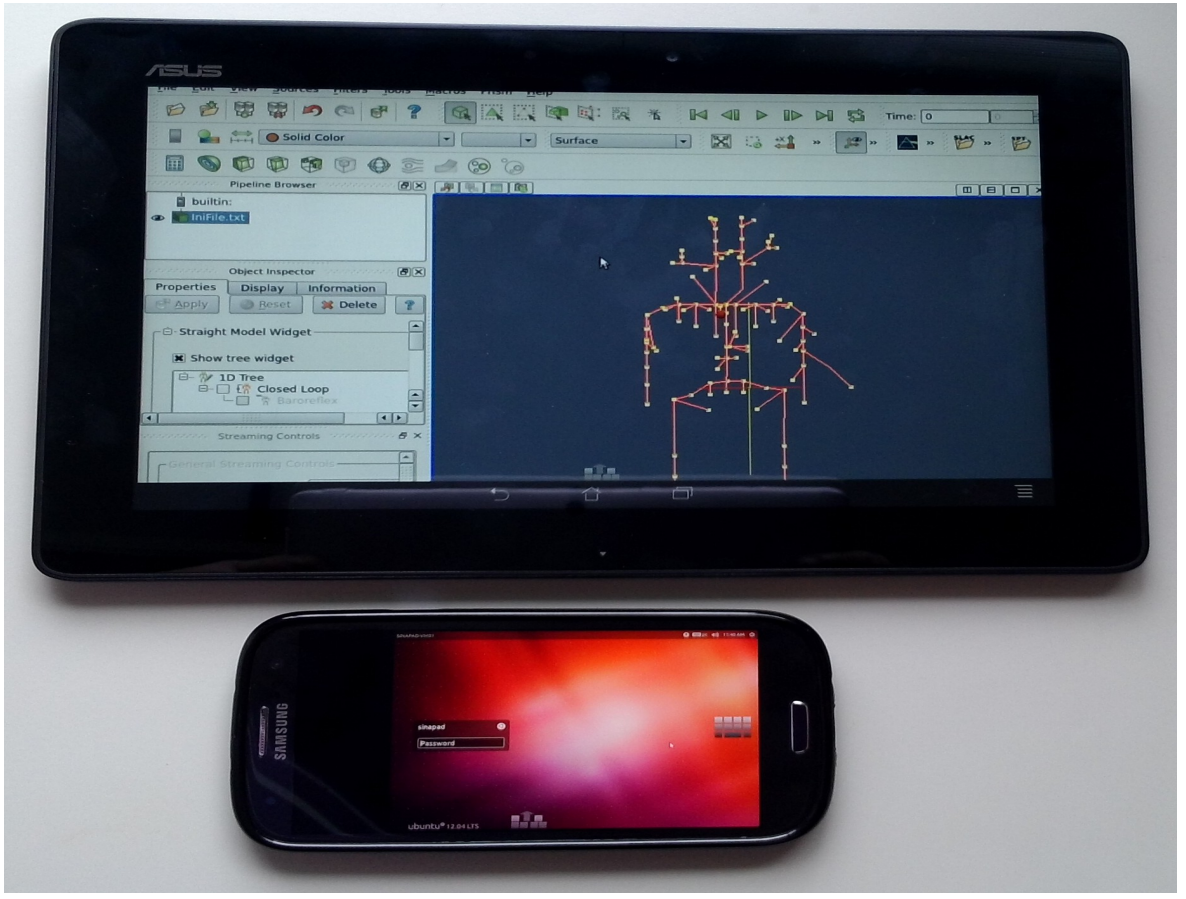
ASUS
Solid Color
Surface
Time: 0
Pipeline Browser
builtin:
IniFile.txt
Object Inspector
Properties
Display
Information
Apply
Reset
Delete
Straight Model Widget
Show tree widget
1D Tree
Closed Loop
Baroreflex
Streaming Controls
SAMSUNG
sinaped
Password
ubuntu 12.04 LTS

Conclusions

- Mixed Infrastructure providing useful resources to scientific research / applications
- Consolidation and usage of many environments
- INCT-MACC – HemoLab / ComCiDis virtual appliances with remote access
- Usage in Testing, Validation and Knowledge Dissemination
- Resource Optimization and Collaboration

Future Work

- Virtualization of many core / multicore new architectures (Nvidia, Xeon Phi, …)
- Deployment of many core / mutlicore based appliances
- Porting of new applications / appliances to Neblina
- Consolidation of Hadoop Virtual Appliances
- Improvements in access control and security

Thanks

schulze@lncc.br

Cloud Computing

- Clouds have emerged as a solution to provide resources on demand;
- Scientific communities viewing or embracing Clouds as an on-demand-infrastructure for their needs;
- Elasticity and almost “unlimited scalability” of Clouds;
- Trade-off between Cost and Performance;
- Performance dictated by Scientific Applications' behavior;

GPUs in Clouds PCI Passthrough

- Passthrough PCI-E GPU device to DomU
- Use Nvidia Tesla CUDA programming model
- Intel VT-d or AMD IOMMU extensions
- Xen / VMware
- ComCiDis has 12x 24GB memory nodes each with 2x GPU's (Tesla C2050)