

Virtualization at the University of Reims Champagne-Ardenne in an HPC context

M. Krajecki, F. Nolot and L. Lucas



Outline

- Introduction
- Virtualization for training and teaching
- Virtualization for remote visualization and HPC
- Future works



University of Reims

Université de Reims Champagne-Ardenne (URCA)

Multidisciplinary university

- about 25 000 students
- a wide initial undergraduate studies program
- graduate studies and PhD program linked with research labs




The ROMEO HPC Center is a platform hosted by URCA

- Funded by European Community, The french government, Champagne-Ardenne Council and the city of Reims
- high performance computing resources
- for both industrial and academic researchers in the region
- an in-depth expertise in different engineering fields: HPC, applied mathematics, physics, biophysics and chemistry.
- first *Cuda Research Center* in France (2012)



Integrated in the European HPC ecosystem

- link between large hardly accessible national centers and small research laboratories and SMEs of the region.
- member of the French Tier 1.5 network *equip@meso* managed by *GENCI*
- member of the European Platform *ETP4HPC*



Romeo HPC Tesla Cluster

Computing



Centre de Calcul ROMEO



Displaying



5th 3131 MFLOPS/W
Bull Cool Cabinet Door



151th 254.9 Tflops
Linpack




260 NVIDIA Tesla
K20X accelerators



130 Bull servers
bullx R421 E3 – Bull AE & MPI

260 INTEL Ivy Bridge E5-2650 v2 Processor, non-blocking **Mellanox Infiniband**, Slurm, 88 To Lustre (NetApp), 57 To home, 100 To Storage



Big Data, on-demand and remote

VirtualGL technology servers
Quadro 6000 & 5800

NVIDIA GRID + Citrix Virtualisation
NVIDIA VGX K2

Scalable Graphics 3D cloud solution
NVIDIA K6000



Virtualization on ROMEO cluster : what are the needs ?

- There are some situations where our HPC cluster's environment is not suitable for user needs:
 - Application requires different operating system (e.g Windows)
 - Application requires different versions of base system libraries and tools
 - Application requires specific setup (installation, configuration) of complex software stack



Virtualization on ROMEO cluster : what are the needs ?

- Application requires special network configuration
- Application requires privileged access to operating system
- ... and combinations of above cases



Outline

- Introduction
- **Virtualization for training and teaching**
- Virtualization for remote visualization and HPC
- Future works



Virtualization : a new world

Container

Isolation

Type 1 Hypervisor

Path Through

Type 2 Hypervisor

Baremetal

Software Defined Network

Network Function Virtualization

PV

Hardware Virtualization

Paravirtualization

HVM



Advantages for training and teaching

- Deploy the right configuration
- Resources available 24h/24h
- Scalability experimentations is easier
- Remote or local access to resources
 - Laptop can have many VM
- Automated correction



How we use virtualization ?

- Starting in 2008 : Remotelabz project
 - Remote control of real network equipment
 - Remote practice exercises for students
 - Booking system and availability 24h/24h
 - Group working is possible on the same resource



Réservation


Fichiers

Groupes

Bugs



Fin à 23:00:00, reste 00:34:54



Remotelabz | Contrôle de Routeur1 - Mozilla Firefox

remotelabz.univ-reims.fr:8080/student/lab/view?nom=Routeur1&hote=10.0.1.21&port=2074&type=telnet

```

Remote Access Server of Reims Cisco Metacad

User Access Verification

Username: m2info
Password:
Password OK


You are connected on line 1/8

% Please answer 'yes' or 'no'.
Would you like to enter the initial configuration dialog? [yes/no]: no

Press RETURN to get started!
  
```

Connected to 10.0.1.21 2074 online

Fin à 23:00:00, reste 00:32:51



And now – Remotelabz v2

- Virtualization of network equipment
 - Evolving tools : Dynamips/Dynagen, IOU then vIOS, VyOS, vASA, ...
- Configuration with Ansible
- Use of HTML5, Websocket , virtual switch, http proxy, ...



Test de configuration réseau sous linux

- Connectez-vous avec les identifiants indiqués sur le terminal
- Lister les interfaces réseaux avec la commande

- Identifier la carte réseau ethX (la valeur de X dépend de votre machine)
- Tapez

afin de récupérer une IP via le serveur DHCP configuré sur le réseau

- Trouvez l'adresse IP que le serveur DHCP vous a attribuée en tapant la commande

- Testez votre connexion réseau en effectuant un

Connected (unencrypted) to: QEMU (Test-VNC-Florent-50)

Send CtrlAltDel

```
root@vm:/home/toto# ip addr show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        inet6 ::1/128 scope host
            valid_lft forever preferred_lft forever
2: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP ql
en 1000
    link/ether 00:18:de:66:66:50 brd ff:ff:ff:ff:ff:ff
    inet 10.22.9.199/24 brd 10.22.9.255 scope global eth1
    inet6 2001:660:4601:7008:218:deff:fe66:6650/64 scope global dynamic
        valid_lft 2591997sec preferred_lft 604797sec
    inet6 fe80::218:deff:fe66:6650/64 scope link
        valid_lft forever preferred_lft forever
root@vm:/home/toto# ping www.google.fr
PING www.google.fr (216.58.208.195) 56(84) bytes of data.
64 bytes from par10s21-in-f195.1e100.net (216.58.208.195): icmp_req=1 ttl=54 tim
e=5.00 ms
64 bytes from par10s21-in-f3.1e100.net (216.58.208.195): icmp_req=2 ttl=54 time=
4.86 ms
^C
--- www.google.fr ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 4.862/4.933/5.005/0.100 ms
root@vm:/home/toto#
```

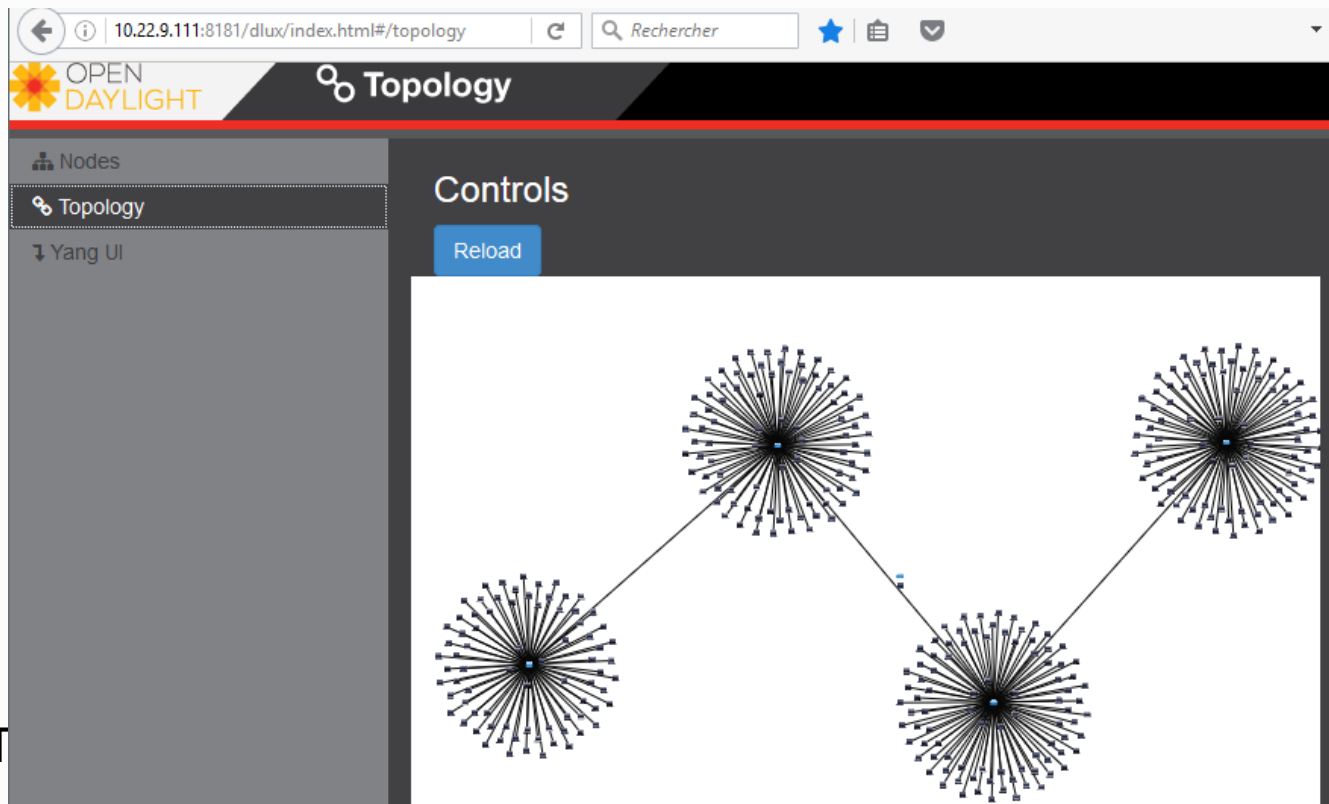
Some examples

- In network course
 - 11 network equipment for each student
 - Some 200 equipment using simultaneously
 - Full automated correction
- In system administration course
 - 4 VM for each student
 - LXC container (for Linux administration)
- In research
 - To build our testbed



Our full virtualized testbed

- To design a new concept for IoT
 - 500 things, 4 OVS, 1 SDN controller
 - Only 128 Go RAM for the host

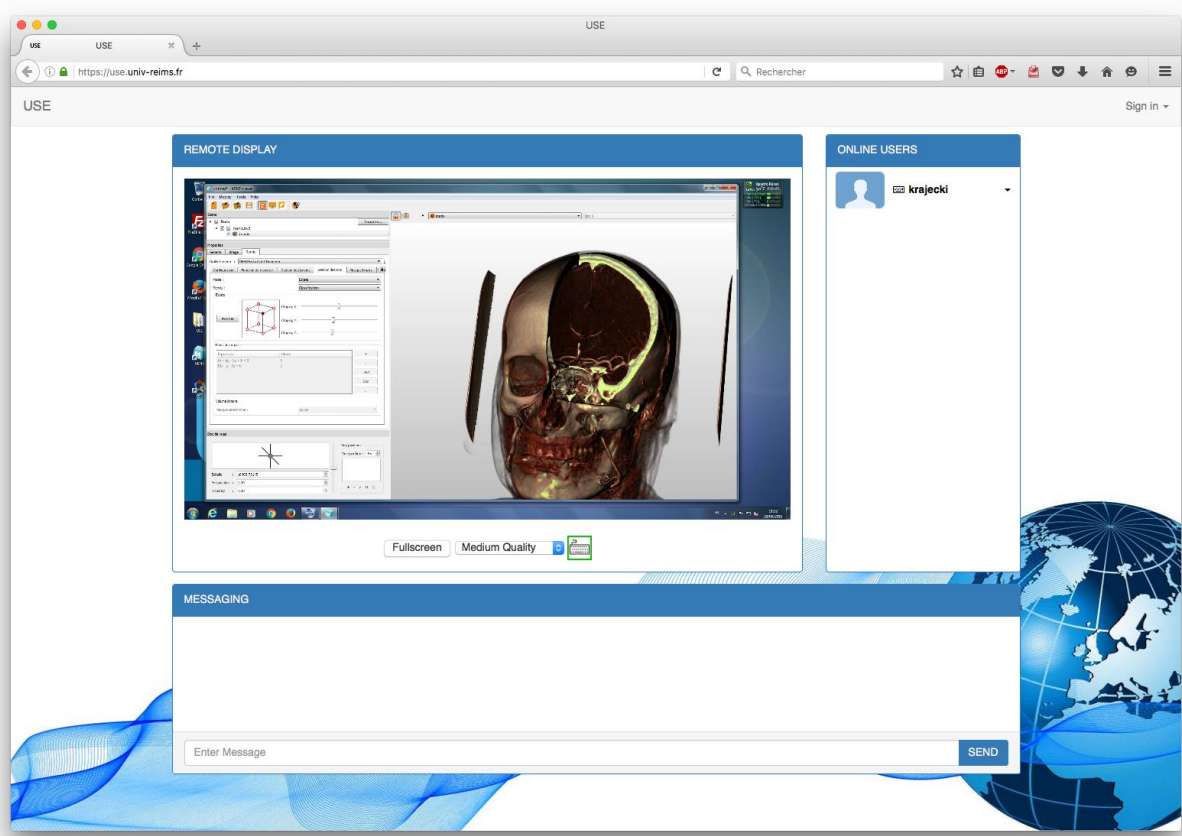


Outline

- Introduction
- Virtualization for training and teaching
- **Virtualization for remote visualization and HPC**
- Future works



USE



Outline

- Introduction
- Virtualization for training and teaching
- Virtualization for remote visualization and HPC
- **Future works**



Future Works

- Move forward in the integration of virtualization in the ROMEO HPC environment
 - Integration in the job scheduler (SLURM)
- Offer virtualization services for all users (not only for remote visualization)
- A key point to be attractive for non specialist users
 - <https://romeo.univ-reims.fr/SIMSEO>

