Provisioning Cloud Services to Academic Users in the Czech Republic

Datacenter IaaS workshop 2014 presentation

Filip Hubík (MUNI), Boris Parák (CESNET), Zdeněk Šustr (CESNET)

contact: hubik@ics.muni.cz

Helsinki, Finland, 11. September 2014

Who Are We (1/2)

MUNI (Masaryk University)

- Institute of Computer Science (ICS)
- Resp. for development of information and communication technologies
- Research conduct (multimedia, supercomputing, . . .)
- Academic connection (communities, support, hiring)
- CERIT Scientific Cloud project
  - One of cloud project participant
  - National center
  - HPC cloud and data storage
  - KVM virtualization

Who Are We (2/2)

CESNET

- Established by academic institutions
- MUNI one of them
- NREN, NGI operator, EGI member
- Open source technology leader (networking, multimedia, . . .)
- MetaCentrum activity
  - Years of experience as grid provider
  - Cloud started alongside grid 3 years ago
  - Primarily Xen virtualization platform
- Close cooperation with MUNI

General cloud computing

Why?

- Users need their own environment (grid)
- Flexibility
- Pay only for the resources you use
- Quick deployment
- No need to own infrastructure

Why not?

- Security issues (third party storage)
- Prone to attacks (IaaS)

Cloud Computing

Joint activity of MetaCentrum (CESNET) and CERIT-SC (MUNI)

Two production clouds

- MetaCloud (MetaCentrum's cloud)
  - Shared resources with CERIT-SC
- FedCloud (EGI federated cloud)

- Infrastructure-as-a-Service (IaaS)
- Both OpenNebula based (C12G labs)
  - Variously integrated into our environment
- Xen and KVM virtualization platform
- More than 3 years of experience
- Support and active coop. with communities

Closer Look at MetaCloud 1/2

Resources

- 396 VCPU, 2 191 GB RAM
- CERIT-SC shared resources (2208 VCPU + 23TB RAM)
- More than 44 TB of shared storage space (GPFS)

Activities

- Research and academic groups in Czech Republic
- Physics (Academy of Sciences, CERN coop.)
- Genetics and biology (Big data, Galaxy, Elixir)
- Advanced networking (Cybernetic polygon project)
- Cloud research groups (Hadoop and NetFlow)
- Education (students, employees)
- Internal testbed and development (SDI4Apps)

Closer Look at MetaCloud 2/2

Architecture

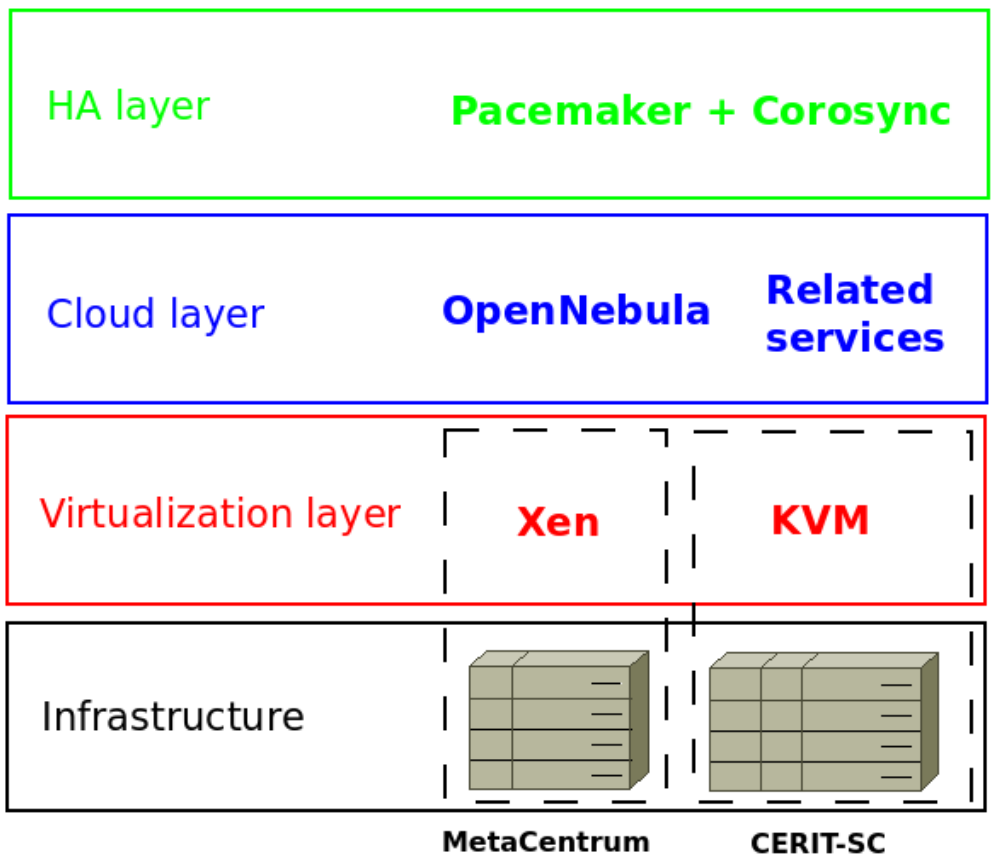

MetaCloud Related Services

- GridFTP upload interface
- Cumulus S3 object storage (REST, fast, metadata)
- High availability (pacemaker + corosync)
  - Critical services virtualized 1:1
  - OpenNebula, MySQL
  - Xen based
  - Not every service can be
- OpenNebula testbed (Puppet on-demand)
- Owncloud testbed (StaaS)

Closer look on CERIT-SC activities

- Infrastructure for academic HPC cloud
- Supercomputing center transformation
- Same OpenNebula front-end
- KVM virtualization platform
- Virtualization of grid nodes 1:1
- Biochemistry, Mathematical models, Physics, Security (KYPO)
- OpenNebula development
  - Cloud-init contextualization connector
  - Red Hat Cluster Suite HA addon

MetaCloud Workflow

1. Connect with us

- cloud@metacentrum.cz or own VO's (EduID)

2. Use OpenNebula (Wiki doc)

- Web – OpenNebula Sunstone
- CLI – One tools or rOCCI

2. Upload own or use our image

- GridFTP
- Web upload (browser limit)
- OpenNebula one tools

3. Configure, deploy and use

- Xen or KVM hosting

MetaCloud Integration and Development

- Perun – our own centralized identity and user management (national e-infrastructure)
- rOCCI – open cloud computing interface (Ruby)
- Kerberos web auth – security
- Nagios probes – monitoring
- OpenNebula – upstream contributions
- Clouditor – accounting
- Cloud-init
- Own cloud images (Puppet, Vagrant, . . .)

Closer Look at FedCloud

Infrastructure

- Part of European Grid Infrastructure (EGI) project
- 240 CPU + 960 GB RAM
- More than 44 TB of shared storage space

Active EGI groups

- Biological data processing (WeNMR)
- Sound and music analysis (PeachNote)
- Astrometry and physics (GAIA, CERN, DIRAC, LHC)
- Development
  - Science portals and GW's – WS-PGRADE, CSGF
  - Databases – BNCWeb (English dictionary)

Another Storage Services

- CESNET's storage department
  - Hierarchical, geo-distributed
  - Backups, archiving
  - NFSv4, FTP, SCP, SSH, rsync
  - GlobusOnline, FileSender
  - Owncloud StaaS
    - Production stage
    - 100GB disk space
    - Share and sync across various devices
    - Browser limit
- MUNI's storage experiments
  - Cloud filesystem benchmark (Ceph, GlusterFS, ...)
  - KVM vs. Xen testing

Future Work

- Support of PaaS experiments (Hadoop)
- Infrastructure growth
- Network virtualization (OpenFlow, Open vSwitch)
- Experiments (SR-IOV, HA)
- rOCCI development
  - Additional backends (EC2 Azure, VMWare, . . .)
  - Emerging OCCI standard specifications (PaaS, XML and JSON rendering, Billing and Monitoring . . .)
- New communities, challenges and opportunities

Thank you for your attention.

References:

CESNET – offical website, http://www.cesnet.cz
MUNI – offical website, http://www.muni.cz
Metacentrum – offical website, http://www.metacentrum.cz/en

Filip Hubík (MUNI), Boris Parák (CESNET), Zdeněk Šustr (CESNET)
contact: hubik@ics.muni.cz