Accelerating Understanding Summit 2016

Using a Hybrid Cloud approach to achieve efficiency, flexibility & high utilisation of your HPC

Antonio Cisternino, Professor of Computer Science - University of Pisa
Maurizio Davini, CTO, IT Center - University of Pisa

Using a Hybrid Cloud approach to achieve efficiency, flexibility and high utilisation of your HPC

Antonio Cisternino, Maurizio Davini

Università di Pisa

What is HPC today?

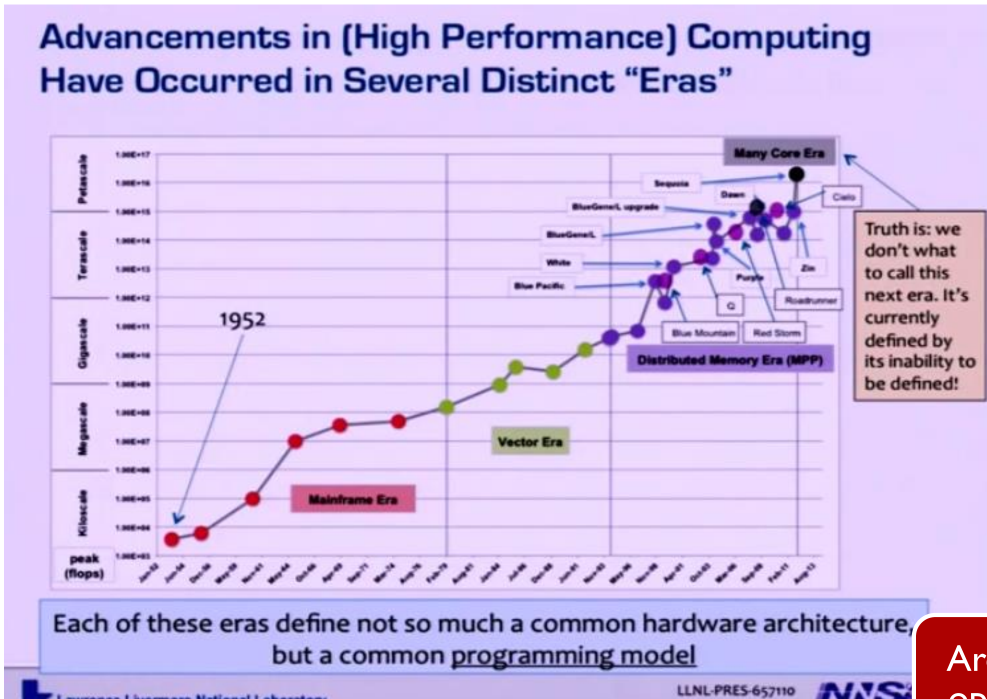

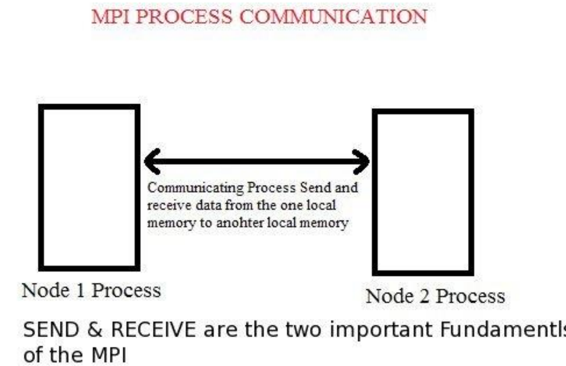

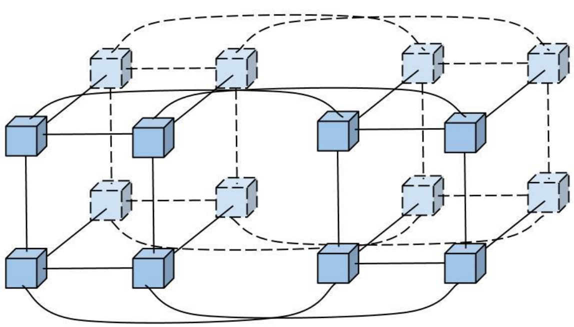

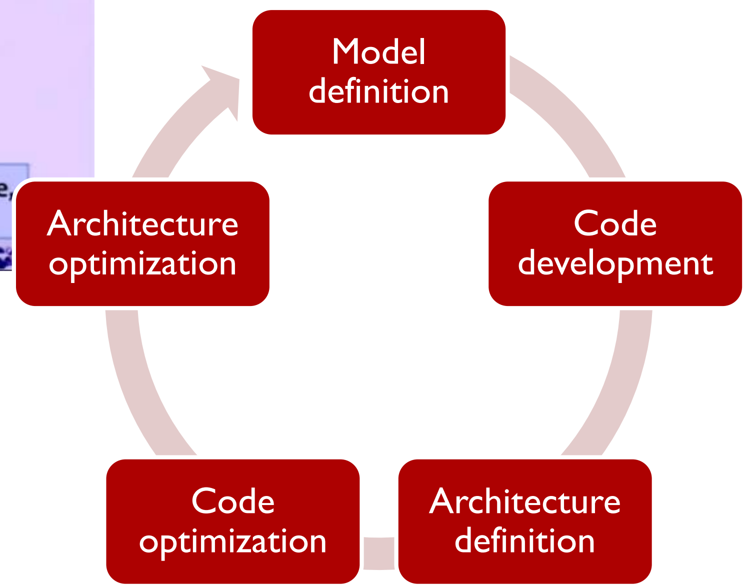

HPC frequently hidden assumptions

- Linear/Polynomial complexity of algorithms
- Numerical codes
- Compute–Communicate ratio
- Problem bigger than a single system
- Extreme custom code optimization

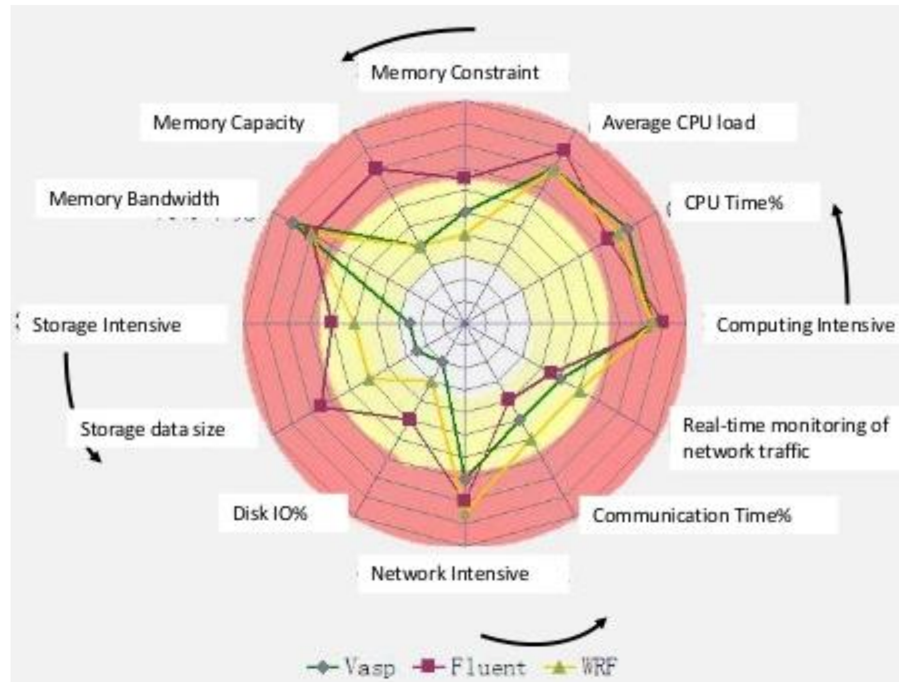

Cloud, HPC and the problem size

HPC

- Problem requires multiple nodes
- Homogeneous structure (jobs and schedulers…)
- Mainly in memory
- Frequent sync (order of latency in communications)

Cloud

- Single node addressing multiple problems
- Heterogeneous structure (just x86…)
- Any memory and disk access patterns
- Communication may happen even through L2 incapsulated in L3

Big data

- Problemrequires multiple nodes
- Not-so-homogeneous structure
- Mainly disk bound
- Communication is rare

Flexibility IS often more important…

Python 3 used what fraction? used how many times more?							
Time-used	|-	|---	25%	median	75%	---|	-|
(Elapsed secs)	±	±	6×	43×	62×	106×	106×

Java used what fraction? used how many times more?							
Time-used	|-	|---	25%	median	75%	---|	-|
(Elapsed secs)	±	±	2×	2×	2×	4×	4×

Critical aspects for achieving flexible HPC

Software stack

- Library dependencies
- Package management (a nightmare): yum, apt–get, npm, nuget
- Debug/run cycle, staging
- Code and data versioning
- Open source development model is worsening this scenario
- Security hard to ensure

Code (in)–efficiency

- Experiment:
  - Take a 4CPU server S with RAM
  - Run CFD workload W on S with MPI multithread
  - Run W on S with MPI multiprocess
  - Run W on a virtual cluster on S with same resources and MPI on vSwitch
- Outcome (2010 and 2015):
  - Virtual execution lead to a speedup up to 30%

Algorithm complexity

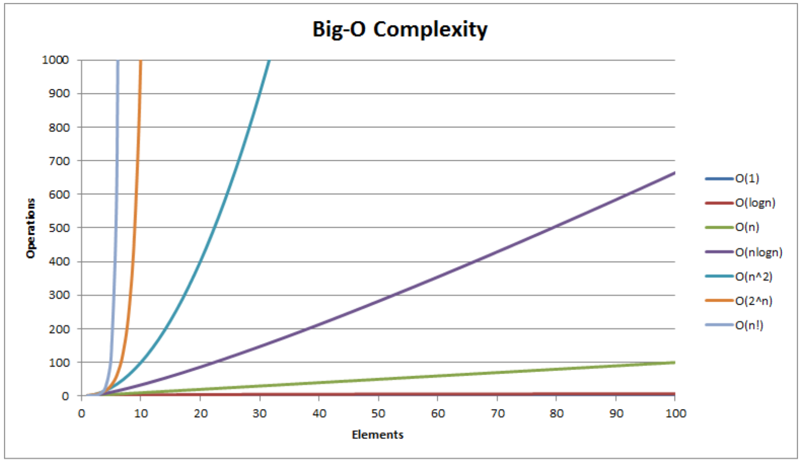

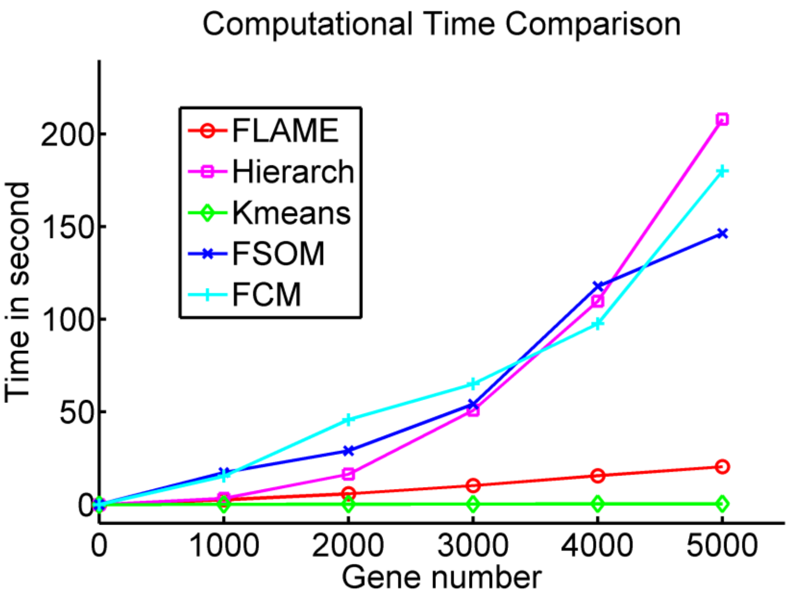

Networking

- East–West traffic
- Low latency
- Topology
- IP Addressing (overlay?)
- Containers
- Data size (i.e. 300GB for a genome, remember?)

Disk size AND performance

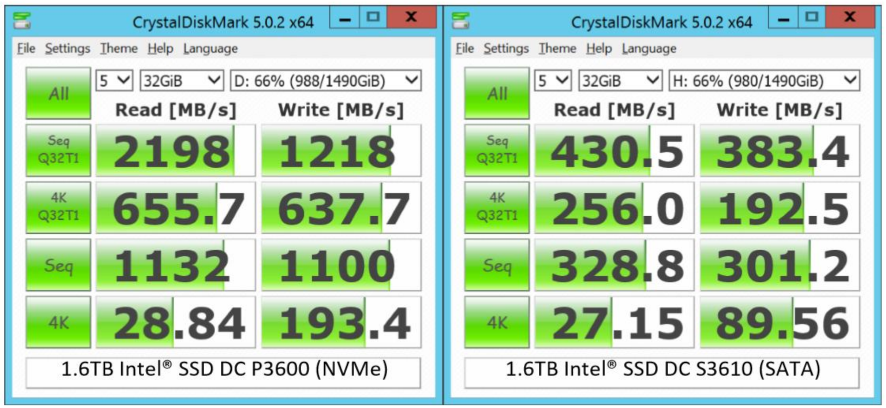

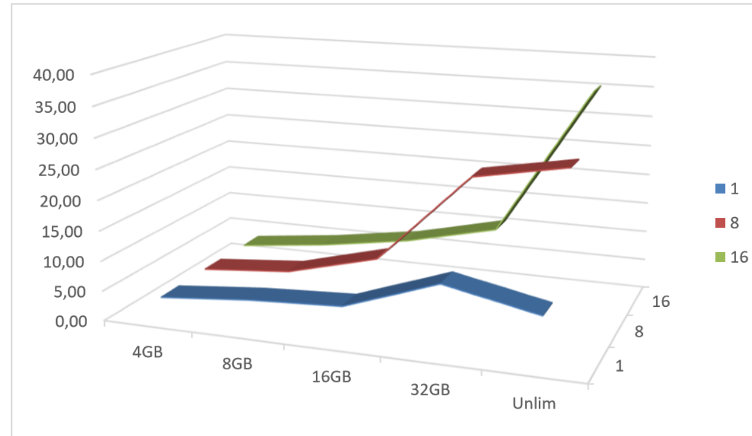

Avg Queue length ratio	Virtual users		
Memory	1	8	16
4GB	3,36	3,47	3,38
8GB	4,71	4,74	5,08
16GB	5,49	8,63	7,40
32GB	10,89	23,70	10,91
Unlim	7,72	26,44	35,55

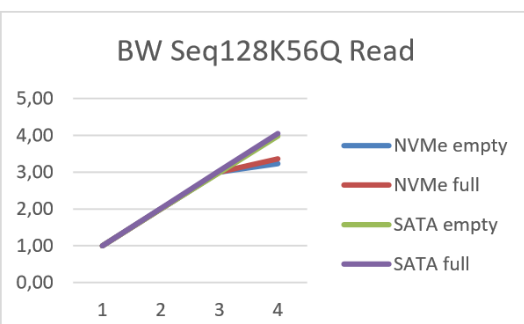

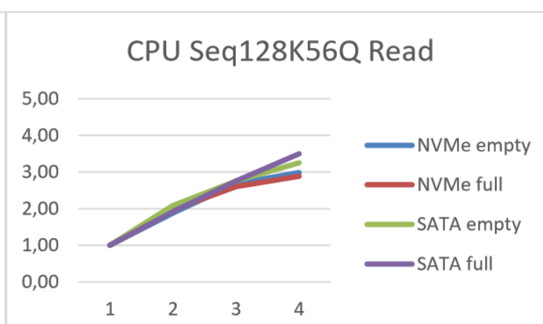

Problem size vs. system size

«Problems usually don't scale up as hardware, and so you start as enterprise to become a small business. Only then you may realize that the beautiful parallel something you devised become useless but it complicates the architecture»

A.C.

(hey, it's me :)

HPC and Cloud

Public Cloud HPC (Azure)

True HPC capabilities in the cloud, on demand

The performance and scalability of a world-class supercomputing center is now available to everyone, on demand in the cloud. Run your Windows and Linux HPC applications using high performance A8 and A9 compute instances on Azure, and take advantage of a backend network with MPI latency under 3 microseconds and non-blocking 32 Gbps throughput. This backend network includes remote direct memory access (RDMA) technology on Windows and Linux that enables parallel applications to scale to thousands of cores. Azure provides you with high memory and HPC-class CPUs to help you get results fast. Scale up and down based upon what you need and pay only for what you use to reduce costs.

Learn more about high performance A8 and A9 compute instances ▸

Read about the performance improvements ANEO achieved with Azure ▸

Public Cloud HPC (AWS)

Try out the demo framework, cfncluster, to see how to run HPC clusters in AWS

High performance computing in the cloud.

High Performance Computing (HPC) on Amazon Elastic Compute Cloud (EC2) is enabled by the Cluster Compute-optimized and GPU instance types (virtual machines). You can use them just like other EC2 instances, but they also have been specifically engineered to provide high performance networking and you can scale to tens of thousands of instances on-demand.

Get an overview of Amazon EC2 »

Announcing the Amazon EC2 C4 instance type.

You can now launch C4 instances, the latest generation of Amazon EC2 Compute-optimized instances. C4 instances are designed for compute-bound workloads, such as high-traffic front-end fleets, MMO gaming, media processing, transcoding, and High Performance Computing (HPC) applications.

C4 instances are available in five sizes, offering up to 36 vCPUs. C4 instances are based on Intel Xeon E5-2666 v3 (codename Haswell) processors that run at a base frequency of 2.9 GHz, and can deliver clock speeds as high as 3.5 GHz with Intel ® Turbo Boost. Each C4 instance type is EBS-optimized by default and at no additional cost. This feature provides 500 Mbps to 4,000 Mbps of dedicated throughput to EBS above and beyond the general purpose network throughput provided to the instance.

Learn more about Amazon EC2 instance types »

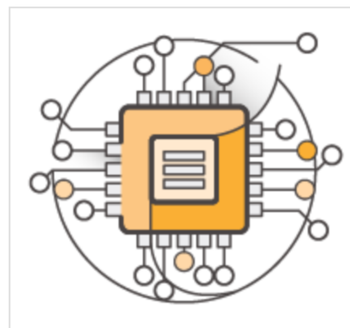

Learn more about EC2 Cluster Compute instance types

Private Cloud (DHCS on prem)

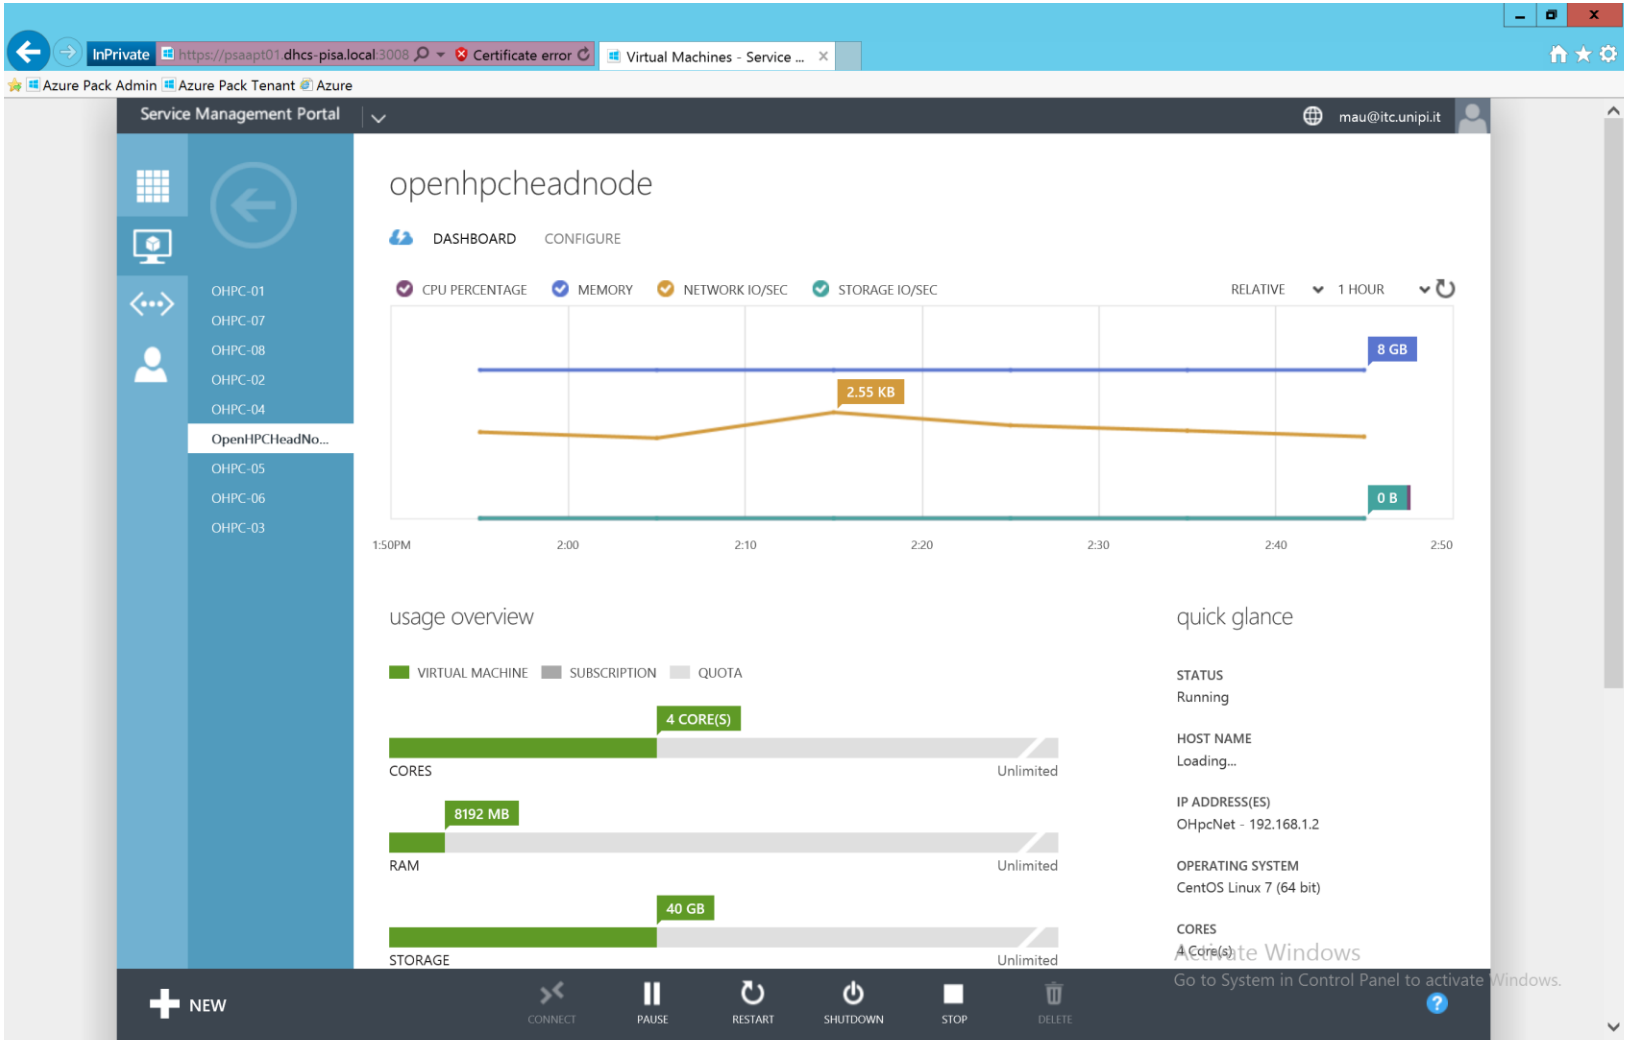

OpenStack?

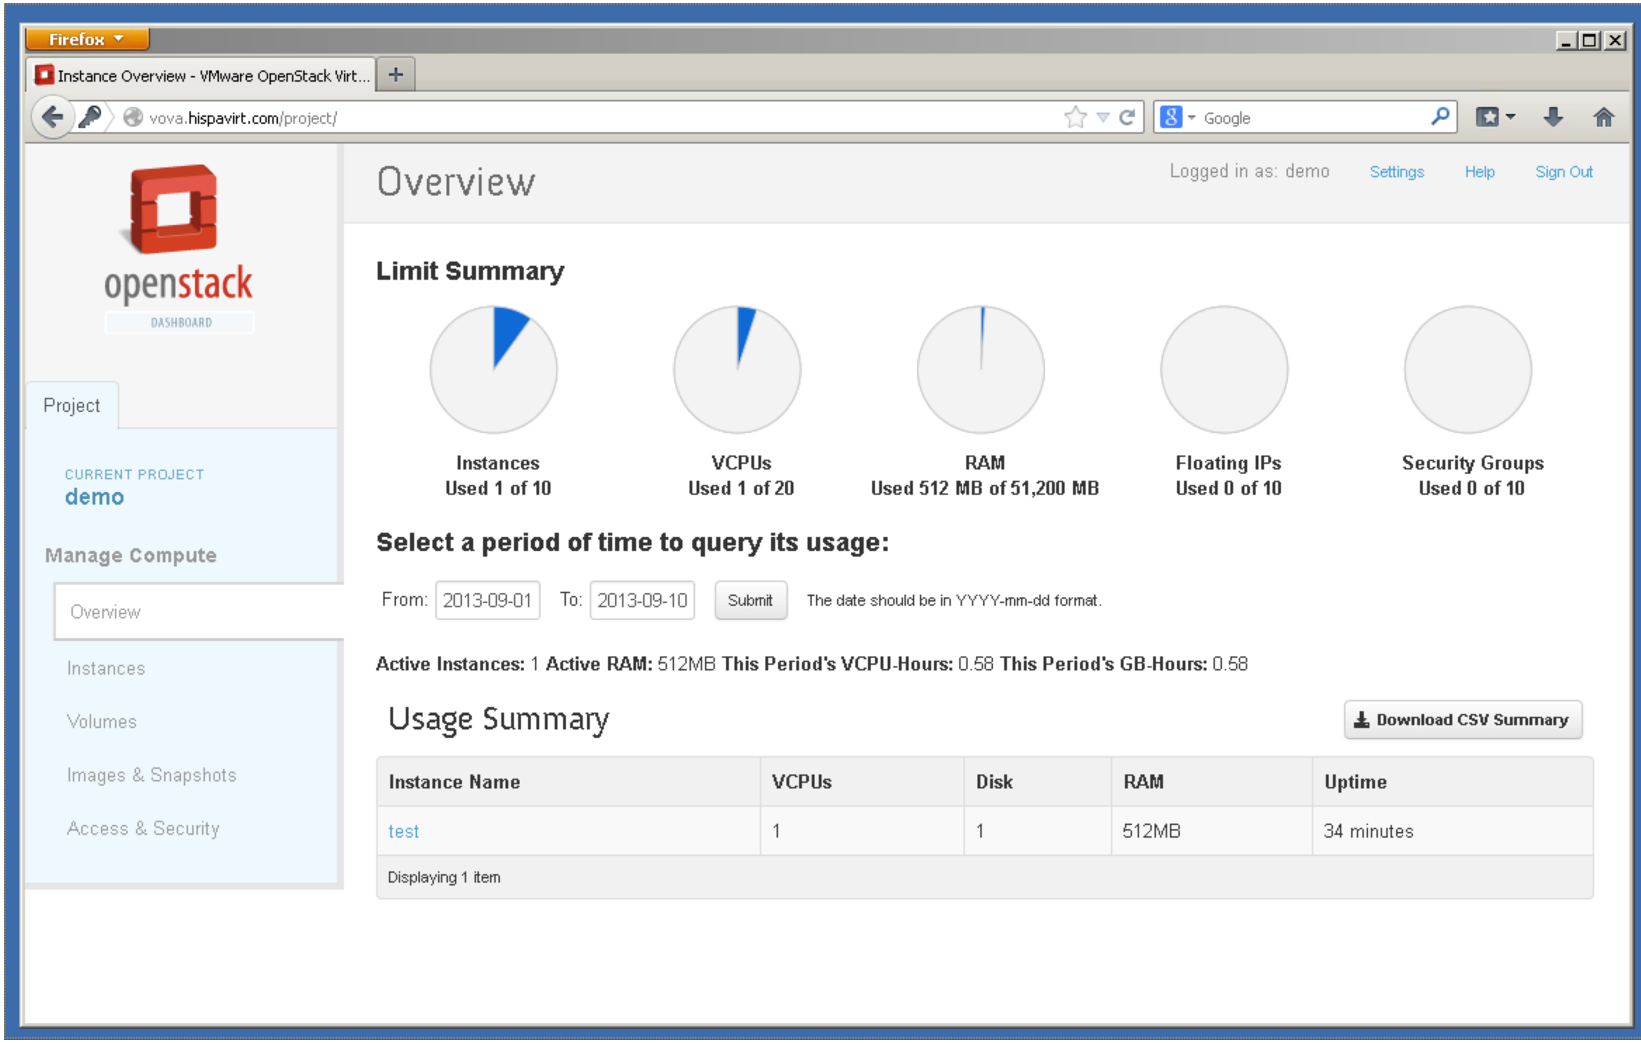
Firefox
Instance Overview - VMware OpenStack Virt...
vova.hispavirt.com/project/
Google
openstack
DASHBOARD
Overview
Logged in as: demo
Settings
Help
Sign Out
Project
CURRENT PROJECT
demo
Manage Compute
Overview
Instances
Volumes
Images & Snapshots
Access & Security
Limit Summary
Instances
Used 1 of 10
VCPUs
Used 1 of 20
RAM
Used 512 MB of 51,200 MB
Floating IPs
Used 0 of 10
Security Groups
Used 0 of 10
Select a period of time to query its usage:
From: 2013-09-01 To: 2013-09-10 Submit The date should be in YYYY-mm-dd format.
Active Instances: 1 Active RAM: 512MB This Period's VCPU-Hours: 0.58 This Period's GB-Hours: 0.58
Usage Summary
Download CSV Summary
Instance Name
VCPUs
Disk
RAM
Uptime
test
1
1
512MB
34 minutes
Displaying 1 item

Containers (Docker)

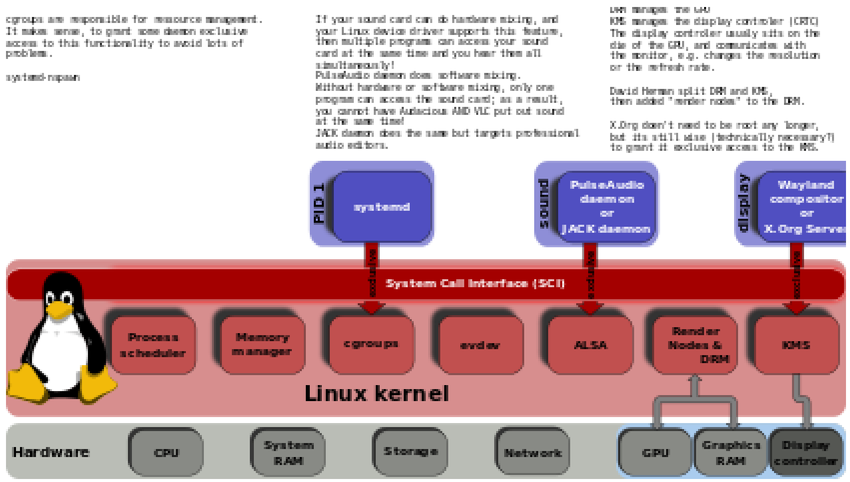

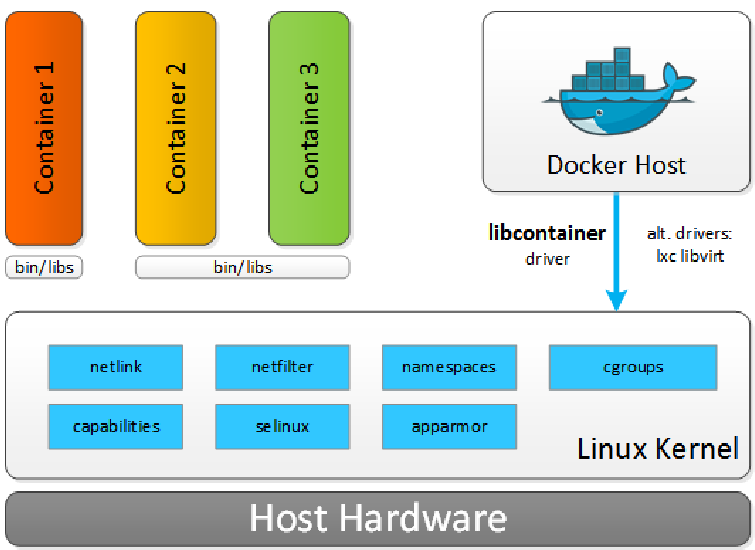

- Use of container for deploying different (potentially incompatible) software
- Containers come with less overhead but networking can be a nightmare (latency)
- It's the dawn of containers stay tuned!

Designing for specific workloads (lot of fun!)

- Design for the workload (DN, BD, traditional HPC…)
- It's not just a matter of CPU and interconnectivity
- Dimension appropriately the resource usage (bus, memory, disk, network)

New horizons

IMAGINATION
INTUITION
VISION
CENTER
SPEED OF EXECUTION

Emerging HUGE problems

- Life sciences
- Big data analytics
- Real time big data analytics
- IoT

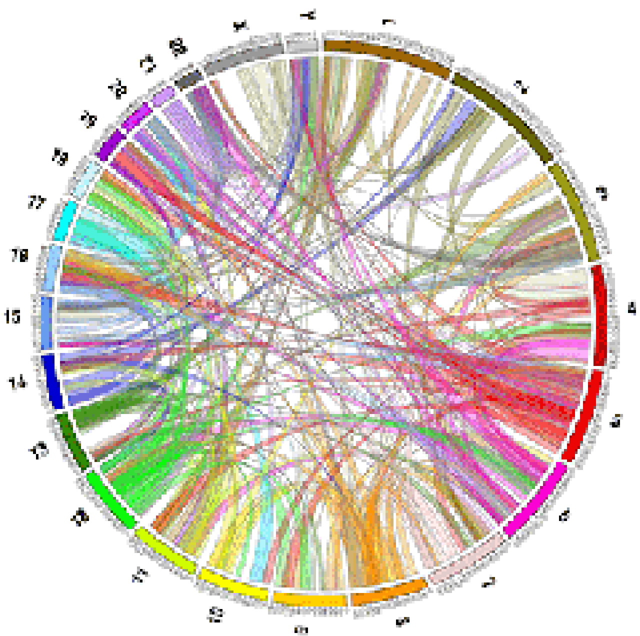

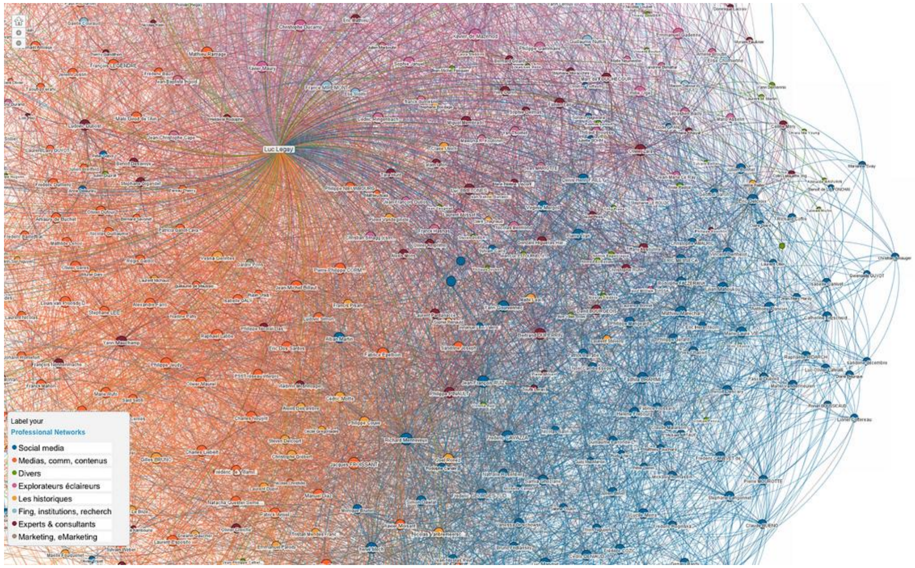

Continuous problem solving

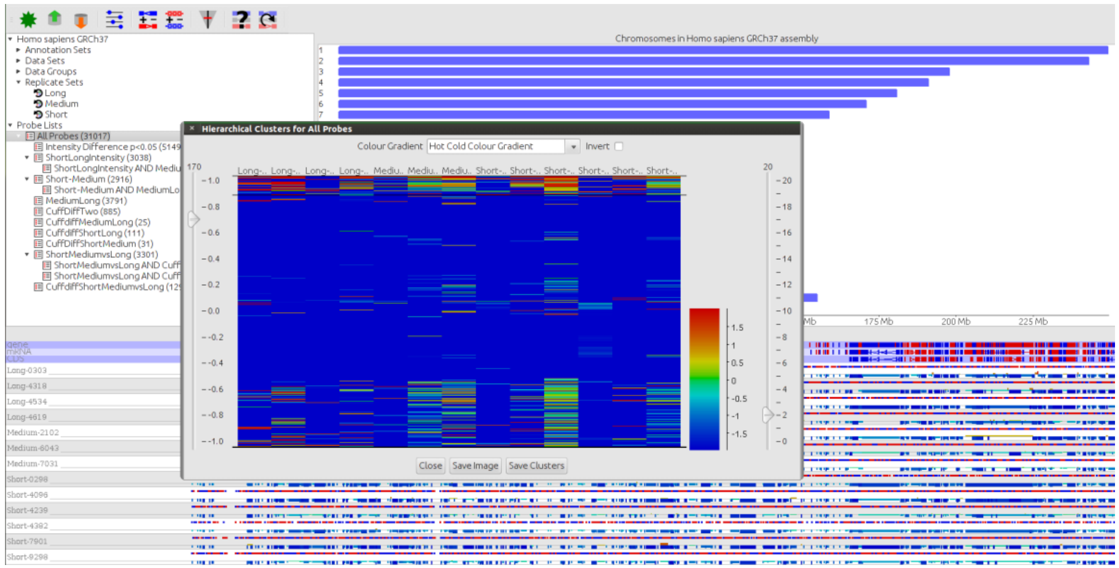

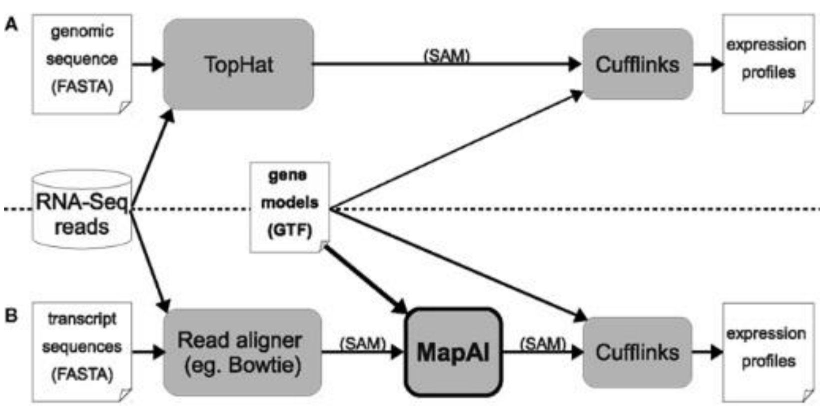

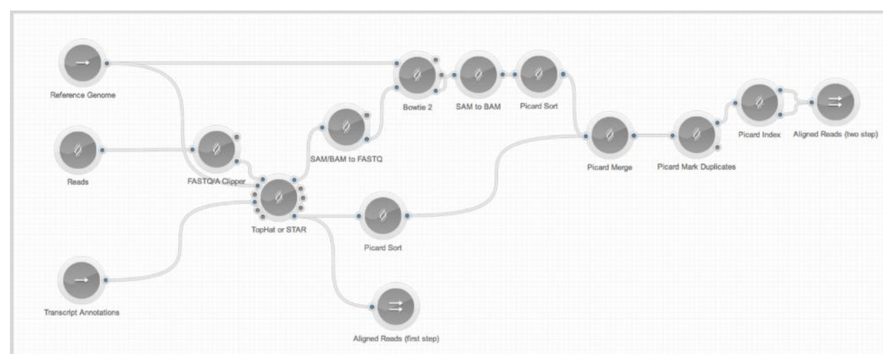

Democratizing HPC

- Data centers in a box
- Predefined, precooked solutions easily deployable
- Some scenarios (Fog, IoT) may benefit from a more distributed information processing

Conclusions

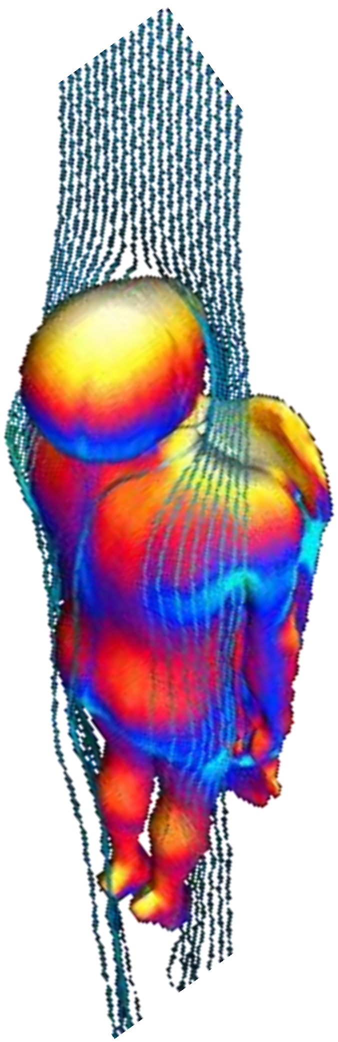

- HPC is redefining itself (again)
- Cloud (public or private or hybrid) + containers offer viable solutions for HPC–kind of problems
- Overlap with non–HPC architectures (i.e. microservices, distributed IoT)
- Flexibility will be an important pillar of the HPC in the future
- We accepted that assembly was low level, it's time to move on

OpenHPC @ University of Pisa Dell Solution Center

Antonio Cisternino
Maurizio Davini
IT Center University of Pisa

UNIVERSITÀ DI PISA

OpenHPC

- **Create a stable environment for testing and validation:** The community will benefit from a shared, continuous integration environment, which will feature a build environment and source control; bug tracking; user and developer forums; collaboration tools; and a validation environment.
- **Reduce Costs:** By providing an open source framework for HPC environments, the overall expense of implementing and operating HPC installations will be reduced.
- **Provide a robust and diverse open source software stack:** OpenHPC members will work together on the stability of the software stack, allowing for ongoing testing and validation across a diverse range of use cases.
- **Develop a flexible framework for configuration:** The OpenHPC stack will provide a group of stable and compatible software components that are continually tested for optimal performance. Developers and end users will be able to use any or all of these components depending on their performance needs, and may substitute their own preferred components to fit their own use cases.

Why OpenHPC

- Many sites spend considerable effort aggregating a large suite of open-source projects to provide a capable HPC environment for their users:
  - necessary to obtain HPC focused packages that are either absent or do not keep pace from Linux distro providers
  - local packaging or customization frequently tries to give software versioning access to users (e.g. via modules or similar equivalent)
- They frequently leverage a mix of external and in-house tools for:
  - provisioning, software installations/upgrades, config management schemes, and system diagnostics mechanisms.
  - although the functionality is similar, the implementations across sites is often different which can lead to duplication of effort

OpenHPC: Early Community Members

OEMs

HPC Sites

ISV-OSVs

Component Overview

SW Block Diagram of Typical HPC System

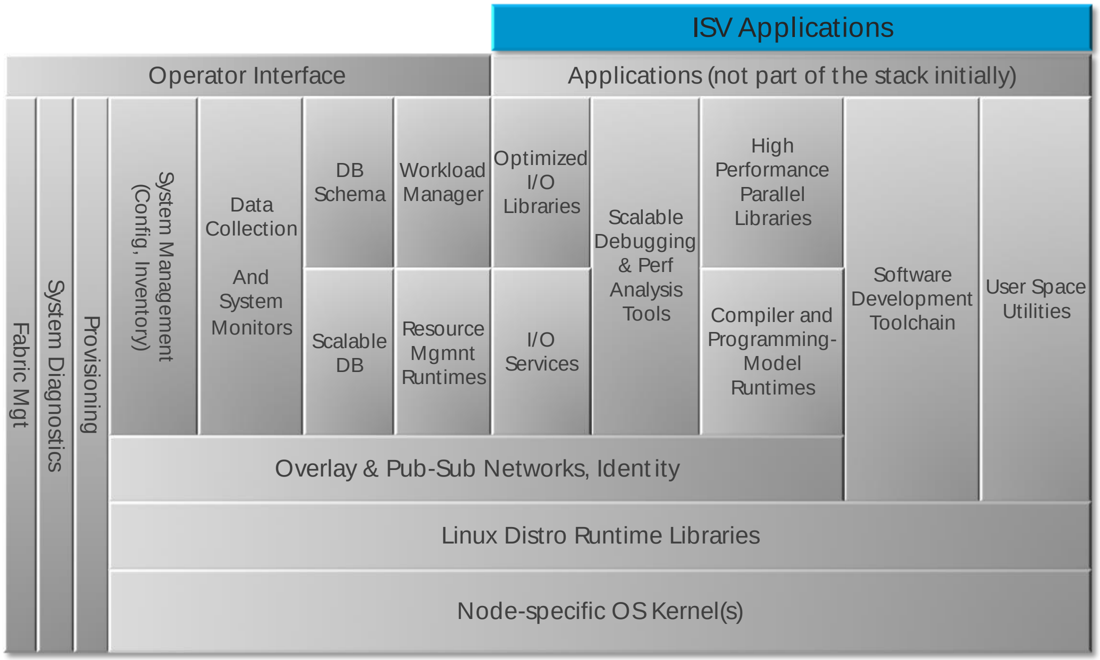

System Architecture

Intention over time is to provide reference design(s) with use cases for big and small system designs:

- usual designation of node types by function
- current recipe example is “diskless”
- intend to include diskfull examples as well

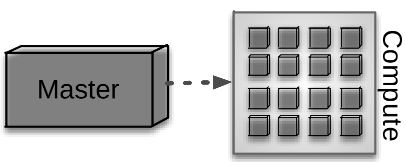

[Flat]

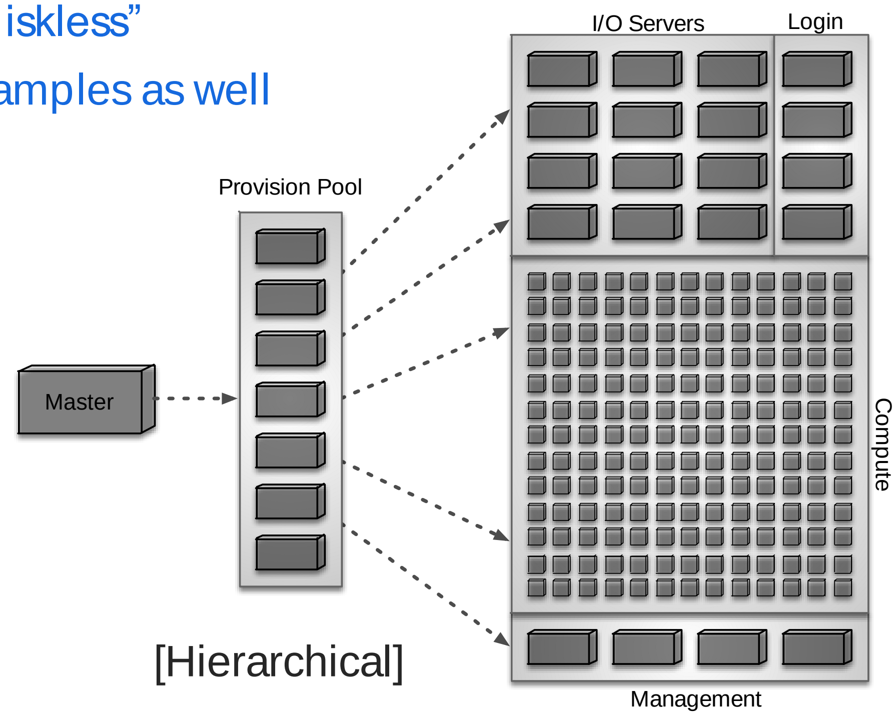

[Hierarchical]

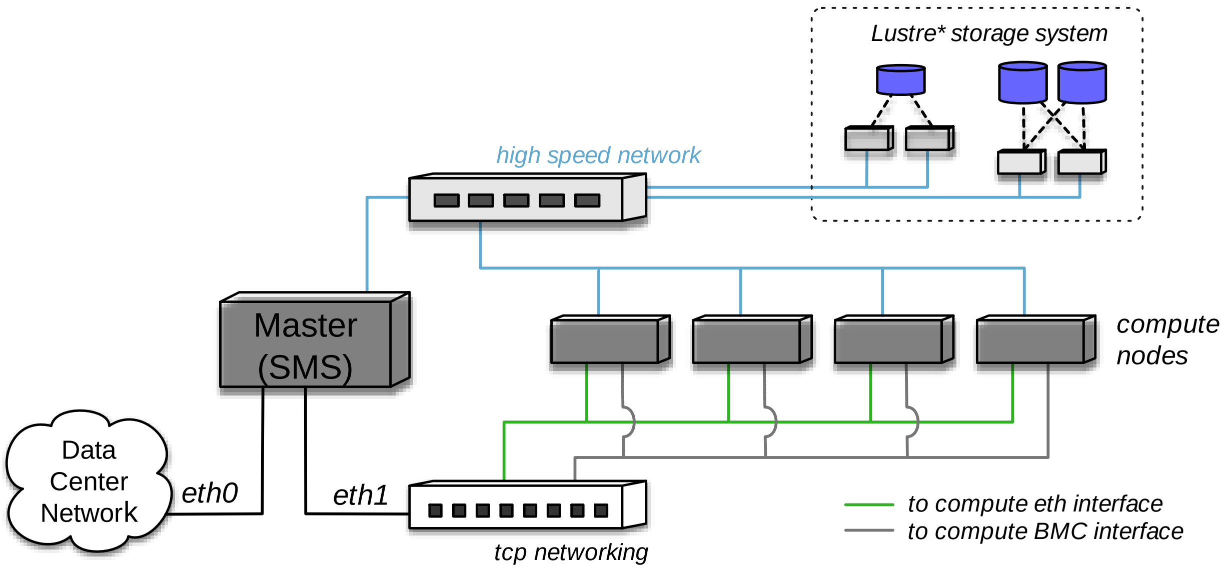

Figure 1: Overview of physical cluster architecture.

OpenHPC 1.0 - Initial starting components

Functional Areas	Components
Base OS	CentOS 7.1
Administrative Tools	Conman, Ganglia, Intel Cluster Checker**, Lmod, LosF, Nagios, pdsh, prun, EasyBuild
Provisioning	Warewulf
Resource Mgmt.	SLURM, Munge
I/O Services	Lustre client (community version)
Numerical/ Scientific Libraries	Boost, GSL, FFTW, Metis, PETSc, Trilinos, Hypre, SuperLU, Mumps, Intel MKL**
I/O Libraries	HDF5 (pHDF5), NetCDF (including C++ and Fortran interfaces), Adios
Compiler Families	GNU (gcc, g++, gfortran), Intel Parallel Studio XE (icc,icpc,ifort)**
MPI Families	MVAPICH2, OpenMPI, Intel MPI**
Development Tools	Autotools (autoconf, automake, libtool), Valgrind,R, SciPy/NumPy, Intel Inspector **
Performance Tools	PAPI, IMB, mpiP, pdtoolkit TAU, Intel Advisor**, Trace Analyzer and Collector**, Vtune Amplifier**

Notes:
- Additional dependencies that are not provided by the BaseOS or community repos (e.g. EPEL) are also included
- 3rd Party libraries are built for each compiler/MPI family (6 combinations typically)
- Resulting repository currently comprised of ~250 RPMs

** Bring your own license model

OpenHPC++ - Potential future efforts

Functional Areas	Components	Contributions by:
Base OS	CentOS 7.1, McKernel, Kitten, mOS	RIKEN, Sandia, Intel
AdminTools	Conman, Ganglia, Intel Cluster Checker**, Lmod, LosF, Nagios, pdsh, prun, EasyBuild, ORCM	Intel
Provisioning	Warewulf, xCAT	Community
Resource Mgmt.	SLURM, Munge, ParaStation mgmt, PMIx, PBS Pro	ParTec, community, Altair
Cross Cutting	OpenStack HPC suitable components	Cray
Runtimes	OpenMP, OmpSs, OCR	BSC, Intel
I/O Services	Lustre client (community version)	
Numerical/ Scientific Libs	Boost, GSL, FFTW, Metis, PETSc, Trilinos, Hypre, SuperLU, Mumps, Intel MKL**	
I/O Libraries	HDF5 (pHDF5), NetCDF (including C++ and Fortran interfaces), Adios	
Compiler Families	GNU (gcc, g++, gfortran), Intel Parallel Studio XE (icc,icpc,ifort)**	
MPI Families	MVAPICH2, OpenMPI, Intel MPI**, MPICH, ParaStation MPI	Argonne, ParTec
Development Tools	Autotools (autoconf, automake, libtool), Valgrind,R, SciPy/NumPy, Intel Inspector **	
Performance Tools	PAPI, Intel IMB, mpiP, pdtoolkit TAU, Intel Advisor**, Intel Trace Analyzer and Collector**, Intel Vtune Amplifier**, Paraver, Scalasca	BSC, Jülich

** Bring your own license model

Installation brief

OpenHPC 1.1 (18 April 2016)

Binary downloads are presently available in the form of RPMs. These RPMs are organized into repositories that can be accessed via standard package manager utilities (e.g. yum, zypper). OpenHPC provides builds that are compatible and tested against CentOS 7.2 as well as SUSE Linux Enterprise Server 12 SP1. A typical deployment on a new system will begin with the installation of the base operating system on a chosen *master* host identified as the system management server (SMS), followed by enabling access to a compatible OpenHPC repository.

The OpenHPC repository is created and maintained using a dedicated instance of the Open Build Service (OBS) that is available here. In addition to serving as the build server, this OBS instance also provides an RPM repository. You can scan the RPM packages that are available via this repository by browsing the `x86_64/` and `noarch/` subdirectories for the 1.1 release at: http://build.openhpc.community/OpenHPC:/1.1.

To get started, you can enable an OpenHPC repository locally through installation of an `ohpc-release` RPM which includes gpg keys for package signing and defines the URL locations for [base] and [update] package repositories. A copy of the ohpc-release file is available for download here:

- ohpc-release-centos7.2-1.1-1.x86_64.rpm (md5sum:f9349b2c2b117a4e3efdac8cd59cc327)
- ohpc-release-sles12sp1-1.1-1.x86_64.rpm (md5sum:a99904b08c90548faaedf7201d60e101)

Alternatively, you can use the package manager to install the RPM directly from the network as in the following examples:

```
# yum install https://github.com/openhpc/ohpc/releases/download/v1.1.GA/ohpc-release-centos7.2-1.1-1.x86_64.rpm
```

or

```
# zypper in https://github.com/openhpc/ohpc/releases/download/v1.1.GA/ohpc-release-sles12sp1-1.1-1.x86_64.rpm
```

Install Recipe(s)

To aid in the installation of OpenHPC packaged components, a companion installation recipe is available. This can be obtained via installation of the `docs-ohpc` RPM after the OpenHPC repository has been enabled locally. Alternatively, copies of the documentation are also provided below:

- CentOS 7.2 Install guide (PDF)
- SLE 12 SP1 Install guide (PDF)

The intent of the guide is to present a simple cluster installation procedure using components from the OpenHPC software stack. The documentation is intended to be reasonably generic, but uses the underlying motivation of a small, stateless cluster installation to define a step-by-step process. Several optional customizations are included and the intent is that these collective instructions can be modified as needed for local site use cases. Please consult the install guide for more detail and discussion regarding a companion template install script.

OpenHPC Build Service

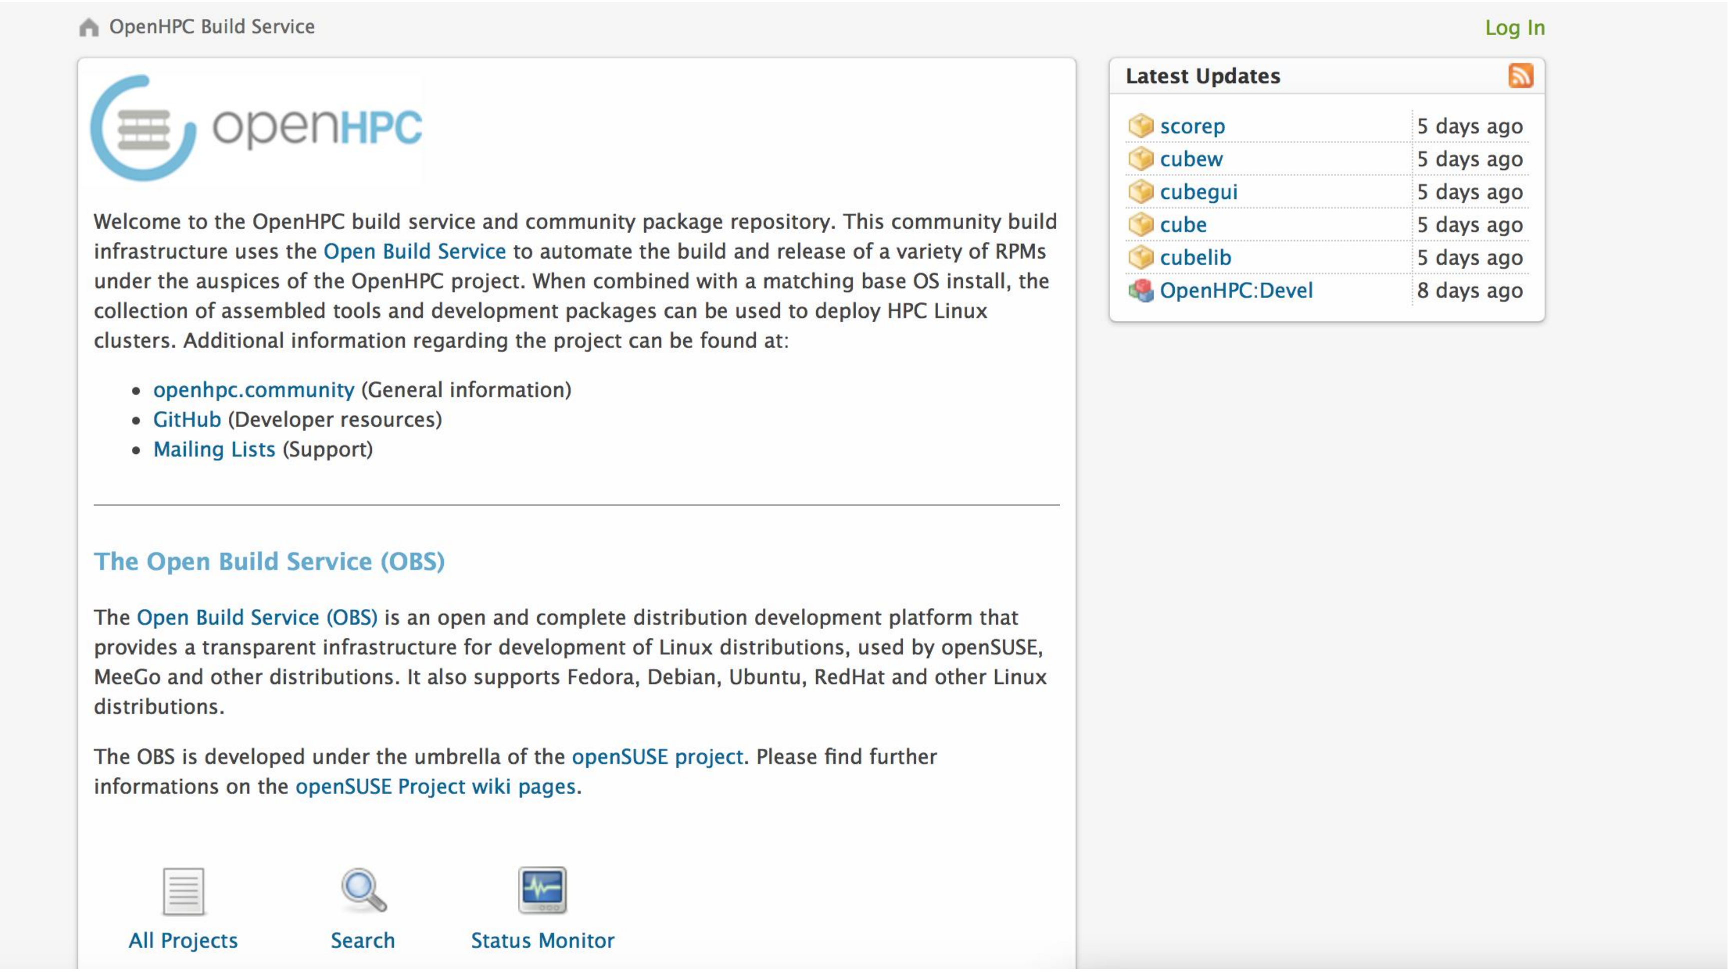

OpenHPC Build Service

Log In

openHPC

Welcome to the OpenHPC build service and community package repository. This community build infrastructure uses the Open Build Service to automate the build and release of a variety of RPMs under the auspices of the OpenHPC project. When combined with a matching base OS install, the collection of assembled tools and development packages can be used to deploy HPC Linux clusters. Additional information regarding the project can be found at:

- openhpc.community (General information)
- GitHub (Developer resources)
- Mailing Lists (Support)

The Open Build Service (OBS)

The Open Build Service (OBS) is an open and complete distribution development platform that provides a transparent infrastructure for development of Linux distributions, used by openSUSE, MeeGo and other distributions. It also supports Fedora, Debian, Ubuntu, RedHat and other Linux distributions.

The OBS is developed under the umbrella of the openSUSE project. Please find further informations on the openSUSE Project wiki pages.

All Projects | Search | Status Monitor

Latest Updates

scorep	5 days ago
cubew	5 days ago
cubegui	5 days ago
cube	5 days ago
cubelib	5 days ago
OpenHPC:Devel	8 days ago

OpenHPC on Microsoft Cloud OS

Solution Brief

Dell Hybrid Cloud System for Microsoft

IT organizations have a journey to take to get to the cloud. Success with cloud requires an understanding that it is not just about the technology stack. Fundamentally, cloud is about applying technology to streamline IT operations and enable new business outcomes.

Business needs must be aligned with technology, and with cloud this relationship is more critical than ever. Dell understands these alignment challenges, having worked with our customers to build and manage some of the largest clouds on the planet.

Introducing Dell Hybrid Cloud System for Microsoft

Dell Hybrid Cloud System for Microsoft is not simply data center technology. The Dell Hybrid Cloud System for Microsoft is the industry's first integrated, cloud solution validated with Microsoft Hybrid Cloud Platform System Standard[1].

Dell Hybrid Cloud System for Microsoft is a turn-key integrated system, with powerful multi-cloud management, that uniquely adapts to and

Business needs and Cloud aligned

The Dell Hybrid Cloud System for Microsoft will make your IT services agile and more efficient. The system is designed to leverage Azure public cloud capabilities for backup and recovery to extend your data center capacity. The solution is also designed to enhance business flexibility with Dell Scale Ready Payment Solutions.

Scale Ready Payment Solutions include Cloud Flex Pay, Pay as you Grow, Provision and Pay and Scale On Demand offers, backed by Dell Financial Services. The Dell Hybrid Cloud System for Microsoft is offered with these highly flexible payment solutions so you can maximize your investments and align the acquisition of cloud technologies with your business needs.

Dell is the only tier one vendor with financing options for an end solution spanning hardware, software and services.

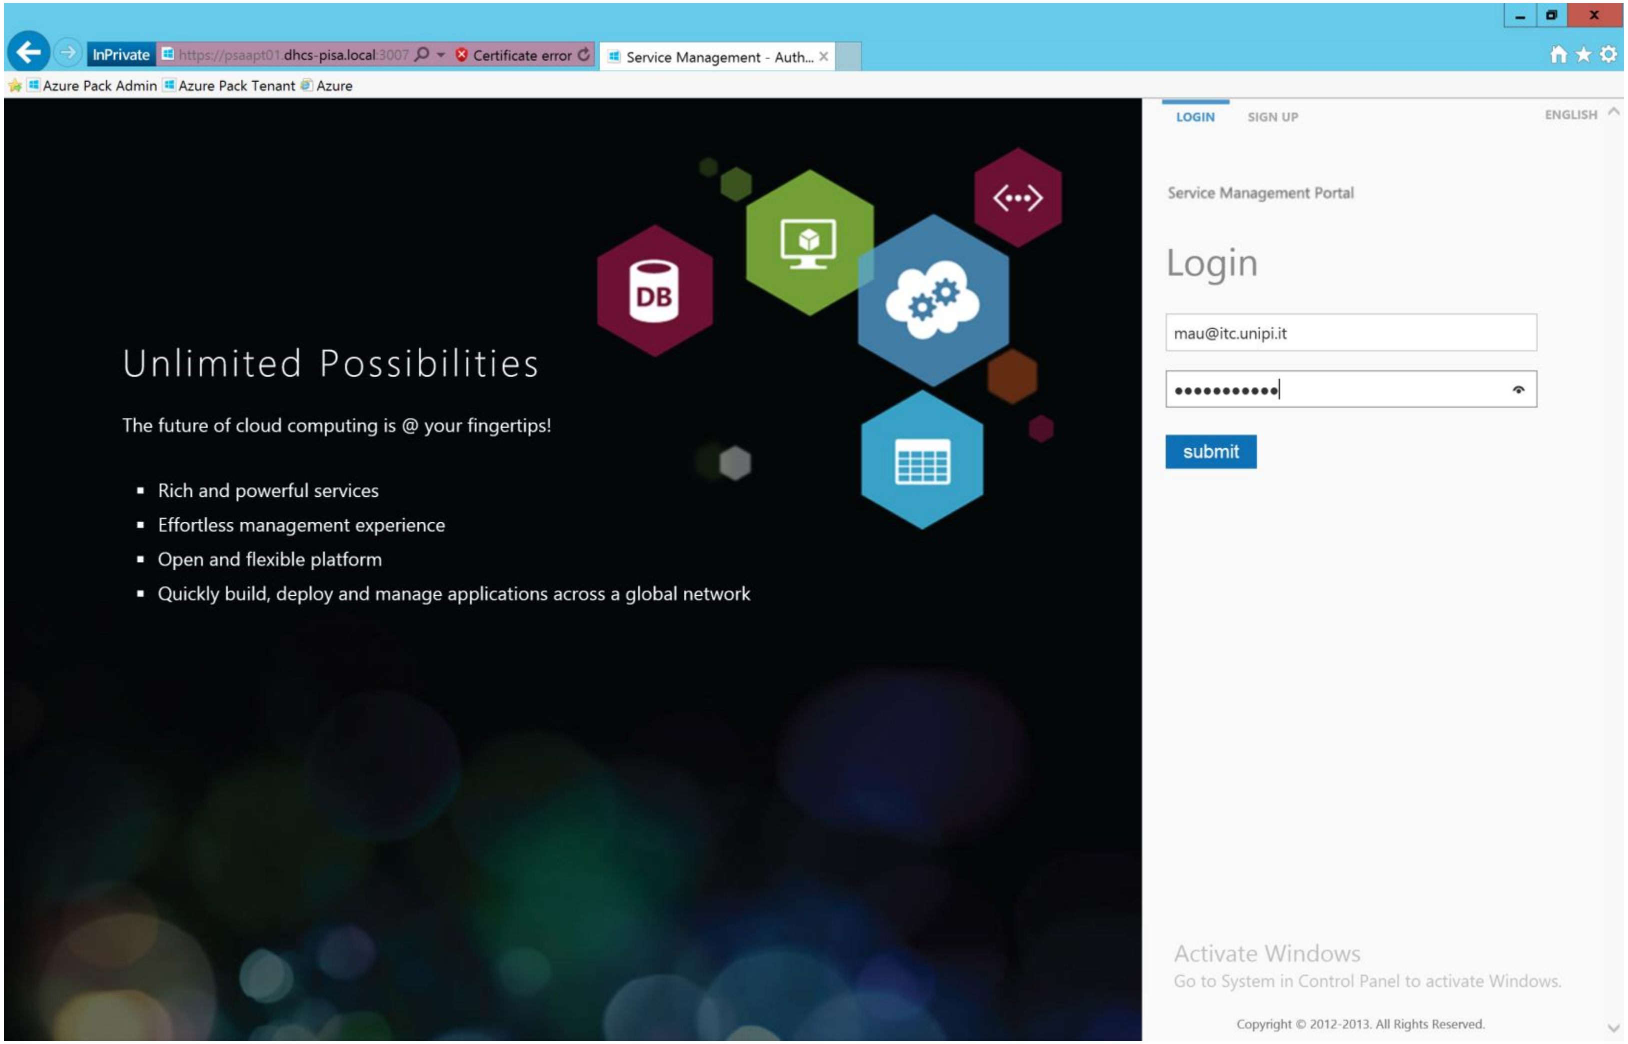
InPrivate
https://psaapt01.dhcs-pisa.local:3007
Certificate error
Service Management - Auth...
Azure Pack Admin
Azure Pack Tenant
Azure
Unlimited Possibilities
The future of cloud computing is @ your fingertips!
Rich and powerful services
Effortless management experience
Open and flexible platform
Quickly build, deploy and manage applications across a global network
DB
LOGIN
SIGN UP
ENGLISH
Service Management Portal
Login
mau@itc.unipi.it
submit
Activate Windows
Go to System in Control Panel to activate Windows.
Copyright © 2012-2013. All Rights Reserved.

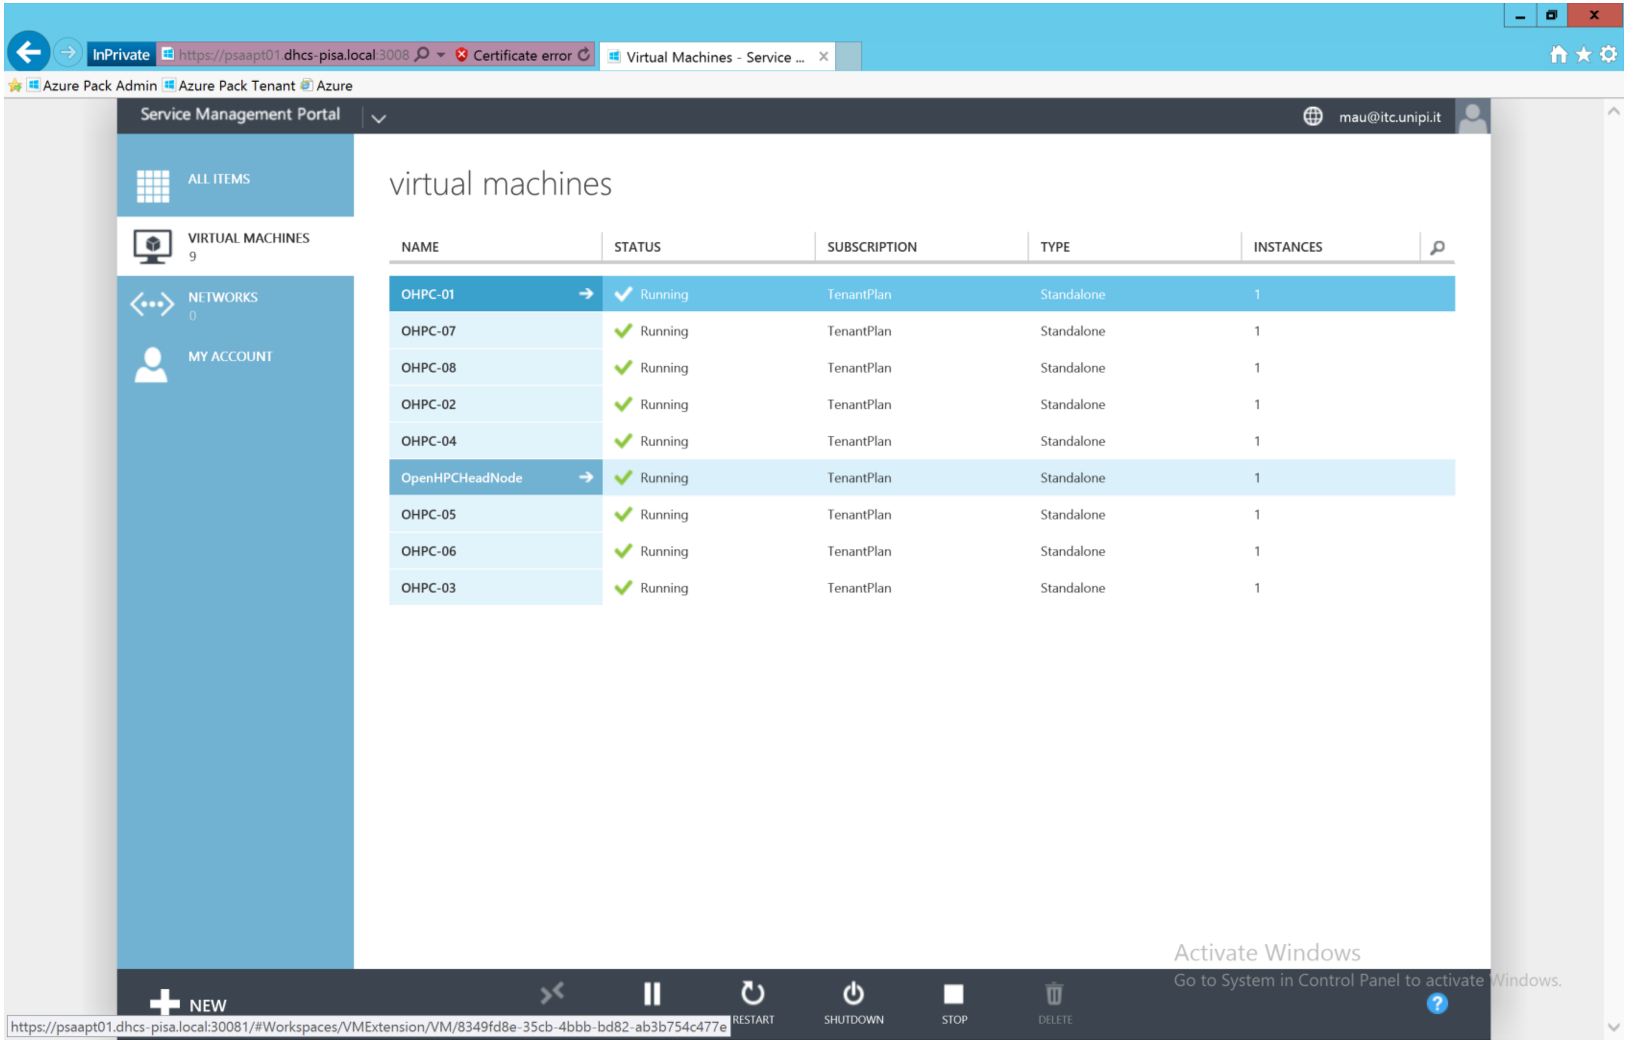
InPrivate
https://psaapt01.dhcs-pisa.local:3008
Certificate error
Virtual Machines - Service ...
Azure Pack Admin
Azure Pack Tenant
Azure
Service Management Portal
mau@itc.unipi.it
ALL ITEMS
VIRTUAL MACHINES
9
NETWORKS
0
MY ACCOUNT
virtual machines
NAME
STATUS
SUBSCRIPTION
TYPE
INSTANCES
OHPC-01 Running TenantPlan Standalone 1
OHPC-07 Running TenantPlan Standalone 1
OHPC-08 Running TenantPlan Standalone 1
OHPC-02 Running TenantPlan Standalone 1
OHPC-04 Running TenantPlan Standalone 1
OpenHPCHeadNode Running TenantPlan Standalone 1
OHPC-05 Running TenantPlan Standalone 1
OHPC-06 Running TenantPlan Standalone 1
OHPC-03 Running TenantPlan Standalone 1
Activate Windows
Go to System in Control Panel to activate Windows.
NEW
RESTART
SHUTDOWN
STOP
DELETE
https://psaapt01.dhcs-pisa.local:30081/#Workspaces/VMExtension/VM/8349fd8e-35cb-4bbb-bd82-ab3b754c477e

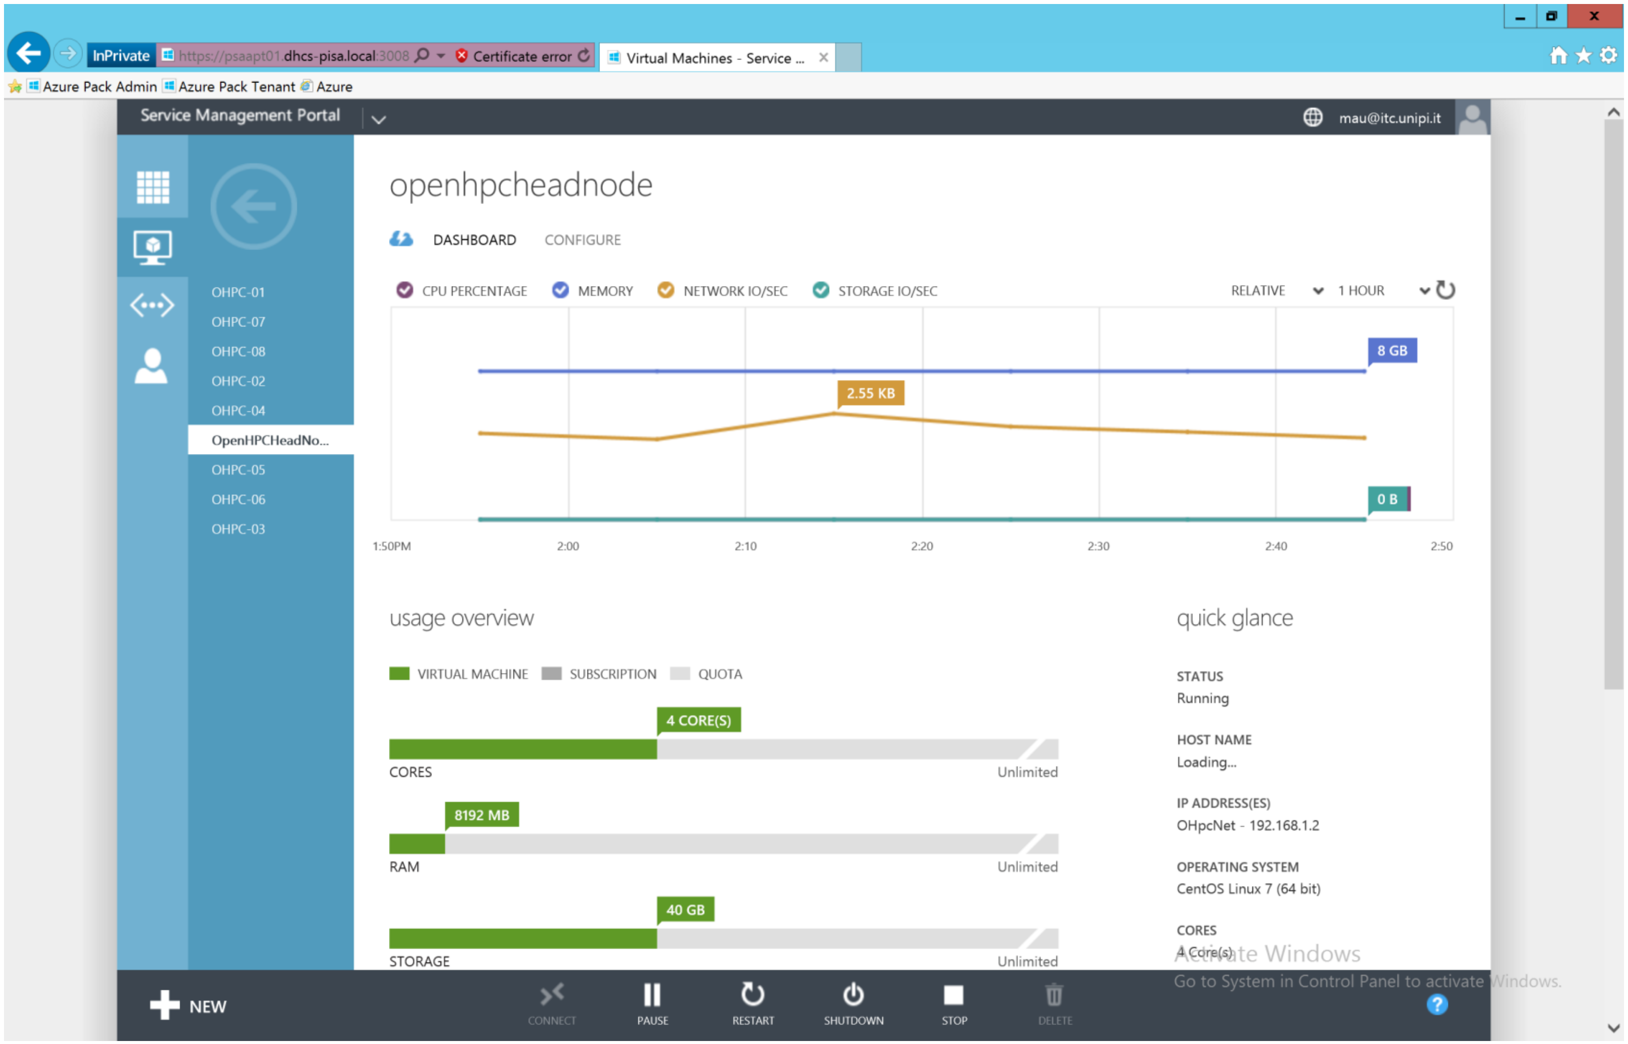

Service Management Portal
mau@itc.unipi.it
openhpcheadnode
DASHBOARD
CONFIGURE
CPU PERCENTAGE
MEMORY
NETWORK IO/SEC
STORAGE IO/SEC
RELATIVE
1 HOUR
OHPC-01
OHPC-07
OHPC-08
OHPC-02
OHPC-04
OpenHPCHeadNo...
OHPC-05
OHPC-06
OHPC-03
8 GB
2.55 KB
0 B
1:50PM
2:00
2:10
2:20
2:30
2:40
2:50
usage overview
VIRTUAL MACHINE
SUBSCRIPTION
QUOTA
4 CORE(S)
CORES
Unlimited
8192 MB
RAM
Unlimited
40 GB
STORAGE
Unlimited
quick glance
STATUS
Running
HOST NAME
Loading...
IP ADDRESS(ES)
OHpcNet - 192.168.1.2
OPERATING SYSTEM
CentOS Linux 7 (64 bit)
CORES
4 Core(s)
NEW
CONNECT
PAUSE
RESTART
SHUTDOWN
STOP
DELETE

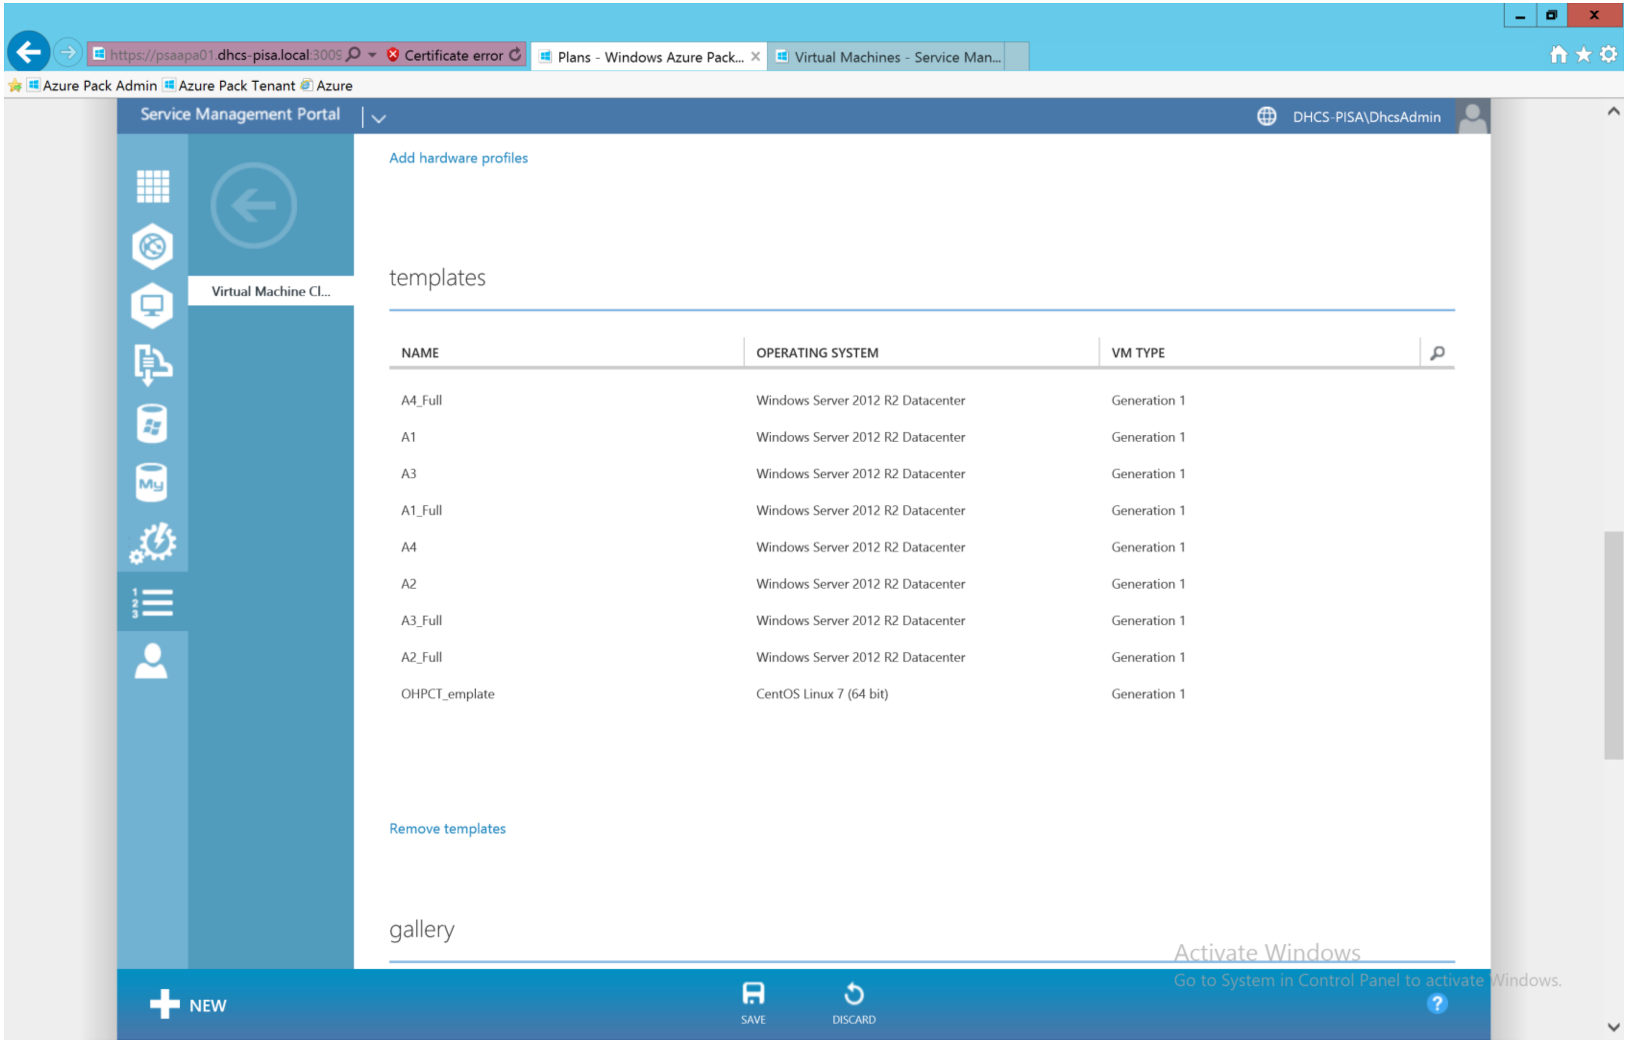

https://psaapa01.dhcs-pisa.local:3009
Certificate error
Plans - Windows Azure Pack...
Virtual Machines - Service Man...
Azure Pack Admin
Azure Pack Tenant
Azure
Service Management Portal
DHCS-PISA\DhcsAdmin
Add hardware profiles
Virtual Machine Cl...
templates
NAME
OPERATING SYSTEM
VM TYPE
A4_Full
Windows Server 2012 R2 Datacenter
Generation 1
A1
Windows Server 2012 R2 Datacenter
Generation 1
A3
Windows Server 2012 R2 Datacenter
Generation 1
A1_Full
Windows Server 2012 R2 Datacenter
Generation 1
A4
Windows Server 2012 R2 Datacenter
Generation 1
A2
Windows Server 2012 R2 Datacenter
Generation 1
A3_Full
Windows Server 2012 R2 Datacenter
Generation 1
A2_Full
Windows Server 2012 R2 Datacenter
Generation 1
OHPCT_emplate
CentOS Linux 7 (64 bit)
Generation 1
Remove templates
gallery
Activate Windows
Go to System in Control Panel to activate Windows.
NEW
SAVE
DISCARD

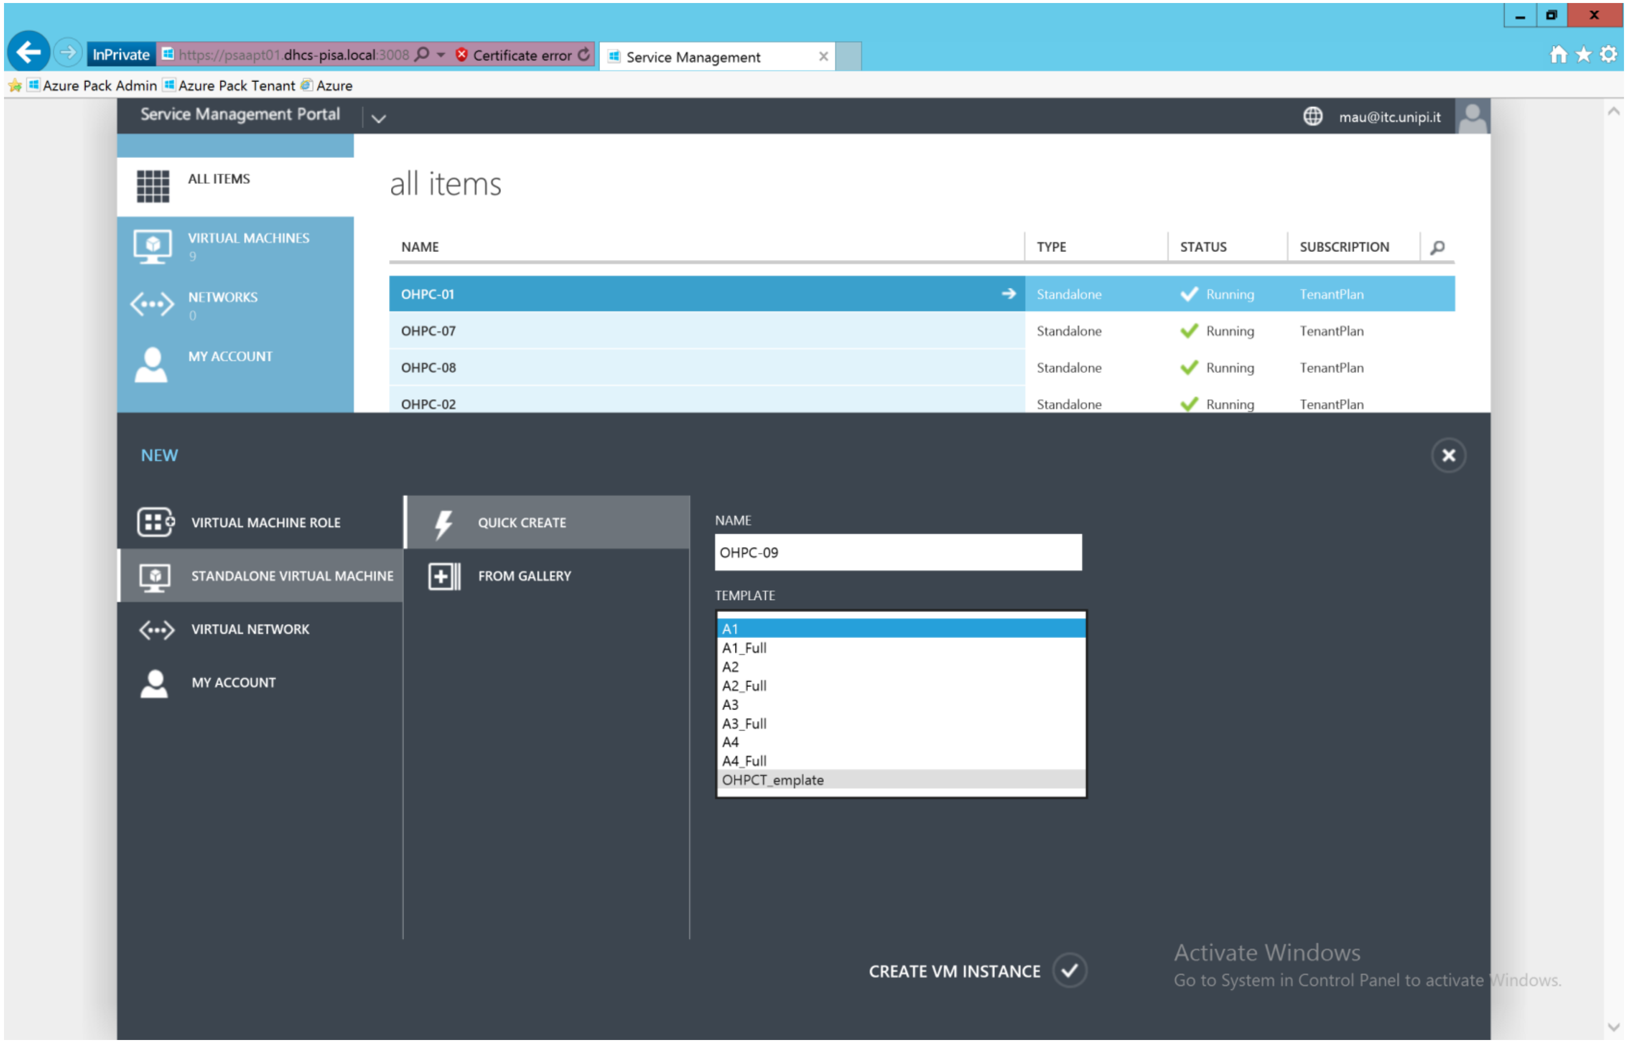
InPrivate
https://psaapt01.dhcs-pisa.local:3008
Certificate error
Service Management
Azure Pack Admin
Azure Pack Tenant
Azure
Service Management Portal
mau@itc.unipi.it
ALL ITEMS
VIRTUAL MACHINES
9
NETWORKS
0
MY ACCOUNT
all items
NAME
TYPE
STATUS
SUBSCRIPTION
OHPC-01 Standalone Running TenantPlan
OHPC-07 Standalone Running TenantPlan
OHPC-08 Standalone Running TenantPlan
OHPC-02 Standalone Running TenantPlan
NEW
VIRTUAL MACHINE ROLE
STANDALONE VIRTUAL MACHINE
VIRTUAL NETWORK
MY ACCOUNT
QUICK CREATE
FROM GALLERY
NAME
OHPC-09
TEMPLATE
A1
A1_Full
A2
A2_Full
A3
A3_Full
A4
A4_Full
OHPCT_emplate
CREATE VM INSTANCE
Activate Windows
Go to System in Control Panel to activate Windows.

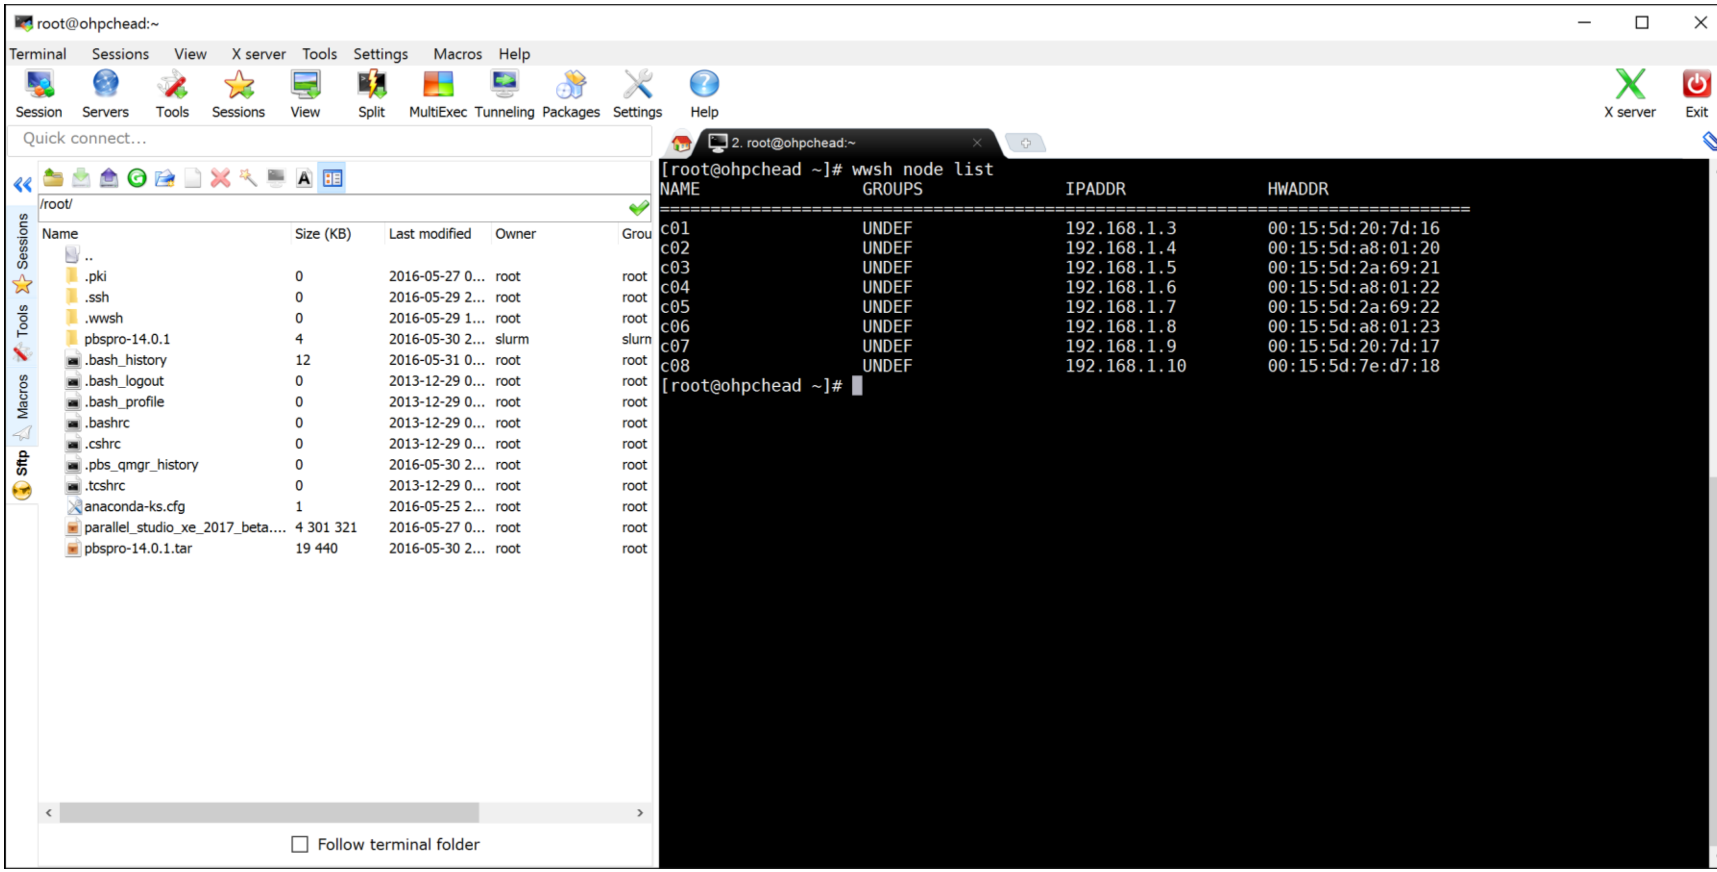
root@ohpchead:~
Terminal Sessions View X server Tools Settings Macros Help
Session Servers Tools Sessions View Split MultiExec Tunneling Packages Settings Help
X server Exit
Quick connect...
2. root@ohpchead:~
/root/
Name Size (KB) Last modified Owner Grou
.pki 0 2016-05-27 0... root root
.ssh 0 2016-05-29 2... root root
.wwsh 0 2016-05-29 1... root root
pbspro-14.0.1 4 2016-05-30 2... slurm slurm
.bash_history 12 2016-05-31 0... root root
.bash_logout 0 2013-12-29 0... root root
.bash_profile 0 2013-12-29 0... root root
.bashrc 0 2013-12-29 0... root root
.cshrc 0 2013-12-29 0... root root
.pbs_qmgr_history 0 2016-05-30 2... root root
.tcshrc 0 2013-12-29 0... root root
anaconda-ks.cfg 1 2016-05-25 2... root root
parallel_studio_xe_2017_beta.... 4 301 321 2016-05-27 0... root root
pbspro-14.0.1.tar 19 440 2016-05-30 2... root root
Follow terminal folder
Sessions Tools Macros Sftp
[root@ohpchead ~]# wwsh node list
NAME GROUPS IPADDR HWADDR
c01 UNDEF 192.168.1.3 00:15:5d:20:7d:16
c02 UNDEF 192.168.1.4 00:15:5d:a8:01:20
c03 UNDEF 192.168.1.5 00:15:5d:2a:69:21
c04 UNDEF 192.168.1.6 00:15:5d:a8:01:22
c05 UNDEF 192.168.1.7 00:15:5d:2a:69:22
c06 UNDEF 192.168.1.8 00:15:5d:a8:01:23
c07 UNDEF 192.168.1.9 00:15:5d:20:7d:17
c08 UNDEF 192.168.1.10 00:15:5d:7e:d7:18
[root@ohpchead ~]#

One Last thing….

PBS Pro: HPC Workload Management & Job Scheduling

Faster time-to-results, better throughput and utilization

- EAL3+ security certification and SELinux support
- Policy-driven, topology-aware scheduling
- Accelerator/Co-processor scheduling
- Green Provisioning™ for power management
- Cgroups and fairshare available
- Proven to run millions of jobs per day
- Backfill, sharing, and shrink-to-fit jobs maximize usage
- Extensible plugin framework
- Open architecture to implement virtually any policy

NASA's Workload Manager of Choice for All NAS HPC Resources

~200k cores scheduled by PBS Professional

PBS Pro Open Source – mid-2016

- Marry community innovations & enterprise expertise
  - Accelerate innovatic n for both
- Founding member c
- Committed to openness and longevity
- Dual licensing: OSI-approved & commercial
  - No change for commercial users
- A single, common "core" for PBS Pro
- Accepted OSS practices, e.g., GitHub
- More info: **pbsworks.com/opensource**

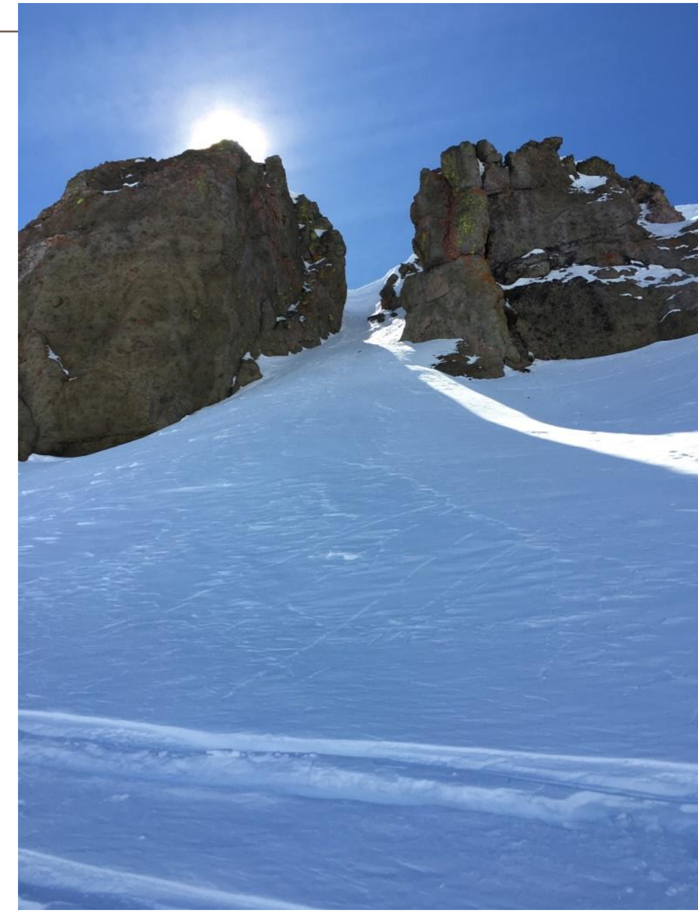

PBS Pro on OpenHPC

mau — root@ohpchead:~/pbspro-14.0.1 — ssh root@131.114.137.127 — 80×24

```
[root@ohpchead pbspro-14.0.1]# ls
aclocal.m4   config.status  doc       m4           README
autogen.sh   configure      INSTALL   Makefile     src
buildutils   configure.ac   libtool   Makefile.am  test
config.log   COPYRIGHT      LICENSE   Makefile.in
[root@ohpchead pbspro-14.0.1]# qstat --version
pbs_version = 14.0.1
[root@ohpchead pbspro-14.0.1]# 
```

PBS and Docker

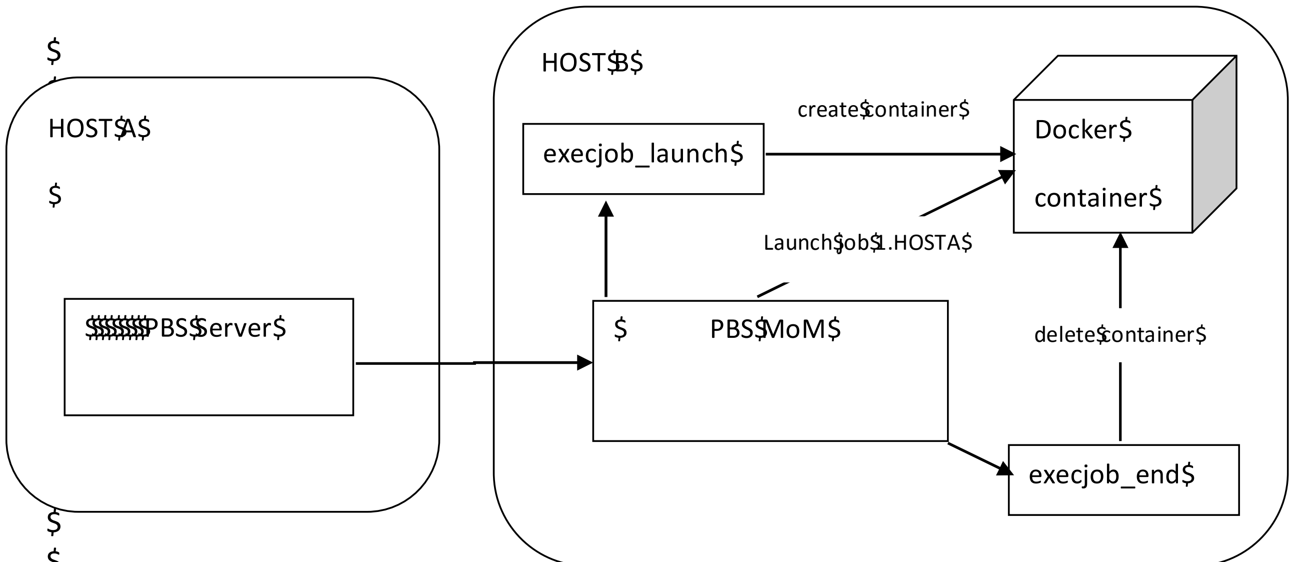

PBS and Docker

- PBS starts a separate Docker container to run each job; the scope of this container ends with the job.
- PBS provides the exact same environment to a job in a container that it provides to any other job when staging files in and out, managing a job's output and error files, and exporting environment variables.
- When PBS starts a Docker container, it configures the container so that the job cannot use more resources than it explicitly requested.
- For multi-vnode jobs, PBS configures any Docker containers created on sister MoMs to be network-linked to the container running on the mother superior, allowing communication between job tasks on separate execution hosts.
- When a job ends, PBS manages shutdown of the Docker containers it started for that particular job.

… to finish

- OpenHPC can be successfully used on flexible clouds platform
- A lot of scientific software will available on this platform

- We have to thanks a lot of people but a special thanks: