CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

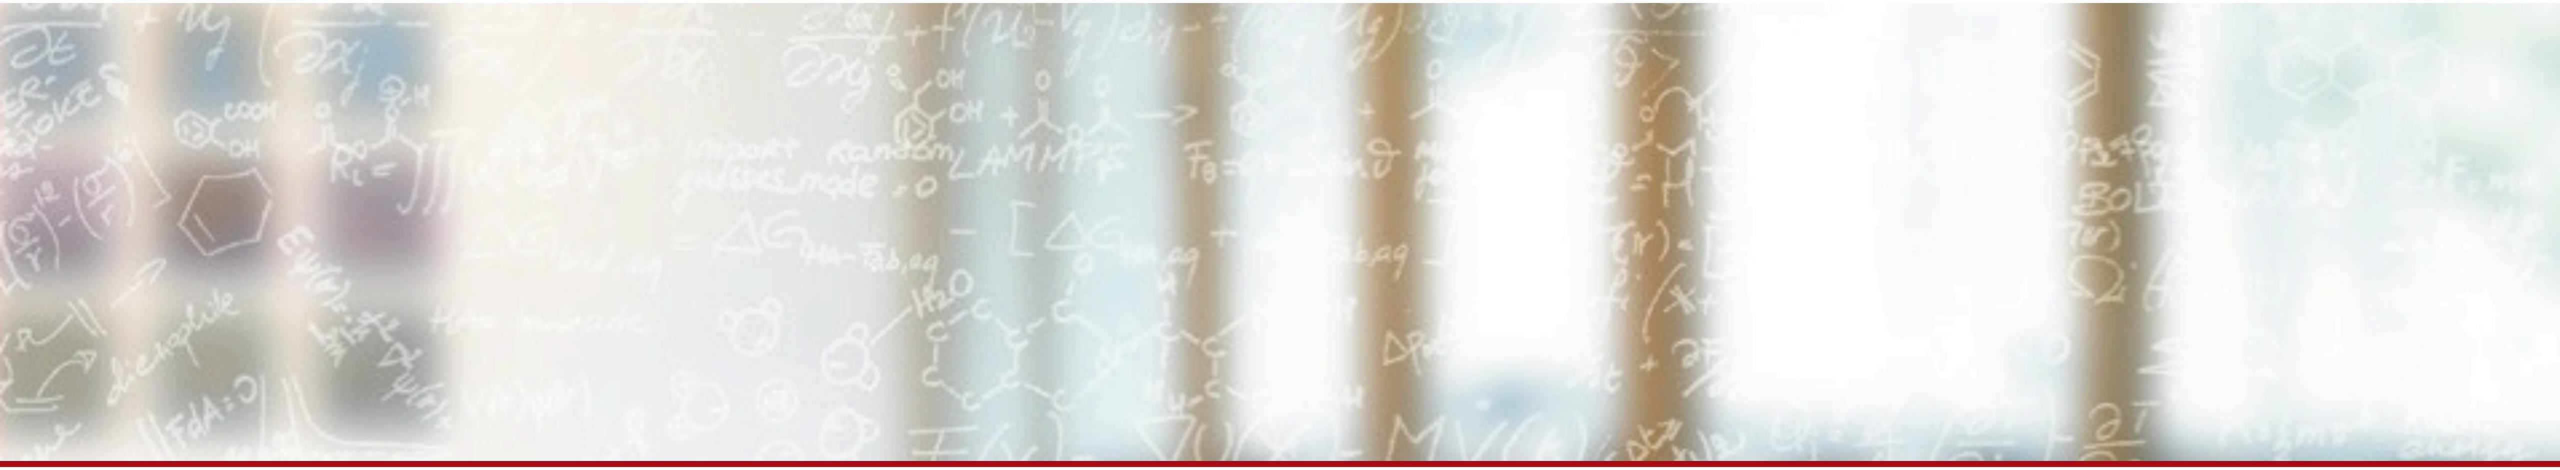

Bringing Cloud Accessibility to HPC Technology, a Use Case Driven Approach

Maxime Martinasso
HPC-CH Forum
May 16, 2019

Presentation overview

1. Cloud and HPC
2. Use cases and requirements
3. HPC and Cloud

What are the benefits of Cloud technology?

- Infrastructure as a Service (IaaS)
  - User can decide of their hardware resource
  - User can deploy their own platform on the infrastructure

- Functionality
  - Identity access management
  - Access in and out to the web
  - Simple interface to create platforms and select resources
  - Clear and straight forward payments or credit allocation

- Drawbacks of Public Cloud providers
  - Commercial focus and not academic
  - Shared and virtualized resource
  - No HPC-like Performance guarantee

Example of Cloud providers using HPC technology

- GPU for compute
  - GPU as an accelerator was introduced by the HPC community
  - Amazon AWS, Google Cloud, Nvidia cloud, …
  - AI and ML workloads
- Easy on-demand HPC cluster on Google Cloud
  - Configure number of nodes
  - Deploy Slurm
- Amazon AWS: Lustre on-demand
  - Lustre is a complex parallel file-system developed by the HPC community
  - Very difficult configuration to guarantee performance
- Microsoft Azure has acquired a Cray system
  - Cray XC + Lustre (thousands of nodes)

Mission of an HPC Centre

- Provide High-Performance compute capability
  - On node: fast CPU, high memory, accelerator
  - Scalability: fast interconnect
- Provide efficient storage service
  - High bandwidth parallel file-system
  - Near-compute fast SSDs
  - Systems-wide file system
  - Archival storage
- Goal: enable scientist to increase scientific discoveries by reducing time-to-solution

Swiss Data Science Centre

- SDSC Mission
  - Accelerate the use of data science and machine learning techniques within academic disciplines
  - RENKU allows scientists to securely *Share/Reproduce/Reuse* Data, Metadata and Code

- Interactivity
  - Spawns Jupyter notebooks to compute nodes
  - Workflow steps are stored in git

- Manage data with git and gitLFS (large file storage)
  - Create a git repository for every project
  - Small files are stored in the git repository
  - Large files are referenced inside a gitLFS repository (pointer to location of data)

- Today: run on Cloud (OpenStack), future: connectivity to HPC (compute+data)

Materials Cloud

MARVEL

NATIONAL CENTRE OF COMPETENCE IN RESEARCH

MATERIALS CLOUD

- NCCR MARVEL Mission
  - Accelerate the design and discovery of novel materials
  - Use quantum-mechanical simulations

- Materials Cloud
  - A platform for open science with educational, research, and archiving tools, simulation software and services, curated and raw data
  - Web portal: uses Jupyter to manage workflows by using “App”
  - All sequence of steps are stored

- Integration with CSCS
  - Web portal running on OpenStack at CSCS
  - Capable of executing jobs on Piz Daint by using SSH from the web portal
  - Use a custom data broker to move data between VM and HPC (small size data)

Paul Scherrer Institute

- PSI Mission
  - Study the internal structure of a wide range of different materials
  - Research facilities: the Swiss Light Source (SLS), the free-electron X-ray laser SwissFEL, the SINQ neutron source and the SμS muon source
- PSI facility users reserve a scientific device for a period of time
  - Compute power should also be available
  - Storage and archive availability during the experiment
  - Data retrievable after experiment by the users of PSI facilities (not at PSI)
- Proposal to interface Piz Daint with their workflow
  - Access to compute and data services (job scheduler, data mover)
  - Create a reservation service to reserve computation nodes (before experiment)
  - Provide a portal running on OpenStack to let PSI users download experiment data

Use cases: common requirements

- Provide an infrastructure to execute dedicated platforms
  - On HPC system if performance required
  - Cloud infrastructure
  - Close to the data

- User identity provided by the platform
  - User are not known at the HPC center (service or temporary account)
  - Use standard authentication and authorization mechanism

- Enable programmable access to HPC resources
  - Compute: dedicated nodes + containerized jobs
  - Data: move data across center filesystem (PFS, archive, on-node)
  - Network: ??

- Connectivity
  - HPC resources ←→ Dedicated platforms
  - Web access from compute resources

Cloud technology adapted to HPC

- Identity access management
  - Leverage Cloud technology: OIDC, OAuth2, …
  - Adapt HPC services for authorization
- Containers and microservices
  - Enable user-defined software stack to run on HPC
  - Suitable for HPC: no privilege escalation, PFS friendly, diskless, WLM
  - Sarus developed at CSCS
- Programmable interface for services
  - Many HPC service are not programmable: resource allocation, data movement
  - FirecREST: a RESTful API, developed at CSCS
- Interactivity
  - Jupyter notebook and JupyterHub are developed for Cloud-like infrastructure
  - Adapt service to HPC: batch jobs vs interactive access
  - Jupyter.cscs.ch: new service, access node of Piz Daint interactively

Cloud technology adapted to HPC

- On-demand data manager versus shared PFS
  - Investigation: use container to deploy data manager on near-to-compute SSD
  - Early stage of a Proof of Concept

- On-demand network resource versus global High-Performance network
  - Requires an hardware change at NIC and switch level
  - Virtualization of interfaces, port configuration and traffic classes

- Elasticity of resource allocations
  - Clustering resources: compute, network and data
  - Needs on-demand resources capability
  - Scheduling resource access in a cluster
  - Scheduling resource allocation among clusters (meta-scheduler)
  - Maximize resource utilization

Convergence of HPC and Cloud

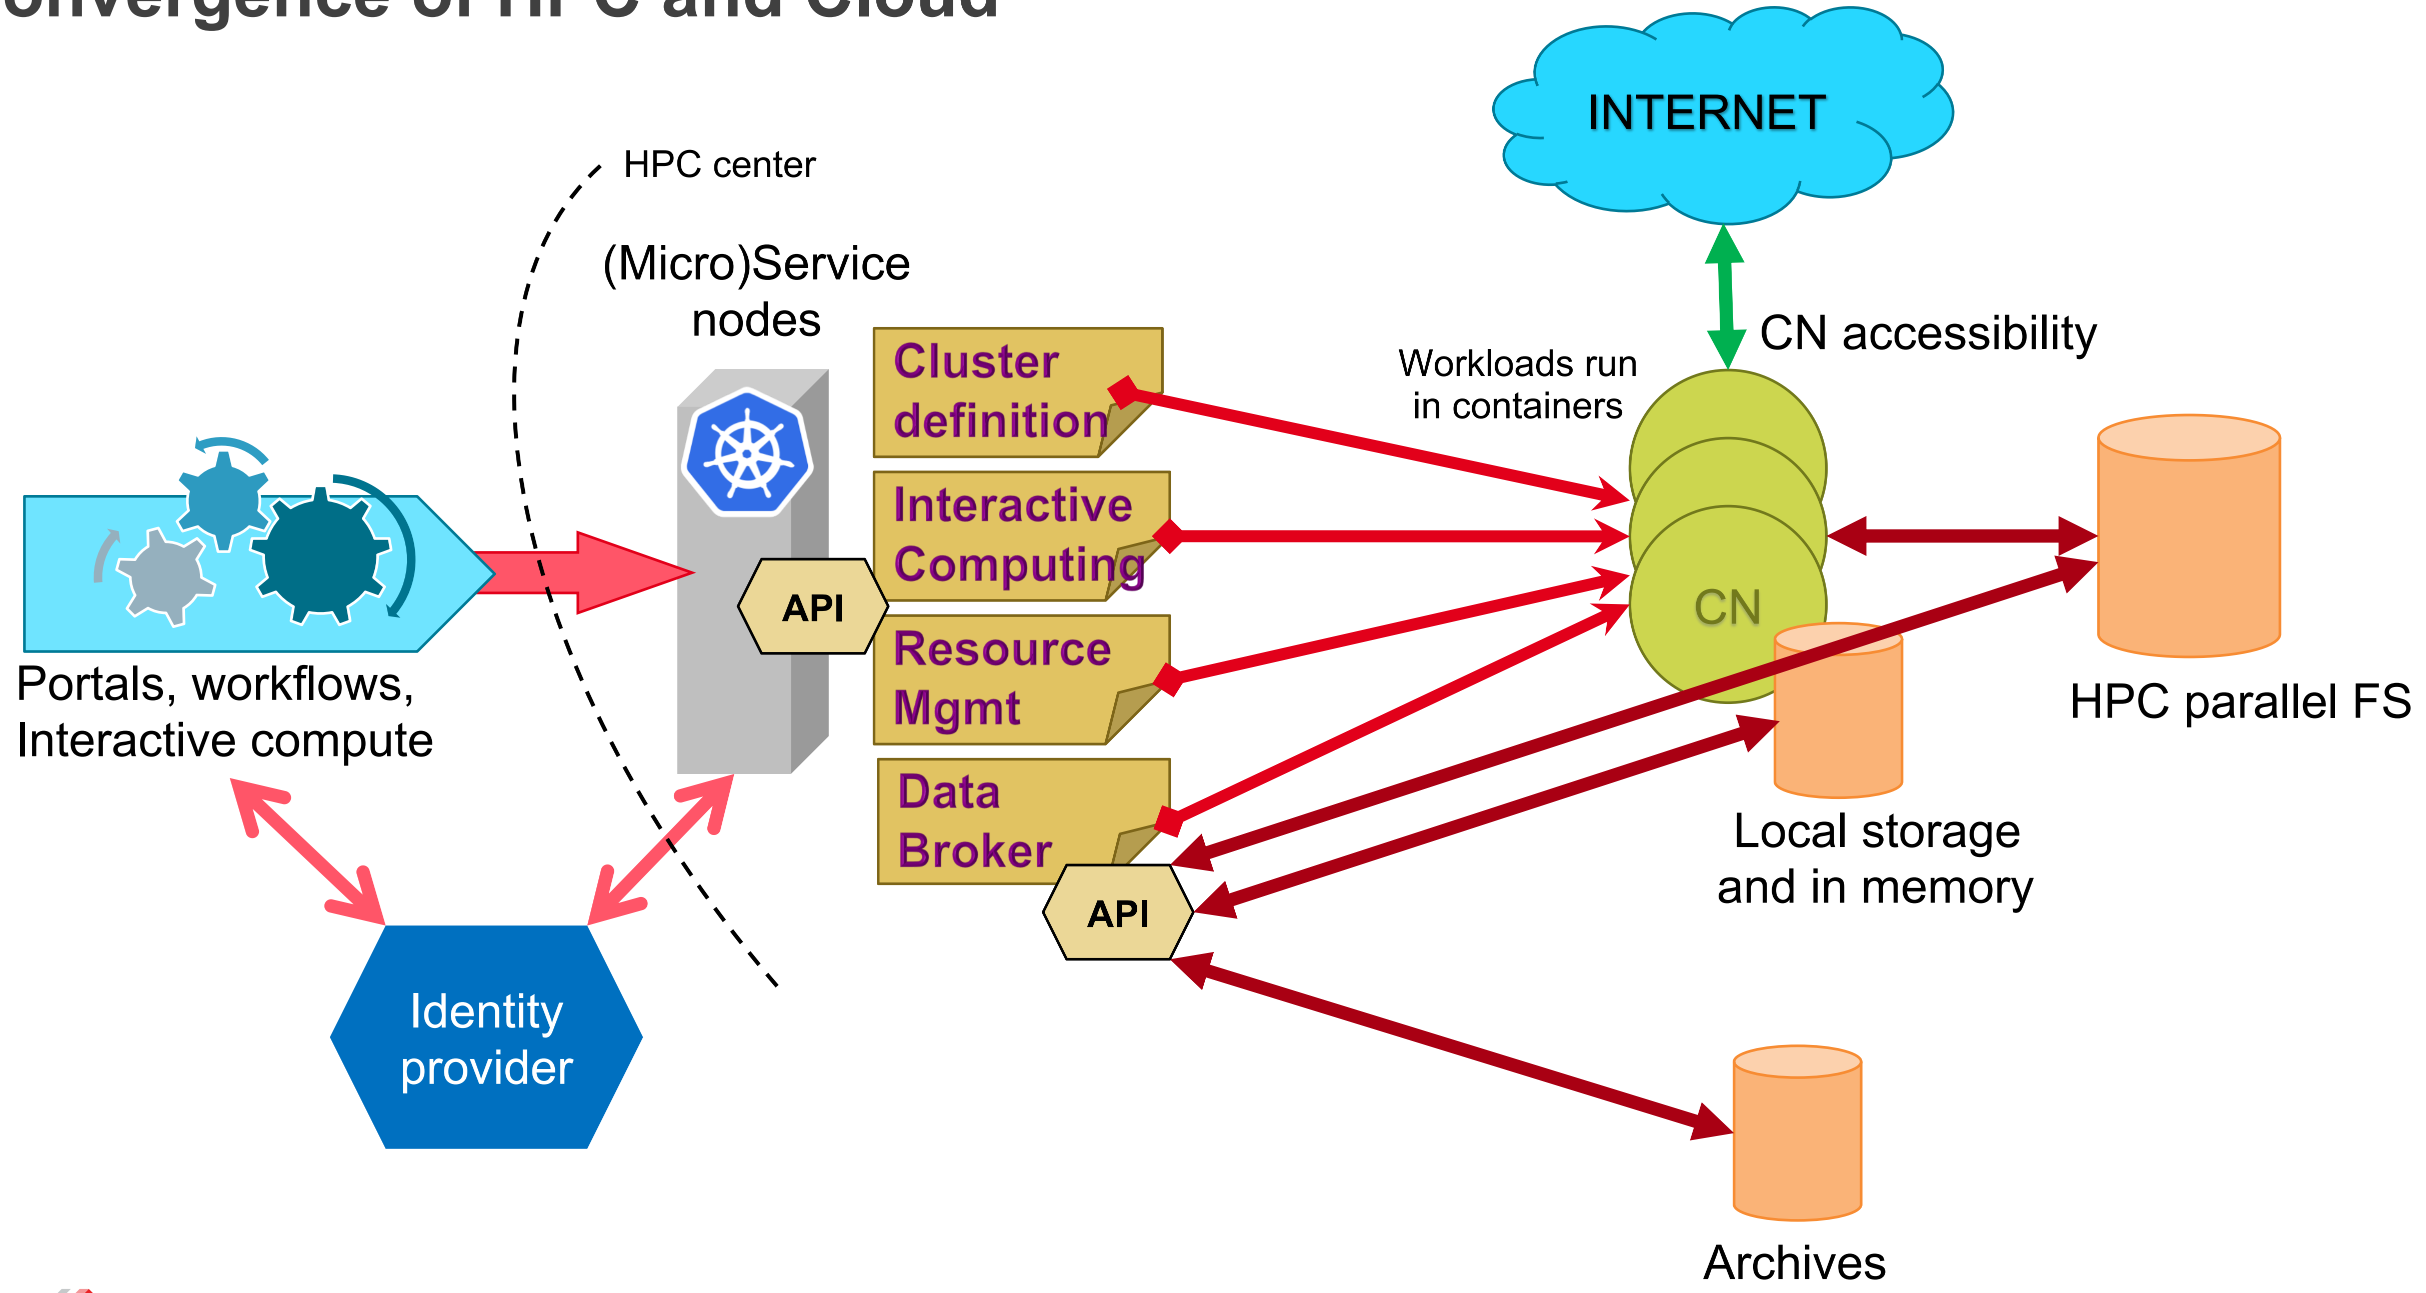

HPC and Cloud convergence

HPC + Cloud Tech ⇔ Cloud + HPC Tech

- Bring selected Cloud technologies to HPC
- Enable scientific use cases to get Cloud-like flexibility
- New management of systems and services at the HPC centre

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

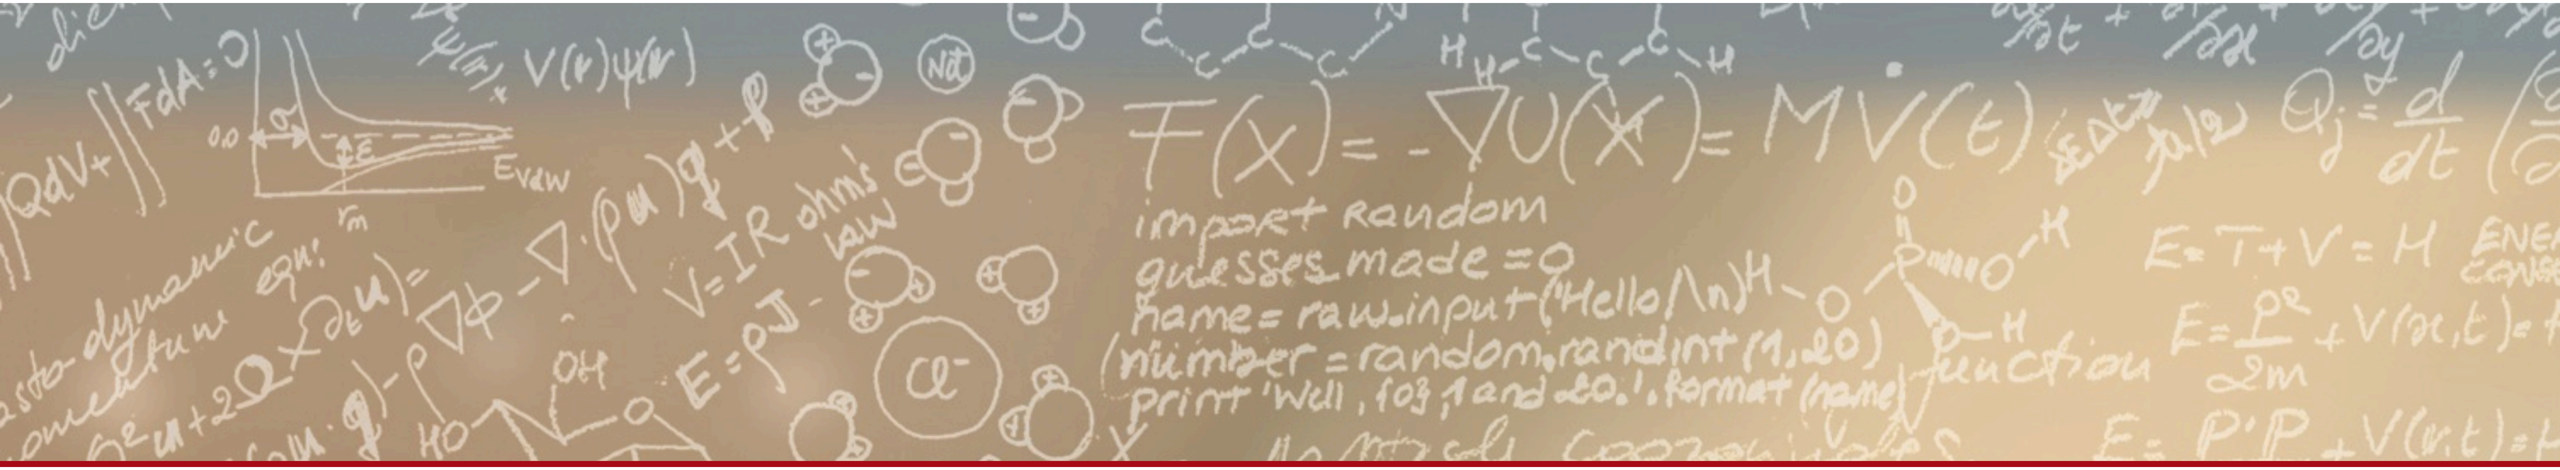

Thank you for your attention.