CERN Experience

D. Giordano / CERN-IT

WLCG Workshop
8-9 October 2016

OpenStack Clouds at CERN

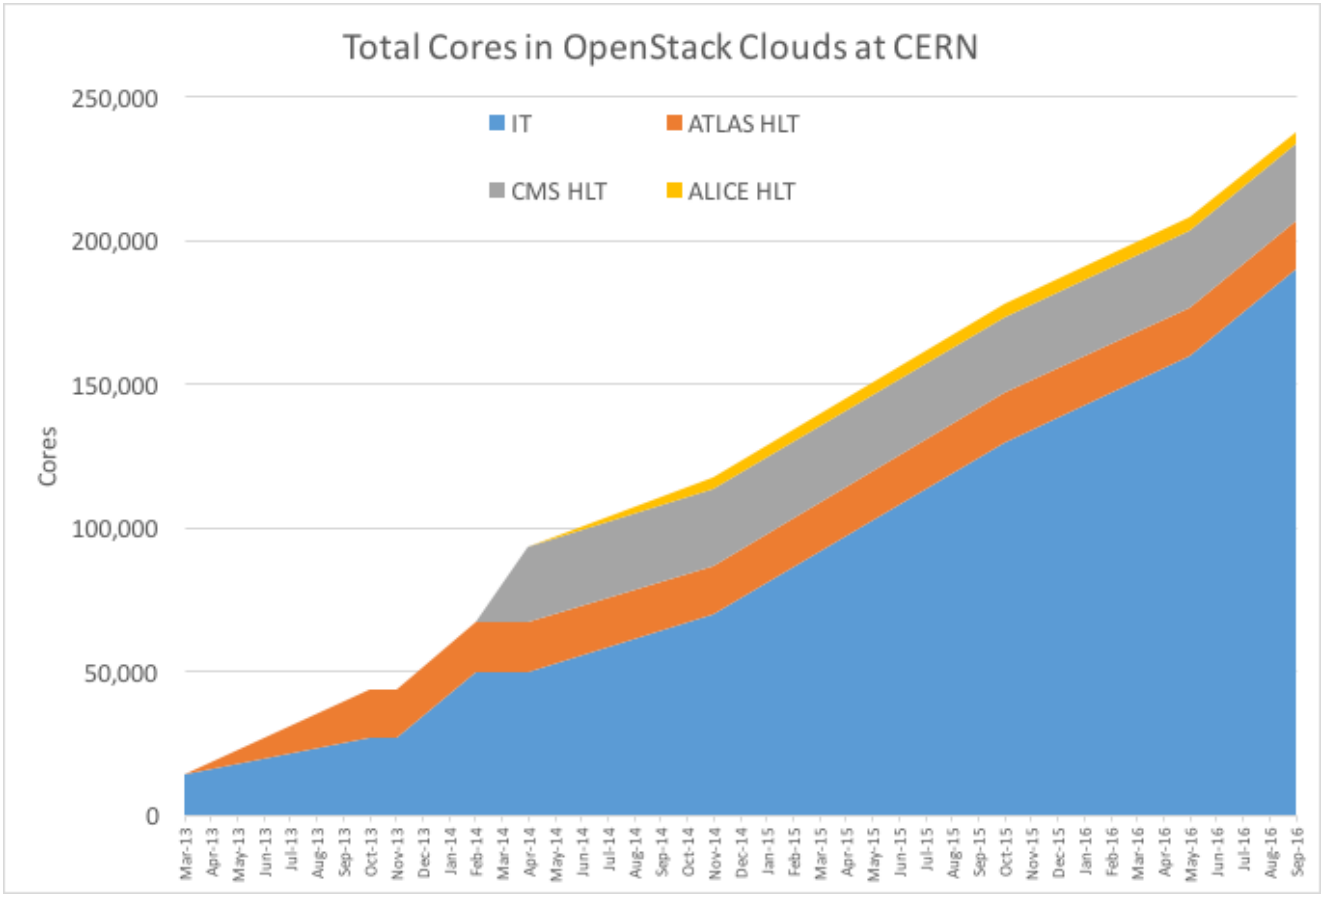

In production:

- ❑ 4 clouds
- ❑ >230K cores
- ❑ >8,000 hypervisors

90% of CERN's compute resources are now delivered on top of OpenStack

A further 42K cores to be installed in next few months

Tier-0 Compute Services 2017

- Universal resource provisioning layer for bare metal, containers and VMs
- HTCondor as the single end user interface with LSF retirement by LS2
- Continue investing in automation and other communities for scaling with fixed staff
- Self service for end users within the policies and allocations

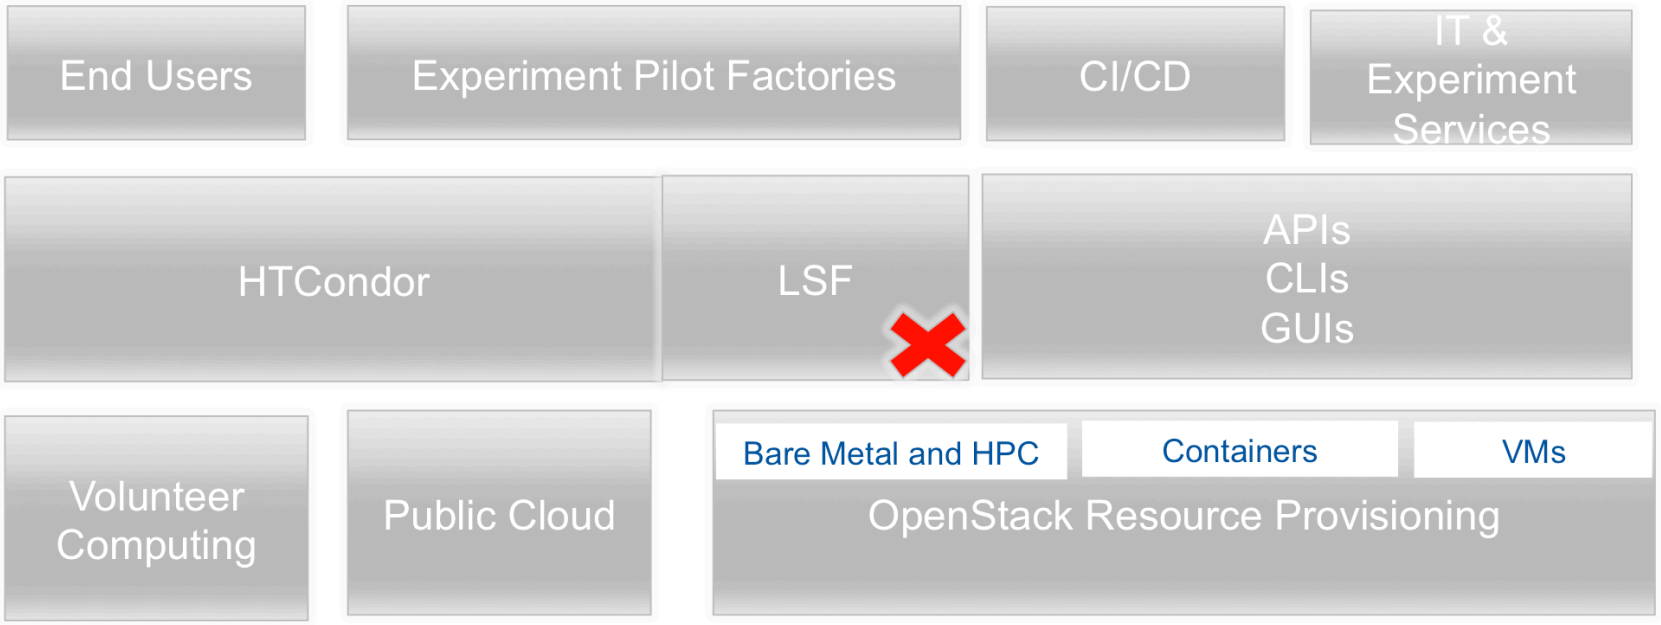

courtesy of T. Bell

Evaluation of
Public Cloud Service Providers (CSP)

Roadmap

Started in 2011 with the EC funded project Helix-Nebula

Since 2015, series of short CERN procurement projects of increasing size and complexity

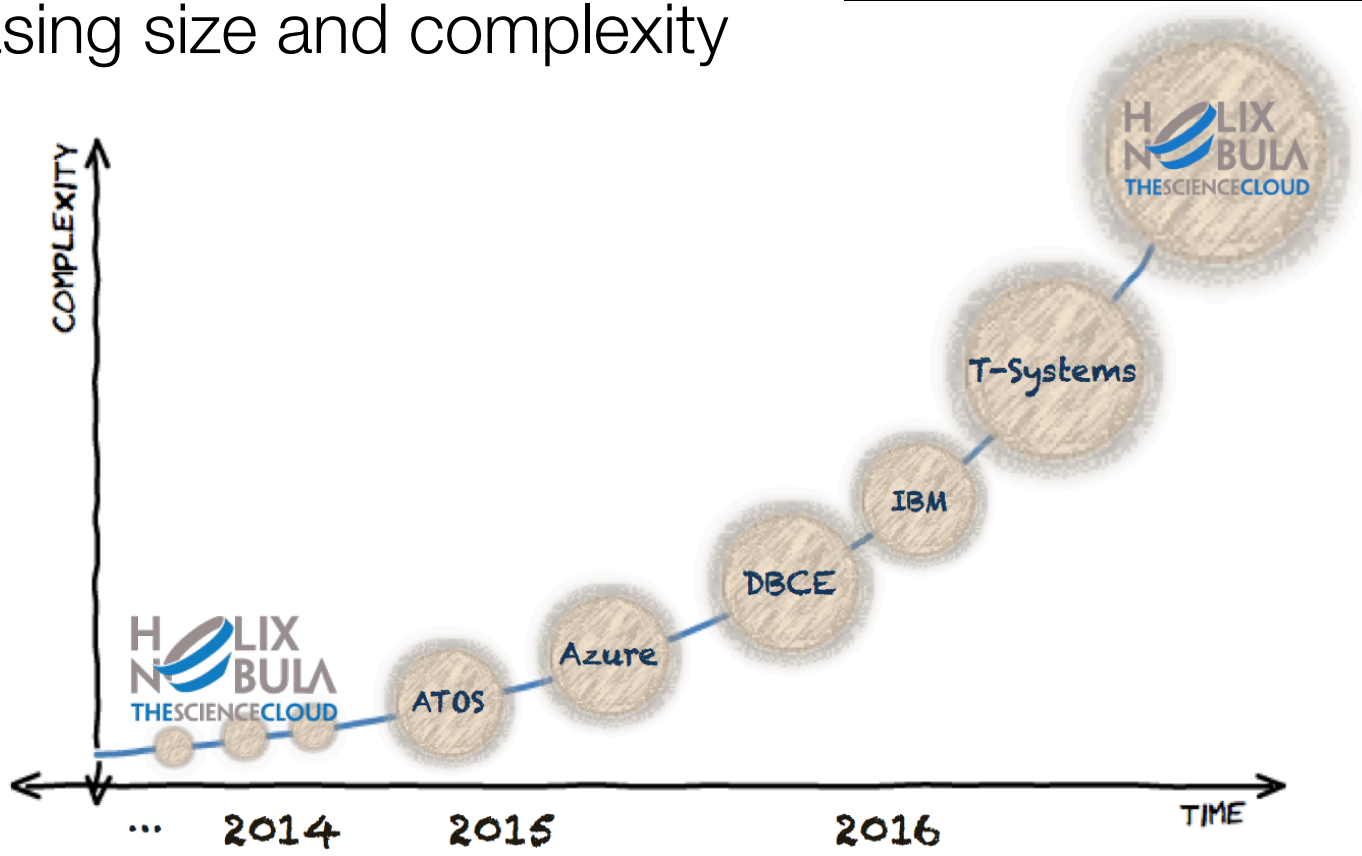

In a nutshell

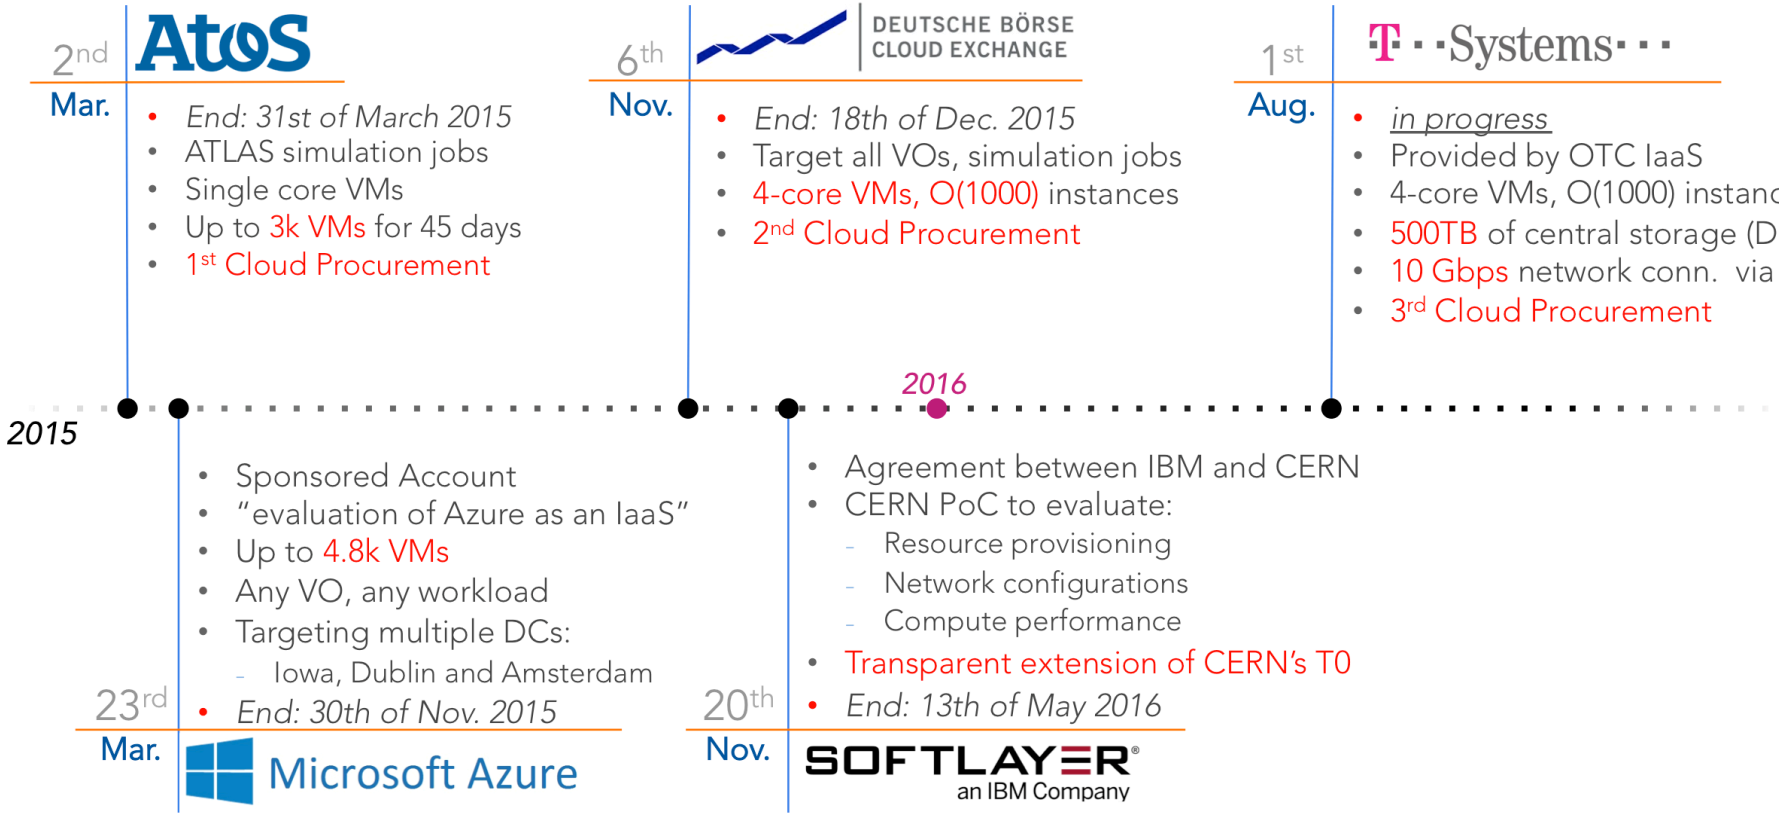

[see CHEP o-26]

Some Lessons learned

Cloud Brokerage not as effective as publicized

Overhead in managing several, small public CSP

Client-side accounting, monitoring and benchmarking are crucial

Cannot avoid preparation and commissioning phases

- Ramp-up speed is often connected to the maturity and stability of the CSP

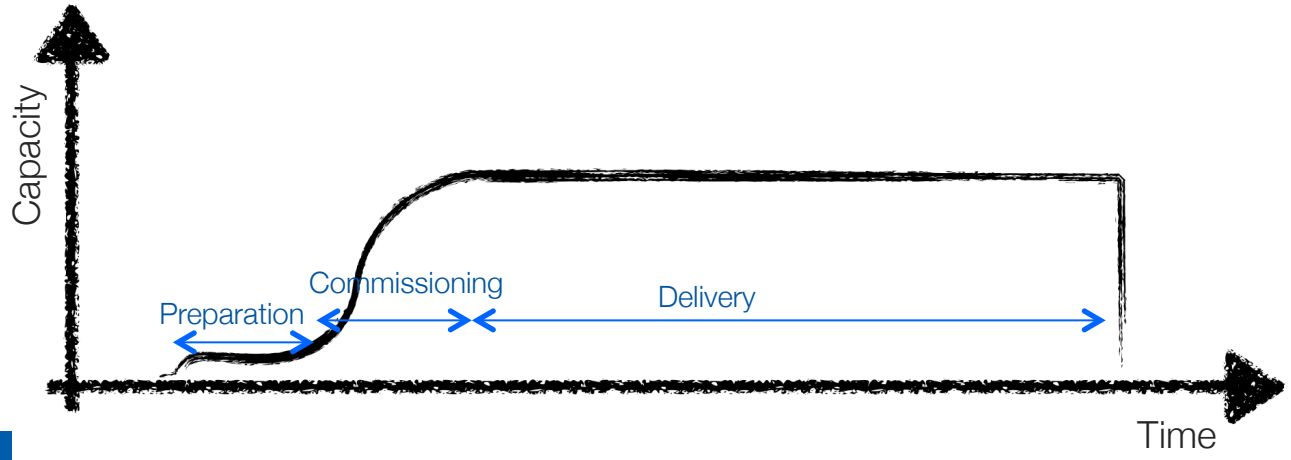

Provisioning, Monitoring, Exploitation

Transparent Extension of CERN Resources

Consolidate the strategies adopted in the past cloud activities

- Manage and exploit external resources using same toolset and entry points as CERN on premises resources
  - ***Puppet*** configuration
  - **HTCondor** for scheduling and match-making
  - Infrastructure **monitoring**

[see CHEP p-22]

- Adopted ***Terraform*** for VM lifecycle management (N.B.: looking for long VM lifetime)
  - Open source toolkit, supports several cloud providers

2015 approaches

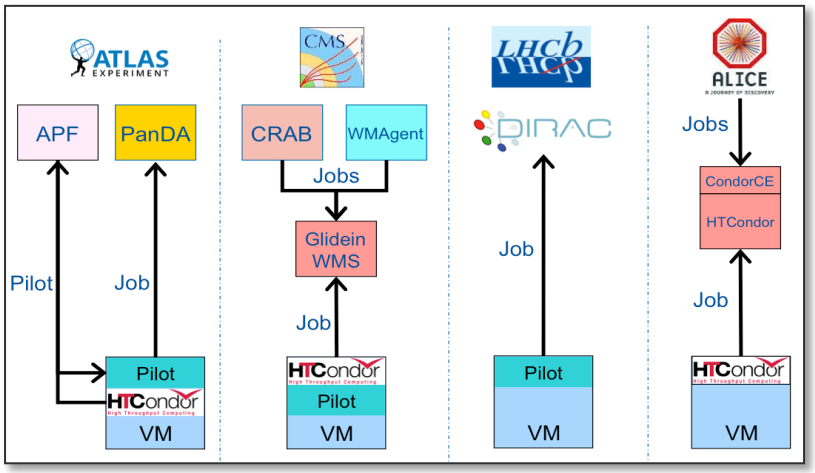

Now

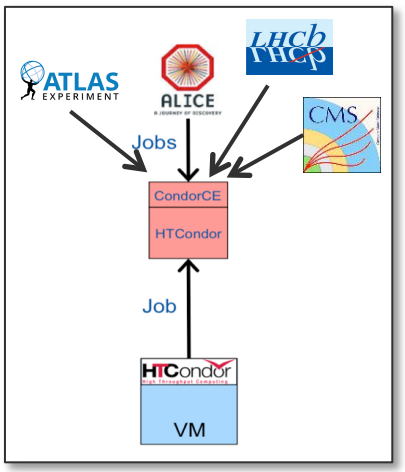

Provisioning tools:

Vcycle; custom scripts for HNX-SlipStream, Azure RM

Terraform

Transparent Extension of CERN Resources

Consolidate the strategies adopted in the past cloud activities

- Manage and exploit external resources using same toolset and entry points as CERN on premises resources
  - ***Puppet*** configuration
  - **HTCondor** for scheduling and match-making
  - Infrastructure **monitoring** *[see CHEP p-22]*
- Adopted ***Terraform*** for VM lifecycle management (N.B.: looking for long VM lifetime)
  - Open source toolkit, supports several cloud providers

2015 approaches

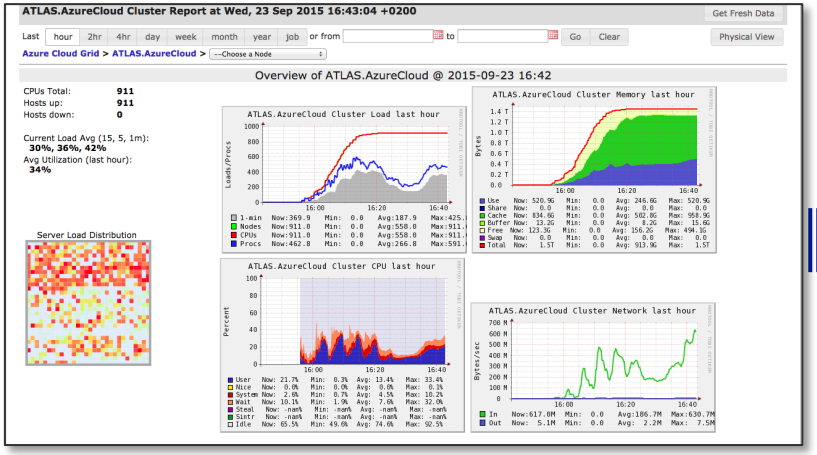

Monitoring:
Ganglia

Now

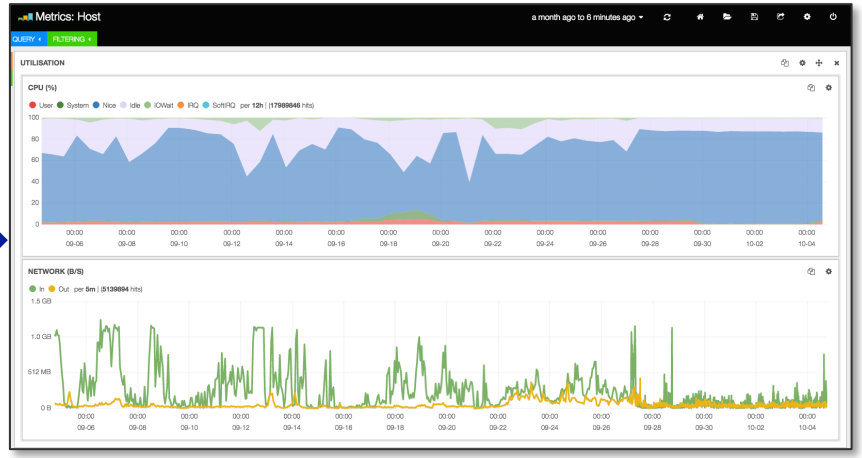

AI Monitoring

Recent activity: T-Systems

- Batch resources fully loaded
  - shared among VOs

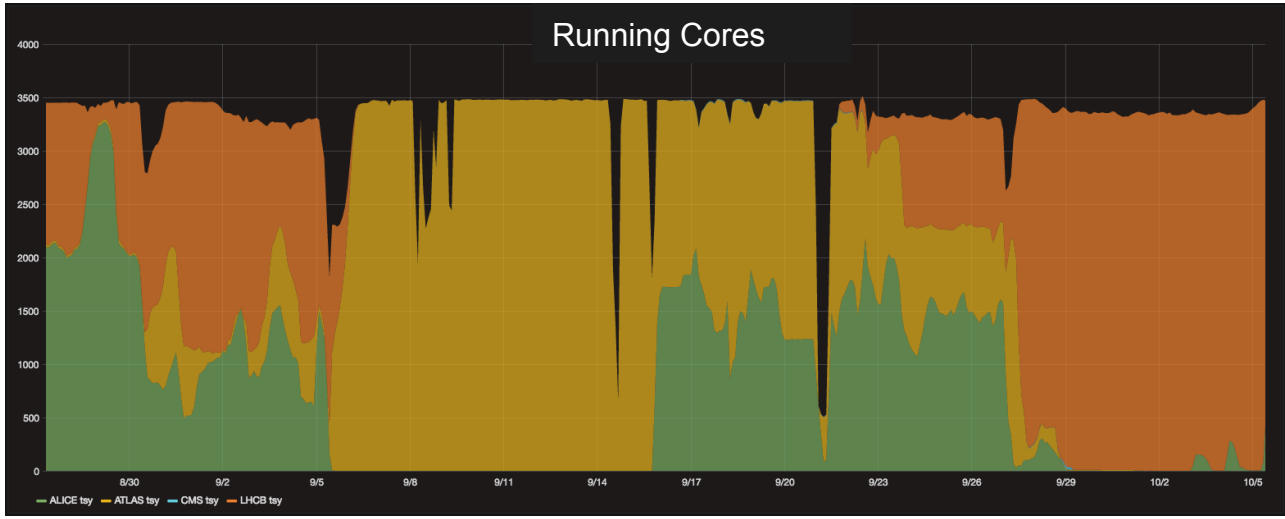

- Mixture of "CPU-intensive" and "network-intensive" tasks
  - MC workloads tend to dominate: easier to manage?

	Max	Avg
LHCB tsy	99.05	85.04
ALICE tsy	93.83	75.98
ATLAS tsy	100.00	64.13

- WAN largely used
  - Sometimes even saturated

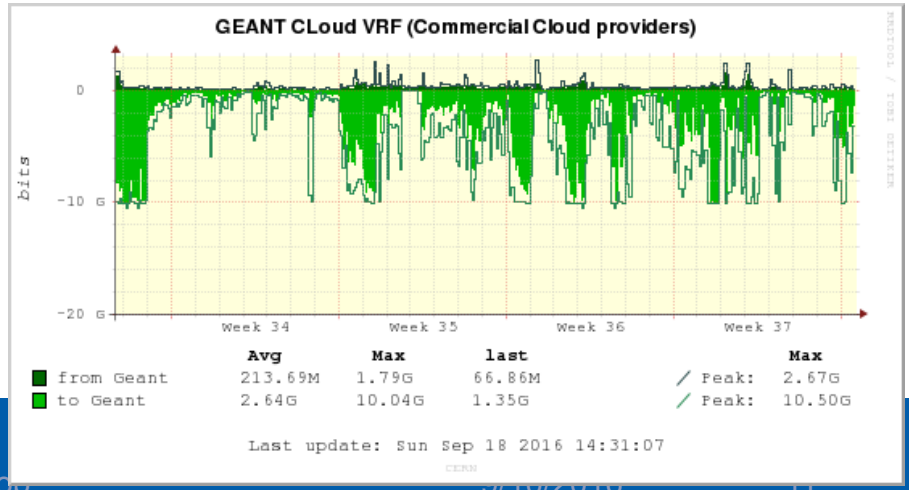

Network Connectivity

WAN connectivity over GÉANT network

Requirement for CSP since the first procurement (early ’15)

GÉANT Cloud VRF is currently connecting CERN and T-Systems (via DFN)

- 10 Gbps of total reserved peak bandwidth available
- The VRF is configured to only allow traffic between CSPs and NRENs; no CSP-CSP traffic is allowed

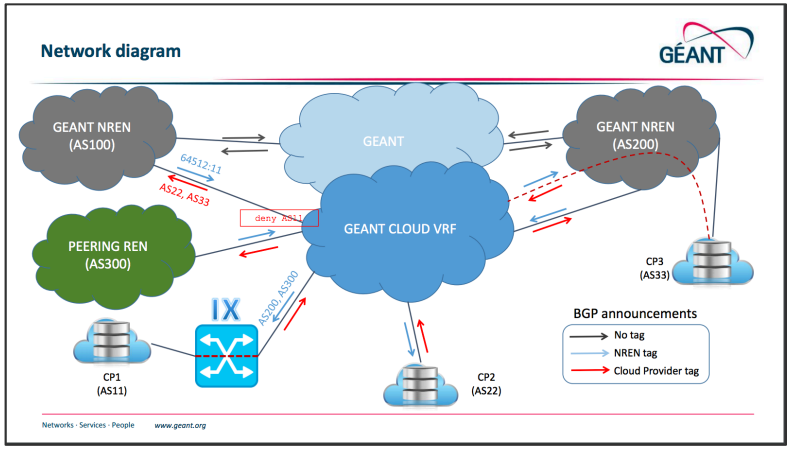

source V. Capone: GÉANT approach to Cloud R&E traffic

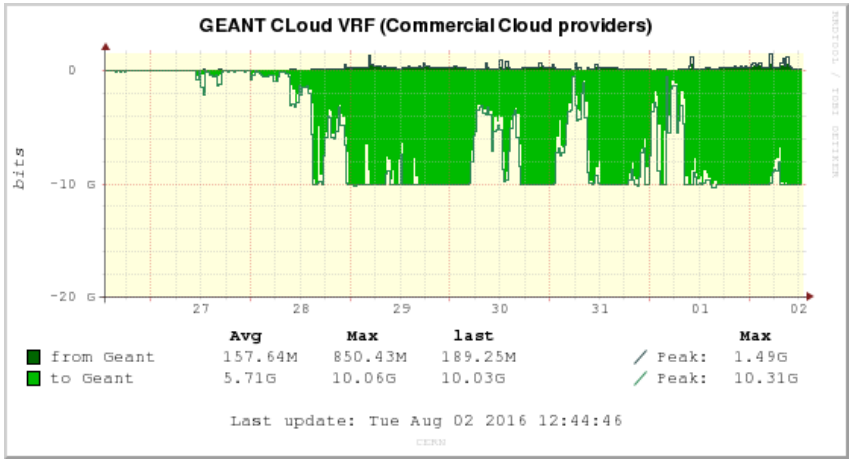

VPN Evaluation

Performed during the IBM-SoftLayer activity over GÉANT

- Established connectivity @ Amsterdam in co-located PoP
- Performance measured using **perfSONAR** (throughput, latency, loss, routing)

VPN evaluation outcome

- Performance: not necessary the best option for maximum throughput
- Management: overhead due to additional configuration
- Security: VPN implies full access to CERN from outside
  - On the contrary w/o VPN going through the CERN firewall. No access to full CERN from outside. Only ports open and to be secured are for Puppet CA.

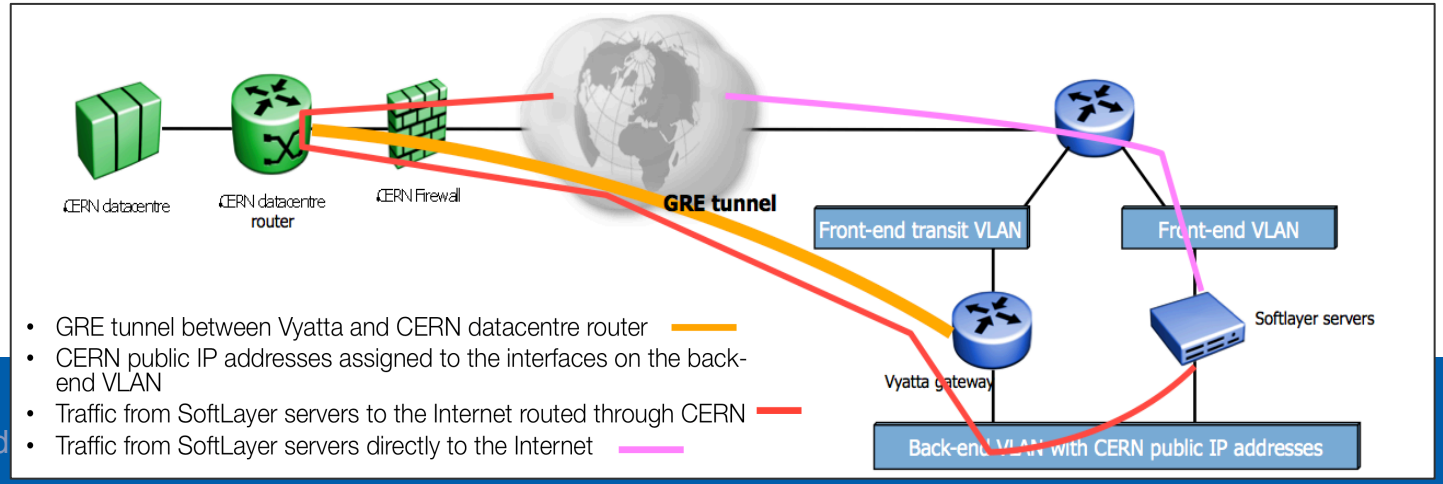

Cloud Storage

Cloud Storage in addition to Compute

Opportunity to study requirements and performance of "data-intensive" workloads, i.e. more than MC simulation

500 TB of Central Storage included in the last contract, awarded to T-Systems

- Block storage, managed via **DPM** (Tier-2 like model)
- Nodes registered in tsy.cern.ch subdomain
- Local DNS service needed for performance reasons with 1:1 NAT

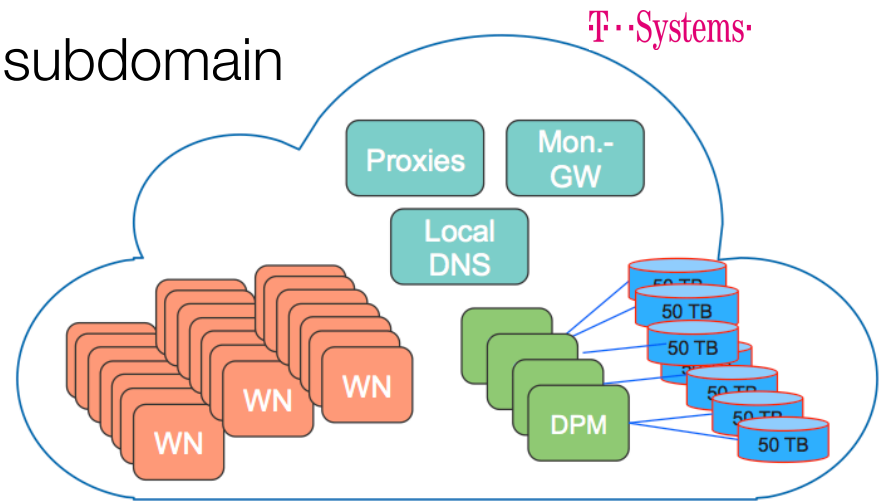

Storage Vs Network
"Block" vs "Object" Storage

Two open questions

1. Is storage (mainly a cache) needed in public cloud?
   - Or is it better to buy more WAN bandwidth?
   - Different requirements from the VOs:
     - Disk-less resources (LHCb)
     - Cache for Minimum Bias (ATLAS and initially CMS the before pre-mixed PU approach)
     - Tier-2 like storage for Analysis (ALICE)

2. Status of adoption of Object Storage in WLCG workloads
   - CSPs prefer to leverage Object Storage
     - More scalable, reliable and cost effective than block storage

Procurement Actions

CERN Procurements

Procurement rules apply

- No credit card to buy cloud services
- Request for Tender with upfront detailed Technical Specifications
- Award to the lowest compliant bid

Compare quality & cost

- Requiring firms to execute benchmarks during the tender procedure
- Adopt CPU benchmarks representative of the experiments' workloads *[see CHEP p-28]*

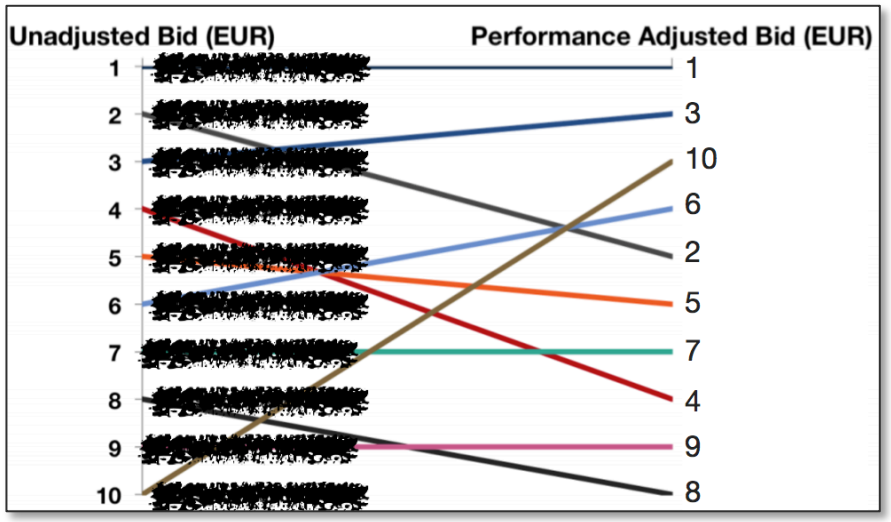

Ranking changes significantly when price is adjusted for expected performance (1 is cheapest)

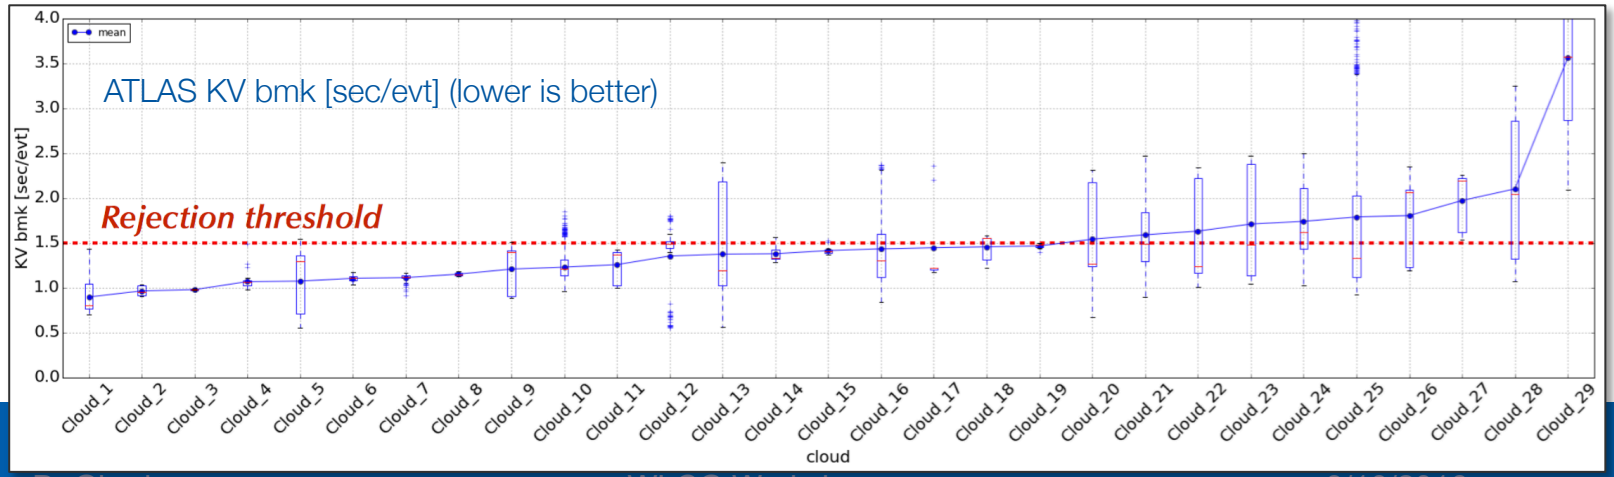

Cost Evaluation

Main focus on the procurement process and features of cloud deployments: functionality, stability, maintainability, …

- Prefer duration (several weeks) to high peak capacity

Side effect: economy of scale not reached

- Shown by the large cost spread per CSP and per service (compute, storage, network, etc)

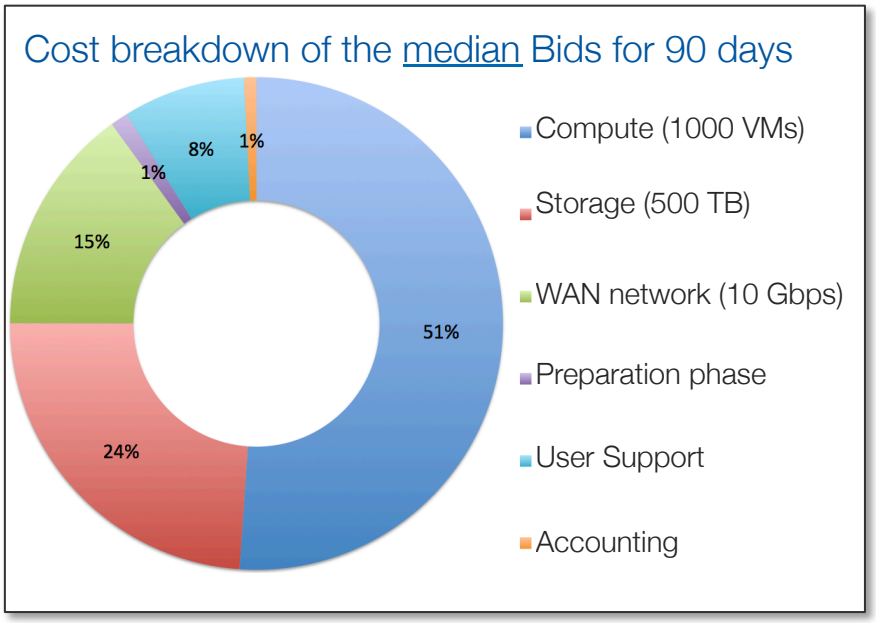

Median Price for compute

- 0.02 CHF/core/hour
- Including RAM (2 GB/core) and Disk (25 GB/core)

Storage (500 TB) represents 24% of the cost

- If not needed can invest more in network

HNSciCloud Pre-Commercial Procurement

- 5.3 MEur procurement of R&D services *[see CHEP o-397,o-399, p-401]*
- Build a Hybrid Cloud platform for the European research community
  - Combining services at IaaS to be integrated with in-house resources
- HNSciCloud challenges, triggering some of the next questions:

- ***Compute and Storage***
  - support a range of virtual machine and container configurations working with datasets in the petabyte range
- ***Network Connectivity and Federated Identity Management***
  - provide high-end network capacity for the whole platform with common identity and access management
- ***Service Payment Models***
  - explore a range of purchasing options to determine the most appropriate ones for the scientific application workloads that will be deployed

www.cern.ch

DNS setup

VMs are generated with private IP addresses

- Public Elastic IP bound to the VM via 1:1 NAT
- VMs assigned to CERN subdomain tsy.cern.ch

NAT

- Different behavior for internal/external clients
- Requires cloud-local DNS

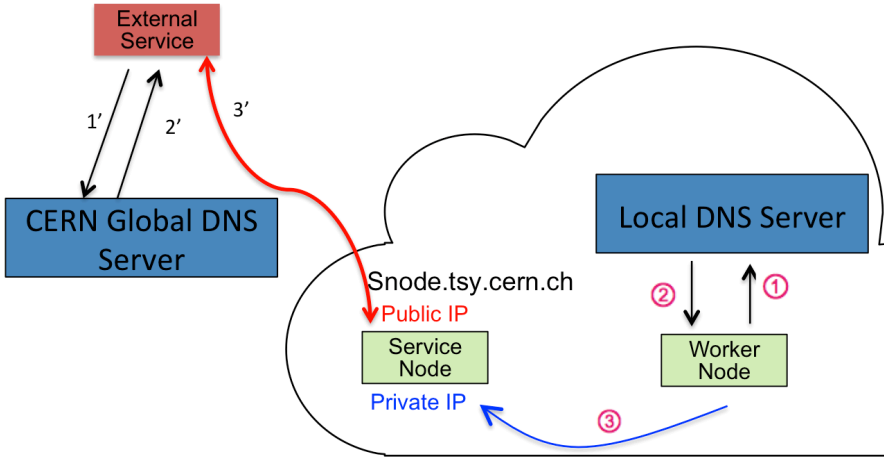

DNS reverse lookups

- Grid clients need this for both public and private IPs
- T-Systems, as authoritative entity for reverse DNS, has implemented the Global DNS reverse function for the CERN-OTC service nodes