Network experiences with Public Cloud Services

TNC2017

30th of May 2017

edoardo.martelli@cern.ch

Content

- CERN experiences with Public Clouds
- HNSciCloud
- RENs’ connectivity solutions

CERN Rationale for Public Clouds

- Additional, temporary resources to fit burst of computing demand
- Potentially replacing on-premise capacity

Potential benefits:
- Economy of scale
- Elastic, adapts to changing demands
- Delegation of hardware procurement and operations

Potential challenges:
- Integration with on-premise resources (orchestration)
- Compute-heavy tasks relatively new to cloud market (doesn't fit overbooking)
- Cloud provider's business models vs. procurement rules and procedures of public organizations

CERN past experiences with Public Clouds

CERN approach to procurement

Series of short procurement projects of increasing size and complexity

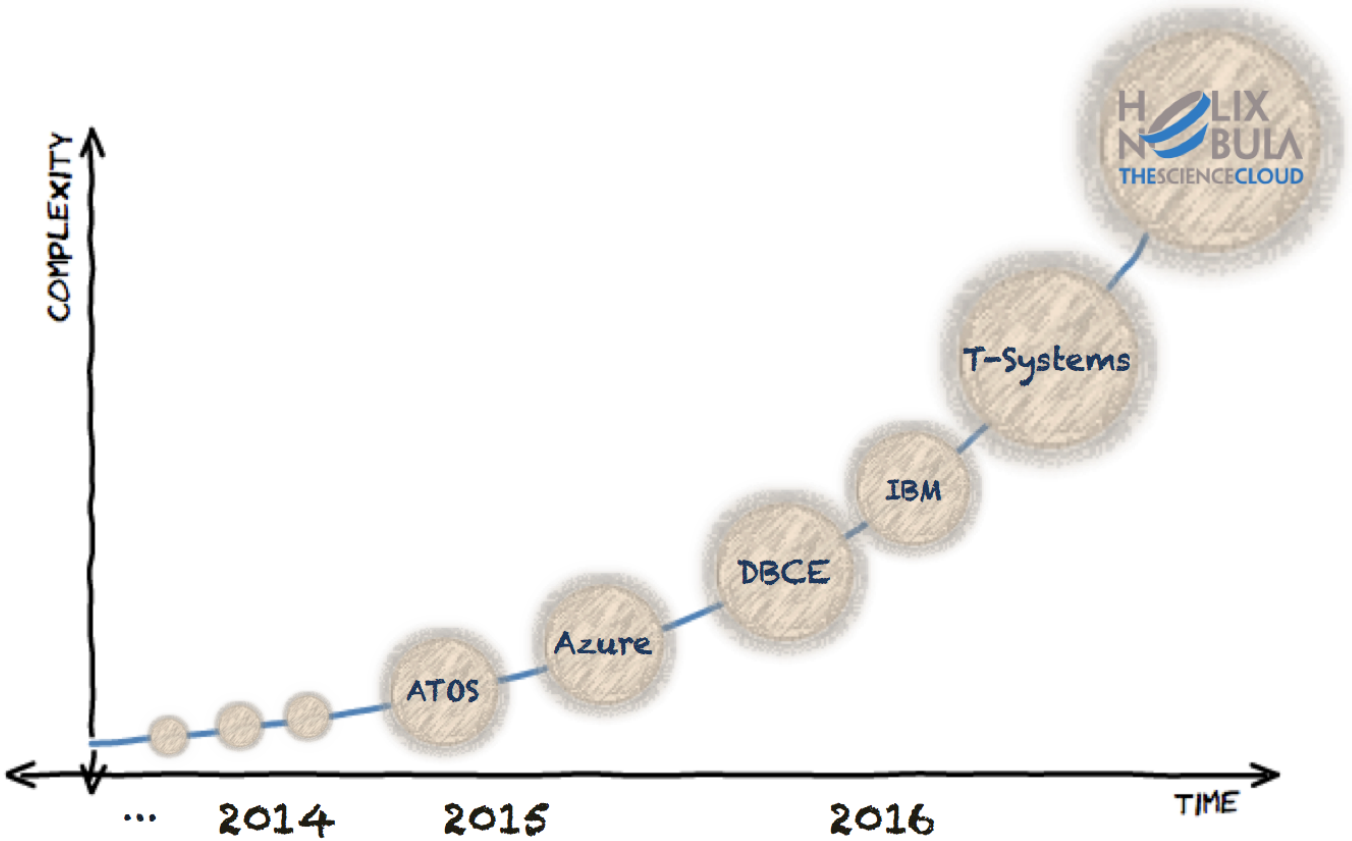

Last procurement (2016)

CERN awarded a contract to T-Systems for resources to be used for physics data processing

Size of procured resources for 3 months (+ 1 month of data extraction)
- 1’000 simultaneous VMs (4’000 cores),
- RAM: 8 GB/VM, Local Disk 100 GB/VM
- 500 TB of Central Storage
- 10Gbps connectivity to GEANT
- 1’000 public IPv4 addresses (resources accessible from all WLCG)

WAN connectivity

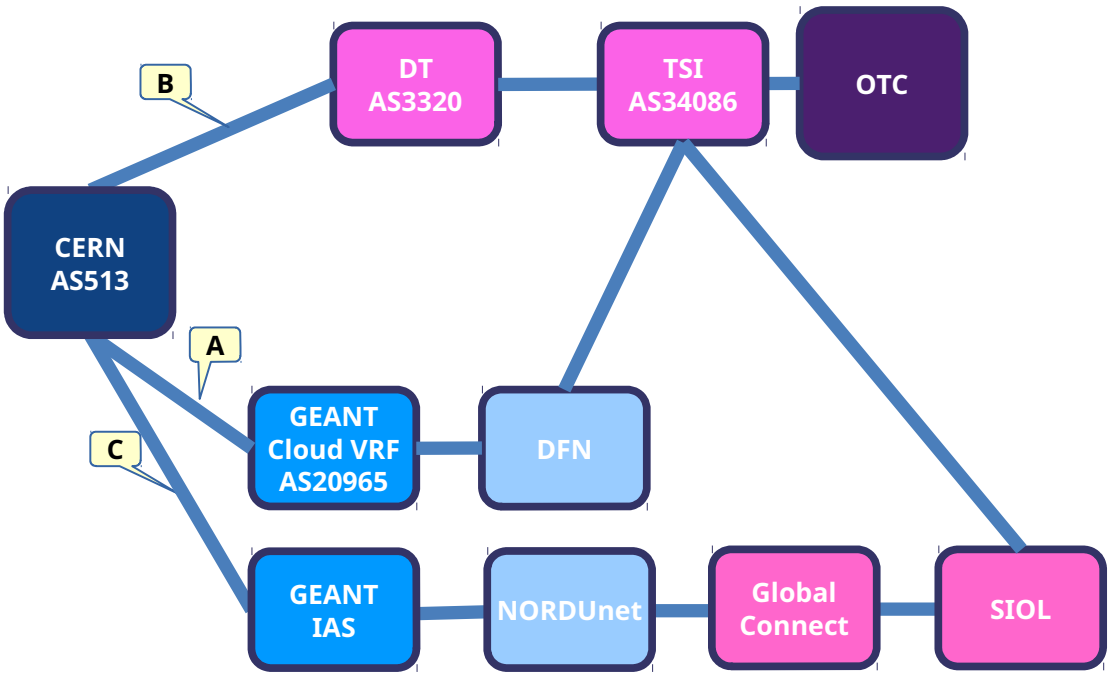

Three options available:

a) Peering with GEANT Cloud VRF through DFN

b) Existing 7Gbps Internet upstream

c) Existing alternative path via GEANT IAS and NORDUnet

Performance over GEANT Cloud VRF

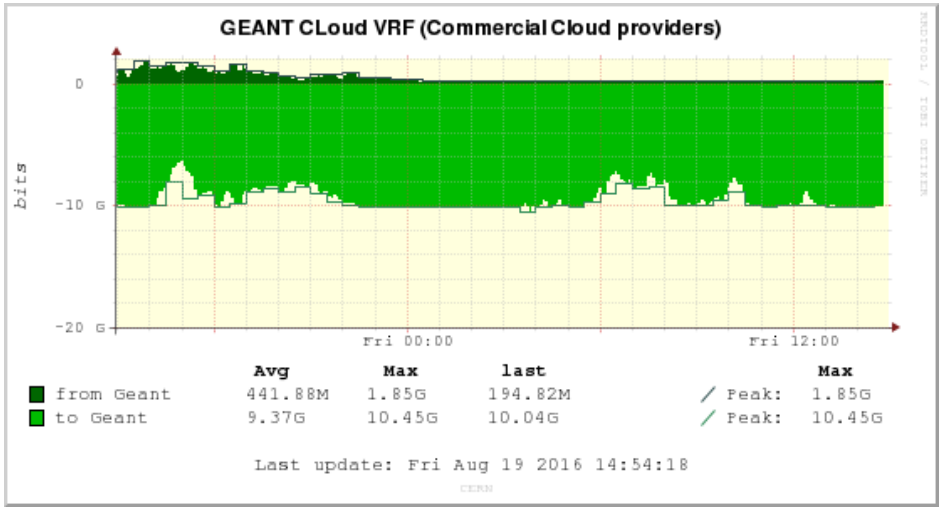

On a day

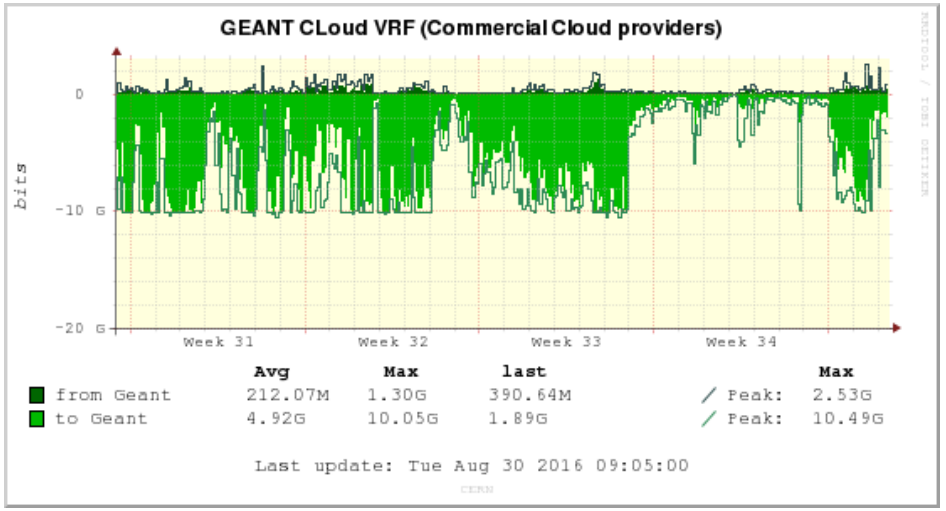

On a month

Network throughput vs number of jobs

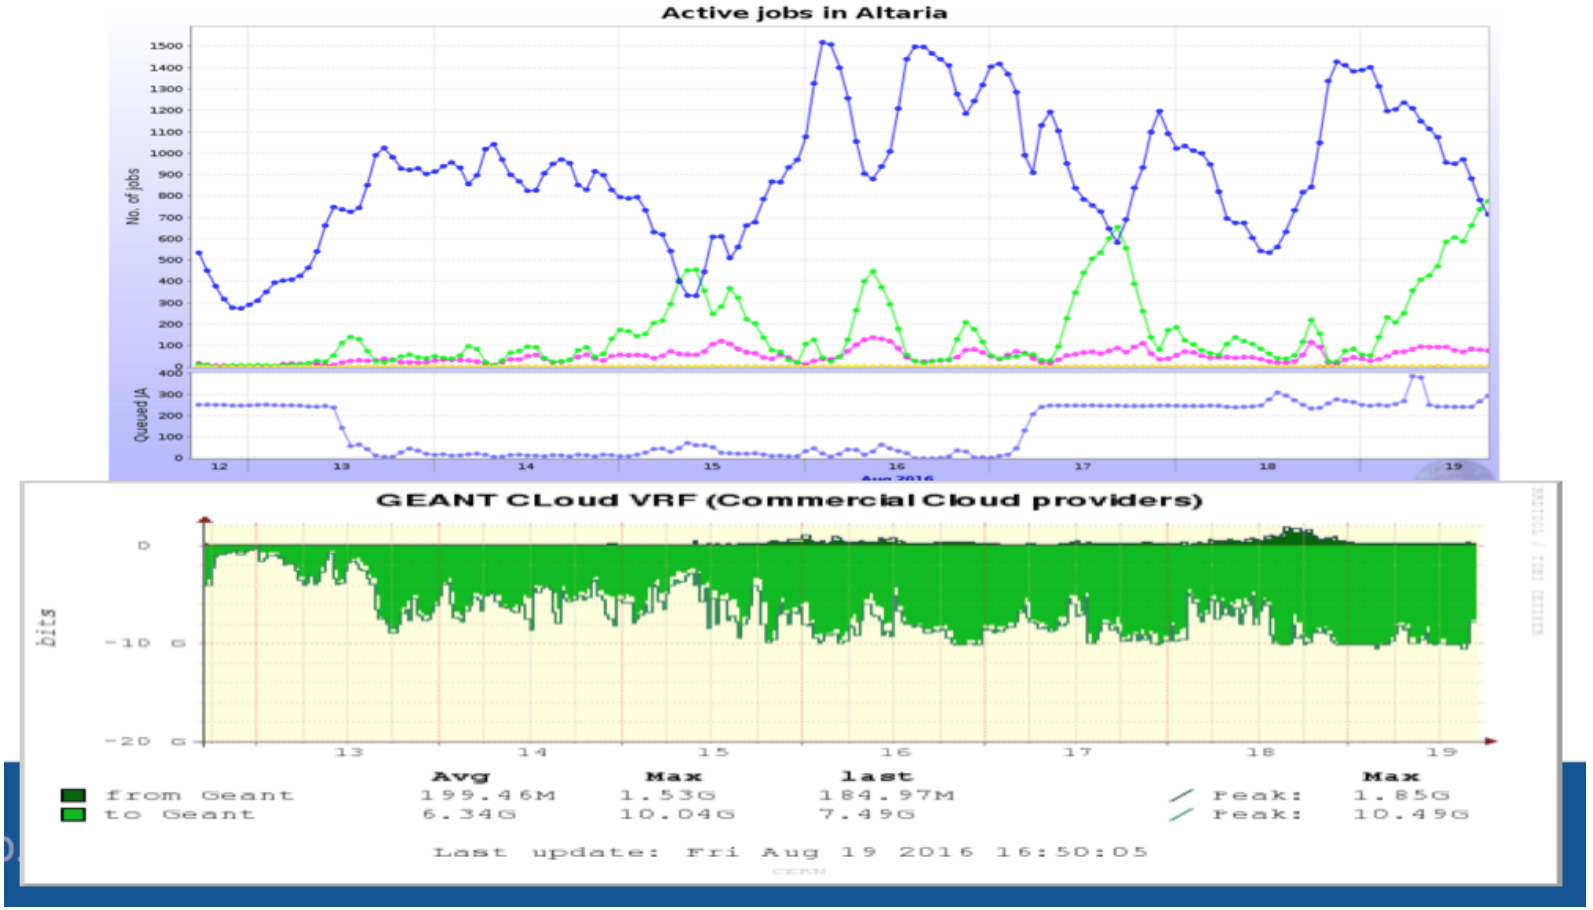

Paths from WLCG Tier1s

Most of the sites would send data through their Commercial Internet upstream (or those of their NRENs).

Only few NRENs had peerings with T-Systems at IXPs

```
CA-TRIUMF->     BCnet, HE, DECIX, T-Systems
CH-CERN->       GEANT-VRF, DFN, T-Systems
DE-KIT->        Belwue, T-Systems
ES-PIC->        CESCA, REDIRIS, NTT,  Netuse, T-Systems
FR-CCIN2P3->    Renater, C&W,  DTAG, T-Systems [from Renater LG]
IT-INFN-CNAF->  Cogent, NTT,  Netuse, T-Systems
NDGF->          NORDUnet, DECIX, Global-Connect, Netuse, T-Systems [from NORDUnet LG]
NL-T1->         SURFnet, NORDUnet, DECIX, Global-Connect, Netuse, T-Systems
RRC-KI-T1->     RETN, Netuse, T-Systems
TW-ASGC->       Chunghwa, Sprint, DTAG, T-Systems
UK-T1-RAL->     Janet, Telia, Netuse, T-Systems
US-FNAL-CMS->   ESnet, DTAG, T-Systems
US-T1-BNL->     ESnet, DTAG, T-Systems
```

Helix Nebula Science Cloud

HELIX NEBULA the Science Cloud

Procurers: CERN, CNRS, DESY, EMBL-EBI, ESRF, IFAE, INFN, KIT, SURFSara, STFC
- Procurers have committed funds (>1.6M€), manpower, use-cases with applications & data, in-house IT resources

Objective: procure innovative IaaS level cloud services
- Fully and seamlessly integrating commercial cloud (Iaas) resources with in-house resources and European e-Infrastructures
- To form a hybrid cloud platform for science

Services will be made available to end-users from many research communities: High-energy physics, astronomy, life sciences, neutron/photon sciences, long tail of science

Co-funded via H2020 (Jan'16-Jun'18) as a Pre-Commercial Procurement (PCP) project: Grant Agreement 687614, total procurement volume: >5M€

HNSciCloud challenges

Innovative IaaS level cloud services integrated with procurers in-house resources and public e-infrastructure to support a range of scientific workloads

Compute and Storage
support a range of virtual machine and container configurations including HPC working with datasets in the petabyte range

Network Connectivity and Federated Identity Management
provide high-end network capacity via GEANT for the whole platform with common identity and access management

Service Payment Models
explore a range of purchasing options to determine those most appropriate for the scientific application workloads to be deployed

High Level Architecture of the Hybrid Cloud Platform including the R&D challenges

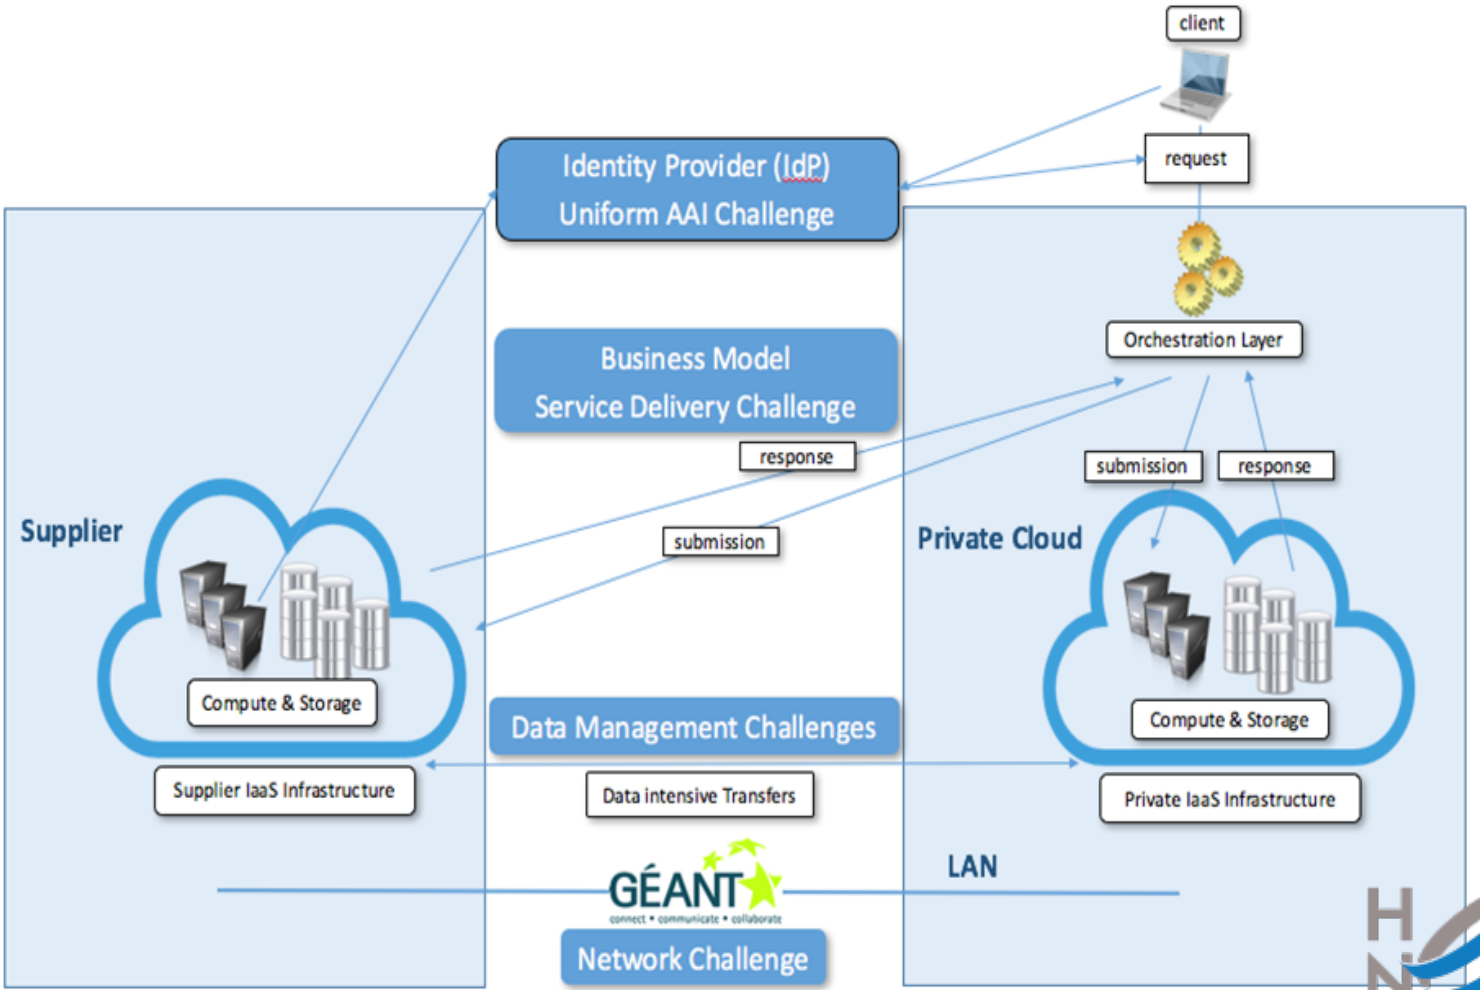

Network Requirements (LAN and WAN)

- Storage nodes globally accessible via public IP addresses (1:1 NAT accepted)
- Computing nodes can reside behind NAT
- At least 1Gbps capacity between pairs of local nodes
- At least 10Gbps of aggregated capacity between Supplier and Customer nodes
- WAN connectivity via one of the following
  - a. an NREN with transit to the customer;
  - b. GEANT or NORDUnet;
  - c. an IXP which the customer is also connected to
- Support of identity federation (eduGAIN) for IT managers

Project phases

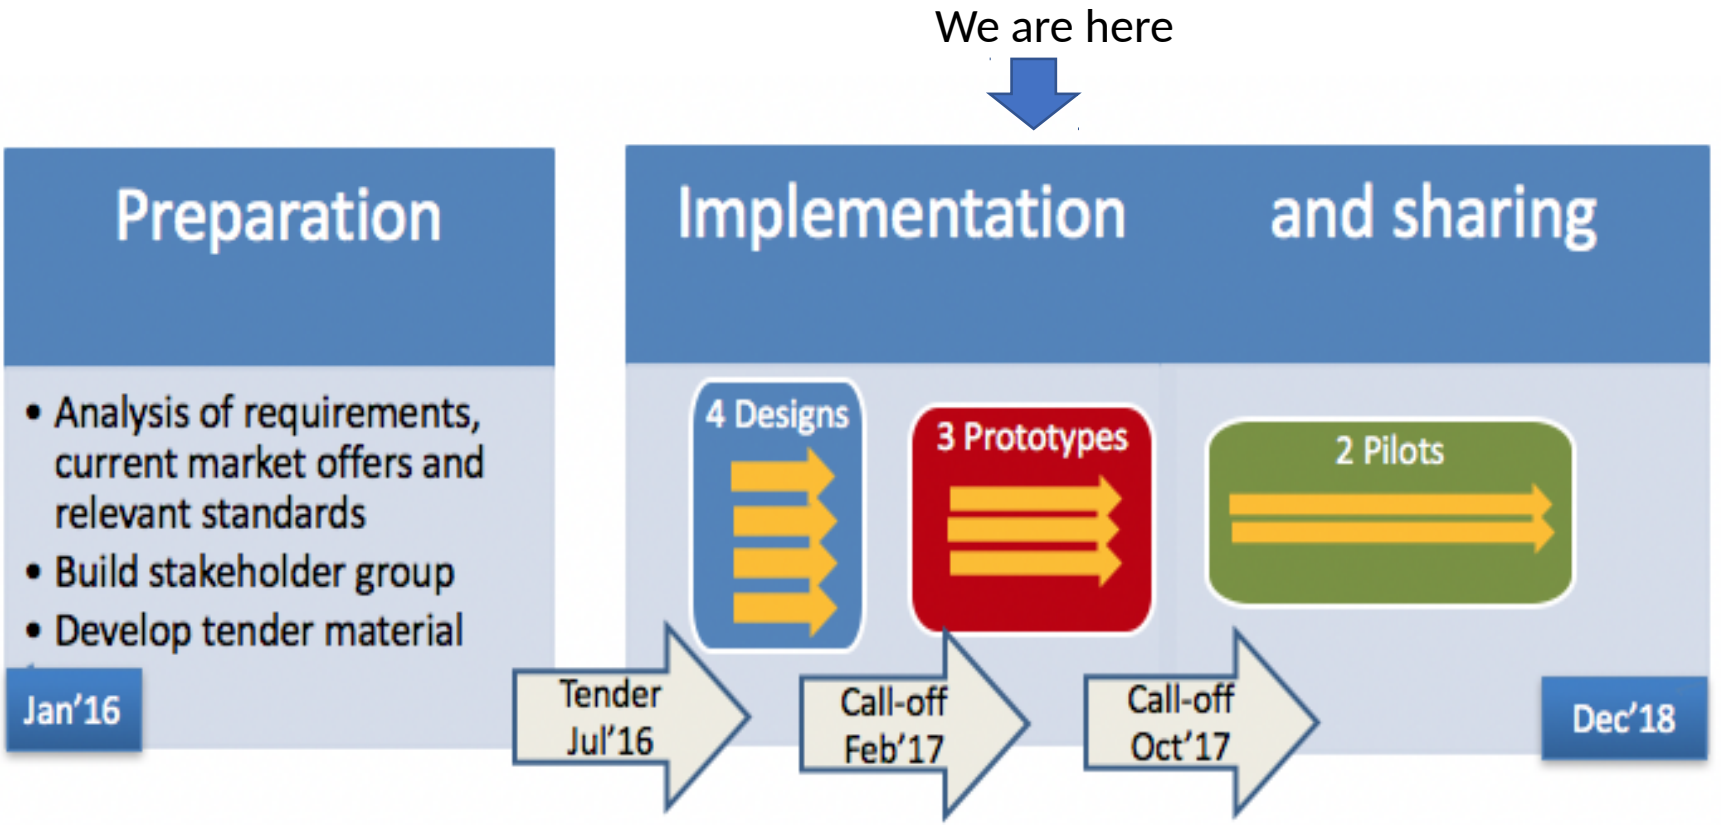

Each step is **competitive** - only contractors that successfully complete the previous step can bid in the next

Selected bidders

Admitted to the Design phase:

- T-Systems, Huawei, Cyfronet, Divia
- IBM
- RHEA Group, T-Systems, exoscale, SixSq
- Indra, HPE, Advania, SixSq

Other major players not interested or dropped out just before tender submission

Admitted to the Prototype phase:

- T-Systems, Huawei, Cyfronet, Divia
- IBM
- RHEA Group, T-Systems, exoscale, SixSq

Current Status: Evaluation of bids in the Prototype Phase

Prototype phase

Collect a number of tests from the Procurers
- to test/validate the design proposed for the several project challenges
- to be able to assess results of the tests as one of the main criteria to be eligible to the next phase (Pilot)

Tests consist in applications and some selected workloads

Next phase (Pilot) will consist in the assessment of the prototypes proposed on:
- performance
- scalability
- security

RENs’ connectivity solutions

R&E connectivity to Public Clouds

Challenges

- Bridge R&N and Commercial Networks with enough bandwidth
- Reduce cost of data transfers
- RENs policies limiting transport of “commercial” traffic (i.e. how to move R&E data between two cloud storage centres)

GEANT: dedicated VRF

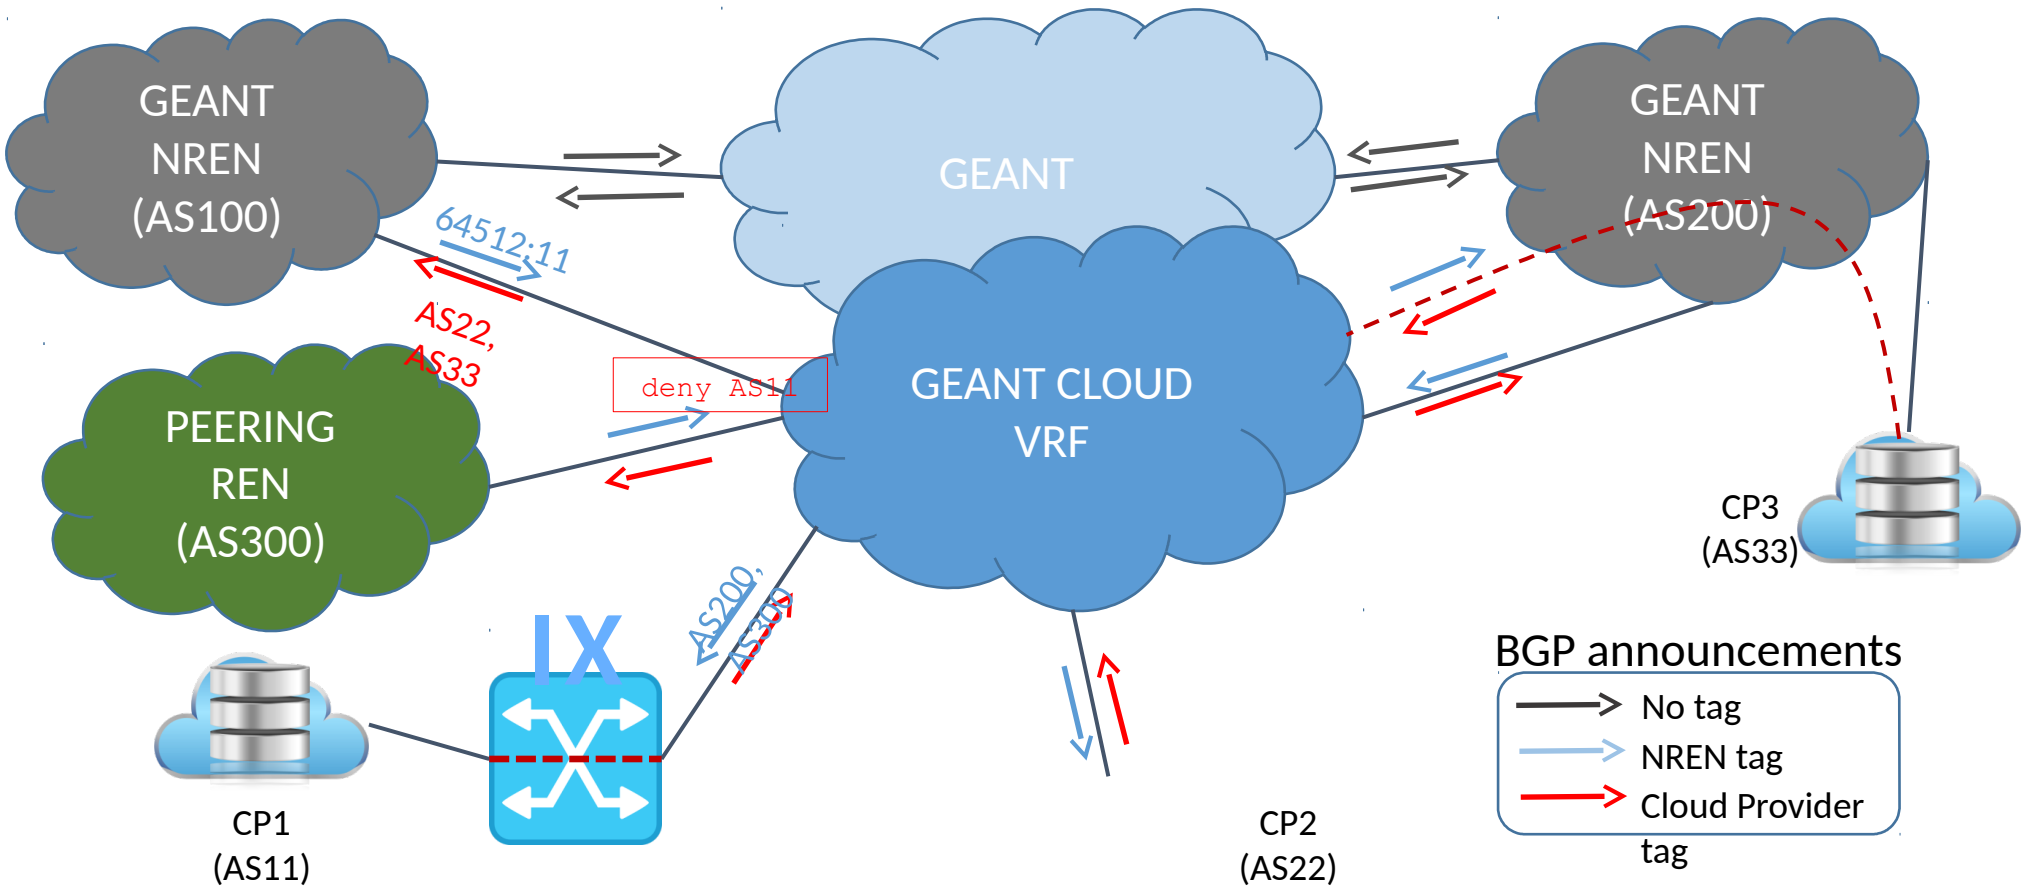

Slide credit: GEANT

ESnet: on-net VPN termination

Virtual “Site Router” (VRTR) Service

At the edge of the cloud

Virtual Site Router at AWS Exchange Point

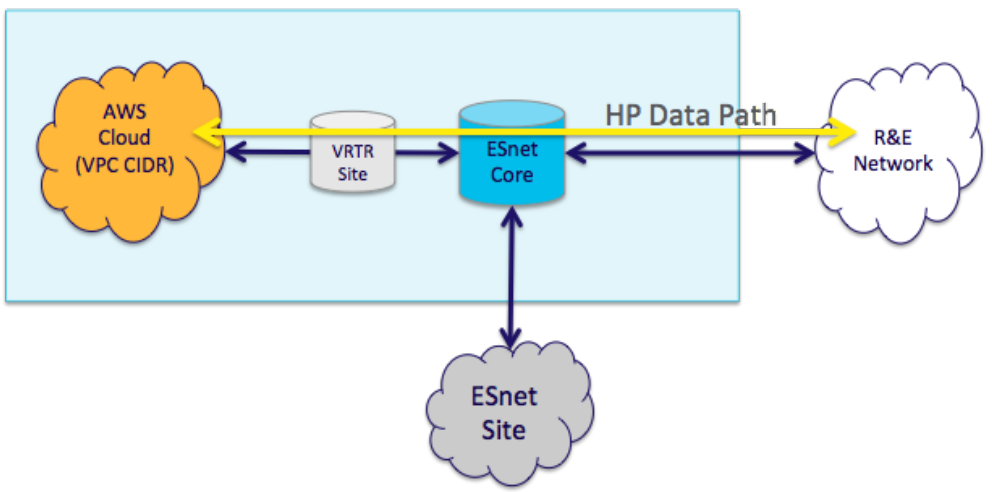

Virtual “Site Router” improves path efficiency and takes pressure off of the site local-loop.

Slide credit: ESnet

NORDUnet: transport to eXchange Points

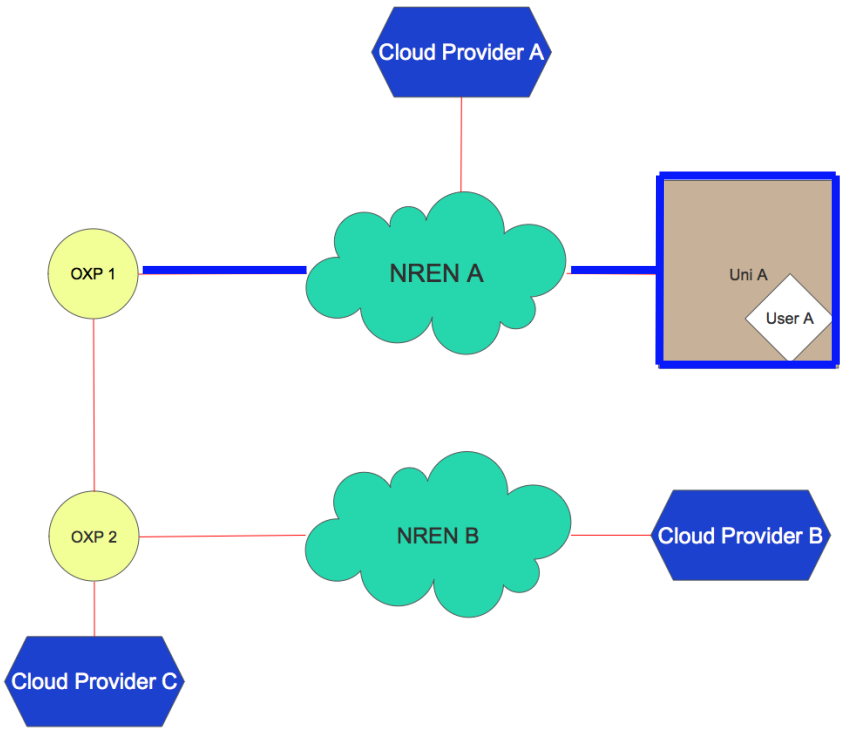

Slide credit: NORDUnet

Best practices for CSP connectivity

NORDUnet presented a document which suggest a way of connecting CSPs and their customers by using Open eXchange Points and connectivity provided by the REN

The document aims to simplify the procurement of the connectivity to CSPs

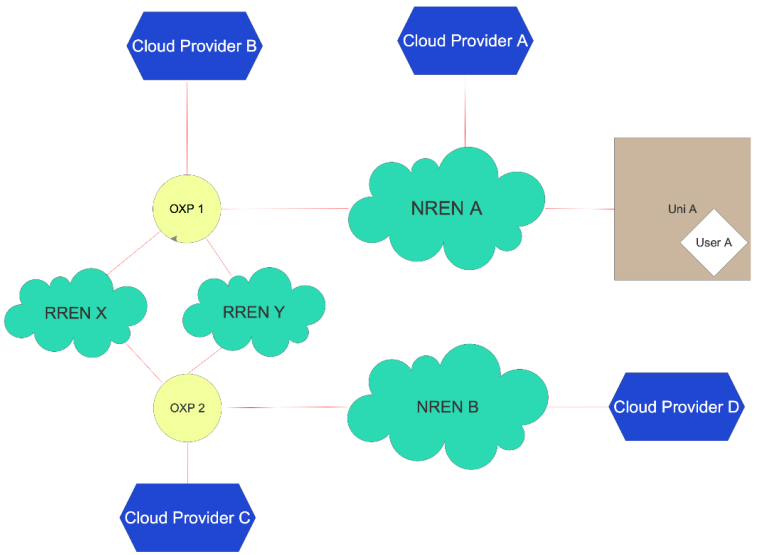

Network connectivity options for CSPs

CERN IT-CS is writing a document which compare the different connectivity options proposed by the RENs

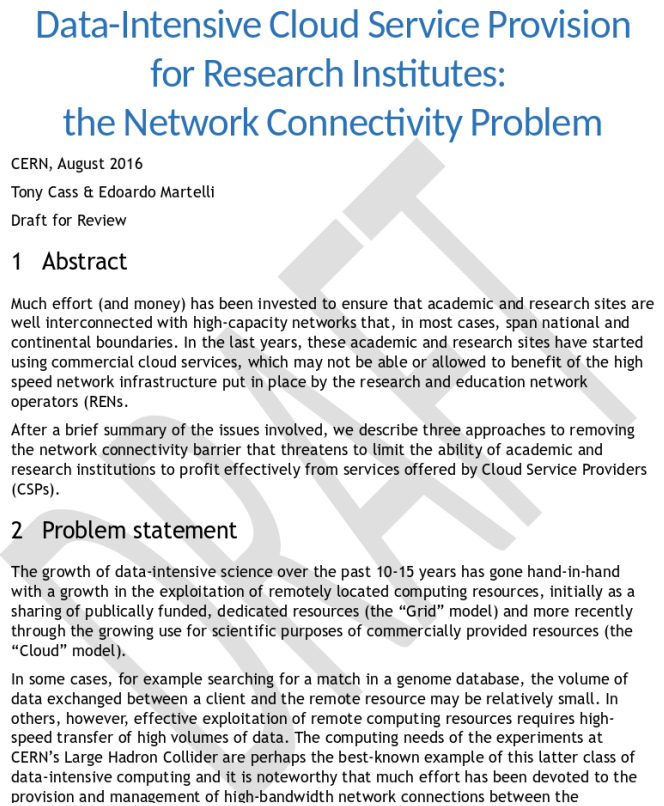

Data-Intensive Cloud Service Provision for Research Institutes: the Network Connectivity Problem

CERN, August 2016

Tony Cass & Edoardo Martelli

Draft for Review

1 Abstract

Much effort (and money) has been invested to ensure that academic and research sites are well interconnected with high-capacity networks that, in most cases, span national and continental boundaries. In the last years, these academic and research sites have started using commercial cloud services, which may not be able or allowed to benefit of the high speed network infrastructure put in place by the research and education network operators (RENs.

After a brief summary of the issues involved, we describe three approaches to removing the network connectivity barrier that threatens to limit the ability of academic and research institutions to profit effectively from services offered by Cloud Service Providers (CSPs).

2 Problem statement

The growth of data-intensive science over the past 10-15 years has gone hand-in-hand with a growth in the exploitation of remotely located computing resources, initially as a sharing of publically funded, dedicated resources (the “Grid” model) and more recently through the growing use for scientific purposes of commercially provided resources (the “Cloud” model).

In some cases, for example searching for a match in a genome database, the volume of data exchanged between a client and the remote resource may be relatively small. In others, however, effective exploitation of remote computing resources requires high-speed transfer of high volumes of data. The computing needs of the experiments at CERN’s Large Hadron Collider are perhaps the best-known example of this latter class of data-intensive computing and it is noteworthy that much effort has been devoted to the provision and management of high-bandwidth network connections between the

Document https://indico.cern.ch/event/527372/contributions/2236895/subcontributions/208050/attachments/1338702/2015050/connectivity-options-for-clouds-draft.pdf

Conclusions

Summary

CERN had positive preliminary tests with Commercial Cloud Providers. Ad-hoc connectivity through GEANT gave much better throughput

HNSciCloud has reached the Prototype phase. Parties are progressing on connectivity through GEANT and NRENs

RENs are working with institutes and Cloud providers to best fit connectivity needs

Questions?

edoardo.martelli@cern.ch