

# Technical challenges of the HNSciCloud H2020 project

more precise – the expected technical challenges

Martin Gasthuber DESY/IT  
Helge Meinhard CERN/IT  
CHEP16, San Francisco

- HNSciCloud – Intro just done before by Helge Meinhard
  
- unlike classical EU projects, we are not the carrier of the R&D activities
  - we ‚just‘ specify, audit and finally do the assessment of what has been developed (and run) by commercial providers – a little different to the classical rules of the game.
  - put *private* money on the table – enlarged by EU funding – forming the *buyers group*
  - total volume is about 5.3M€ (for all 3 phases) – at least 50% of that being spent on R&D activities

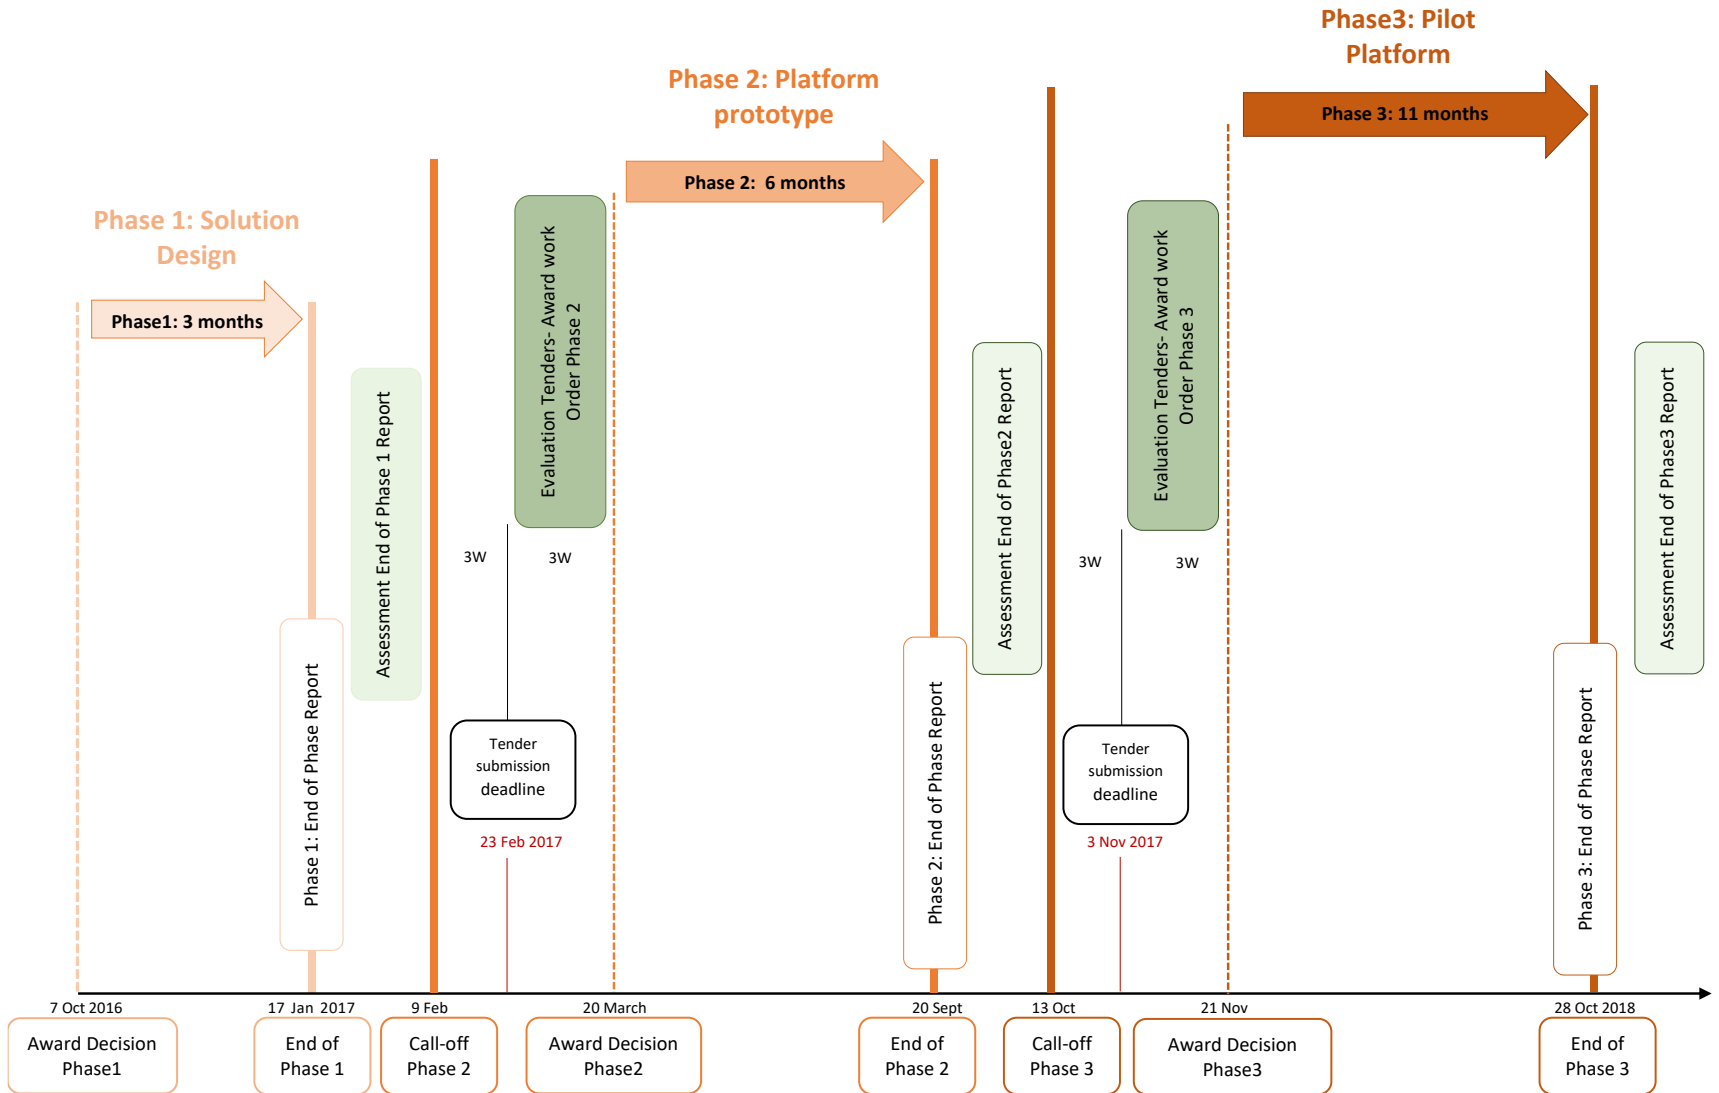
## > challenges outlined in the call – what the buyers see as challenges

- data access – no change of local application (near POSIX expected)
- authentication, identity federation with eduGAIN (SAML2)
- container support – not only for deploying services
- hybrid – N x buyers group + M x public
- HPCaaS – small HPC use-cases (mostly photon science)

## > challenges from .... instinct

- WAN – bulk data transfer - SDN, GEANT link,...
- operations – hybrid operations
- over commitment – hyperthreading – (real) mem per core
- container environment – local development & testing, remote run @massive scale

# fixed timeline & phases



## > cultural clash

- between buyers communities, between buyer and commercial providers
- nomenclature – find common wording
- efficiency in communication (i.e. call handling) – time is money

## > none of the listed technical challenges are technology challenges

- technology exists, known to run in dedicated environment, ...
- make is usable – at large scale, on and off-prem, independent of off-prem provider

# storage & data access

- > supplier would like to leverage object stores
  - best known limitations
- > primary copy of (all) data stays on buyers side (implicit assumption by all buyers)
  - legal, supplier lock-in, ... a lot of nervous feeling
- > no change of user application, POSIX like, auto managed (not self managed)
  - POSIX like – full semantic not required
  - way back into buyers storage – writes
  - namespace forwarding/connection – connect to buyers local system(s)
  - speed of local (off-prem) data access
  - direct remote (to buyer) access + pre fetching
  - auto managed – no end user involvement

- > hybrid scale – scale out on-prem and off-prem
- > separation of responsibilities and work – beside the well known benefits
  - scientist – standard/clear context for development & testing – no need to start with the stuff you've installed on your laptop
  - local admin – focus on local scale and shipping containers off-prem, container definition for scientists, control scale-out to off-prem resources
  - off-prem admin – just scale as large as possible with lowest costs, hidden to user/scientists

## > connection of different worlds

- GEANT to Cloud-Providers networks
- Cloud-Provider X to Cloud-Provider Y
- trans-ocean – i.e. LHCONE – not directly asked for

## > VPN usage and risks

## > identity federation (eduGAIN)

- admins & end-user (scientist)
- not limited to WEB applications (CLI mode)

## > HPCaaS

- driven by photon science community – small and (few) large HPC jobs expected – some are just multi-threaded, some using MPI
- low experience on buyers side – probably mirrored on supply side – high learning curve
- IPC network - probably the most limiting factor (10GE)

# resources to be expected

## > prototype (phase 2)

- ~3500 cores, ~350TB shared storage, 10Gbs GEANT link – per supplier for 3 months

## > pilot (phase 3)

- ~10K cores, 1PB shared storage, 40Gbs GEANT link – per supplier for 8 months

## > VM characteristics (minimum)

- 4-8 vCPU per VM – some use-cases could use more
- $\geq 1.875$  GiB or  $\geq 3.75$  GiB per vCPU
- ~30GB local VM storage
- Windows & Linux

# rough volume of R&D efforts possible

	Max Total Budget	Expected minimum number of Contractors	Maximum budget for each Contractor	Payments	Maximum fixed Duration
PCP budget	5.300.000				
Phase 1: Solution design	520.000	4	130.000	At the end of the Phase	3
Phase 2: Prototype Platform	1.325.000	3	441.667	Interim Payment 30% and final payment	6
Phase 3: Pilot Platform	3.455.000	2	1.727.500	Interim Payment 40% and final payment	11

> questions ?