Helix Nebula – The Science Cloud
Status Update

Helge Meinhard / CERN
With contributions by Martin Gasthuber / DESY
HEPiX Fall 2016 @LBNL – contribution # 34

Presentation at HEPiX Zeuthen: https://indico.cern.ch/event/466991/contributions/1143637/

Scale of LHC Data Tomorrow

Ian Bird, WLCG workshop Oct 2016

Scale of LHC Data Tomorrow

Short term: until end of Run 2

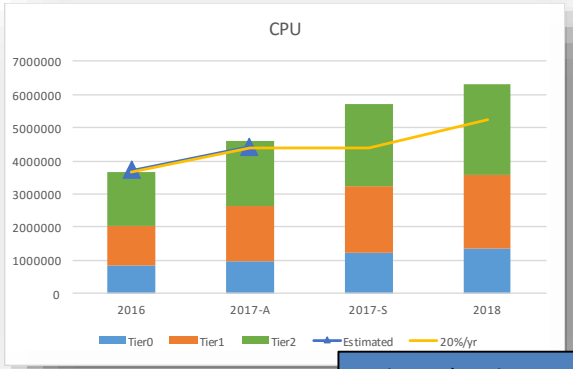

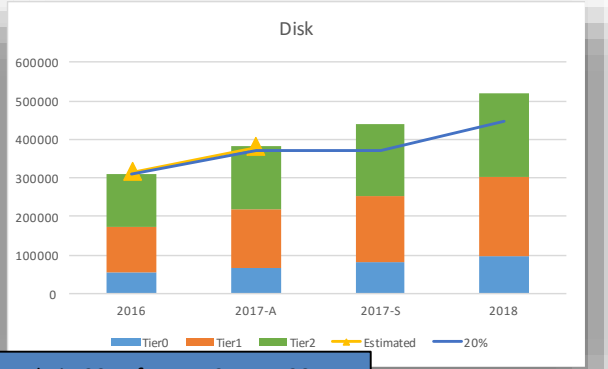

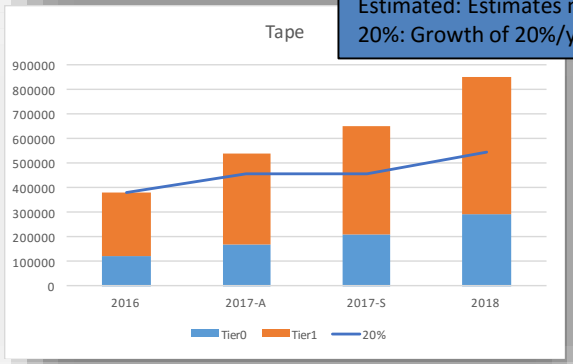

Estimated: Estimates made in 2014 for Run 2 up to 2017
20%: Growth of 20%/yr starting in 2016 ("flat budget")

Ian Bird, WLCG workshop Oct 2016

Scale of LHC Data Tomorrow

Short term: until end of Run 2

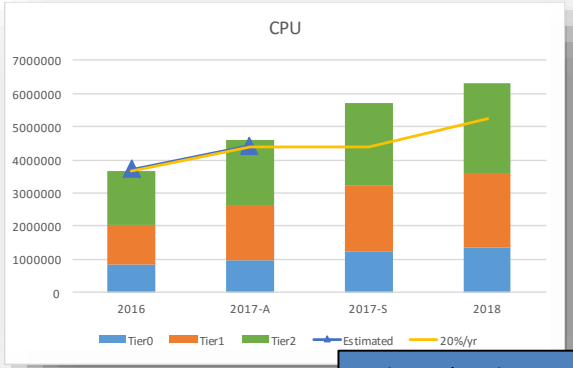

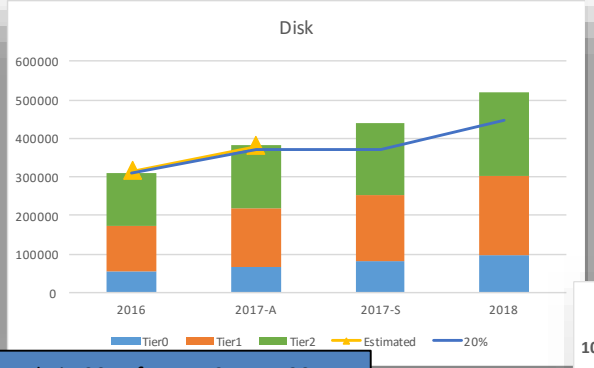

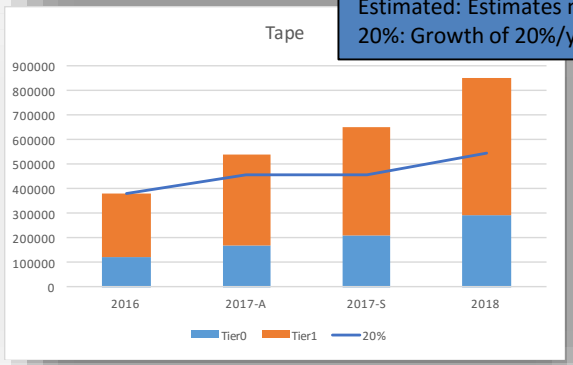

Estimated: Estimates made in 2014 for Run 2 up to 2017
20%: Growth of 20%/yr starting in 2016 ("flat budget")

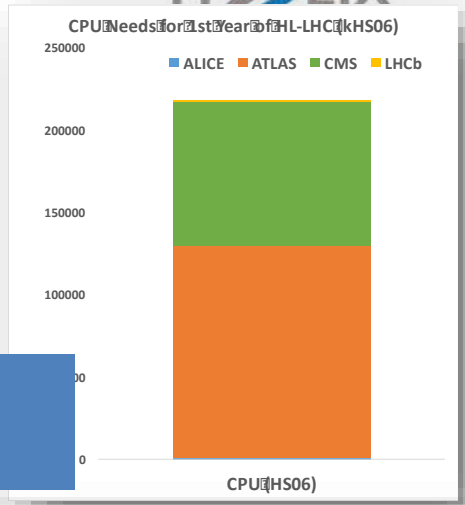

2026: HL-LHC

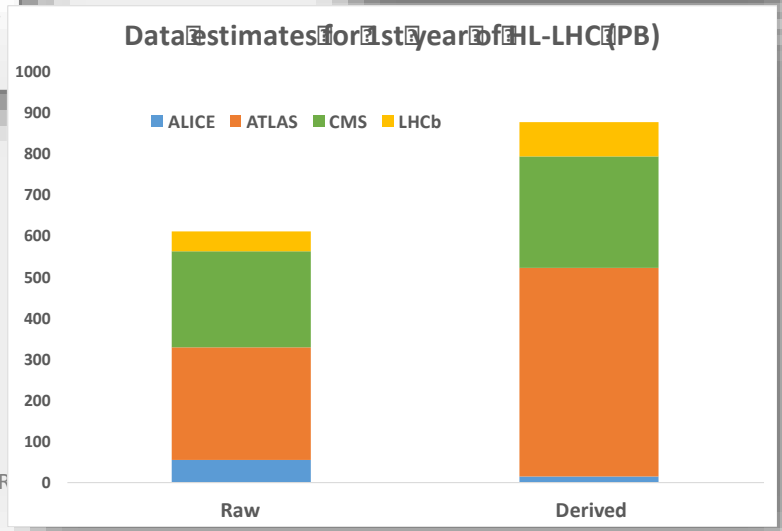

Ian Bird, WLCG workshop Oct 2016

Scale of LHC Data Tomorrow

Short term: until end of Run 2

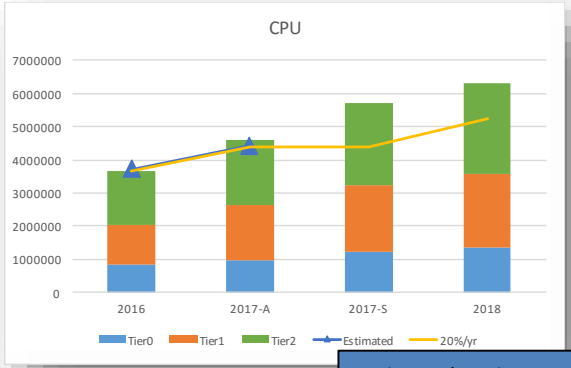

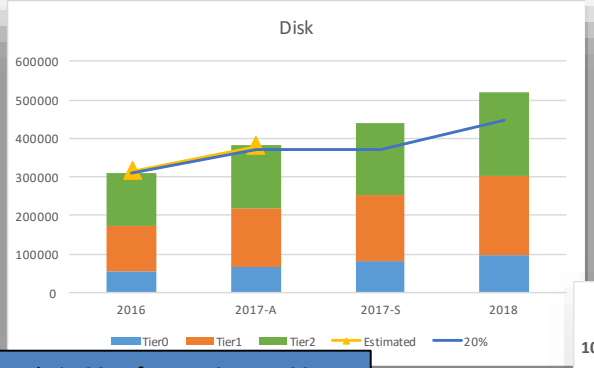

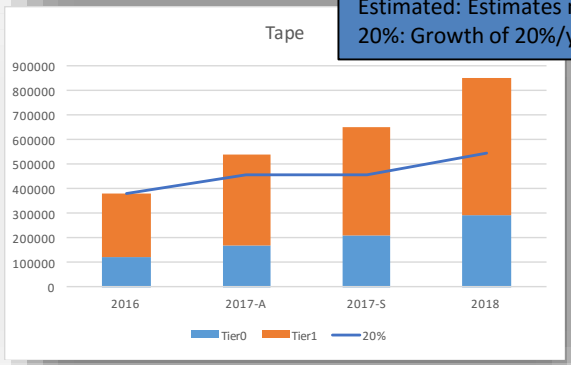

Estimated: Estimates made in 2014 for Run 2 up to 2017
20%: Growth of 20%/yr starting in 2016 ("flat budget")

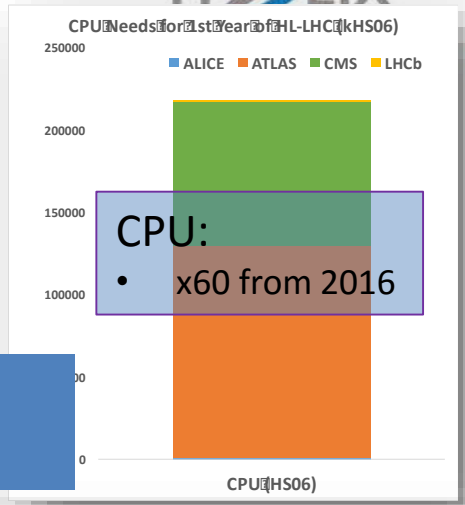

CPU:
- x60 from 2016

2026:
HL-LHC

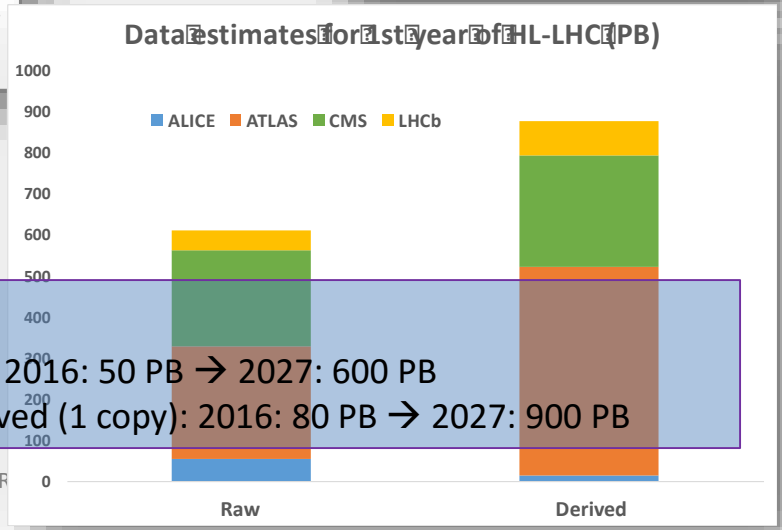

Data:
- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB

Ian Bird, WLCG workshop Oct 2016

Future Requirements

- Not only LHC, but a number of particle physics projects with high data rates
- Not only particle physics, but also other physics fields (e.g. astronomy)
- Not only physics, but also other sciences (e.g. life sciences, material science)
  - E.g. EBI: Data doubles every 12 months

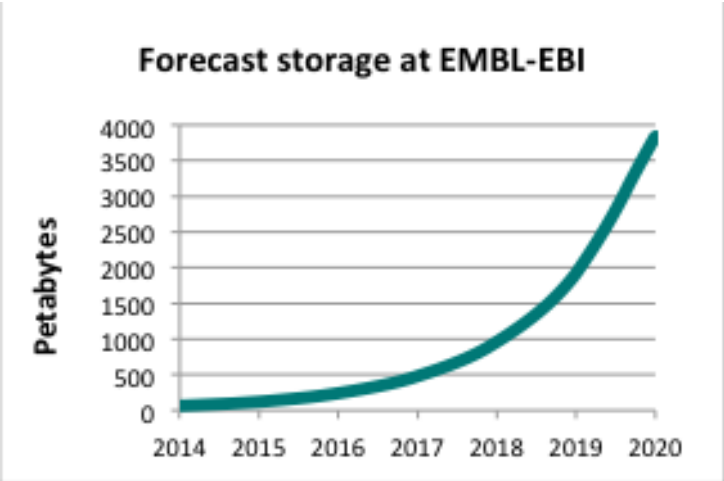

Scaling up: Public Clouds (1)

- Additional resources, perhaps later replacing on-premise capacity
- Potential benefits:
  - Economy of scale
  - More elastic, adapts to changing demands
  - Somebody else worries about machines and infrastructure

Scaling up Further: Public Clouds (2)

- Potential issues:
  - Cloud provider's business models not well adapted to procurement rules and procedures of public organisations
  - Lack of skills for and experience with procurements
  - Market largely not targeting compute-heavy tasks
    - Performance metrics/benchmarks not established
  - Legal impediments
  - Not integrated with on-premise resources and/or publicly funded e-infrastructures

HELIX NEBULA Science Cloud
Joint Pre-Commercial Procurement

Procurers: CERN, CNRS, DESY, EMBL-EBI, ESRF,
IFAE, INFN, KIT, SURFSara, STFC
Experts: Trust-IT & EGI.eu

The group of procurers have committed
- >1.6M€ of procurement funds
- Manpower for testing/evaluation
- Use-cases with applications & data
- In-house IT resources

To procure innovative IaaS level cloud services integrated into a hybrid cloud model
- Commercial cloud services
- European e-Infrastructures

Services will be made available to end-users from many research communities

Co-funded via H2020 (Jan'16-Jun'18)
- Grant Agreement 687614

Total procurement commitment >5M€

User groups to be supported

- High Energy Physics
  - LHC experiments
  - Belle II
  - COMPASS
- Astronomy
  - CTA – Cherenkov Telescope Array
  - MAGIC
  - Pierre Auger Observatory
- Life Sciences
  - ELIXIR
  - Euro-BioImaging
  - Pan-Cancer
  - BBMRI
  - WeNMR
- Photon/Neutron science
  - PETRA III, European XFEL, 3DIX, OCEAN, OSIRIS
- Long tail of science

The HNSciCloud Project

- Unlike “classical” EU projects, we are not the carrier of the R&D activities
- We “just” specify, audit and finally assess what has been developed (and run) by commercial providers
- Total volume is about 5.3 MEUR for the three phases
  - At least 50% to be spent on R&D in each phase

Technical Challenges

- Compute
  - Integration of some HPC requirements
  - Support for containers at massive scale
- Data access
  - Caching at provider's site, if possible automatically (avoid managed storage)
  - Transparent for application; POSIX subset
- Network
  - Connection via GÉANT
  - Support of identity federation (eduGAIN) for IT managers

Semi-Technical Challenges

- Procurement
  - Match of cloud providers' business model with public procurement rules
- Different cultures
  - Between buyers' communities
  - Between buyers and commercial providers
- Nomenclature – common wording required
- Communication – effective and efficient channels to be set up
- Challenge is to make technology usable – at large scale, on- and off-premises, independent of off-premise provider

Storage and Data Access

- Suppliers would like to leverage object stores
  - Best price point, best understood
- Implicit assumption: Primary copy of (all) data stays on buyers' premises
  - Reasons: legal, supplier lock-in, ... a lot of nervous feeling
- No change of user application, POSIX-like access, auto-managed
  - POSIX like – open, close, read, write, seek suffices
  - Way back into buyers' storage – writes
  - Namespace forwarding/connection – connect to buyers local system(s)
  - Speed of local (off-premise) data access
  - Direct remote (to buyer) access + pre fetching
  - Auto-managed – no end user involvement

Container Support

- Hybrid scale – scale out on-premise and off-premise
- 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-premise, container definition for scientists, control scale-out to off-premise resources
  - Off-premise admin – just scale as large as possible with lowest costs, hidden to user/scientists

Network, Identity Federation, HPCaaS

- Connection of different worlds
  - GÉANT to cloud provider networks
  - Cloud provider X to cloud provider Y
  - Trans-ocean – i.e. LHCONE – not directly asked for
- VPN usage and risks
- Identity federation (eduGAIN, SAML 2.0)
  - 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
  - Little experience on buyers' side – probably mirrored on supply side – steep learning curve
  - IPC network - probably the most limiting factor (10GE)

HNSciCloud Project Phases

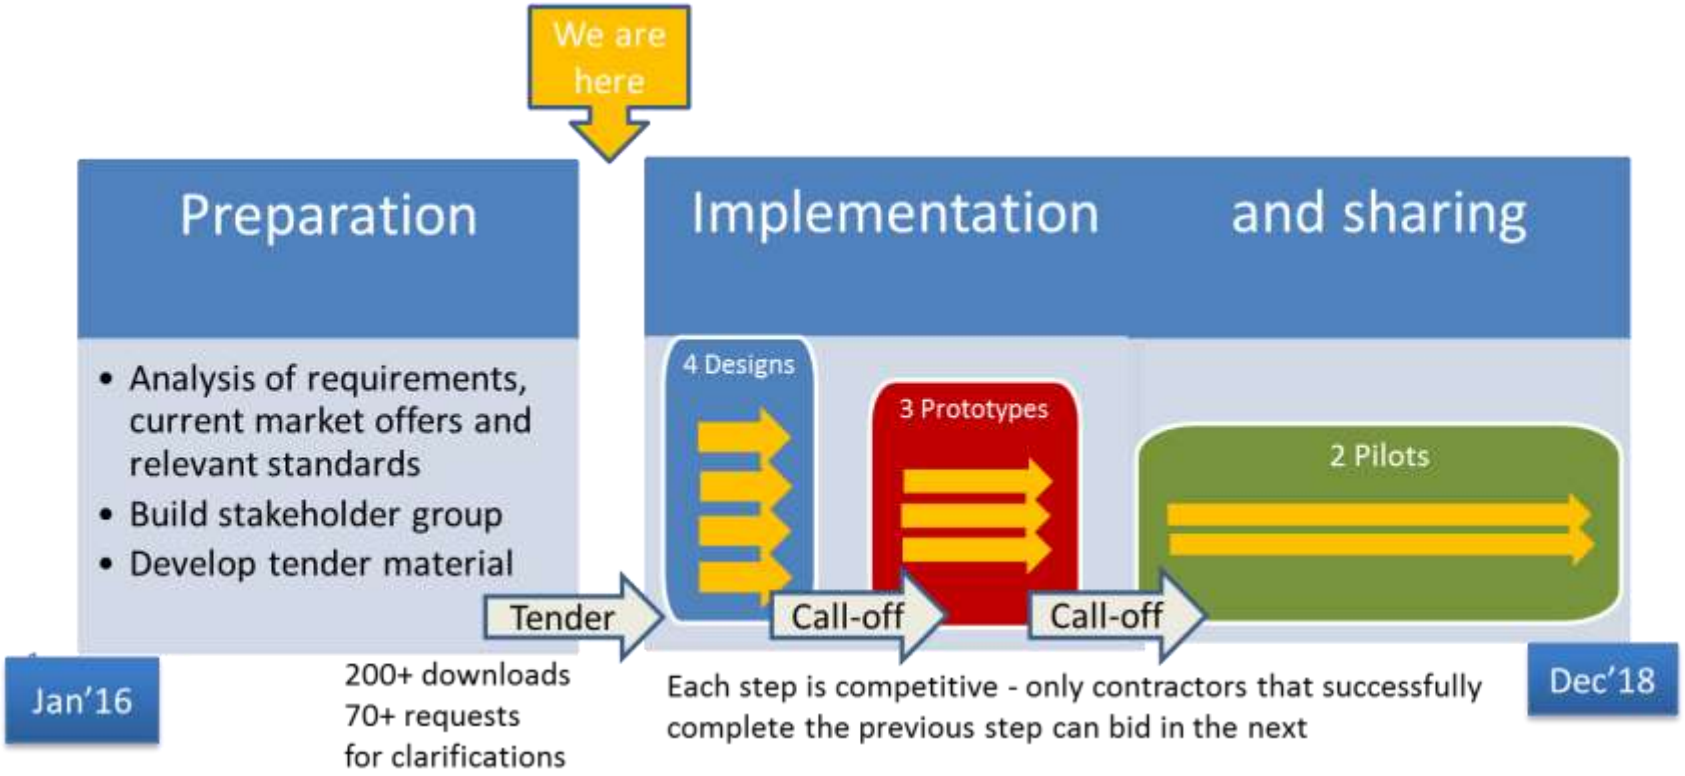

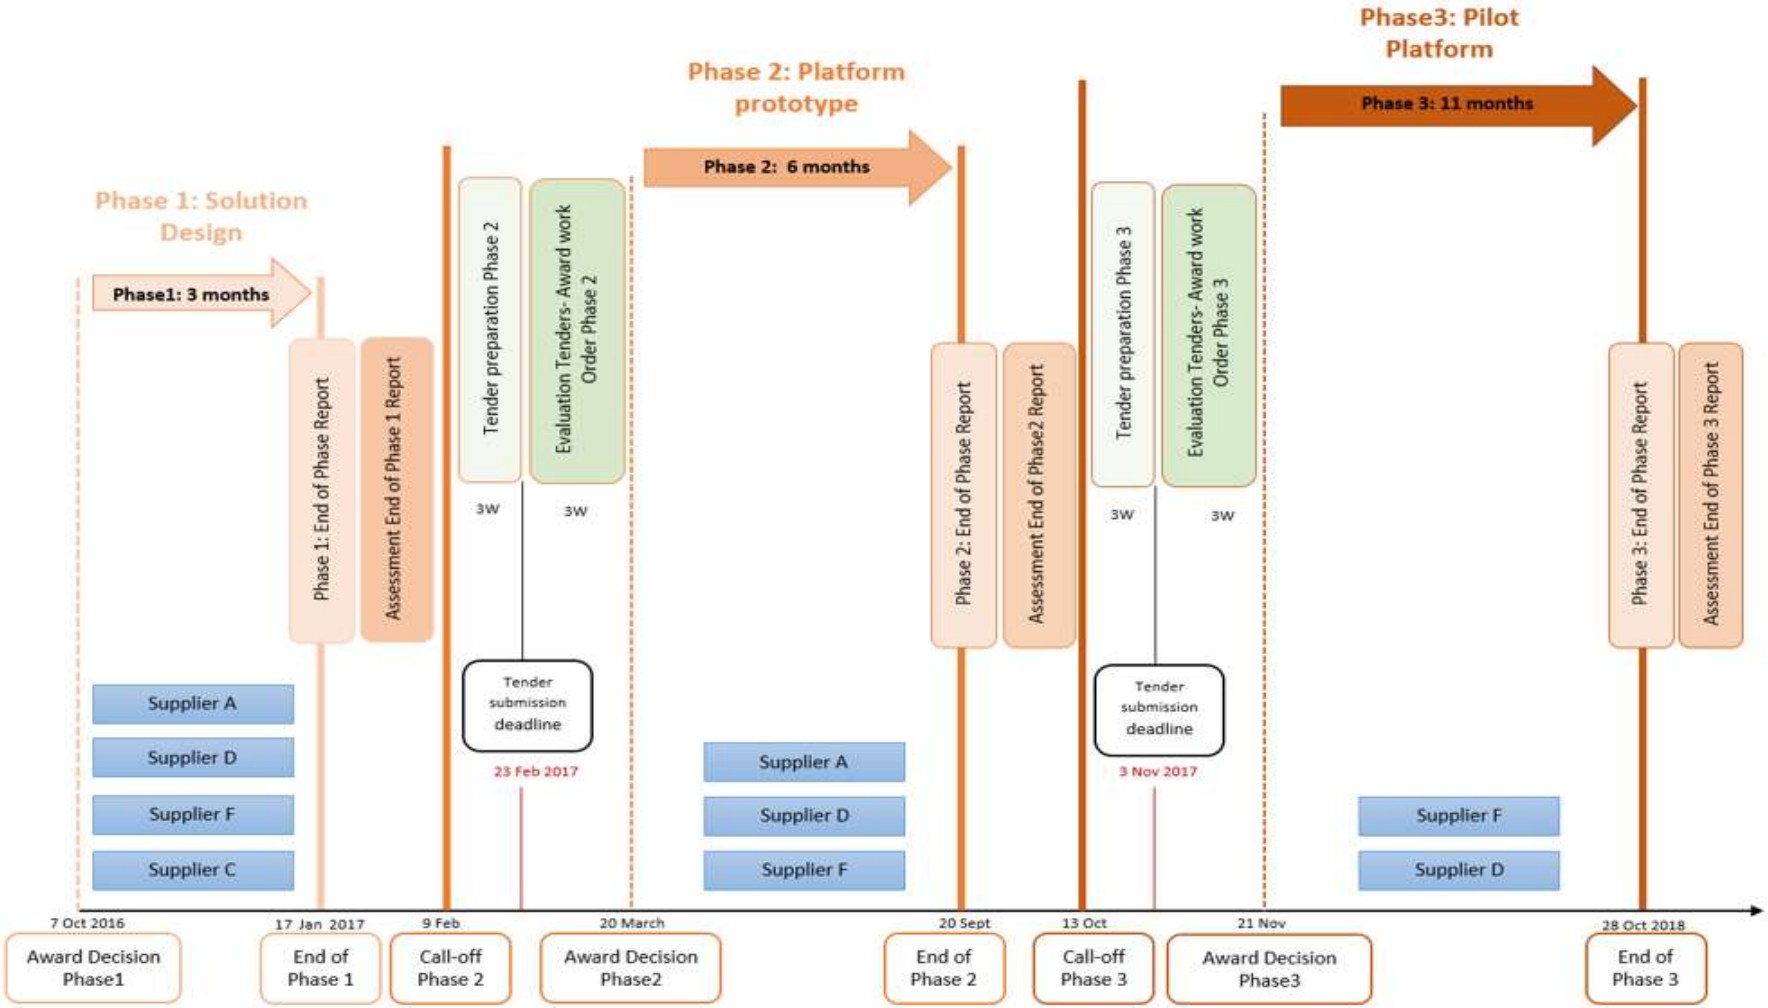
Phase 1: Solution Design
Phase1: 3 months
Phase 1: End of Phase Report
Assessment End of Phase 1 Report
Tender preparation Phase 2
Evaluation Tenders- Award work Order Phase 2
3W
3W
Tender submission deadline
23 Feb 2017
Supplier A
Supplier D
Supplier F
Supplier C
Phase 2: Platform prototype
Phase 2: 6 months
Phase 2: End of Phase Report
Assessment End of Phase2 Report
Supplier A
Supplier D
Supplier F
Tender preparation Phase 3
Evaluation Tenders- Award work Order Phase 3
3W
3W
Tender submission deadline
3 Nov 2017
Phase3: Pilot Platform
Phase 3: 11 months
Supplier F
Supplier D
Phase 3: End of Phase Report
Assessment End of Phase 3 Report
7 Oct 2016
Award Decision Phase1
17 Jan 2017
End of Phase 1
9 Feb
Call-off Phase 2
20 March
Award Decision Phase2
20 Sept
End of Phase 2
13 Oct
Call-off Phase 3
21 Nov
Award Decision Phase3
28 Oct 2018
End of Phase 3

Expected IaaS Resources

- Prototype (phase 2):
  - ~3500 cores, ~350 TB shared storage, 10 Gbps GÉANT link – per supplier for 3 months
- Pilot (phase 3):
  - ~10K cores, 1 PB shared storage, 40 Gbps GÉANT link – per supplier for 8 months

- VM characteristics (minimum)
  - 4 or 8 vCPU per VM – some use-cases could use more
  - >=1.875 GiB or >=3.75 GiB per vCPU
  - ~30 GB local VM storage
  - Linux (and Windows)

HNSciCloud – Current Status

- Official start of project: Jan 2016, duration: 30 months
- Tender announced in Jan 2016
- 17-Mar-2016: Open market consultation
- 21-Jul-2016: Tender issued (> 200 downloads, > 70 requests for clarification)
- 07-Sep-2016: Tender information day – design phase
- 19-Sep-2016: Deadline for tender replies
  - Sufficient number of valid tenders received
  - Evaluation by administrative and technical experts
- 07-Oct-2016: Award decision, contracts
- 02-Nov-2016: Kick-off meeting with Phase 1 contractors

Summary

- Commercial cloud services are expected to play an increasing role in the computing models of scientific Research Infrastructures as part of a hybrid cloud platform
- Such a hybrid cloud platform has the potential to serve many high-profile research projects
- ***Helix Nebula Science Cloud*** is a Pre-Commercial Procurement project with a budget of more than 5M€ that is co-funded by the European Commission
  - The objective is to produce a hybrid cloud platform for the European research community
- Changes to the procurement process in the public research sector are necessary to benefit from a dynamic Digital Single Market and should be supported by the platform
- A hybrid cloud poses a number of technical challenges that are being addressed by ***Helix Nebula Science Cloud***
- ***Helix Nebula Science Cloud*** is the first in a foreseen series of EC co-funded projects which will contribute to the European Cloud Initiative