M-PIL-3.2 Public Session
Next steps

Andrea Chierici

Helix Nebula – The Science Cloud with Grant Agreement 687614 is a Pre-Commercial Procurement Action
funded by H2020 Framework Programme

Ramping-up of the IaaS Resources

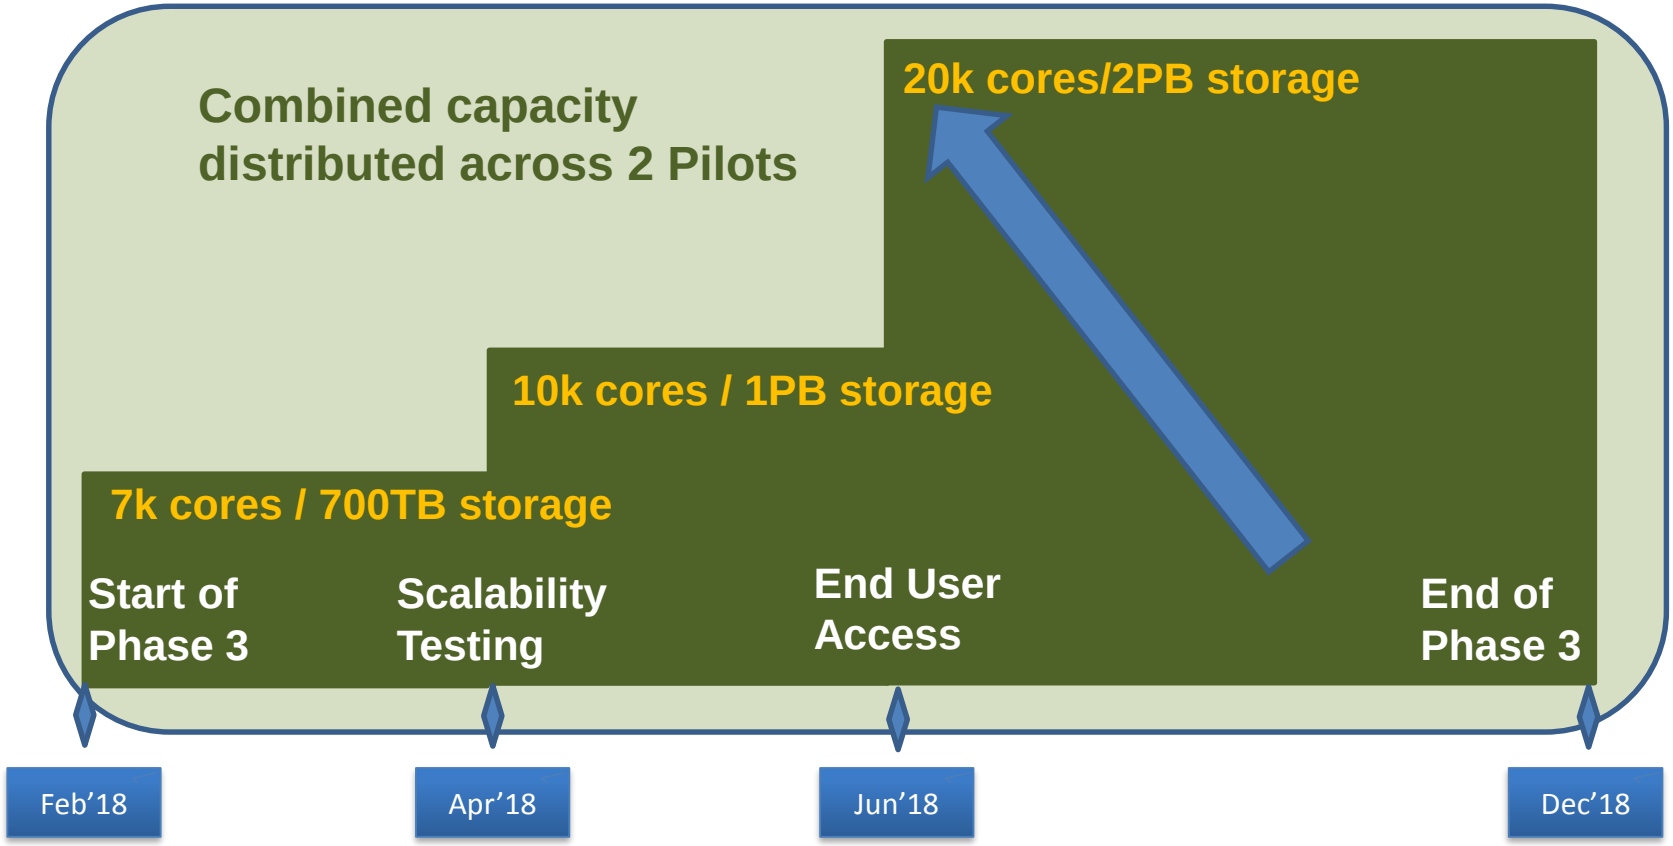

Network aggregated speed: 40Gbps

Quota management

- Procurers shared resources in order to achieve large scale tests
  - Quota limitation may have prevented reaching important results
  - Optimized utilization of resources (reached 80% in one case)
- Consolidated WLCG approach
  - Single access point with shared resources
- New use-cases
  - Each Procurer is responsible to allocate some of their capacity to all use-cases they propose
  - In some cases of high demand the procurers voluntarily donate capacity for a limited time.

Upcoming events

- July 9: CHEP 2018, **Sofia**
- August 28: GridKa School, **Karlsruhe**
  - Hands-On session organised by KIT
- September 11: HNSciCloud meeting, **Amsterdam**
  - Organised by SURFsara
- October 9-11: DI4R 2018, **Lisbon**
- October 24th: **Hamburg**
  - Organised by DESY
- November 28th-30th: HNSciCloud meeting, **CERN**
- December 4-6: ICT 2018, **Vienna**
  - Demonstrations


KIT
Karlsruher Institut für Technologie
GridKa School
GridKa School
Aug. 27.-31. 2018
Computing and Science Fair
Modern Programming
Machine Learning & AI
Cyber Security
Big Data Analytics
KIT – Campus North
Karlsruhe, Germany
August 27th - 31st 2018
More Information and Registration at
http://gridka.school

TCO Study: PanCancer

- Up to 400 VMs, at least 250 VMs concurrently, 8-16 vCPUs, 16G RAM and 30 GB scratch
- 1 PB dataset, more than 4,000 files 5-30GB range, more than 4,000 files 100-500 MB range. Outputs ca. 92,000 files in kb range
- a minimal set of resources will be in constant use, and at periods of incoming data we will ramp up resource usage to a maximum
  - **Continuous deployment (will be running for the full 12 months)**:
    - Compute: 7 VMs, 50 CPUs, 50GB RAM
    - Storage: 0.5PB, not accessed frequently
    - Network: Minimal
  - **During bursts (24-48 hours duration, 1-2x / month)**:
    - Compute: ~250 VMs, 1000 vCPUs, 8,5TB RAM
    - Storage: 0.5PB, random access with high I/O requirement
    - Network: 10GBit/s intra-node traffic, up to 4GBit/s ingress

TCO Study: Alice

- 3 types of jobs (workloads) within the ALICE use-case (each job uses a single core)
  - Monte Carlo (detector simulation)
    - low priority suitable for 'backfilling' unused capacity
    - duration: 6 hours, 2 GB/core, disk 2 GB/core
    - Network: 200MB down / 350MB up/ 0,3 Mbps
  - Raw Data reconstruction
    - process recently recorded data or reprocess older data
  - Analysis Trains

Intro of Vouchers short-term usage

- Initial voucher emission: 10 per procurer
  - Value: 250 euro
  - Available for any service (cpu, disk, gpu)
    - Consumption blocked once value has been reached
    - Possibility to export data even if credit exhausted
  - Feedback provided by early users
- Expected 100 vouchers emitted in total
- Means of paying for IaaS services consumed by long tail of science users that will execute EGI Applications on Demand