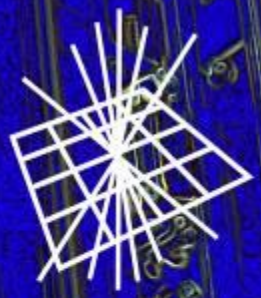

GridPP

UK Computing for Particle Physics

Tier-1 in GRIDPP5

Andrew Sansum

30th April 2015

Science & Technology Facilities Council
Rutherford Appleton Laboratory

Tier-1 in GridPP4+ Strategy

- Our 2014 GridPP4+ Strategy
  - Meet Run-2 requirements
  - Prepare for Run-3
  - Convince funding agency that we are efficient and effective
  - Requested 17.5 FTE effort

- Funding agencies accepted GridPP (not just Tier-1) justification for staffing level but not able to fund beyond FY15.
  - Tier-1 still funded at 19.5 FTE
  - Kicked baked-bean can down road 1 year
- Clearly our 2014 GridPP4+ strategy will not work again in 2015

One Possible GridPP5 Tier-1 Strategy

- Cut manpower – cap effort
  - Stop development
  - Stop quality control
  - Discard efficient infrastructure
  - Find new cheap way of working

- High Risk Strategy
  - Discarding what we have built
  - Losing agility
  - Quality control is efficient
  - Risking data loss
  - Staff retention issues

Staff Effort Analysis (Bottom up)

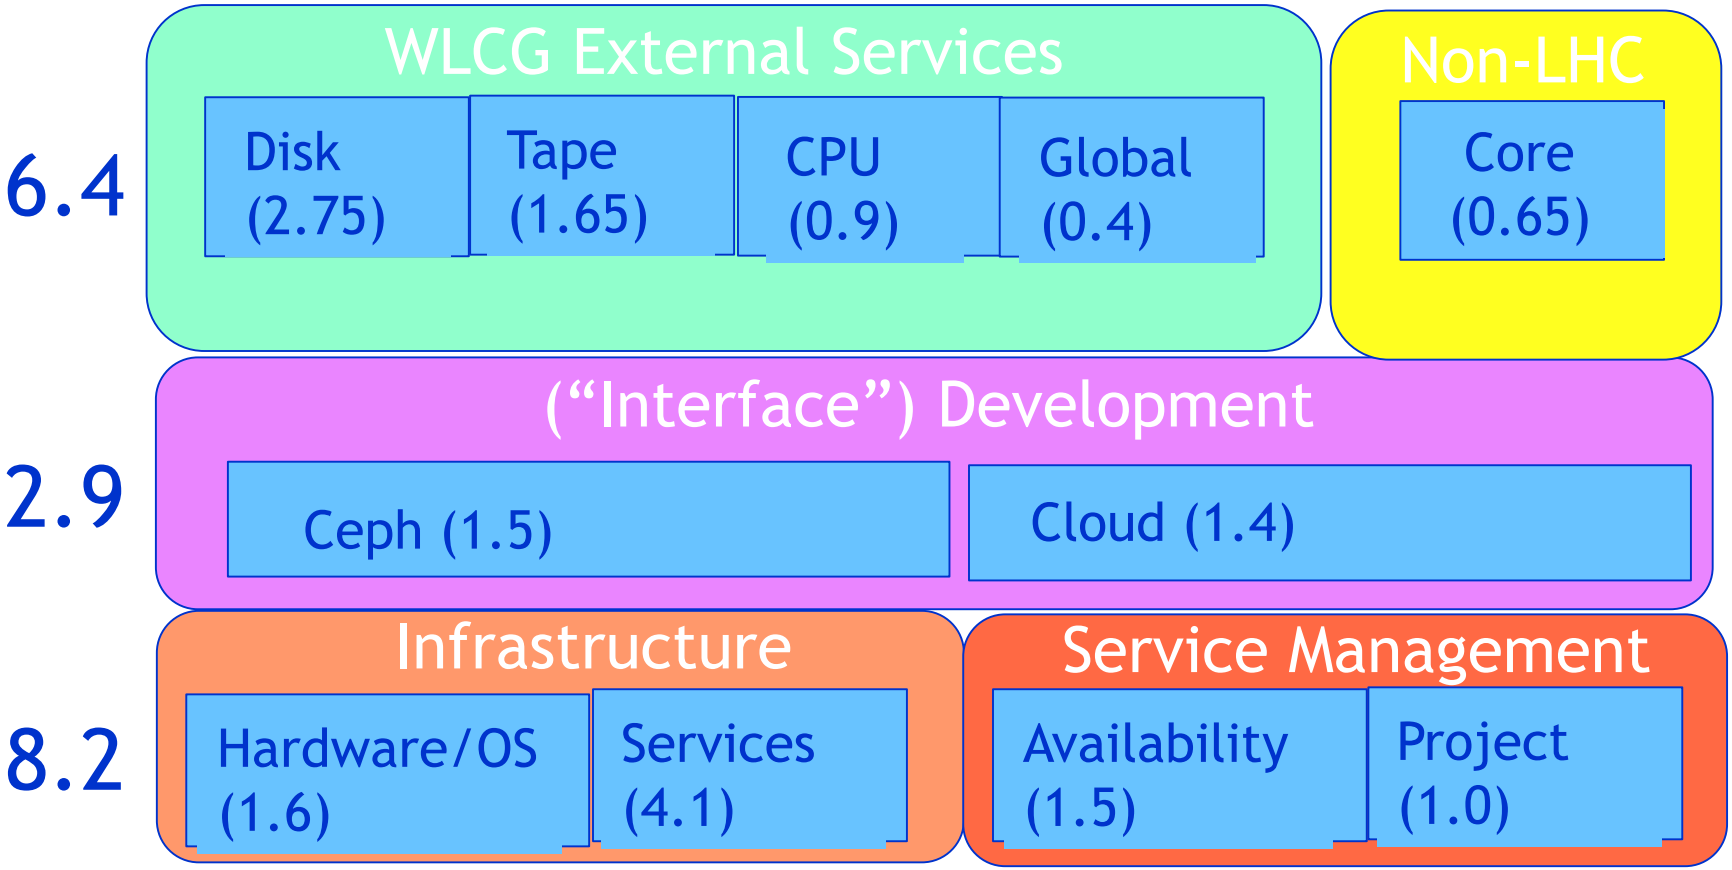

Observations

- Based on a snapshot of the Tier-1 from November 2014
  - Based on detailed lower level service decomposition – provides protection from salami slicing cuts.
  - Plus a few plans
  - But adapted to avoid problems during proposal process
  - Recognise there will be service evolution to and through GRIDPP5 – some services will close but some new services will emerge
  - Most development encapsulated invisibly within components – but make two big lumps visible (CEPH+Cloud)
- Total required effort 17.5 FTE. Decomposition:
  - External services + development = 9.3
    - Deliver to end users (only slight hardware volume scaling = 0.7 FTE)
  - Hidden underpinning = 8.2
    - The bit that allows agility and minimises scaling

The Problem(s)

- Guidance from Dave during GridPP5 modeling
  - 100% Scenario 17.5 ☺
  - 90% Scenario 16.0 FTE ☹
  - 70% Scenario 11.5 FTE ☹
- Cannot just trim everything a bit
  - Cutting bottom layer of stack will reduce our efficiency and raise costs
  - Don't want to stop doing anything at top of stack

The Importance of Narrative

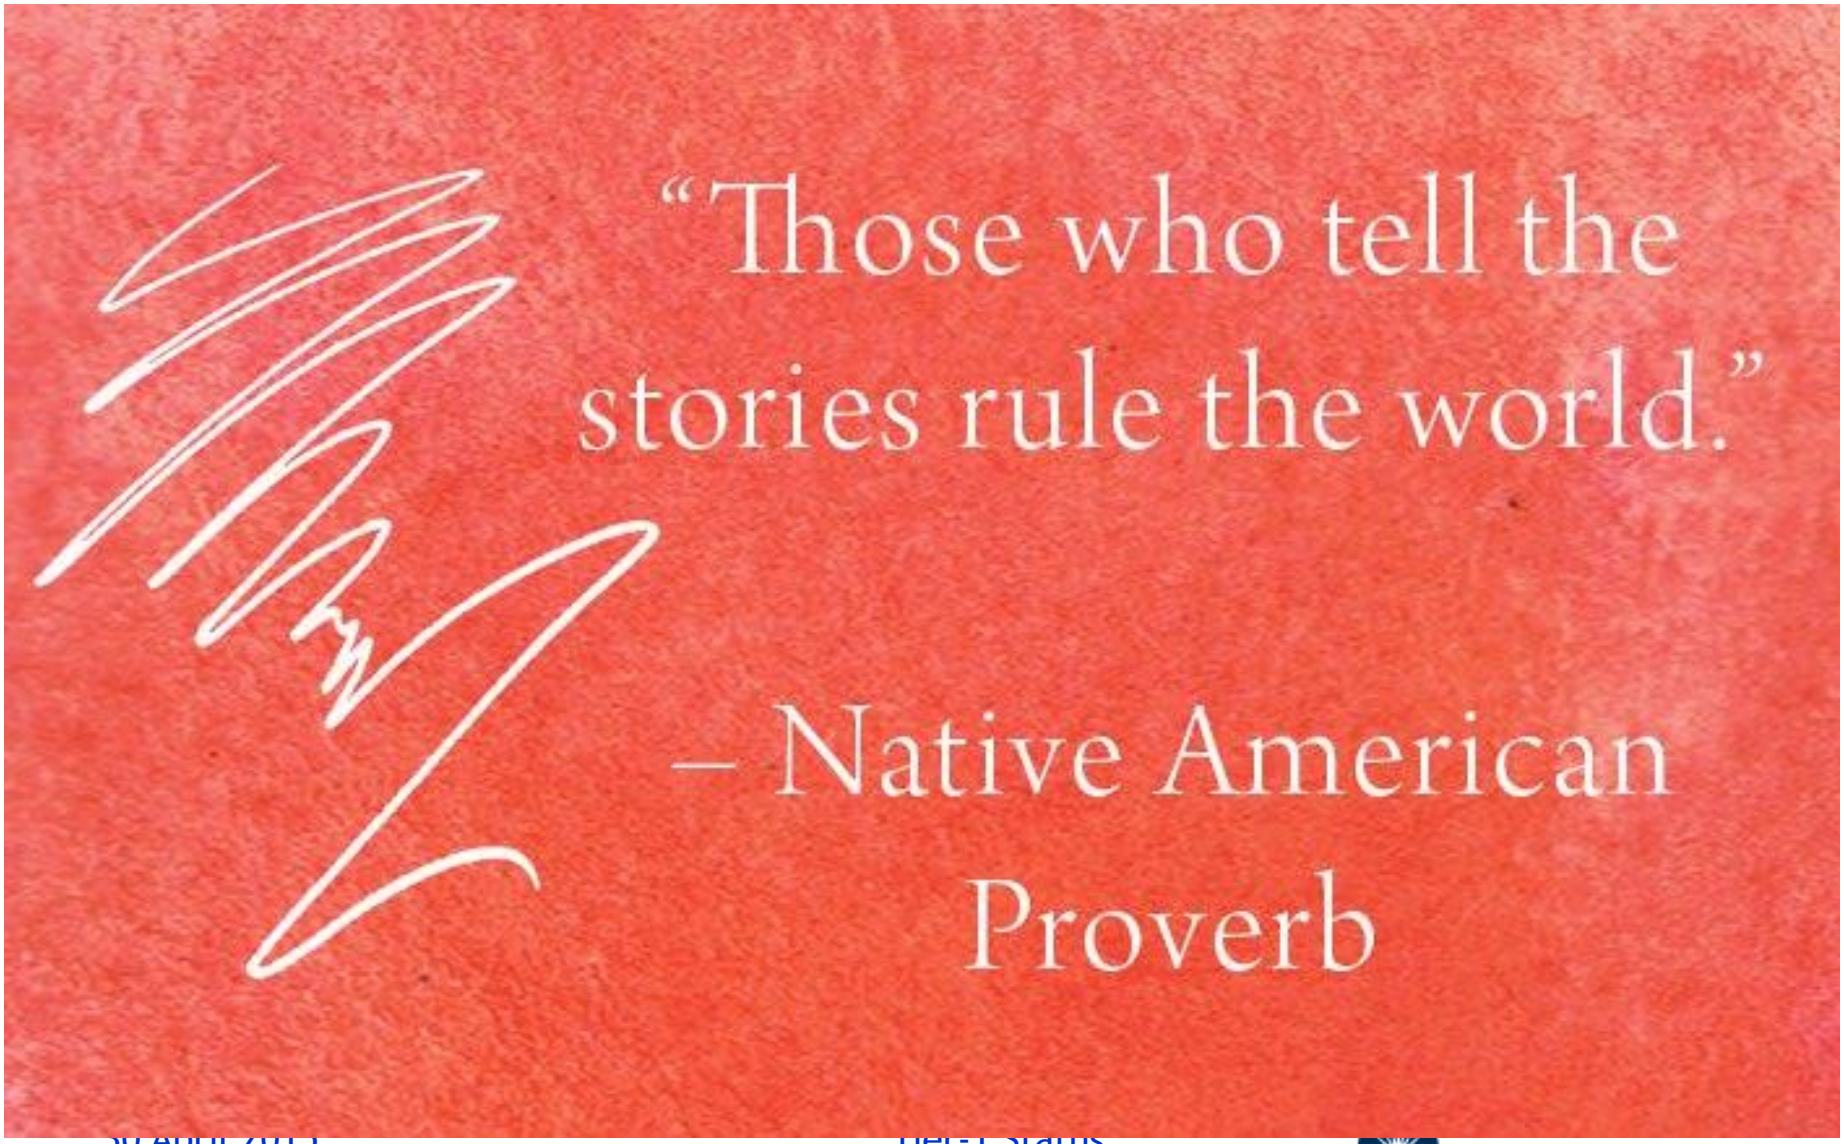

Or alternativly

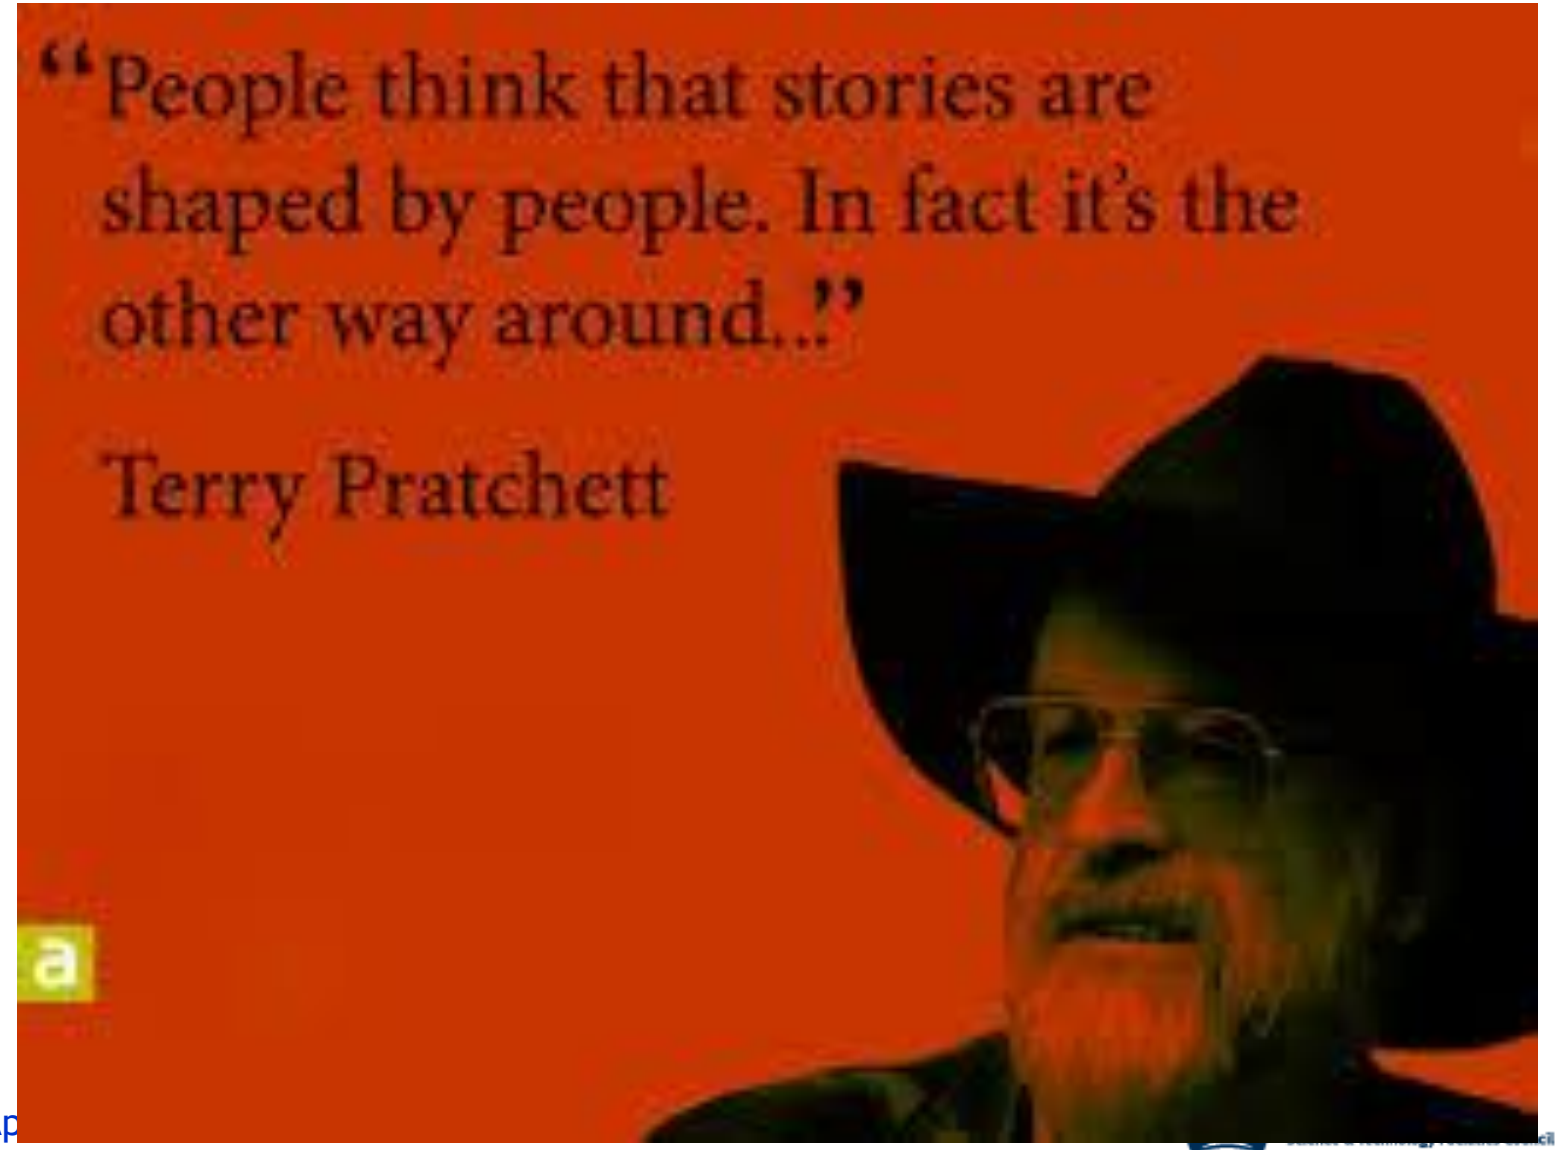

What Narrative to present to STFC ?

- Our strategy must be to cut STFC LHC computing cost.
  - No more major efficiency savings at the Tier-1 (did it in GridPP3/4) so cost reduction must come from elsewhere
  - Cannot stop development as LHC Run-3 and Run-4 around corner
- Only solution for 90% and lower scenarios is to share costs with other customers
  - Solve other STFC problems
  - Support other communities
- Cannot do it now – so must show a trajectory for staff effort funded by GridPP: 17.5, 17.5, 14.5, 14.5
  - Lost effort must be replaced by other projects in order to sustain required development in 2nd half of GridPP5
- Meet strategic objective do same in 100% scenario

Obstacles to Cost Sharing

- Almost everything we have built is an obstacle – particularly:
  - The Storage element – eg: SRM/xrootd
  - The Compute element
  - Authentication (x509 certificates)and the authorisation framework
  - And the reputation of Grid Computing
- Also the Tier-1's myopic focus on LHC computing and the need to meet the MoU
- And the lack of effort/funding by our target communities
- And the lack of (perceived) requirement for resource

But currently opportunities

- Many experiments are (or soon will be generating large data volumes. Eg: SKA, LSST, LOFAR, ITER etc etc. Also STFC facilities like ISIS, also DIAMOND.
  - Most don't have solutions (or even realise they have a problem)
  - We are one of the largest Data-Intense academic infrastrictures in the country
- STFC's shortage of funds for ICT for "big science" favors incumbents
  - STFC actively seeking a joined up UK/EU e-infrastructure (EU-T0 and UK-T0)
  - Also success of JASMIN service at RAL for earth observation demonstrates benefits of pooling infrastructure
- In order to plug 3 FTE GridPP gap, RAL probably need to get a 6 FTE project – 50% bespoke, 50% (3 FTE) contribution to existing shared infrastructure.
  - But to succeed, must have a marketable product (ie not Grid).

CEPH

- Tier-1 needs a new storage system (to replace CASTOR disk). Must work for LHC, but also have resonance with new potential customers:
  - Big technical challenge to make work for LHC (we understand risks)
- Clearly hugely successful (just now) in attracting interest from other funding streams
  - ISIS very interested (and already have a CEPH instance) can we run theirs or replace theirs with ours
  - EUDAT have H2020 funding to interface IRODS to object store (CEPH)
  - SAGE (RAL/Diamond/CCFE/SEAGATE) collaboration funded in H2020 – RAL role will be to compare SEAGATE h/w solution with Tiered CEPH instance
  - JASMIN will test CEPH object store
  - ZEPHYR project unfunded in H2020 but would build on CEPH – try again
  - CCFE interested in interfacing their data server on top of our CEPH object store – precursor to ITER.

Cloud

- Tier-1 does not need a new batch computing interface. Cloud is all about engagement with other communities.
  - Except insofar as it might allow us to place a load of Grid infrastructure (but think cloud will cost more)
  - Probably will cost more than current Grid interfaces will cost
- Possible customers:
  - SCD (already fund 0.5 FTE)
  - STFC facilities (eg ISIS and Diamond) – ISIS might fund 0.5 FTE for ISIS specific tailoring
  - Federated infrastructures like EGI or INDIGO-Datacloud
  - The EU-T0 and UK-T0 (resonance with STFC)
  - Specific STFC projects – eg LOFAR, SKA etc etc
  - H20202 (eg INDIGO-Datacloud)) 0.3 FTE

Tape

- Can we make the Tier-1 tape robots available to STFC storage projects.
- Eg DiRAC need to store 5PB on tape – can they use Tier-1
  - No money but crucial to join up UK-T0

Summary

- Reductions in funding at Tier-1 difficult to address without new funding streams
- If we don't plug the gap Tier-1 likely begin to become dysfunctional in FY18
- Need to shift our focus from solely LHC computing to STFC HTC computing in order to share costs
  - Also hugely important to STFC science. We know stuff that will help these projects.
- Need to fund a usable (marketable) product set in order to attract new customers. Bespoke middleware simply will not provide the narrative we need.
- Many opportunities but also huge risk. No plan B