



Computing Infrastructure for PP (and PPAN) Science

Pete Clarke

PPAP Town Meeting 26/27th July 2016



Computing Infrastructure

- **HTC computing and storage**
 - LHC
 - Non-LHC
 - Future requirements across PPAN
- **HPC computing - DiRAC**
- **Consolidation across STFC**
 - UKT0
 - Making case for government investment in infrastructure



HTC Computing & Storage

LHC Support

What exists today: GridPP5

18 Tier2 sites



~10% of GridPP4 resources for non-LHC activities



Tier1 RAL Computer Centre R89

~ 60k logical CPU cores
~ 32 PB Disk
~ 14 PB Tape

~ 10% of the Worldwide LHC Computing Grid (WLCG)


LHC computing support: UK share of WLCG

Developed by CESGA 'EGI View': / sum_normcpu / 2015:4-2016:3 / TIER1-VO / lhc (x) / GRBAR-LIN / I

2016-04-06 02:06

TIER1 per TIER1

UK Tier1 share is ~10%



LHC computing support: process


- LHC Experiments estimate requirements annually
 - Firm request are made for year N+1,
 - Plus estimates for year N+2..
 - Documents submitted to CRSG (computing resources scrutiny group)
- Experiment requests scrutinised by CRSG
 - Scrutiny/meetings/adjustments.....
 - Eventual approval by RRB
 - Approved official experiment requirements appear in system called “REBUS”
- This is an international process - its not a UK thing
- The WLCG then requests fair share “pledges” from all countries
- UK (GridPP) then pledges exactly its share - proportional to author fractions.
- Projected UK fair share requirements are requested in each GridPP funding cycle
- So hardware support for LHC experiments is “sort of” OK until 2019/2020
- But severe shortage of computing staff in the experiments

LHC computing support: actual usage

The total histogram (envelope) shows the actual CPU used in 2015/16 by experiments

LHC experiments get fair share support from UK funded by STFC

LHC experiments use more
This is provided in UK using leveraged resources (not funded by STFC)




This is possible because the Tier2 sites actually provide ~double that which they are funded for (+ fund all of the electricity)



Non-LHC Computing Support

Non-LHC computing support

- Non-LHC activities supported are shown in this [log plot](#)



- These are supported through:
 - Trying to maintain 10% of GridPP resources reserved for non-LHC activities
 - Local leverage at Tier2 sites

Non-LHC computing support

- Currently supported PP activities include:
 - ATLAS,CMS, LHCb,ALICE,
 - T2K,
 - NA62,
 - ILC,
 - PhenoGrid,
 - SNO,
 -other smaller users.....
- New major activities on horizon in next 5 years:
 - Lux-Zeplin [already in production]
 - HyperK, DUNE
 - LSST
- Every effort is made to support any new PP activity within existing resources
- But as more and more activities arise then eventually unitarity will be violated
 - marginal cost of physical hardware resources
 - spreading staff even more thinly

Non-LHC computing support

- Policy published on GridPP web site: new activities are encouraged to:
 - liaise with GridPP when preparing any requests for funding
 - at least make their computing resource costs manifest when seeking approval
 - where these are “large” then to request these costs where possible
 - this is particularly important if a large commitment (pledge in LHC terms) is required to an international collaboration.
 - Each new activity should consider the complete costs of computing:
 - Marginal hardware (CPU, storage)
 - Staff:
 - operations
 - generic services
 - user support
 - activity specific services
- ↑ Economies of scale increase
- Of course, if it is not “timely” to obtain costs, then best efforts access remains

Astro-Particle Computing Support

- Lux-Zeplin
 - LZ is already a mainstream GridPP computing activity - centred at Imperial
- Advanced-LIGO
 - A-LIGO already has a small footprint at the RAL Tier1
 - This could be developed further as required by LIGO
- CTA
 - No request for computing to the UK yet - but GridPP is expecting to support this
 - CTA UK management will address this later



HPC computing



DiRAC

HPC computing for theory

- HTC : for embarrassingly parallel work (e.g. event processing)
 - cheap commodity “x86” clusters
 - ~ 2GByte/core
 - no fancy interconnect
 - no fancy fast file system
- HPC : for true highly parallel work (e.g. lattice QCD, cosmological simulations)
 - can be x86 but also more specialist very many-core processors
 - high speed interconnect, can be clever topology
 - large memory per core / large coherent distributed memory / shared memory
 - often fancy fast file system
- The theory community relies upon HPC facilities
 - these are their “accelerators”
 - produce very large simulated data sets for analysis
- DiRAC is the STFC HPC facility.

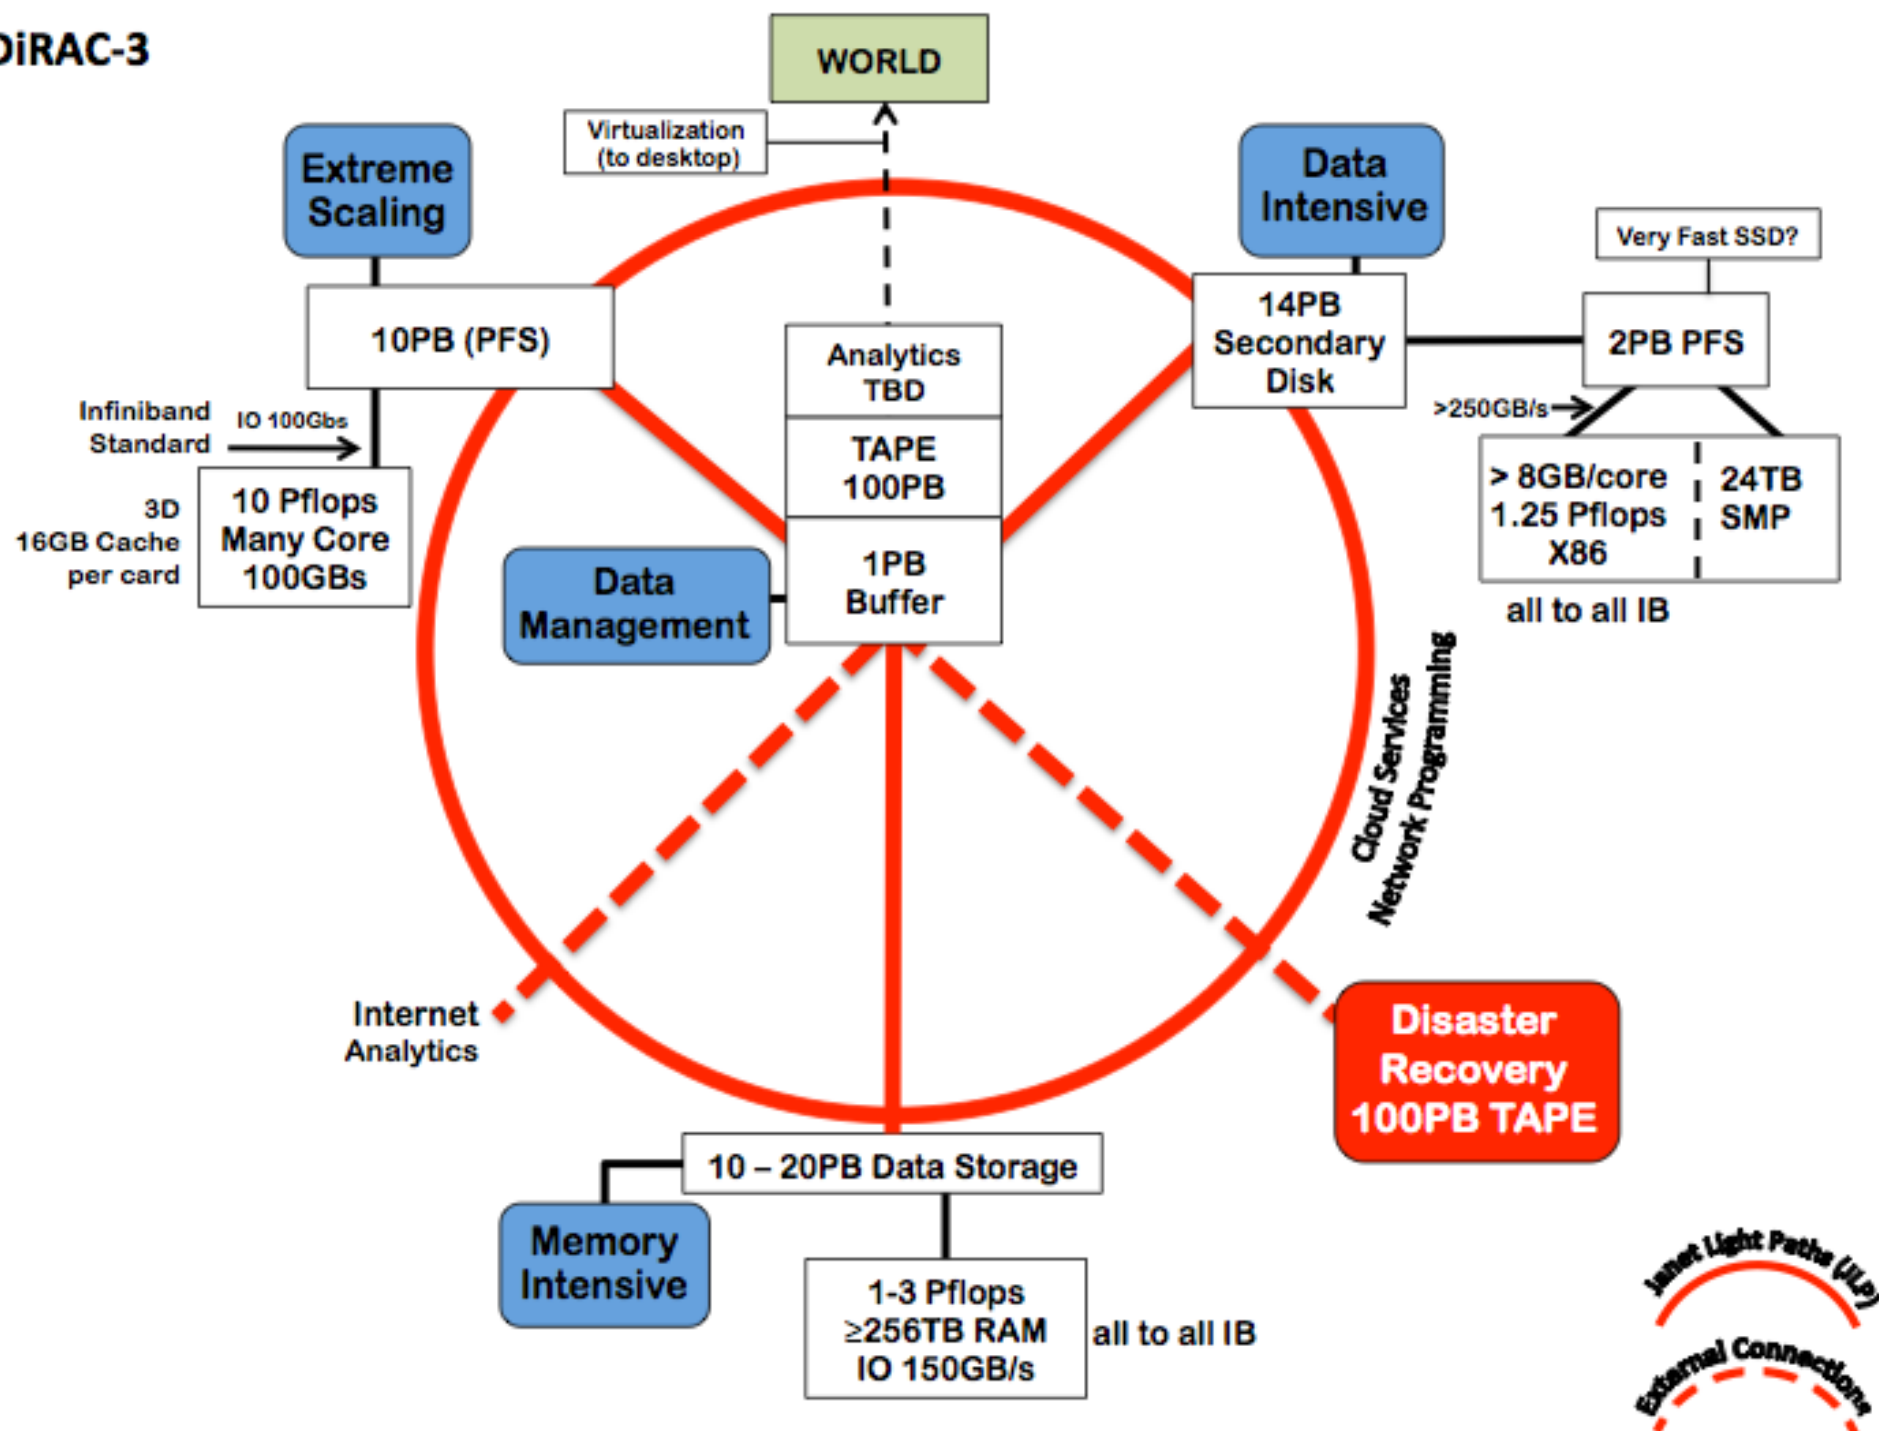
HPC computing for theory

- DiRAC-2
 - 5 machines at Edinburgh, Durham, Leicester, Cambridge
 - ~2 PetaFlops/s
 - Excellent performance - has given UK an advantage
 - In production > 5 years. Now end of life
- DiRAC-2 sticking plaster
 - Ex-Hartree Centre *Blue Wonder* machine going to Durham
 - Ex-Hartree Centre *Blue Gene* going to Edinburgh for spare parts.
- DiRAC-3 is needed by the theory communities across PPAN
 - The scientific and technical case has been made ~ 2 years ago
 - ~15 PetaFlops/s + 100 PB storage
 - Funding line request of ~ £20-30M
 - But no known funding route at present !
- Situation is again very serious for the PPAN Theory Community !



121,000 cores
1.8 Pflop/s
3.7 PB

DiRAC-3





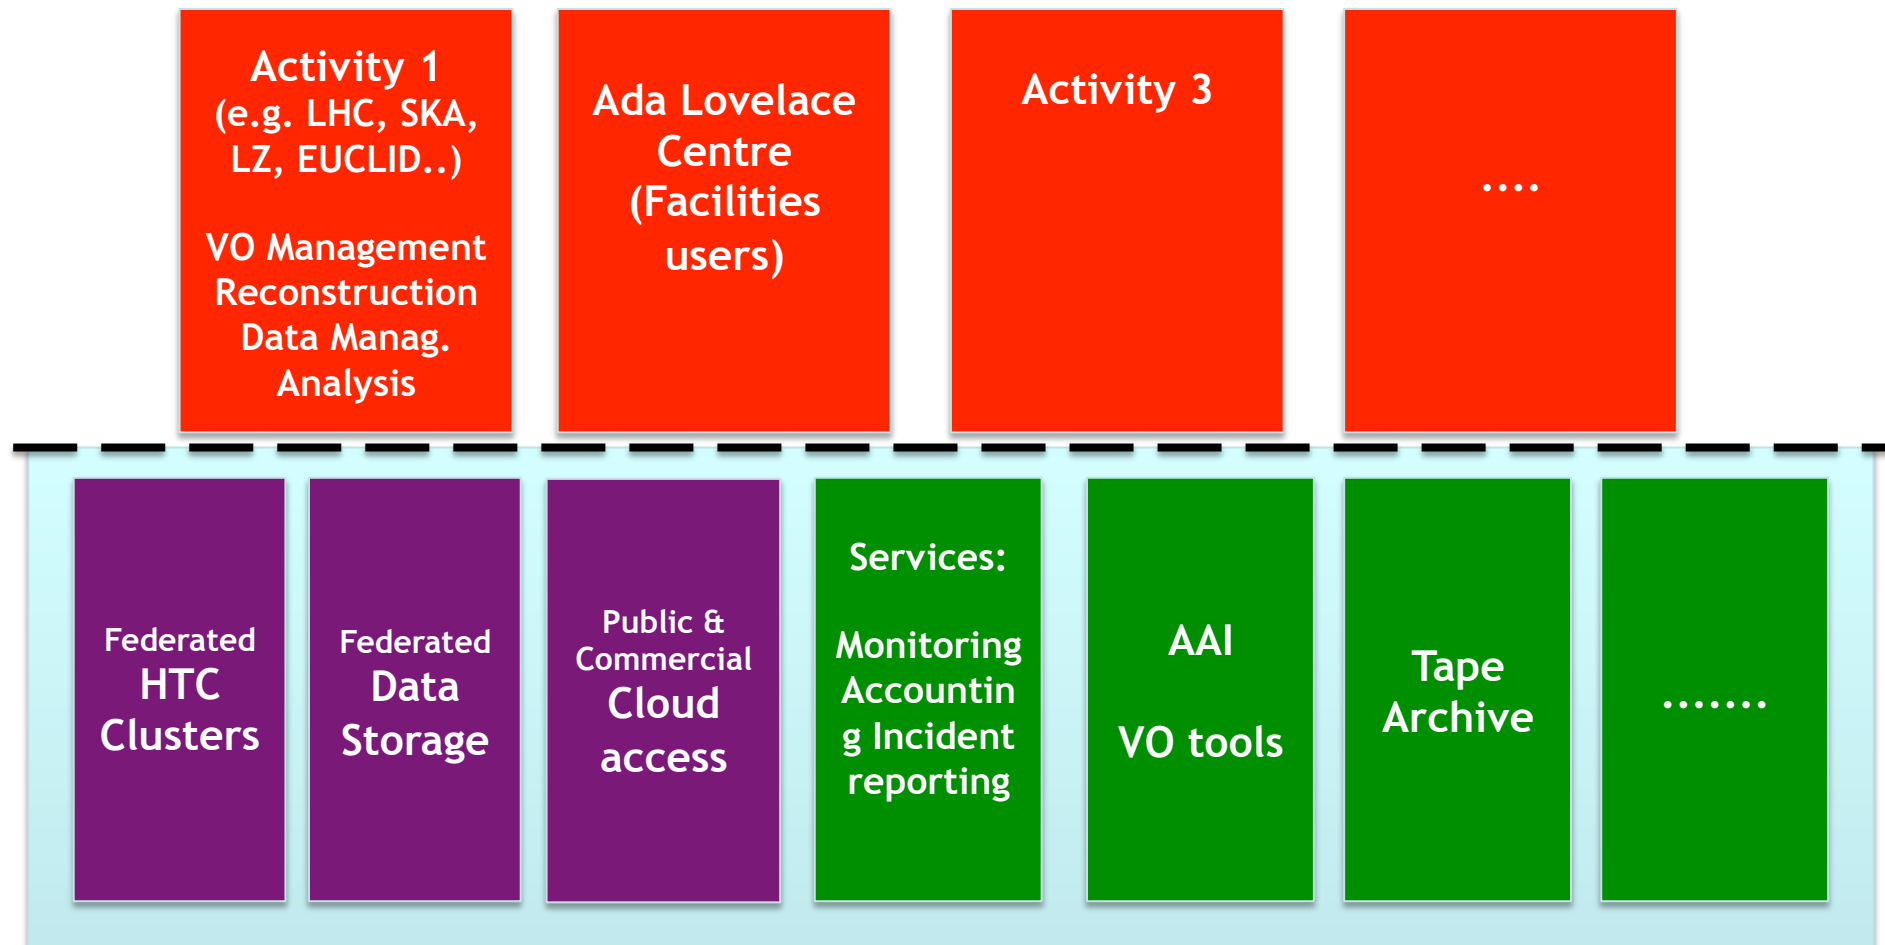
Consolidation across STFC

Consolidation across STFC: UKT0

- There are many good reasons to consolidate and share infrastructure
 - European level: in concert with partner funding agencies
 - UK level: BIS and UKRI
 - STFC level: it makes no sense to duplicate silos
 - Scientist level: shared interests and common sense
- An initiative was taken in 2015 to form an association of peer interests across STFC - this called UKT0
- So far:
 - Particle Physics: LHC + other PP experiments
 - Astro: LOFAR, LSST, EUCLID, SKA
 - Astro-particle: LZ, Advanced-LIGO
 - DiRAC (for storage)
 - STFC Scientific Computing Dept (SCD),
 - National Facilities: Diamond Light Source, ISIS
 - CCFE (Culham Fusion)
- Aim to
 - share/harmonise/consolidate
 - avoid duplication
 - achieve economies of scale where possible

Consolidation: ethos

Science Domains remain “sovereign” where appropriate



Share in common where it makes sense to do so

Consolidation: PP ⇔ Astro links


- Already strong links between PP ⇔ Astronomy
- LSST
 - PP groups at Edinburgh, Lancaster, Manchester, Liverpool, Oxford, UCL, Imperial are involved
 - Proof of principle resources used by LSST@GridPP to do galaxy shear analysis
 - Joint PP/LSST computing post in place to share expertise (Edinburgh)
 - Recent commitment made from GridPP to support DESC (Dark Energy Science Consortium)
[relying mainly upon local resources at participating groups]
- EUCLID
 - EUCLID is a CERN recognised activity - particularly to use CERNVM technology
 - EUCLID has been enabled on GridPP and has carried out piloting work which was a success
- SKA
 - SKA is a major high profile activity for the UK
 - Many synergies with LHC computing to be exploited
 - Joint PP/SKA computing post in place (Cambridge)
 - RAL Tier1 are involved in SKA H2020 project
 - Joint GridPP ⇔ SKA meeting planned for November 2016

PPAN wide HTC requirement

- PP requirements grow towards LHC Run-III
- Astronomy requirements are growing fast
 - Advanced LIGO
 - LSST
 - EUCLID
 - SKA

- Figure shows CPU requirements (2015 cores)
 - GridPP5 funded
 - PP requirements
 - PPAN requirements




[some of difference between green and purple is currently made up of leverage]



- Similar plots for storage
- PPAN requirements are approximately double the known funded resources

Consolidation: reminder of reality

- Obvious but: co-ordinating activities and consolidation means:
 - cost per unit hardware resource to each activity will reduce
 - operations and common service staff can be shared - reducing cost per activity, avoiding duplication
- But it does not actually make operating costs go down in absolute terms when the required capacity is over doubling
- Its just that costs scale less-than-linearly with required capacity (logarithmically?)





Case for BIS investment in eInfrastructure for RCUK

RCUK/BIS spreadsheet for e-Infrastructure investment

- The landscape is changing
 - BIS → DBEIS
 - RCUK → UKRI
- An RCUK wide group has been working for > 2 years to make case to BIS to invest in eInfrastructure across RCUK
- STFC is represented on this BIS/RCUK group
- A funding case was submitted to BIS via RCUK in Jan 2016
 - This contained a substantive lines for 5 years for:
 - PPAN
 - National Facilities (DLS,ISIS,CLF)
 - DIRAC
 - Included staff element
- Case was last reviewed May 6th
 - discussions are still going on
 - some hope for next autumn statement ??

LHC cat and non-LHC cat have to share



It was not possible to fund all hardware costs in GridPP5 for all LHC and non-LHC requirements



local leverage and determination



next 5 years ... we have to work as UKT0



DBEIS invest in bigger basket ?



DBEIS invest in bigger basket ?



..and a high performance basket

- Computing for LHC is approximately OK until 2019/20 through GridPP5
- Requirements from non-LHC activities are growing. Up to a point this can be handled using leveraged resources at Tier2 sites - but at some point unitarity is violated for both hardware and staff
- New activities are encouraged to liaise early with GridPP, to work within UKT0 framework, and to request and contribute marginal costs as part of the “collective”
- DiRAC HPC provision for theory is at end of life - the situation is rapidly becoming serious.
- **Severe shortage of computing staff in experiments**
- There is much good will across STFC (PPAN and National Facilities) to work together to minimise costs
- Case being made strongly to BEIS for elnfrastructure investment