

Tier-1 Procurement

Alastair Dewhurst

Introduction

- FY19/20 is the final year of GridPP5.
- Capacity purchased this year will go into production for FY20/21 the first year of GridPP6.
- Key GridPP6 differences:
 - GridPP6 is not yet approved.
 - The fractional share of each LHC VO is changing.
 - The hardware model is reducing how much spare capacity we need.

$$\text{Capacity} = [\text{Requirement}] * [\text{GridPP Allocation}] * [\text{UK Fraction}] * [1 + \text{Hardware Overhead}]$$



Requirements

LHC VO Requirements

- We are currently in a Long Shutdown so the LHC VO requirements are only increasing slightly.
- The Table shows the **total** LHC VO requirements for 2019 and 2020 as taken from Rebus.

FY19/20	ALICE	ATLAS	CMS	LHCb	SUM
CPU (HEP-SPEC06)	365000	1057000	650000	271000	2343000
Disk (Tbytes)	37900	88000	68000	27900	221800
Tape (Tbytes)	37700	221000	220000	50900	529600
FY20/21	ALICE	ATLAS	CMS	LHCb	SUM
CPU (HEP-SPEC06)	365000	1057000	650000	328000	2400000
Disk (Tbytes)	44000	88000	68000	33200	233200
Tape (Tbytes)	37700	221000	220000	55500	534200
Difference	ALICE	ATLAS	CMS	LHCb	SUM
CPU (HEP-SPEC06)	0	0	0	57000	57000
Disk (Tbytes)	6100	0	0	5300	11400
Tape (Tbytes)	0	0	0	4600	4600



GridPP Fraction

- Currently we should assume that GridPP will be expected to provide the LHC VOs with 100% of the UK commitment.
- The Table shows how each LHC VOs fraction changes between GridPP5 and 6.
- Some experiment requirements drop significantly meaning they face a reduction in resources.

Experiment	GridPP5	GridPP6
ALICE	2.0%	3.0%
ATLAS	12.5%	14.8%
CMS	8.0%	6.8%
LHCb	30.0%	22.5%



Disk Requirement

- CMS and LHCb capacity drops.
 - CMS have used all their current capacity.
 - LHCb haven't yet, but are requesting a significant increase next year for their upgrade work.
- Proposal to keep CMS and LHCb allocations constant for FY20/21.
 - We predict that by FY21/22 their allocation will grow again.

Tier-I Disk Allocation	FY19/20 (TB)	FY20/21 Current (TB)	FY20/21 Proposed
ALICE	758	1476	1476
ATLAS	11000	13468	13468
CMS	5440	4624	5440
LHCb	8370	7065	8370
IRIS	5346	6600	6600
Total	30914	33133	35354



CPU Requirement

- Non-LHC increases from 10% to 13% in GridPP6
- CMS and LHCb capacity drops
- CPU fair share can easily be modified but not having enough CPU to exploit data stored on disk would be counter productive.
- Proposal to keep CMS and LHCb allocations constant for FY20/21.

Tier-I CPU Allocation	FY19/20 (HS06)	FY20/21 Current	FY20/21 Proposed
ALICE	7300	10950	10950
ATLAS	132125	159692	159692
CMS	52000	44200	52000
LHCb	81300	73800	81300
Other	27273	37523	37523
Total	300000	326165	341465



Tape Requirement

- VO have traditionally under utilized Tape so their allocation is often a lot lower than their pledge.
- It is unlikely that CMS will exceed this F20/21 pledge before next year, so their allocation can be their pledge.
- LHCb are likely to exceed their F20/21 pledge and so will be allowed to keep their current allocation.
- ATLAS growth is probably double what is indicated here as they ran a deletion campaign just before May.

	Usage in May 2019 (TB)	Growth in previous 6 months	FY19/20 Pledge	FY19/20 Allocation	FY20/21 Pledge	FY20/21 Proposed Allocation
ALICE	882	0	754	882	1131	1131
ATLAS	18375	1394	27625	23000	32708	29000
CMS	12585	983	17600	16000	14960	14960
LHCb	12011	440	15150	13000	12578	13000
Other	4425	394	N/A	8000	N/A	8000
Total	48278	3211	53193	52882	61249	66091



Procurement

Disk Procurement

- HY20/21 Disk Requirement: 35354 TB
- Current capacity: (5 years lifetime) 35374 TB
- Capacity to be decommissioned: 1318 TB
- For LHC VOs, 95% of pledge goes into Echo (5% Tape buffer).
- GridPP allocates 5% of LHC capacity to non-LHC VOs. i.e. buying the LHC pledge for Echo gives us enough capacity for all GridPP users. Just don't forget to buy disk for the tape buffer!
- GridPP6 assumes 12% overhead
 - In GridPP5 were using 20% overhead.
- Proposed to continue using larger overhead for now.
 - Need to purchase 7.3PB raw capacity in Echo to meet requirement.
 - Budget of £350k for £48.15 / raw TB
 - **£250k from GridPP and £100k from IRIS.**



CPU Procurement

- FY20/21 CPU Requirement: 341465 HS06
- Current capacity: 335201 HS06 (6 years lifetime)
 - Some capacity works into 7th year but that is just a bonus.
- Capacity to be decommissioned: 45966 HS06
- Assume 4% capacity unavailable:
 - 3% for first 5 years and 10% in 6th year.
- Need to purchase 50500 HS06 to meet requirement.
 - Budget of £450k for £8.91 / HS06
 - Includes increase cost due to SSD.



Tape Procurement

- FY20/21 Tape Requirement: 66091 TB.
- Current Capacity: 58956 TB (10TKD media).
 - We can borrow 13000 TB from IRIS (IBM/LTO media).
- Required overhead: 400 Tapes (3400 TB)
 - This is point at which we need to order more tapes or risk running out.
- We do not need to purchase any tapes this financial year.
 - Current capacity should last till 2021.
 - Do not intend to ever buy any more T10KD media.



2019 Procurement

- Tier-I Capital:
 - Initially £782k.
 - ‘Almost Certain’ extra £150k for SSD for CPU.
 - Might get more as part of GridPP underspend.
- Spend:
 - Minimum £450k on CPU
 - Minimum £250k on Disk
 - £232k not-committed. Can split between disk and CPU or spend elsewhere (e.g. new Tape stuff).
- Anything we buy this year will last for all of GridPP6. In the long run we need a lot more of both, so will just try and keep disk:CPU ratio the same.



Procurement batches

- Traditionally in GridPP5, the Tier-1 purchased two batches of each type of hardware each year.
 - This is a significant overhead to the procurement (effort and cost).
- Last year we agreed to buy one batch if total size was less than 15% of the current capacity.
 - Both minimum spends are well under this 15% limit.
- Proposal: For GridPP6 we should buy one procurement each year.
 - Larger spends give better value for money (for a relatively small increased risk).
 - We have setup tender so we can purchase extra capacity later in the year (opportunistic capital).
 - If there are problems, we should be aware of them before spending more money.



Summary

Allocations FY20/21	Tape (TB)	Disk (TB)	CPU (HS06)
ALICE	1132	1476	10950
ATLAS	29000	13468	159692
CMS	14960	5440	52000
LHCb	13000	8370	81300
Other	8000	6600 (IRIS)	37523

- Procurements will be from a single vendor this year.
 - CPU Tender should go live next week!



Backup

GridPP6 Proposal

	Tier-1 Resource Requirements				Tier-2 Resource Requirements			
	2021	2022	2023	2024	2021	2022	2023	2024
CPU [KHS06]	417	540	661	709	402	512	626	657
Disk [PB]	39	55	70	72	32	45	56	60

Above is the required amount of capacity needed at the Tier-1 and spread across the Tier-2s.

Below is the estimated costs based on our conservative pricing.


Hardware		FY20	FY21	FY22	FY23	Sub-Total	Totals	Grand Total
Tier-1	CPU	£1,990,648	£1,136,201	£1,422,270	£697,288	£5,246,407	Capital	£11,465,651
	Disk	£257,348	£1,390,496	£1,031,012	£653,585	£3,332,440		
	Tape	£552,780	£539,889	£690,115	£264,020	£2,046,804		
	Ops (C)	£35,000	£35,000	£35,000	£35,000	£140,000	Non-Capital	
	Ops (R)	£175,000	£175,000	£175,000	£175,000	£700,000		
	Netwk(R)	£150,000	£100,000	£100,000	£100,000	£450,000		
Tier-2	CPU	£1,158,984	£1,350,373	£1,439,378	£941,346	£4,890,080	Capital	£7,984,502
	Disk	£591,856	£982,109	£922,669	£597,788	£3,094,422		
	Network	£0	£0	£0	£0	£0	£7,984,502	
	Ops	£0	£0	£0	£0	£0		
Total		£4,911,615	£5,709,068	£5,815,443	£3,464,026	£19,900,153		£19,900,153



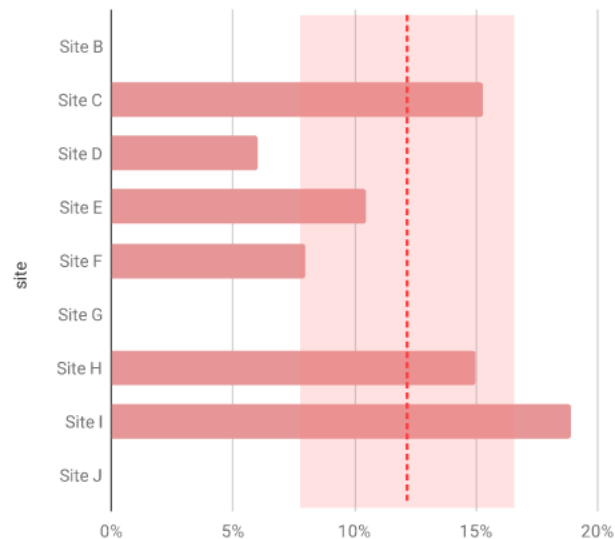
WLCG CPU costs

- Last year's price (from XMA) was £308574.70 (inc VAT) for 34000 HS06.
- £9.08 / HS06
- Had to use direct award rather than mini tender due to external time constraints

CPU cost (2018)



CPU cost yearly decrease rate



10.3 €/HS06

-12%/year

anaud.vernet@cern.ch

2019-03-20

RR

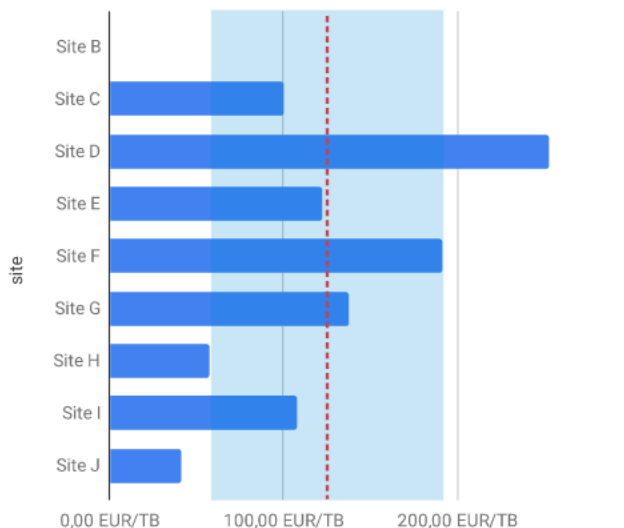
[1] https://indico.cern.ch/event/759388/contributions/3337465/attachments/1814794/2965614/Vernet_CostModelling.pdf



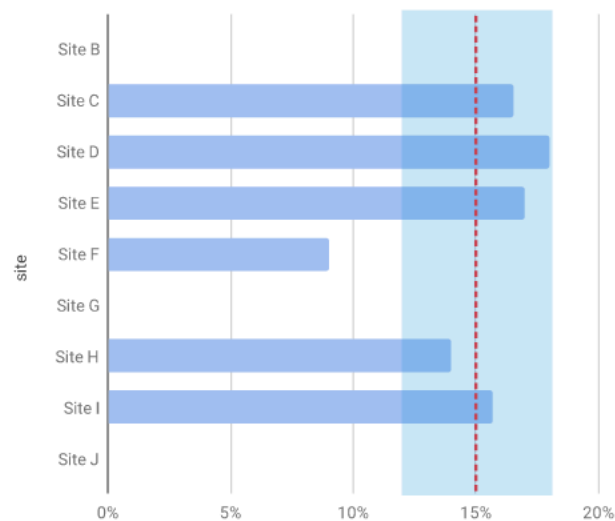
WLCG Disk costs

- Last year's price (from DELL) was £299,263.20 (inc VAT) for 6480TB storage
 - £46 / TB Raw or £80 / TB usable.

Disk cost (2018)



Disk cost yearly decrease rate



126 €/TB

-15 %/year

[1] https://indico.cern.ch/event/759388/contributions/3337465/attachments/1814794/2965614/Vernet_CostModelling.pdf



Cost modelling

- GridPP has monitored effective cost per unit of resource since 2001 at the Tier-1.
- We are seeing much slower price improvements.
- Other countries are seeing worse than 20% but slightly better than us.

