IRIS Computing Capacity Requirements 2020-2024

Presented on behalf of the IRIS Consortium.
Contact editor: Pete Clarke (peter.clarke@ed.ac.uk)

August 27, 2020

1 Introduction

This report shows the computing requirements for STFC for CPU, Disk, GPU and Tape.These are aggregated across Programmes, National Facilities and SCD.

The figures are compiled from the IRIS RSAP[1] process carried out in Q4-2019 and Q1-2020. Each science Activity submitted a **Resource Request Document**. (RRD) This contained firm requests for 2020 and forward look estimations for 2021-2024. Each RRD was scrutinised by the RSAP to produce final allocations. The 2020 RSAP full report is available in a separate document.

In addition to this

- LSST provided a new estimate for their proposed pipeline processing contribution in lieu of subscription;
- SKA provided a revised estimate to include the SRC;
- GridPP provided an estimate of the shortfall not funded in the GridPP6 award. this is actual during the GridPP6 funding period 2020-2023 and extrapolated for 2024;

[1]The RSAP is the Resource Scrutiny and Allocation Panel. It is an independent body constituted within IRIS. Panel members are drawn from independent scientists in scientific communities. Resource Request Documents are scrutinised in a formative process to ensure that requests are commensurate with the approved science programmes.

- CASU and WFAU have been added explicitly.

The requested and allocated resources are all contained in JSON files. These JSON files have been used to auto-generate the tables, plots and costings in this document.

2 Requirements

The first Table 1 and Figure 1 show the annual incremental costs required to build the required capacity for all of STFC. The final line shows the cumulative total.

The subsequent pairs of Tables show, for each science Activity, the (i) annual cumulative requirements for each of CPU, Disk and Tape and GPUs and the total and (ii) the cost estimate. These are followed by figures showing the cumulative requirement and annual incremental cost for each.

An additional line labelled as "other" is included to allow headroom for activities not yet known or approved and include some costs for non-capacity infrastructure (racks, switches, head nodes,....etc).

3 Notes on the calculations

- The cost-per-unit-resource figures have been compiled from past experience in IRIS as well as actual costs for recent purchases.
- For the forward projections the incremental shortfall for a subsequent year-n is calculated as: (Net requirement for year n) - (Net requirement for year n-1 assumed to now exist)*(1- the obsolescence factor given in the table).
- The existing resources are "retired" at the obsolescence factor. Currently an approximate value of 10% obsolescence per annum is used. This is certainly an underestimate when hardware is in steady state, it is more like 15-20%. The lower figure of 10% is used as IRIS hardware is very new.

Table 1: Overall cost summary table in kPounds

Cost (kPounds)	2020	2021	2022	2023	2024
CPU cost	2036	3048	4664	4358	4192
Disk cost	59	1338	2519	2635	2820
Tape cost	0	206	635	968	906
GPU cost	1867	1176	1495	1152	1447
Total cost	3962	5768	9313	9113	9365
Cuml. cost	3962	9730	19043	28156	37521

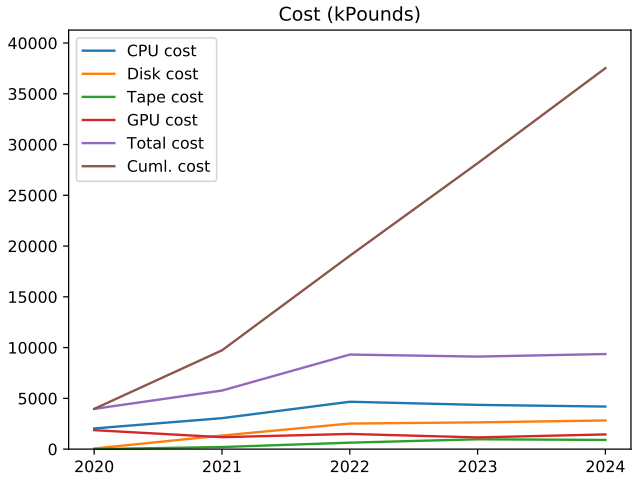

Figure 1: All lines show the annual incremental cost summary, apart from the line labelled as Cuml. cost which shows the cumulative total.

All of the following large Tables show the cumulative capacity required by each year. I.e. the column for year-n shows the total capacity which must be installed by that time.

The associated costing Tables then calculate the incremental spend each year needed to achieve the above. The figures in the last line of each are then copied into Table 1.

Table 2: CPU-cores-NetRequestTable

CPU-cores	2020	2021	2022	2023	2024
CCFE	1000	1100	1200	1300	1400
CCP4	300	400	500	500	500
CLFOCTOPUS	3072	3840	4608	5376	6144
CTA	570	570	570	570	570
DAKOTA	2000	2000	2000	3000	3000
DLS	4000	6000	8000	10000	10000
DUNE	1197	1918	2032	2641	3434
EUCLID	4600	6000	8000	8000	8000
GAIA	383	1000	1000	1000	1000
ISIS	1024	2048	3072	4096	5120
JBCA	32	42	52	62	72
JINTRAC	200	600	900	1000	1200
JLab	1000	1140	1308	1210	1452
LIGO	270	400	5000	5000	5000
LSST	1328	1351	488	435	720
LZ	300	300	1500	2000	2400
PANOSC	0	0	0	0	0
RAYSECT	500	500	1000	1000	1000
SOLID	0	0	0	0	0
CASU	384	768	1152	1152	1536
WFAU	80	80	16	48	80
LSST-Pipeline	0	0	11000	16000	21000
SKA-SRC	640	960	1920	2240	3520
GRIDPP-shortfall	103	9654	14844	27368	39891
Other at 20.0%	4596	8134	14032	18799	23407
Total IRIS requirement	27579	48805	84194	112797	140446

Table 3: CPU-cores - Net shortfall summary with estimated cost in kPounds

CPU-cores	2020	2021	2022	2023	2024
Total IRIS Requirement	27579	48805	84194	112797	140446
IRIS existing	14000	14000	14000	10000	6000
Net IRIS Requirement	13579	34805	70194	102797	134446
Inc.shortfall (10.0% obs.p.a.)	13579	22583	38869	39622	41928
Unit cost (Pounds)	150.0	135.0	120.0	110.0	100.0
Cost (kPounds)	2036	3048	4664	4358	4192

Table 4: Disk-TB-NetRequestTable

Disk-TB	2020	2021	2022	2023	2024
CCFE	30	30	30	30	30
CCP4	0	0	0	0	0
CLFOCTOPUS	140	190	270	350	430
CTA	228	228	228	228	228
DAKOTA	20	20	50	100	100
DLS	2000	3000	3000	4000	4000
DUNE	2323	3799	5004	6505	8456
EUCLID	250	300	300	300	300
GAIA	24	100	100	100	1000
ISIS	1000	2000	3000	4000	5000
JBCA	196	296	396	496	600
JINTRAC	1	10	20	50	100
JLab	0	0	0	0	0
LIGO	0	0	100	100	100
LSST	1894	2406	5529	4249	8345
LZ	300	300	3500	5300	6800
PANOSC	2000	3000	4000	5000	6000
RAYSECT	10	10	20	20	20
SOLID	270	410	0	0	0
CASU	400	1500	2000	2500	3000
WFAU	1296	1000	750	1000	1296
LSST-Pipeline	0	0	20000	31600	43200
SKA-SRC	85	350	1000	2900	4500
GRIDPP-shortfall	0	13000	21000	39000	57000
Other at 20.0%	2493	6389	14059	21565	30101
Total IRIS requirement	14960	38338	84356	129393	180606

Table 5: Disk-TB - Net shortfall summary with estimated cost in kPounds

Disk-TB	2020	2021	2022	2023	2024
Total IRIS Requirement	14960	38338	84356	129393	180606
IRIS existing	14000	14000	14000	10000	6000
Net IRIS Requirement	960	24338	70356	119393	174606
Inc.shortfall (10.0% obs.p.a.)	960	23474	48451	56072	67152
Unit cost (Pounds)	62.0	57.0	52.0	47.0	42.0
Cost (kPounds)	59	1338	2519	2635	2820

Table 6: Tape-TB-NetRequestTable

Tape-TB	2020	2021	2022	2023	2024
CCFE	0	0	0	0	0
CCP4	0	0	0	0	0
CLFOCTOPUS	140	190	270	350	430
CTA	0	0	0	0	0
DAKOTA	0	0	0	0	0
DLS	0	0	0	0	0
DUNE	0	0	0	0	0
EUCLID	0	0	0	0	0
GAIA	0	0	0	0	0
ISIS	1000	2000	3000	4000	5000
JBCA	0	0	0	0	0
JINTRAC	0	0	0	0	0
JLab	0	0	0	0	0
LIGO	0	0	0	0	0
LSST	0	0	2048	8192	12288
LZ	0	0	0	0	0
PANOSC	100	200	400	600	600
RAYSECT	0	0	0	0	0
SOLID	0	0	0	0	0
CASU	0	0	0	0	0
WFAU	0	0	0	0	0
LSST-Pipeline	0	0	9000	20200	31400
SKA-SRC	0	0	0	0	0
GRIDPP-shortfall	0	5000	11000	21000	30000
Other at 20.0%	248	1478	5143	10868	15943
Total IRIS requirement	1488	8868	30861	65210	95661

Table 7: Tape-TB - Net shortfall summary with estimated cost in kPounds

Tape-TB	2020	2021	2022	2023	2024
Total IRIS Requirement	1488	8868	30861	65210	95661
IRIS existing	2000	2000	2000	2000	1000
Net IRIS Requirement	0	6868	28861	63210	94661
Inc.shortfall (10.0% obs.p.a.)	0	6868	22679	37235	37772
Unit cost (Pounds)	32.0	30.0	28.0	26.0	24.0
Cost (kPounds)	0	206	635	968	906

Table 8: GPU-cards-NetRequestTable

GPU-cards	2020	2021	2022	2023	2024
CCFE	16	16	16	16	16
CCP4	0	0	0	0	0
CLFOCTOPUS	48	60	72	84	96
CTA	10	10	10	10	10
DAKOTA	0	0	0	0	0
DLS	20	40	60	80	80
DUNE	0	0	0	0	0
EUCLID	0	0	0	0	0
GAIA	0	0	0	0	0
ISIS	73	130	187	244	301
JBCA	0	0	0	0	0
JINTRAC	0	0	0	0	0
JLab	0	0	0	0	0
LIGO	0	2	4	6	10
LSST	1	2	2	0	0
LZ	0	0	0	0	0
PANOSC	0	0	0	0	0
RAYSECT	0	0	0	0	0
SOLID	0	0	0	0	0
CASU	0	8	16	16	16
WFAU	0	0	0	0	0
LSST-Pipeline	0	0	0	0	0
SKA-SRC	40	60	120	140	220
GRIDPP-shortfall	0	0	0	0	0
Other at 20.0%	41	65	97	119	149
Total IRIS requirement	249	393	584	715	898

Table 9: GPU-cards - Net shortfall summary with estimated cost in kPounds

GPU-cards	2020	2021	2022	2023	2024
Total IRIS Requirement	249	393	584	715	898
IRIS existing	0	0	0	0	0
Net IRIS Requirement	249	393	584	715	898
Inc.shortfall (10.0% obs.p.a.)	249	168	230	189	254
Unit cost (Pounds)	7500.0	7000.0	6500.0	6100.0	5700.0
Cost (kPounds)	1867	1176	1495	1152	1447

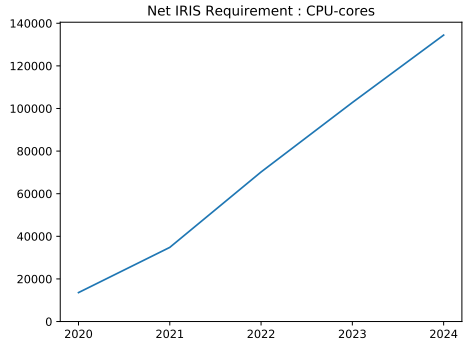

(a) CPU 2020 cores (cumulative).

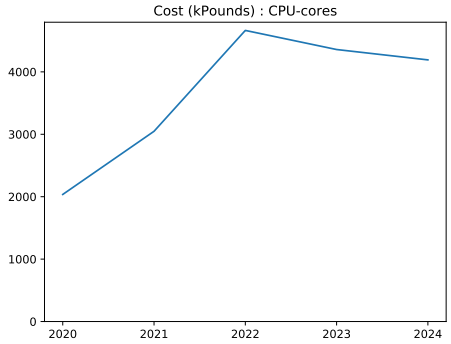

(b) Annual incremental cost £k

Figure 2: CPU requirements

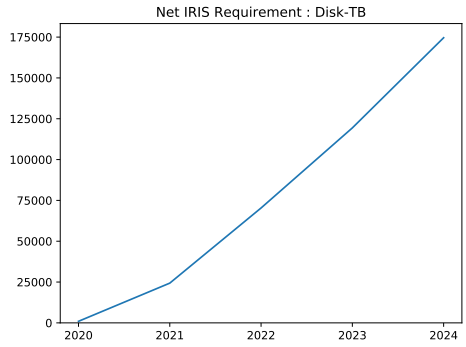

(a) Disk TB (cumulative).

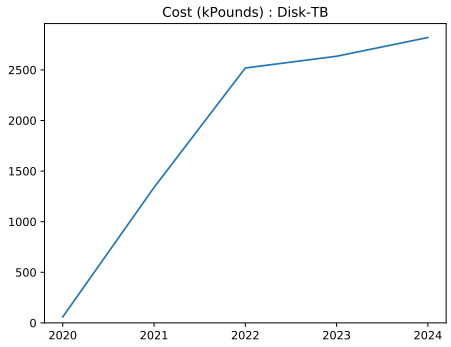

(b) Annual incremental cost £k

Figure 3: Disk requirement

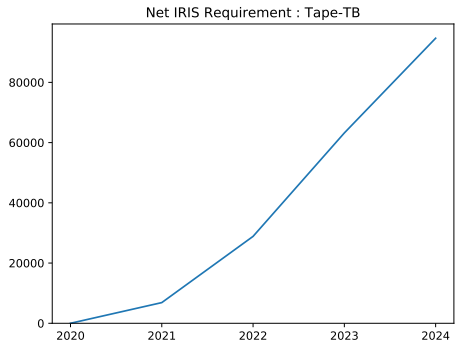

(a) Tape TB (cumulative).

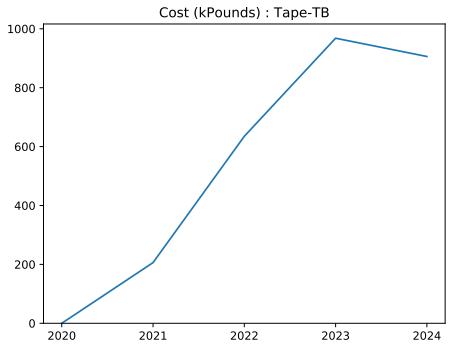

(b) Annual incremental cost £k

Figure 4: Tape requirement

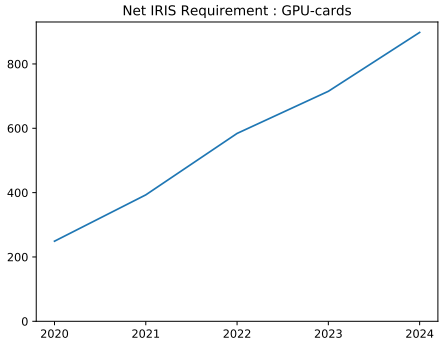

(a) GPU V100 equivalent (cumulative).

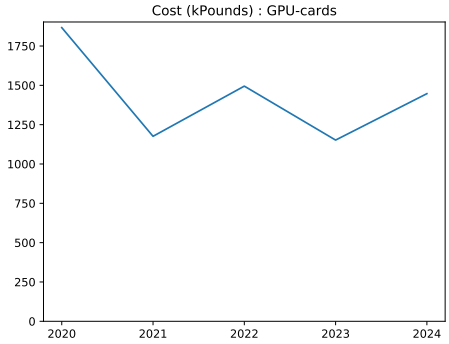

(b) Annual incremental cost £k

Figure 5: GPU requirement