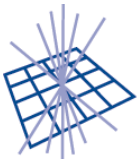

GridPP
UK Computing for Particle Physics

GridPP Project Management Board

GridPP5 Hardware Requirements and Costs

Document identifier:	**GridPP-PMB-166-GridPP5_HW.doc**
Date:	**25/1/2014**
Version:	**1.0**
Document status:	**Final**
Author	**DB**

Introduction

This background paper presents details of the hardware resource requirements used in the GridPP5 proposal.

Computing Resource Requirements

LHC Experiments

The UK contribution to the global resource requirements is based on the size of the UK involvement in each experiment. For the Tier-1 we take the number of UK authors divided by the total number of authors at Tier-1 hosting countries leading to UK contributions, of 2%, 12.5%, 8% and 31.5% for ALICE, ATLAS, CMS and LHCb respectively. For CMS, the number should be 7% but in previous phases of GridPP the UK has agreed to provide 8% as this is the minimum viable size for a CMS Tier-1 site. For the Tier-2s, the denominator is the total number of authors[1] leading to UK pledges of 2%, 10%, 4% and 21.5% of the global resource requirements. However, GridPP provides additional Tier-2 resources to ATLAS and CMS for local (i.e. UK) analysis (LHCb use the Tier-1) and the Tier-2 fractions actually provided are 2%, 12.5%, 5% and 21.5% respectively.
The LHC experiments have presented their computing requirements for 2014 and 2015 to the Computing Resource Scrutiny Group (C-RSG). The numbers have been endorsed and included in the REBUS system[2]; however, following an agreed schedule change to solve an approval/pledge timing problem, the new CRRB cycle will confirm (possibly updated) 2015 requests in spring 2014. For 2016 and 2017, preliminary estimates of the resource requirements by the LHC experiments are documented[3] in a review of the LHC Experiment Computing Models, presented to the LHC Computing Committee (LHCC) in December 2013. The numbers contained in Version-2.5 of that document have been used in the preparation of the GridPP5 resource requirements, together with a more recent update from LHCb containing numbers that will be submitted to the C-RSG in 2014.
The LHC Run-2 starts in the spring of 2015 and extends to mid-2018. However, the GridPP5 proposal covers the period up to March 2019 and must, therefore, provision the hardware that will need to be procured in 2018 for use in 2019. To estimate the hardware requirements for 2018, we have linearly extrapolated the growth estimated from 2015 to 2017 and assumed a 25% increment is needed. Very roughly, this is a flat-cash increment, if one assumes a 2-year Moore's Law growth; and is consistent with the increments between 2015-16 and 2016-17. For 2019 when the LHC will not be running, we have assumed a 5% increment.

Non-LHC Virtual Organisations

In addition to supporting the four LHC experiments, GridPP has traditionally provided computing resources to other particle physics experiments and, contributing to STFC's impact agenda, to other sciences and to commercial entities. Historically, we have been funded to provide 5-10% of overall resources to non-LHC groups and this has proved to be a good match to demand: Over the last 10-years, for example, 9.8% of CPU in the UK went to non-LHC VOs. Storage usage has generally been lower but has been increasing more recently. In estimating the hardware for GridPP5, we contacted all the non LHC-VOs and received some explicit guidance on resource requirements from T2K, ILC and NA62. These have been included explicitly in the planning. For other users, and

[1] Excluding CERN authors for LHCb; this is logically correct, as there is no CERN Tier-2; but a small effect for ATLAS/CMS.
[2] The WLCG REsource Balance and USage system is the official repository of requirements, capacities, and pledges.
[3] https://indico.cern.ch/getFile.py/access?contribId=0&resId=1&materialId=1&confId=212502

potential users, we have assumed that resources at the level of 10% of the LHC Tier-1 and Tier-2 resources should be provisioned.

		From Draft Planning Doc			Extrapolations			From Draft Planning Doc			Extrapolations	
		Tier-1 Resources			1.25	1.05		Tier-2 Resources			1.25	1.05
		2015	**2016**	**2017**	**2018**	**2019**		**2015**	**2016**	**2017**	**2018**	**2019**
ALICE (UK)	CPU [KHS06]	2.4	3.2	4.2	5.3	5.5		4.0	4.8	5.4	6.8	7.1
	Disk [PB]	0.4	0.4	0.4	0.5	0.6		0.4	0.5	0.6	0.8	0.8
	Tape [PB]	0.2	0.3	0.4	0.5	0.5						
ATLAS (UK)	CPU [KHS06]	56.3	69.0	86.4	108.0	113.4		65.0	76.0	91.5	114.4	120.1
	Disk [PB]	4.6	6.1	7.3	9.1	9.5		6.5	9.4	12.3	15.3	16.1
	Tape [PB]	8.1	10.5	13.5	16.9	17.7						
CMS (UK)	CPU [KHS06]	24.0	32.0	42.0	52.5	55.1		25.0	35.0	40.0	50.0	52.5
	Disk [PB]	2.1	2.8	3.6	4.5	4.7		1.5	2.0	2.4	3.0	3.2
	Tape [PB]	5.9	8.0	10.8	13.5	14.2						
LHCb (UK)	CPU [KHS06]	41.6	49.8	59.9	74.8	78.6		15.9	19.1	23.0	28.8	30.2
	Disk [PB]	3.7	4.7	5.1	6.4	6.7		0.5	0.9	1.2	1.5	1.6
	Tape [PB]	8.5	15.7	22.1	27.6	29.0						
LHC Total	CPU [KHS06]	124.2	154.0	192.4	240.5	252.6		109.9	134.9	159.9	199.9	209.9
	Disk [PB]	10.7	14.0	16.4	20.5	21.6		8.9	12.8	16.5	20.6	21.6
	Tape [PB]	22.8	34.5	46.8	58.5	61.4						
T2K	CPU [KHS06]	0.7	0.8	1.0	1.5	2.0		0.1	0.3	0.5	0.7	0.8
	Disk [PB]	0.7	1.0	1.5	2.5	3.5		0.1	0.3	0.5	0.7	0.8
	Tape [PB]	1.0	1.2	1.5	2.0	2.5						
ILC	CPU [KHS06]	0.3	0.5	0.7	0.9	1.0		0.3	0.6	1.0	2.0	3.0
	Disk [PB]	0.1	0.2	0.3	0.4	0.4		0.1	0.2	0.3	0.4	0.5
	Tape [PB]	0.2	0.3	0.4	0.5	0.6						
NA62	CPU [KHS06]	0.5	0.7	0.9	1.1	1.1		1.0	1.2	1.5	1.8	2.0
	Disk [PB]	0.1	0.2	0.2	0.2	0.2		0.1	0.2	0.2	0.3	0.3
	Tape [PB]	0.1	0.2	0.3	0.4	0.5						
Others	CPU [KHS06]	12.4	15.4	19.2	24.1	25.3		11.0	13.5	16.0	20.0	21.0
	Disk [PB]	1.1	1.4	1.6	2.1	2.2		0.9	1.3	1.6	2.1	2.2
	Tape [PB]	2.3	3.5	4.7	5.8	6.1						
Operation	CPU [KHS06]	4.1	5.1	6.4	8.0	8.5		3.7	4.5	5.4	6.7	7.1
	Disk [PB]	1.3	1.7	2.0	2.6	2.8		1.0	1.5	1.9	2.4	2.5
	Tape [PB]	1.3	2.0	2.7	3.4	3.6		0.0	0.0	0.0	0.0	0.0
Grand Total	CPU [KHS06]	142	177	221	276	290		126	155	184	231	244
	Disk [PB]	14	19	22	28	31		11	16	21	26	28
	Tape [PB]	28	42	56	71	75						

Table-1: Preliminary estimates of Tier-1 (RHS) and Tier-2 (LHS) hardware requirements. The resources in 2016-2019 would need to be funded in a GridPP5 project running from 2015-2018.

Computing Resource Costing

The WLCG planning document[4] that contains the experiment hardware estimates also contains a forward look at technology and estimates that there will be a 25% price/performance decrease per year for CPU and a 20% price/performance decrease for disk over the next few years. We have used these assumptions in our calculation of costs, though we have some concerns that there will be a net decrease in the efficiency with which the new many-core CPU architectures can be used.

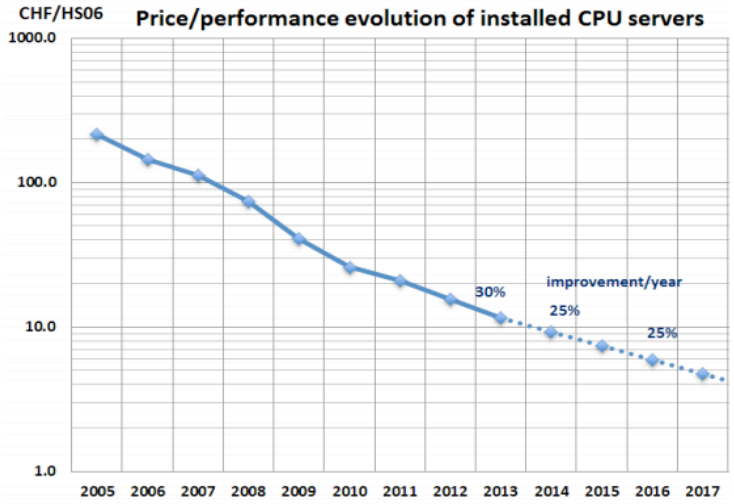

Figure 40: Evolution of price/performance of server systems, example of the CERN computer centre

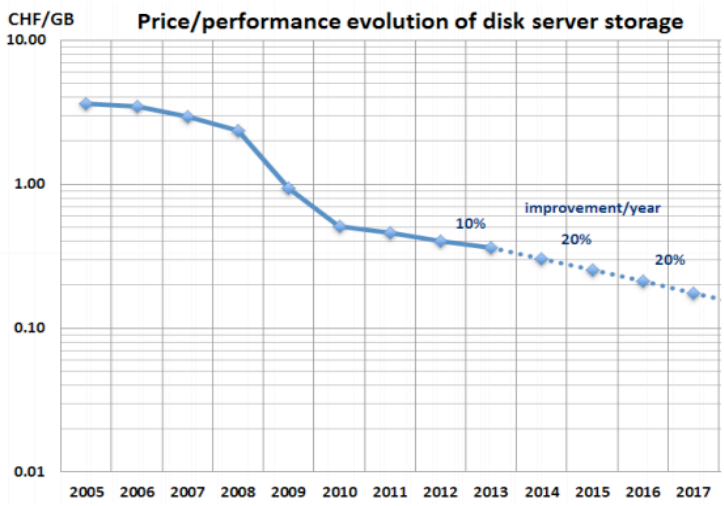

Figure 43: Evolution of price/performance of disk, example of CERN computer centre

The above plots are based on CERN computer centre purchase and a careful look at the technology evolution trends. The plots below show similar data accumulated by GridPP from purchases at RAL:

[4] https://indico.cern.ch/getFile.py/access?contribId=0&resId=1&materialId=1&confId=212502

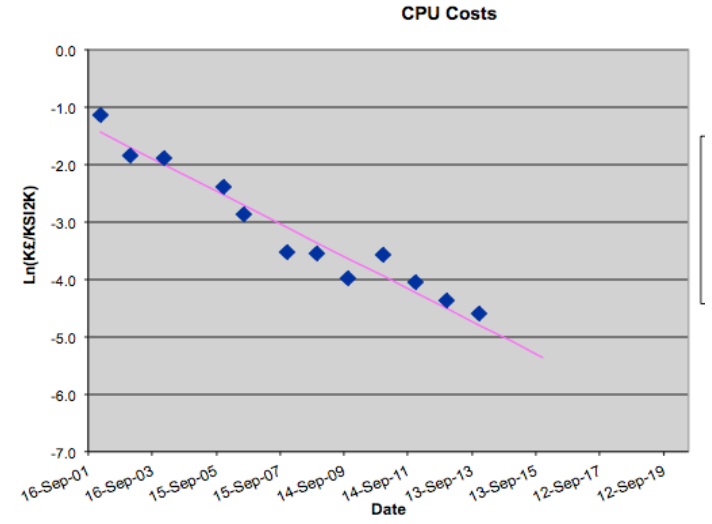

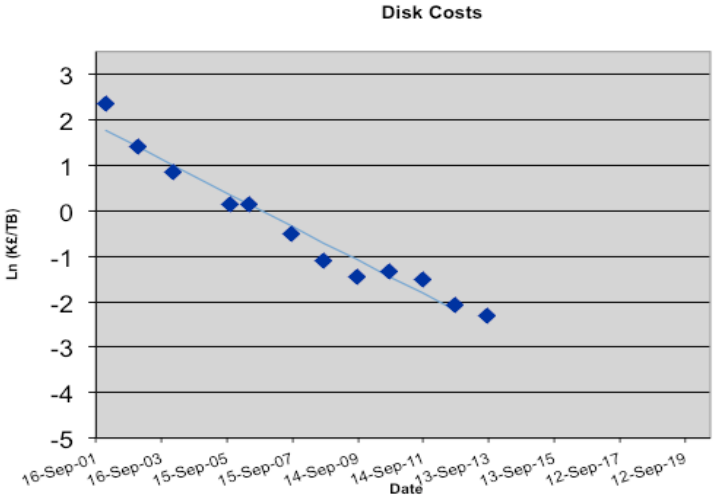

The kink in the CERN disk cost between 2008 and 2010 is probably related to the change in £/CHF exchange rates in the period, though it is interesting that it is not so evident in the CPU cost plot.

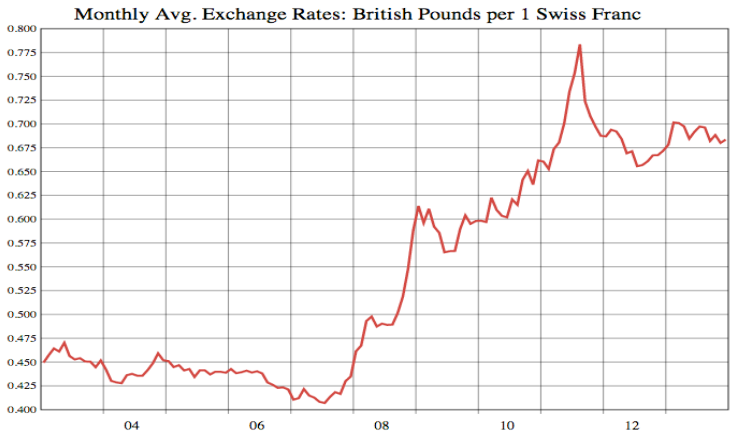

The GridPP data predicts an identical cost evolution for CPU but is somewhat more optimistic about the disk costs. However, the CERN data is not a straight fit, but also considers the likely technological evolution, so we use the CERN numbers in our predictions.

The resources presented in Table-1 are costed in Table-2. For disk and CPU we start from our planned 2014 costs and use the CERN assumption that these will decrease by 20%/year for disk and 25%/year for CPU. We also assume a resource lifetime of 4-years for disk and 5-years for CPU (i.e. slightly longer lifetimes than currently, but consistent with the trend we observe). For tape, we assume, based on publically available information together with some information obtained under non-disclosure agreements, that we can manage, for the likely period of GridPP5, with the current robot installation at the Tier1 and thus that the major cost will be for media. The indicative costing of the tape plan is £1.8m.

Table-2 also contains a line for non-capacity Tier-1 hardware. These are the systems that provide the various services and databases, elements of the internal network and other such overheads. Finally, there is a networking line that covers the cost of the OPN (optical-private-network) that connects RAL to CERN. We assume that we will saturate the current links at some point in GridPP5 and that there will be an increase in cost for a significantly larger bandwidth link. At the Tier-2s, the "DRI" funding in 2012 has allowed us to future-proof most sites until about 2015/6. At that point we assume that some investment in upgrading network links will be needed in the following two years. Similarly, the core network at RAL will need to be upgraded with 100Gb/s-capable routers in 2017.

	FY15	FY16	FY17	FY18	Total
Tier-1 Disk and CPU	£1,203,680	£1,046,402	£1,105,626	£661,311	£4,017,019
Tier-1 Tape	£450,000	£450,000	£450,000	£450,000	£1,800,000
Tier-1 Non-Capacity	£270,000	£280,000	£290,000	£300,000	£1,140,000
Tier-1 OPN Network	£140,000	£140,000	£300,000	£300,000	£880,000
Tier-1 Core Network	£0	£0	£200,000	£0	£200,000
Tier-2 Disk and CPU	£1,133,211	£985,006	£979,094	£565,421	£3,662,731
Tier-2 Networking	£0	£0	£600,000	£0	£600,000
TOTAL	£3,196,891	£2,901,408	£3,924,720	£2,276,732	£12,299,751

Table-2: Estimated hardware costs for GridPP5 (2015 – 2018).

We conclude that the cost of hardware for GridPP5 is about £12.3m, though this number probably has a 20% uncertainty due to the uncertainties in future costs and requirements.