

Modeling PPA Cluster Needs

Introduction

PPA has started buying hardware for the common good across the directorate. This was initiated in 2012 with the first purchase of 185 "Bullet" nodes in early 2013. These are infiniband-connected, with Lustre storage. Historically the cluster was provisioned largely for BABAR, with other experiments riding its coattails. Currently there are three projects of comparable batch allocation size: BABAR, ATLAS and Fermi. BABAR stopped taking data in 2009 and it is presumed that their usage will tail off; Fermi is in routine operations with modest near real time needs and a 1.5-2 year program of intensive work around its "Pass8" reconstruction revamp; ATLAS operates a Tier2 center at SLAC and as such can be viewed as a contractual agreement to provide a certain level of cycles continuously. It is imagined that at some point, LSST will start increasing its needs, but at this time - 8 years from first light - those needs are still unspecified.

The modeling has 3 components:

- inventory of existing hardware
- model for retirement vs time
- model for project needs vs time

A python script has been developed to do the modeling. We are using "CPU factor" as the computing unit to account for differing oomphs of the various node types in the farm.

Table of Contents:

- [Introduction](#)
- [Jump to Summary](#)
- [Needs Estimation](#)
- [Purchase Record of Existing Hardware](#)
 - [Snapshot of Inventory for Modeling - ATLAS hardware removed](#)
- [Retirement Models](#)
- [Summary](#)
- [Purchase Options](#)
 - [Option 1: Expand to 18 fully populated chassis with all-IB and future expansion capability \(revised for increased IB cost \(+6k/chassis\)\)](#)
 - [Option 2: Expand to 15 full IB chassis and 4 non-IB chassis](#)
 - [Notes](#)
 - [Add-on to either option:](#)

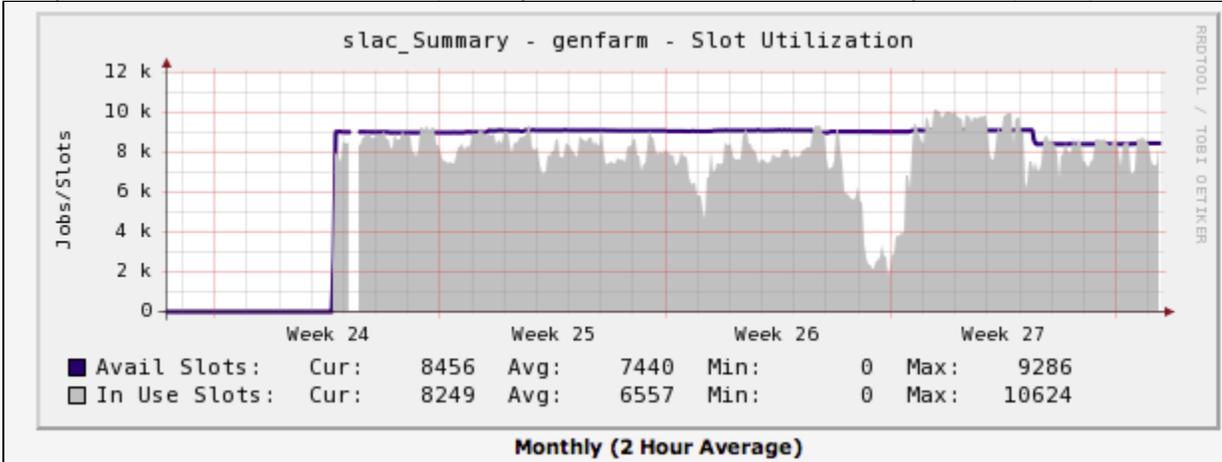
[Jump to Summary](#)

Needs Estimation

PPA projects were polled for their projected needs over the next few years. This is recapped [here](#). Fermi intends to use its allocation fully for the next 2 years; that allocation is sufficient. BABAR should start ramping down, but they do not yet project that ramp. KIPAC sees a need for about 1600 cores for the needs of its users, including MPI and DES. All other PPA projects are small in comparison to these two, perhaps totaling 500 cores. In terms of allocation units, taking Fermi and BABAR at their current levels and an average core being about 13 units, the current needs estimate is:

Project	Need (alloc units)
Fermi	36k
BABAR	26k
KIPAC	21k
other	8k
total	91k

A reasonable model at present is to maintain this level for the next two years as projects such as LSST, LBNE et al start their planning. These are the average needs. We should build in headroom for peak usage - add 20%, to 109k units. Cluster slot usage over the past couple of months is shown here:



illustrating that the cluster is on average heavily used. Week 27 is when Fermi reprocessing began, so we expect full usage for some time to come.

Purchase Record of Existing Hardware

Purchase Year	Node type	Node Count	Cores per node	CPU factor
2006	yili	156	4	8.46
2007	bali	252	4	10.
2007	boer	135	4	10.
2008	fell	164+179	8	11.
2009	hequ	192	8	14.6
2009	orange	96	8	10.
2010	dole	38	12	15.6
2011	kiso	68	24	12.2
2013	bullet	185	16	14.

Of these, ATLAS owns 78 boers, 40 fells, 40 hequs, 38 doles and 68 kiss. Also, note that as of 2013-07-15, the Black Boxes were retired, taking with them all the balis, boers and all but 25 of the yilis.

Snapshot of Inventory for Modeling - ATLAS hardware removed

Purchase Year	Node type	Node Count	Cores per node	CPU factor
2006	yili	25	4	8.46
2008	fell	277	8	11.
2009	hequ	152	8	14.6
2009	orange	96	8	10.
2013	bullet	179	16	14.

Retirement Models

Two models have been considered: strict age cut (eg all machines older than 5 years are retired) and a do not resuscitate model ("DNR" - machines out of Maintech support left to die; up to that point they get repaired). The age cut presumably allows better planning of the physical layout of the data center, as the DNR model would leave holes by happenstance. On the other hand, the DNR model leaves useful hardware in place with minimal effort, but does assume that floor space and power are not factors in the cost.

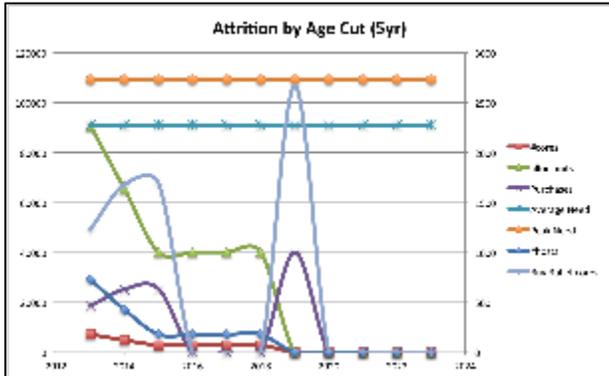
In practice, we may adopt a hybrid of these two, especially since a strict age cutoff would make sudden drops in capacity, given our acquisition history.

There are industry estimates (Fig 2) for survival rates vs time. With a 5 year horizon, machines bought in 2013 only start to retire slowly in 2017. For both tables to follow, the purchases are calculated to maintain the peak leave of 109k allocation units.

Age cut:

Using a strict 5 year cut, here is the survival rate of the existing hardware (in 2019 it is all gone):

Year	#hosts	#cores	SLAC-units	To Buy	Buy "bullets" (cores)	Cost (k\$)
2013	729	7164	90698	18.5k	1.3k	500
2014	427	4848	65476	25.2k	1.8k	700
2015	179	2864	40067	25.4k	1.8k	700
2016	179	2864	40067	-		
2017	179	2864	40067	-		

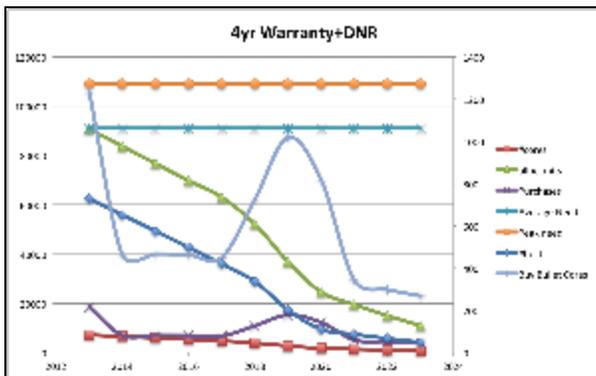


Basically by the end of 2014, all hardware before the bullet purchase would have been retired.

DNR model:

None are allowed after 10 years.

Year	#hosts	#cores	SLAC-units	To Buy	Buy "bullets" (cores)	Cost (k\$)
2013	729	7164	90698	18.5k	1.3k	500
2014	652	6568	83744	7k	0.5k	200
2015	575	5968	76736	7k	0.5k	200
2016	499	5376	69789	7k	0.5k	200
2017	420	4792	63039	7k	0.5k	200
2018	337	3944	52183	11k	0.8k	300



Summary

Projections for PPA's cycles needs for the next few years are flat at 90k allocation units average. We expect the current capacity to be saturated for at least two years due to Fermi reprocessing and simulations needs for Pass8. It is still early for LSST to be needing serious cycles. It would be prudent to have some headroom, hence we recommend a 20% increase for peak bursts, taking us to a need for 109k allocation units.

The installed hardware is already old, except for the 2012-2013 purchase of "bullet" nodes as the first installment of the PPA common cluster purchases. The current capacity just matches the average need; 18.5k units are needed to cover peak usage. Depending on the retirement model, we need to replace the old hardware within two years at 25k units per year, or if via DNR, we need 7k units per year after this year.

If we were to buy the 18.5k units needed, this corresponds to about 1.3k bullet cores. Currently 256 cores cost \$100k, so this could cost \$500k.

We had planned to intersperse storage purchases in with compute nodes. The new cluster architecture is relying on Lustre as a shared file system, and also to provide scratch space for batch. Such an upgrade was anticipated and the 170 TB existing space can be doubled by adding trays to the existing servers for about \$60k.

Purchase Options

This proposal is to expand the bullet cluster with combined funds from PPA, ATLAS, and Theory. This would double our existing parallel file system size (173->346TB) and add either 1649 or 1904 cores depending on which option we choose. The first option is to provision infiniband (IB) in all nodes and add IB switches to allow additional future expansion of the IB network. Because of the IB network topology allowing future expansion implies a jump in the number of core switches from 4 to 8. The second option would split the cluster into IB and non-IB parts with the ATLAS nodes being non-IB. Note the pricing below is based on several different quotes that have been refreshed. Hence the pricing is to be verified but very close to actual. The details are:

Option 1: Expand to 18 fully populated chassis with all-IB and future expansion capability (revised for increased IB cost (+6k/chassis))

- 6 full chassis @97.227k => 583.4k (quote 661664648)
- 7 blades w/IB added to existing empty slots @5.018k => 35.2k (quote 661664256)
- 4 IB switches with cables 50.2k (includes active fiber IB cables for lustre switch) (quote 662008993)
- 2 60x2TB disk trays with controllers @30.3k => 60.6k (quote 661663331)

Total is \$729.3k for 1616 cores and storage expansion.

Gross bullet cluster core count would then be 2960+1648=4608 (all IB)

costs are:

- ATLAS: 60 blades @4.96730k => 296.73k (960c) (Based on quote 659024769 for a non-IB full chassis)
- Theory: 97.227k + 5*5.018k => 122.32k (336c)
- PPA: 310.83k - 2.7k => 308.93 (352c)

Notes:

- PPA cost/core is not meaningful because it includes the storage expansion and subsidizing the IB infrastructure for the Atlas blades

Benefit here is that we have a uniform cluster.

Option 2: ~~Expand to 15 full IB chassis and 4 non-IB chassis~~

- ~~4 full non-IB chassis @79k => 316k~~
- ~~3 full IB chassis @97k => 291k~~
- ~~7 blades w/IB @4.817k => 33.72k~~
- ~~2 60x2TB disk trays with controllers @34k => 68k~~

~~Total is \$709k for 1904 cores and storage.~~

~~Gross bullet IB cores is then 3840.~~

~~Gross bullet non-IB cores is 1024.~~

~~costs are:~~

- ~~ATLAS: 4 @79 => 316k (1024c)~~
- ~~Theory: 97k + 7x4.817k => 130.72k (368c)~~
- ~~PPA: 262k (512c)~~

Notes

- Revised on 8/21 to account for IB price change since the original nodes were purchased. Full IB chassis changed from 91->97 when IB changed from QDR to FDR (increase in performance).
- Need to verify Theory (Hoeche) budget (is 131k too large?)
- Need to verify Atlas budget (option 2 is over 300k)
- Revised 9/3 with latest quotes and to reflect choice of option one and the actual budget estimates.

Add-on to either option:

We could get new GPU servers (kipac's are old!) which are equivalent to bullet blades with ~5000 gpu-cores for ~10k each. ~~So we could top off to 300k if we get 3 of these.~~ We (PPA) do need to replace or existing GPU "system" that is hosted by kipac. A good case for adding some of these is [presented here](#) by Debbie Bard.

Pros and Cons of these options should be discussed.