

Processing and data service requirements: FY 2019

Please brain-dump here requests, requirements and suggestions for moderate-to-large scale processing tasks, storage needs, Science Platform service expansion, etc. that we'll need to undertake during FY2019 (October 2018 through September 2019). These will be used to inform Data Facility procurement.

Summary	Requested by	Estimated compute or storage requirements	Comments
Commissioning team (incl. associated grad students) and Camera team (Science Platform service expansion)	Leanne Guy	Up to ~50 active users Note that stack-club currently has about 8 active users writing notebooks on the lspdev from commissioning and operations	Commissioning and Camera team scientists using lspdev for analysis of test images
Science Collaboration Users (Science Platform service expansion)	Leanne Guy	Up to ~40 active SC users Note that stack-club currently has about 8 active users from science collaborations writing notebooks on the lspdev	LSST science collaboration users are becoming active users of the SUIT notebook aspect on lspdev, especially via stack-club. While we are not scoped to support a large community of users on lspdev, early feedback from a small subset is valuable. I'm thinking about 40 science users from across all the Science Collaborations. We need to distinguish between user access to PDAC (which is intended for phases of formal user testing of full LSP prototypes interspersed with periods of integration work during which user access would be closed), and user access to lsst-lspdev, which is intended to be an ongoing resource for DM and LSST project staff development efforts but has been extended to be accessible to a larger community.
Storing spectrograph test data at LDF	Leanne Guy	~40MB/image, 100 images /day (upper estimate)	Spectrograph test data from Tucson to make available to the wider DM and commissioning teams. Eventually AuxTel data from the summit as well.
DESC 1.2i	Leanne Guy	Subset of ~25 sq deg of the final DESC DC2 dataset. ~2000 images, ~20TB, ~ 1/3 size of HSC PDR1.	NCSA will host this for DESC and DM will use this DC2 subset for testing with LSST-like images and Qserv testing (KPM50)
Gaia DR2	Leanne Guy	~1.2 TB, 1.7 billion rows	Understood that Gaia DR2 was already planned for? http://cdn.gea.esac.esa.int/Gaia/gdr2/
WISE single-epoch	Gregory Dubois-Felsmann	19-20 G rows, or 20-21 TB per year: about 60 TB of raw table data, plus indexes, etc., needed to load Years 2-4.	Once the simplified bulk-loading tools from the Database group are complete, we wish to load Years 2 and beyond (up through Year 4 currently available) of the NEOWISE single-epoch photometry, available here: https://irsa.ipac.caltech.edu/data/download/ . At least Year 2 was supposed to have been accounted for in earlier space requests, but this should be verified. Year 5 may be released during FY19.
HSC RC reprocessing	John Swinbank and the DRP team	Should be possible to get these from Hsin-Fang Chiang based on FY18 activities.	Assuming that this will continue through FY19 at approx. the same cadence as during FY18.
HSC PDR1 reprocessing	John Swinbank , Gregory Dubois-Felsmann	Again, estimate based on FY18	Based on past performance, expect 2-3 full reprocessing runs during this year. We should also plan to load the whole of PDR1 into the PDAC LSP; however, there are concerns about making the data available outside the DM team (we have been asked to limit its use to "engineering" purposes).
HSC PDR2	John Swinbank		I believe that PDR2 will be available in summer 2019, and we'll likely want to reprocess that, but given that that is late in FY19 anyway, and adding some margin for the release date slipping, it may be premature to budget for it.
Storing camera data at LDF	Leanne Guy	Estimate between 60-200 TB to bulk transfer.	Need to get estimates for expected ongoing rate & volume.
SSD for Firefly servers in lspdev	Unknown User (xiuqin)	512GB X2	We want to use SSD for Firefly server cache to improve the file access performance, assuming two Firefly servers, each with 512GB SSD.
Mini-broker testing	Eric Bellm	per discussions at LSST2018 with Unknown User (mbutler) , 3 dedicated K8 nodes	We want to develop the mini-broker architecture and operations concept and understand its performance limits without interfering with other K8 uses.
Regular AP reprocessing	Eric Bellm	Unknown User (emorganson) to estimate based on DES SN dataset	AP analogue of the HSC RC reprocessing.
DAX Web Services	Fritz Mueller	60 cores 4-5GB/core 25-50 GB SSD/core Shared: 10 TB SSD (GPFS/NFS)	Ideally, these requirements will fit with other machines that are deployed in the commons, but in real terms this is roughly two physical machines. Shared disk in GPFS is desired for storing asynchronous outputs from the web services.

PDAC DB node updates	Fritz Mueller	35 nodes, ~14 cores/node; 40 TiB storage/node w/ RAID controllers; at least 384 GiB RAM/node	<p>We need to take PDAC Qserv instance on to next step on the glide-path to production scale; would like to run KPM50 and KPM75 on NCSA with NCSA infrastructure as well as at CC-IN2P3. The plan here would be to get 35 new nodes with contemporary hw and use those for KPM50; later in the year we'd join these with the current nodes for an even more expanded system fo KPM75.</p> <p>Note: no new Qserv czar node is anticipated to be needed in FY19; the two existing ones should be adequate.</p>
Alerts DB	Fritz Mueller	~20 cores; ~5TB storage configurable for object store; budget as capacity in k8s commons?	For prototyping of the Alerts DB (generated alerts, probably noSQL, not the PPDB). Although there are currently no milestones anchoring this, it seems reasonable to expect some development activity and budget some capacity for this.
Summit shared filesystem	Kian-Tat Lim	4 TB SSD; one NFS server machine	Concerned about reliability and uptime, but this should be good enough to start with.
K8 commons expansion to support Science Platform development	Frossie Economou	3x expansion over current usage to about 60 32-core nodes (from the current 20)	<ul style="list-style-type: none"> In the coming year SQuaRE will prototype (and possibly release) ad-hoc dask clusters in jellybean to allow for large catalogue operations from the science platform. The goal is to allow catalogue operations using GAIA DR2 as a data source. Public tutorial sessions will be expanded from ~50 participants to ~250 (eg AAS) The Stack Club, which is being used as an early access for science users to our pipelines and capabilities, is expecting to double its membership by Oct 2019
Object Storage for services running on the Kubernetes commons.	Frossie Economou	5TB storage to back an object store service; object store service OR appropriate privileges for us to deploy a k8s-hosted minio service	<ul style="list-style-type: none"> In order to deploy some of the services we currently deploy on AWS, such as Jenkins agents, in the LDF we need an S3 compatible object store. We need about 1TB of persistent space to back this service We are open to deploying our own object store service on top of k8s (probably using minio) or using an S3 compatible service if that is provided.
Jenkins test system for release manager	Frossie Economou	1 node-equivalent	<ul style="list-style-type: none"> Unknown User (gcomoretto) wishes to deploy Jenkins test instance in the LDF Nebula may be sufficient for this purpose