

Data processing services for FY2020 (Procurement at LDF 2020)

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 FY2020 (October 2019 through September 2020). These will be used to inform Data Facility procurement.

Summary	Requested by	Estimated compute or storage requirements	Comments
<i>What is needed?</i>	<i>Lead person to ask question of</i>	<i>Try to be as accurate as you can.</i>	<i>Why is this needed; where should it go;...</i>
Qserv for AuxTel/Comcam commissioning data connected to lsp-stable – servers with internal disk; How much internal disk in each server (Fritz Mueller Kian-Tat Lim) and I assume a head node? is this 1 or 2? and do you know how much SSD you need in the head node? shall I order one like what is on the lsp-int one today? (same as PDAC?)	Fritz Mueller	servers + internal disks; 1? head node with SSD	This will be for qserv access for the -stable side of LSP for commissioning data and auxtel data.
APDB machines	Fritz Mueller	couple of servers (failover?) with shared disk resources? or internal disks that are replicated between servers? or 1 server for now because it's test?	Alert processing database systems
LSP development (lsp-int)	Simon Krughoff Leanne Guy	Add equivalent of 5 more nodes to the integration cluster.	<ul style="list-style-type: none"> • Currently we have 4 nodes, but because of testing/persistent node problems we rarely have access to all 4 nodes. • We need a larger integration environment to test <code>dash</code> and other user facing technologies. • We had originally planned on 4 nodes. Adding 5 nodes will hopefully give us overhead to have nodes down for testing and such and still almost double the cluster size. • Maybe the integration cluster is the place to do the trade study suggested in the next row. I.e. adding nodes with a different resource profile.
Stack-club / LSP-club support (lsp-stable)	Leanne Guy Simon Krughoff	Add equivalent of 10 more nodes.	<ul style="list-style-type: none"> • We need to take into account that some non-negligible percentage (20%?) of nodes is persistently down because of various reasons. We need to take that into account in our sizing and procurement. • This will bring us back to the original plan of 20 nodes given node downtime overhead. • We are not currently using all of the resources, but activities ongoing in 2020 will increase demand: e.g. AuxTel and ComCam data coming on line, qserv migrating to stable with more interesting data. • We should do a trade study looking at procuring smaller nodes. We currently have 32 core nodes, so when a node goes down, it is a lot of resources. I know a 16 core node is not half the cost, so it's not obviously a zero sum game.
Optimized server pool for Firefly operations UNCONFIRMED	Gregory Dubois-Felsmann	3-4 servers per heavily-used LSP cluster? Probably would request fast (i.e., SSD) local disk.	Experience suggests that Firefly servers run on the existing "vanilla" Kubernetes cluster nodes run significantly more slowly (2-10x slower) than the existing dedicated server on <code>lssst-demo</code> . The reason is not fully understood. Experience at IPAC shows that performance is substantially improved by ensuring that jumbo frames are supported at all layers of the Kubernetes virtualization stack. We have asked for this to be applied at NCSA and are waiting to do further debugging until that has been done. It <i>may</i> turn out that performance is also significantly affected by the availability of fast local disk on the server nodes (as is available on <code>lssst-demo</code>), but this is really difficult to understand until the network performance is improved.

Jenkins	Simon Krughoff	Add equivalent of 2 nodes for dedicated jenkins execution	<ul style="list-style-type: none">• For various reasons, our CI system (Jenkins) imposes requirements on the k8s system that go beyond those of most other execution contexts.• Moving Jenkins to dedicated resources will ease security concerns and allow for a tailored environment for our continuous integration operations.
---------	--------------------------------	---	--