Science Platform Instances - a re-evaluation

This page emerged from a discussion at the 09 Aug 2018 PDAC meeting in which the proliferation of Science Platform instances, and of planned audiences for them, was discussed and raised some concerns from the development teams.

Among the issues to be considered for each use case and each LSP instance which may satisfy a use case, or several, are:

1. Requirements for availability and stability from a user perspective
2. Identity of the intended users (DM / Project / non-Project, etc.)
3. Need for persistent user environments (e.g., software installations, personal data)
4. Datasets to be held
5. Need for Qserv (whether for development or because of the scale of the data to be held)

For the avoidance of doubt: "LSP" here refers to the entire three-Aspect Science Platform, not just to the Notebook Aspect as one sometimes hears in casual conversation.

1. Use Case perspective

1.1. Formal testing of LSP before initial and subsequent deployments to DACs during commissioning and operations

This use case is for carrying out LSP-wide system tests, as well as Aspect-level tests that require near-exact replication of the real DAC deployment conditions. Some of these tests will be associated with Level 2 milestones for DM. After the initial rounds of acceptance testing, it is anticipated that there will be an ongoing need for pre-deployment testing of new LSP software releases. It is likely that in commissioning and/or early operations there will be frequent need for updates to the operational LSP instances, and in order to minimize downtime these must be able to be tested at scale before a public release changeover is made.

1. Must remain stable long enough to carry out prescribed testing.
2. User access only for designated system test personnel
3. Persistent user environments either not needed at all (because all test scripts are held elsewhere) or only for as long as a round of testing takes.
4. Datasets TBD, but must be large enough to stress quantitative requirements, and LSST-like enough to ensure that the test is on-point.
5. Requires Qserv *per se*.

This use case was traditionally intended to be met by the "Integration Cluster", which at present is largely represented by the PDAC hardware.

Even if an operational DAC is not supported during Commissioning, deployments to the Commissioning Cluster will still need to be tested in advance to avoid blocking Commissioning progress.

1.2. Science Platform integration

Because of the limited availability of large-scale hardware platforms for Aspect-level testing, and the current lack of "dummy loads" allowing the interfaces between Aspect to be easily exercised locally by the various Aspects' development teams, the availability of a common integration platform with a significant hardware scale is crucial to continued progress. For this to be useful, it is essential that the integration platform be able to be in a "broken" state from time to time, as major integration challenges are tackled, including a) cross-Aspect integration, b) integration of the LSP components with underlying services such as A&A, and c) R&D in deployment technologies (such as configuration management for Kubernetes-based system deployments).

1. Need remain stable only long enough to verify correct operation of the integration issues at stake at any given time. Maximum availability for testing changes, alternate deployment mechanisms, etc., is useful for developers.
2. User access *during integration itself* is minimal and limited to project staff checking behavior. However, periodically, integration must be verified with larger-scale testing and user exposure as in 1.4 below.
3. Persistent user environments must be present as part of the suite of things being integrated, but there is no long-term requirement for their stability / preservation across rounds of integration work.
4. Datasets must be large enough to challenge the scaling of the components, and LSST-like enough to ensure that the full range of features needed in LSST are explored. Presently the WISE time-domain data and a synthetically enlarged dataset are the ones being used to challenge scaling. The first reasonably LSST-like dataset will arrive with HSC data integration this year (2018).
5. Needs Qserv to ensure integration.

This use case was also traditionally intended to be met by the "Integration Cluster", with pre-deployment testing interleaved with or the result of integration work. Note that increasing the size and fidelity (using mocks/dummy loads) of Aspect-level *development* resources could reduce (but probably not eliminate) the need for this use case.

1.3. Qserv development and test at scale

The further development of Qserv increasingly requires performing testing on platforms with hardware characteristics close to those of the expected final DAC deployments.

1. This stability requirements here are limited to what is needed to support the Qserv group's own work, e.g., it must stay up for long enough to perform challenging performance test series such as those associated with "KPM30".
2. The users are primarily the Qserv team itself.
3. Persistent user environments are not required.
4. Datasets must be large enough to support the scaling and performance testing required. Relatively unrealistic test datasets may still be useful or even desirable, to focus on the specific needs of a test.

5. Qserv is required *per se*.

This use case was traditionally intended to be met by part of the "Development Cluster", although that has never been projected to be large enough to run 30%, 50%, or 100% scale tests. Instead, it had been thought that tests at that scale could be run on temporarily-acquired hardware or on the production hardware prior to deployment in production. (For example, a 100% DR1 Qserv instance was to be procured for use in FY2020 prior to its use for DR Science Validation.)

1.4. Exposure of large scientifically useful datasets through the Science Platform to encourage user evaluation and feedback

This was one of the primary motivations for the allocation of a significant chunk of "Integration Cluster" resources in FY2017 to create the "PDAC" (Prototype Data Access Center) as a means of realizing the FDR-era notions of "three prototype releases" of what was then called the "science user interface" (including "Level 3" support) for science community review prior to the start of operations.

1. A testing environment exposed to science users must remain stable for long enough that meaningful investigations can be carried out. At a minimum we would like our test users to be able to attempt to reproduce analyses they have done on similar data in other ways, to confirm that existing community science use cases can be addressed in the LSP (including all its aspects). Even more desirable would be the possibility that the data provided and the unique capabilities of the LSP may enable a certain amount of new science to be done, as this would really ensure that the testing carried out was need-driven. To enable a new analysis, a period of stability permitting multiple query-analyze-think cycles would be needed.
2. Limited resources for operational support during the present era of construction, before early operations funding becomes available, have always led us to assume that the number of users for this test environment would be small, perhaps O(10), and carefully selected to include a range of scientific perspectives and areas of interest.
3. User data needs to be persistent long enough to permit analysis, but results could be transferred elsewhere and the environment torn down after the end of the testing period.
4. Datasets need to be large and scientifically meaningful.
5. Qserv is essential.

Again, this use case was assigned to the "Integration Cluster", interleaved with other uses, assuming that "Development Cluster" resources were sufficient to enable periods of integration to be short. The increase in demand for use case 1.2, however, seems to be making this plan infeasible.

1.5. Interactive environment for LSST DM developers, especially Science Pipelines, successor to ssh-in lsst-dev systems

A common development system and shared stack are still useful for many purposes, but there are some disadvantages to that model. Some of those limitations were supposed to be handled by per-developer virtual machines obtained via Nebula OpenStack, but that has not (yet) proved to be culturally or operationally adoptable. A notebook environment with pre-deployed stack provides some of the advantages of Nebula with some of the advantages of the shared stack.

This use case was traditionally assigned to a combination of the "Development Cluster" and, to the extent that it is actually use case 1.6, the Science Validation instance of the LSP.

1.6. Science data quality analysis environment

Analyzing the results of simulated productions needs to occur in a production-like environment. This use case may overlap to some extent with 1.5.

This use case was traditionally assigned to the Science Validation instance of the LSP.

1.7. Analysis of AuxTel data in 2019

It is not clear that this is different from 1.6.

1.8. Pop-up demonstrations of the LSP to users at workshops, training sessions, etc.

This is a new use case that was not previously anticipated.

1.9. Commissioning team and other non-DM Project personnel training on LSP and preparation of code for the commissioning era

This is also a new use case that should have been anticipated.

1.10. "Stack Club", DESC, ISSC, etc. access to the LSST Stack to enable preparation for the science era

Unfortunately, this use case, 1.8, and, to some extent, 1.9 are merging and supplanting the carefully-thought-out plans for 1.4 (see LDM-482). It seems that users are being encouraged to use the LSP and Science Pipelines code to experiment with the available precursor datasets or any data that a user brings in on a near-permanent basis. To the extent that the "Stack Club" was actually satisfying 1.9, this was acceptable, but now accounts have been offered to O(100) users, apparently with the expectation that not only their data but also their access to compute resources will continue indefinitely (and, implicitly, with significant stability and support). None of this was budgeted, and it is not clear that it is necessary several years in advance of the availability of data products.

2. Allocated Resources / PMCS perspective

The following makes reference to an FDR-era concept for the division of hardware resources to be acquired during construction. Things have changed *de facto* but here I initially just want to catalog what the picture *in PMCS* is of the scale of resources expected, the number of clusters, and the rough timeline / funding profile for the acquisition and installation of the resources.

For each "cluster" or increment of hardware, we should try to fill in:

1. Funding available for hardware acquisition
2. Funding profile (i.e., in what fiscal year the resources were meant to be acquired)
3. Originally framed purpose
4. Was a Qserv instance meant to be included?

2.1. Integration Cluster

LDM-144 calls for $2M, 10% for FY15, 20% for FY16, 35% for FY17, 35% for FY18 (no refreshes anticipated)

2.2. Development Cluster

LDM-144 calls for $2M, 10% for FY15, 20% for FY16, 35% for FY17, 35% for FY18 (no refreshes anticipated)

2.3. Commissioning Cluster

LDM-144 calls for $200K, all for FY18, with another $200K for FY27 for a refresh

2.4. US DAC

LDM-144 calls for a ComCam/Commissioning purchase for FY19 ($110K), including ~800 TB of Qserv, with a full DR2-sized purchase for FY20 ($1.12M), including ~16 PB of Qserv. Full DR2 production hardware, including things like "Cutout Service Compute", "L3 Community Scratch", the PPDB, etc. was anticipated to be purchased for either FY19 or FY20 (most internal production items in FY19, user-facing items at ComCam scale in FY19 and full scale in FY20).

2.5. Chilean DAC

LDM-144 calls for a full DR2-sized purchase for FY20 ($1.18M)