

LSP Integration and PDAC Meeting 2019-01-10

Date

10 Jan 2019

Attendees

- [Gregory Dubois-Felsmann](#) (apologies today)

Goals

- Advance the re-integration of the LSP components in the merged Kubernetes environment

Discussion items

| Time | Item | Who | Notes |
|------|--|--|--|
| | Status of getting <code>kubectl</code> access to Kubernetes development environment for devs | Michelle Butler | <ul style="list-style-type: none"> • See  IHS-1637 - Jira project doesn't exist or you don't have permission to view it. • Michelle Butler says that nothing should have changed here and that SUIT developers should have <code>kubectl</code> access to the development environment. She has asked for the list of such developers to be explicitly added to that ticket. |
| | Progress in re-establishing the PDAC DAX services and Portal on <code>lsst-lsp-int</code> | Christine Banek Tatiana Goldina | <ul style="list-style-type: none"> • DAX restart completed, see  DM-17059 - Deploy legacy DAX services on <code>lsst-lsp-int.ncsa.illinois.edu</code>  • Starting to work on reconnecting the Portal, partially blocked by IHS-1637, see  DM-17072 - Re-establish the legacy PDAC Portal setup on <code>lsst-lsp-int</code>, using Kubernetes , now assigned to Loi Ly, who should coordinate with Adam Thornton if any updates are needed at the cluster level. • It would be useful to have a discussion about how to actually implement the relocatability of the cross-Aspect connections (so that, by default, e.g., <code>lsst-lsp-stable</code>'s Portal would refer to <code>lsst-lsp-stable</code> DAX services, and <code>lsst-lsp-int</code>'s Portal would refer to <code>lsst-lsp-int</code> DAX services), but allowing for overrides. Adam Thornton has some ideas for this emerging from last week's successful effort to add a configuration override for the Notebook-to-Firefly connection. <ul style="list-style-type: none"> • Notebook need to talk to Firefly; currently handled with override environment variable • Notebook may eventually need to talk to DAX • Portal needs to talk to DAX • DAX shouldn't need to talk to anyone else • Should this be handled with service names? Possibly not if we want users to be able to change it. Also effectively bypasses ingress including authentication; better for internal service connections. |
| | Establishing a basic Firefly service on <code>lsst-lsp-stable</code> | Loi Ly Adam Thornton | <ul style="list-style-type: none"> • What is a good way for the IPAC group to communicate to the LSP "ops team" (currently some combination of NCSA & SQuaRE) what container images are to be used for each instance? (Discuss, e.g., DockerHub image tagging conventions.) • For now, we're trying to have Adam Thornton coordinate, so communication should go through him. As this is likely to prove unscalable, we'll need to come up with official procedures for release and deploy of new versions both in development and in production. |
| | Authentication /authorization progress | Brian Van Klaveren Kian-Tat Lim | <p>It is really urgent to get this functional.</p> <ul style="list-style-type: none"> • Status of rights-checking micro-service • Status of authentication proxy • Both are blocked on getting the ID token through the proxy to services behind it. Most existing services apparently use access tokens instead. Likely need to customize our own fork to pass the ID token or else use the access token to access the <code>userinfo</code> endpoint. ETA still end of this week. • Client Aspects should use JWT authenticators (one exists for JupyterLab) rather than contacting the <code>userinfo</code> endpoint. • Brian Van Klaveren needs new IAM client ids for the OAuth2 proxy; these requests go to CILogon help, not IHS. |

| | | | |
|-------------------|-------------------------------------|--|--|
| (if time permits) | Progress on IVOA-ish DAX services | Fritz Mueller Christine Banek Kenny Lo | <ul style="list-style-type: none"> • Status of SODA service work <ul style="list-style-type: none"> • We would be first to implement SODA and underlying VO services in Python on the server side; have to implement the entire "stack". • Not clear whether calling out from Java to Python is more difficult (wasn't really evaluated; seems to be possible due to Hadoop/Python example, but VO services may need more astronomical concepts in the framework code rather than in external "job" code). • First milestone is DALI; second is images; ETA for completion of both is this cycle (end of S19). • Using existing PDAC datasets as tests. • Status of TAP-QServ connectivity <ul style="list-style-type: none"> • Christine making progress after being blocked on syntax issues with qserv_* pseudo-functions. Expecting other ADQL issues afterwards. Month to 6 weeks before running. Simple queries can run now, but geometry will take a while. Will come up with a list of increasingly-complex queries so that support can be "checked off". • Current portal has multi-cone queries that may be difficult for Qserv. Should be supportable at Qserv layer, although possibly not with good performance (although performance is partly the point). Not a particularly high priority. • Portal can retrieve records for a particular Object or Source, but needs an areaspec? Object should not, as the "secondary index" provides an Object ID to chunk mapping. Source is too big for a "secondary index", but part of its ID could potentially be decoded and used as a spatial hint. We need to prepare for indexing columns within chunks (e.g. magnitudes), once they're identified. |
| (if time permits) | Progress on WebDAV workspace access | Brian Van Klaveren | <ul style="list-style-type: none"> • WebDAV working with SciTokens but not yet ID tokens; should be simple once proxy is fixed. |

Action items

