

PDAC Meeting 2017-02-23

Date

23 Feb 2017

Attendees

- Gregory Dubois-Felsmann
- Fabio Hernandez
- Kian-Tat Lim
- Brian Van Klaveren
- Unknown User (jalt)
- Unknown User (xiuqin)
- Trey Roby
- Igor Gaponenko
- Fritz Mueller
- Kenny Lo

Goals

- Continue progress on PDAC service deployment, discuss user feedback

Discussion items

Time	Item	Who	Notes
	Infrastructure / NCSA	Unknown User (jalt)	<ul style="list-style-type: none">• Discussing how we are going to do OS updates in the PDAC cluster<ul style="list-style-type: none">• How to apply updates without disrupting operations• Gregory Dubois-Felsmann says: we would like 24 hours' notice for unscheduled OS changes and reboots<ul style="list-style-type: none">• For user-facing services we are only promising 15-30• We will start observing NCSA's routine patch day on PDAC: 8-10 Central 3rd Thursday• Trey Roby: please post a summary to #dm-pdac channel<ul style="list-style-type: none">• Particularly concerned to know about reboots and about changes in filesystems• Unknown User (jalt): we can take "exclude lists" of packages not to be updated• Rough prototype of monitoring system running (based on Grafana), writing a document about it<ul style="list-style-type: none">• Possible for users to take a look at it at https://141.142.208.107:3000• Unknown User (jalt) will confirm access/instructions
	imgserv cutout services	Fritz Mueller	<ul style="list-style-type: none">• Kenny Lo's work reviewed and on Github, containers built, working on deployment today<ul style="list-style-type: none">• Brian Van Klaveren will be working on spinning up the new Docker images• Kenny Lo will be writing a Confluence page about the current and new APIs<ul style="list-style-type: none">• Will include "I11" and "I12" as well as some other new work done by John Gates last year• Kian-Tat Lim: please get the build-and-deploy process documented• Fritz Mueller: looking into using Kubernetes for deployment over the next two months, initially testing at IN2P3
	metaserv	Brian Van Klaveren	<ul style="list-style-type: none">• Both old and new versions of metaserv running, but no actual metadata for the PDAC tables.• Waiting for IPAC request for "20 favorite columns"
	WISE database loading	Igor Gaponenko	<ul style="list-style-type: none">• "Object"-like and "ForcedSource"-like table data downloaded and reformatted for MySQL ingest in December, partitioned for Qserv (~150,000 partitions)<ul style="list-style-type: none">• Partitioning was on columns with the plain names "ra" and "dec" in both tables, and on the obvious foreign-key relationship between the tables. Referential integrity & consistency between the two spot-checked.• Overlap and data location will permit spatial joins between the Stripe 82 and WISE Object-like tables out to a 1 arcmin radius<ul style="list-style-type: none">• Some discussion of whether the "substripe" scheme is optimal• Loading "Object"-like table into Qserv - tables loaded, now generating secondary indexes<ul style="list-style-type: none">• Expecting to be able to test access through Qserv by the end of the day• Loading of "ForcedSource"-like tables awaits a restart of Qserv with updated containers (already built)• Igor Gaponenko and Gregory Dubois-Felsmann will recover some work from December to document the IPAC metadata service API
	other database / DAX / SLAC		<ul style="list-style-type: none">• John Gates is working on improved handling / scheduling of Qserv queries with large result sizes

	SUIT / IPAC		<ul style="list-style-type: none">• Trey Roby Early next week we will be ready to deploy our time-series viewer, assuming things work out with the cutout service<ul style="list-style-type: none">• Questions about user table upload (e.g., to support bulk cone-search queries)• Fritz Mueller: Yes, there will be DAX support for this• For now, bulk queries will have to be done by iterating over single queries; should work at the level of 20 simultaneous queries being launched (c.f. underlying <code>websevr</code> 40-process pool)• Note that our new user is already interested in scales more like 10^3 to 10^5 objects in bulk queries.
--	-------------	--	---

Action items

