

Database meeting 2022-03-09

Date

09 Mar 2022

Attendees

- [Fritz Mueller](#)
- [Colin Slater](#)
- [Andy Hanushevsky](#)
- [John Gates](#)
- [Andy Salnikov](#)
- [Fabrice Jammes](#)

Discussion items

Time	Item	Who	Notes
	Project News	Fritz Mueller	<ul style="list-style-type: none">• Refrigeration issues: ongoing meetings this week; project is aiming to resolve a path forward by Thursday.• COVID re-openings proceeding more or less institution by institution. For SLAC folks, our work agreements are already in place and there is no immediate need to change our work patterns, so steady-as-she-goes and we're good for now. Mask and testing mandates becoming more lax at the lab (follow lab emails for latest details.) Use your common sense and stay safe!• Data Facilities planning workshop, three days, first week of April. Fritz will attend full time, Igor may be interested in much of it, perhaps Andy S. for some re. PPDB/APBD/OGAR
	Grand Unified Repo	Fritz Mueller	<ul style="list-style-type: none">• Nate has PR for merging in <code>qserv_testdata</code>; on Fritz Mueller's plate for review• Probably soonest for Fabrice Jammes bringing in <code>qserv-operator</code> would be next week
	DB init	Fritz Mueller Fabrice Jammes	<ul style="list-style-type: none">• On further review, Fabrice Jammes is positive on using mounted volume init technique for db init, but still thinks separating db containers to their own services/pods is not a good plan. One further issue of note: even with affinities set to co-schedule db services and db service clients (e.g. to allow Unix-domain socket), how would stateful set instance indices be coordinated?• Fritz: good point – still thinking on it...
	Subchunk sizes and overlap	John Gates , team	<ul style="list-style-type: none">• John notes full-sky n.n. joins on KPM50 are currently taking, literally, days. With some back-of-envelope calculations, it looks as if moving to ~8x number of subchunks would significantly accelerate this.• Fritz to check on current partitioning parameters to get an idea of overlap storage impact if we did this
	Worker load imbalance and chunk affinity	John Gates , Andy Hanushevsky , team	<ul style="list-style-type: none">• John's n.n. query testing has shown some strange artifacts wrt. load imbalance between workers. Almost linear work distribution from worker with no work to worker with maximal work across all 30 workers soon after query was started. Seems to have leveled out overnight though? Are XRootD chunk affinities acting up, or is this something else?• ABH: need redirector logs to take a look...• Fritz to verify version of XRootD built into instance being tested• Fritz/John to get redirector logs to Andy