Next steps: Packaging a PanDA Distribution

Work will begin to package a minimal PanDA distribution as a Docker stack, so that a working, self-contained PanDA ecosystem can be deployed on a commodity machine as well as a dedicated cluster.

- Build images from latest sources automatically
- Require only Docker to be installed on the host machine(s)
- Include the server, the client, the monitor, and hopefully Harvester
- Install from GitHub-hosted files and/or directly from DockerHub
- Deploy offline (i.e. no calls to AGIS or external endpoints)