Operations Use Cases 2

1. An operator wants to configure and submit a DRP pipeline via the Batch Processing Service using particular inputs for execution on a shared-nothing computing cluster.
2. An operator wants to configure and submit a DRP pipeline via the Batch Processing Service using particular inputs for execution on a computing cluster using a shared filesystem as a staging area.
3. An operator wants to pre-fetch particular DRP input files from tape.
4. An operator wants the log files from a pipeline execution ingested into logstash.
5. An operations staff member wants to find file(s) from the data backbone based upon metadata (e.g., visit, ccd, filter).
6. An operations staff member wants to find file(s) in the data backbone based upon qc values.
7. An operations staff member wants to find file(s) in the data backbone based upon catalog values (e.g., object values).
8. An operations staff member wants to find in the data backbone all output file(s) from a particular pipeline execution.
9. An operations staff member wants to find in the data backbone all input file(s) for a particular SuperTask execution.
10. An operations staff member wants to find in the data backbone all input file(s) for the SuperTask execution that produced a particular output file.
11. An operations staff member wants to monitor current pipeline executions.
12. An operator wants to manually ingest into the data backbone files created externally to the camera or Batch Processing Service to be used by pipelines
13. An operator wants to delete DB rows from bad or test runs that are not needed for historical provenance, e.g. objects.
14. Operations staff member wants to easily run a couple pipeline steps on their own workstation or LSST development cluster on particular inputs.
15. Operations staff member needs to run queries against operational DB to perform quality checks across data.
16. Operations staff member (e.g., Michelle Gower) needs to run queries against operational DB to check for non-science issues such as “bad” machine, dramatic increase in run times, memory, failures, etc.
17. An L1 Archiver at the Base wants to ingest raw file(s) into the data backbone.
18. The L1 Prompt Processing Enclave wants to pre-gather input data from the Data Backbone in order to not affect L1 processing.
19. The L1 Prompt Processing Enclave wants to copy data into the Data Backbone at a cadence that does not affect L1 processing.
20. The Observatory Operations staff wants access to raw files at the base, in particular from their computing cluster.
21. The Science Release Team(?) wants a particular set of data released.
22. Operations wants data backed up for disaster recovery in a manner that makes the data retrievable in a timely manner to continue operations.
23. Operations wants data copied to the LSST DAC in Chile.
24. Operations wants raw files and other DRP input files copied to IN2P3.
25. LSST DAC in IN2P3 wants to serve some subset of the release data.
26. A non-project DAC wants to serve some subset of the release data.