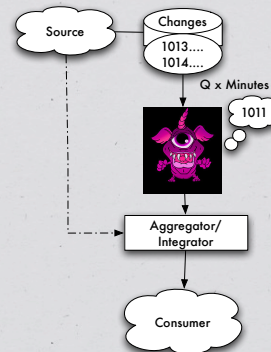


Changelog Consumption [2]

- * Source generates a changelog, assigning monotonically increasing IDs to changes and recording them sequentially. A purple changelog eater tracks its last successful change ID and periodically “rolls forward” changes to aggregator. Changelog entries may be “heavy” (with all associated data) or “light” (aggregator has to retrieve changed data)



Changelog Consumption [2]

* Advantages

- * Reduced latency
- * Relatively cheap runtime
- * Potentially self-healing (if the PCE memory is good)
- * Very auditable
- * Tx integrity assurable

* Disadvantages

- * Still, guaranteed latency
- * Somewhat brittle/policy
- * We don't make no steenkin' changelog
- * Impedance mismatch - atomic vs. composite ops source vs. consumer