LIGO Data Replication Status and Plans for S6

Stuart Anderson, Scott Koranda, Dan Kozak,
Ed Maros, Greg Mendell

LIGO

September 24, 2008

On the question of low latency data replication for S6

- What we can do today
- What we can do by S6
- What we can do by S6 if relieved of other tasks

On the question of low latency data replication for S6

What we can do today...

Following the bits

1. Framebuilder writes 32 s raw frames onto shared disk (32 s)
   - CDS reads back, validates before declaring valid (42 s)
   - Adds ≈ 10 seconds
2. Diskcache detects frames, records metadata & path (2 m 42 s)
   - Latency ≈ 2 minutes
3. Disk2Disk queries Diskcache & copies to `/archive` (3 m 42 s)
   - Latency ≈ 1 minute
4. Diskcache detects new location on `/archive` (5 m 42 s)
   - Latency ≈ 2 minutes

Following the bits

5. LDAS queries Diskcache for new raw frames `(6 m 42 s)`
    - Latency ≈ 1 minute
    - First runs FrCheck and md5sum, adds ≈ 10 + 10 sec `(7 m 2 s)`
6. LDAS computes h(t), RDS 1,3 write to `/archive` `(9 m 10 s)`
    - For S5 used 128 seconds of raw to compute h(t)
    - Note computation goes ≈ real time on $>$ 2 GHz core
    - Question: how many seconds raw frames needed for S6?
7. Diskcache detects new h(t) and RDS frames `(11 m 10 s)`
    - Latency ≈ 2 minutes
8. Pub scripts query Diskcache, publish to local LDR `(12 m 10 s)`
    - Latency ≈ 1 min
    - Checksums computed in parallel (only for RDS)

Following the bits

9. Local RLS updates remote RLS (13 m 10 s)
    - Latency can be dialed down to 1 minute?
10. Remote LDR pulls new metadata
    - Latency can be dialed down to 1 minute?
    - Asynchronous with RLS update above
11. Remote LDR schedules transfer (14 m 10 s)
    - Latency can be dialed down to 1 minute?
12. Remote LDR pulls file (15 m 10 s)
    - Latency can be dialed down to 1 minute?
    - Latency is in the daemon
    - Bandwidth is not issue
13. Remote LDR computes checksums (15 m 20 s)
    - Computed in parallel
    - Only after pass checksum is file available in `/archive`

Following the bits

14. Diskcache detects transferred h(t) and RDS frames (17 m 20 s)
    - Latency ≈ 2 minutes
15. Diskcache hash dumped to ASCII, ≈ 1 minute (18 m 20 s)
16. Dumps mirrored to different nodes, ≈ 30 seconds (18 m 50 s)
17. Dumps loaded into LDRdataFindServer, ≈ 1 minute (19 m 50 s)

At this point frame files available via LDRdataFindServer to analysis pipelines. (20 m)

On the question of low latency data replication for S6

What we can do by S6...

Make Diskcache synchronous

- Enable Diskcache to receive signal indicating new files
- Latency at each step reduced from $\approx$ 2 m to $\approx$ 10 s
- Overall decrease: 7 m 20 s
- New total: 12 m 30 s

Make publishing script synchronous

- Enable publishing script to receive signal indicating new files
- Latency reduced from ≈ 1 min to ≈ 10 seconds
- Overal decrease: 50 s
- New total: 11 m 40 s

LDAS read from /frames and bypass Disk2disk

- Have LDAS read raw frames from `/frames`
- Bypass need for Disk2Disk for h(t)
- Overall decrease: 1 m 10 s
- New total: 10 m 30 s

On the question of low latency data replication for S6

What we can do by S6 if relieved of other tasks

Scalable LDRdataFindServer

- Scalable LDRdataFindServer is being coded in C
- Remove need to dump ASCII hash of Diskcache and ingest
- Overall decrease: 2 m 30 s
- New total: 8 m