

Sprint Review Meeting Minutes

Attendees: Tuntufye Mwakalasya, Deryn Hurst, Myckheal Hosang, Nicholas Merjech, Maksim Kharutski

Start time: 6:00 PM

End time: 7:00 PM

After a show and tell presentation, the expansion upon the following user stories were accepted by the product owners: All.

1. Logic Edge-Case Validation: Completed a full "Time-Travel" simulation in the staging environment, verifying that the 7/4/12 retention logic correctly identifies and protects the specific yearly and monthly snapshots while purging redundant daily files without a single "false positive" deletion.
2. Dashboard Accuracy Audit: Cross-referenced the automated dashboard's projected cost savings against actual cloud billing APIs and storage metrics; confirmed that the financial reporting is accurate within a 1% margin of error before going live.
3. Alerting Resilience Test: Conducted a "Silent Failure" drill by intentionally interrupting the dry-run job's network connection; verified the system correctly identifies the timeout and successfully posts a "Fail: Connectivity Issue" status to the team chat rather than failing silently.
4. Cold Storage Retrieval Drill: Validated the secondary cloud region's transition rule by manually forcing a subset of test data into the Archive tier, then successfully performed a "Restore and Read" operation to ensure data remains recoverable despite the lower-cost storage state.
5. Security Signature Stress-Test: Executed a "Bit-Flip" simulation on a sample backup manifest to intentionally invalidate its HMAC signature; confirmed the security trigger successfully identified the tampering and pushed a high-priority alert to the security channel in under 15 seconds.

The following ones were rejected and moved back to the product backlog to be assigned to a future sprint at a future Spring Planning meeting.

1. I want the pipeline to support multiple database types so I can verify backups across heterogeneous systems.
2. There is need to fix the SHA-256 Hashing algorithm system which is still not as we expect it to be working to meet our expectations
3. I want cryptographic integrity proofs so I can be confident that backups are bit-for-bit consistent.
4. I want performance metrics so I can trend backup/restore performance and spot regressions.
5. Slack notifications are enabled, when a DAG fails.

6. As a compliance officer, I want automated retention rules applied to backup sets (e.g. keep 7 daily, 4 weekly, 12 monthly), so that storage costs are controlled and audit requirements are met without manual cleanup.
7. As a DR engineer, I want to execute a restore in dry-run mode that validates all files and checksums without writing to disk, so that I can rehearse recoveries confidently without touching production.
8. As a resilience engineer, I want backup sets automatically replicated to a secondary cloud region or storage tier, so that a single-region outage does not result in data loss.
9. As a security engineer, I want every backup manifest signed with an HMAC key and written to an immutable audit log, so that tampering with backup records can be detected and investigated.
10. Demonstrated a new automated dashboard/report that calculates exactly how much storage space and money is being saved each week by the 7/4/12 retention cleanup.
11. howcased a scheduled job that automatically runs the zero-impact dry-run validation every Sunday night and posts a simple "Pass/Fail" health summary directly to the engineering team chat.
12. Confirmed a new lifecycle rule where backups in the secondary cloud region automatically transition to a cheaper "cold storage" tier (like AWS Glacier or Azure Archive) after 30 days, further optimizing disaster recovery costs.
13. Validated a new security trigger that instantly sends a high-priority alert to the security team channel if a backup manifest's HMAC signature ever fails validation, ensuring immediate response to potential tampering.