

# Run the drill again.

A 90-day restore drill for ZFS environments. Run it, time it, write the numbers down.

**Why every quarter now.** AI-enabled adversaries increased operations 89% in 2025 (CrowdStrike 2026 Global Threat Report). A frontier AI model has found zero-days in every major OS (Anthropic, April 2026), so patch cycles shorten and restore paths change more often. Inside your own walls, AI agents with production credentials have already deleted live data (Replit, July 2025). An annual test was sized for neither.

## 1 BEFORE THE DRILL

- ☐ Pick one production dataset. Write down its RTO and RPO targets.
- ☐ Stand up an isolated restore host with no route back to production.
- ☐ Confirm encryption keys for raw sends are stored outside the pool you protect.

## 2 RESTORE

- ☐ Find the newest replicated snapshot on the target. Is it inside your RPO?  
`zfs list -t snapshot -o name,creation -s creation`
- ☐ Send it to the isolated host without mounting.  
`zfs send -w pool/data@snap | ssh drhost zfs receive -u drpool/restore`
- ☐ Load keys, mount, and start the application against the restored data.
- ☐ Stop the clock when the application serves real requests. Compare to RTO.

## 3 VERIFY PROTECTION

- ☐ Scrub the DR pool and check for errors.  
`zpool scrub drpool; zpool status -v drpool`
- ☐ Critical snapshots carry holds.  
`zfs holds pool/data@snap`
- ☐ The replication account cannot destroy data.  
`zfs allow pool/data`
- ☐ AI agents and automation accounts have the narrowest permissions that still work.

## 4 AFTER

- ☐ Write down what broke and how long the fix took.
- ☐ Put the next drill on the calendar, 90 days out.

Date

Dataset

RTO target / actual

RPO target / actual

**Want Klara engineers in the room for the next one?**

The Quarterly ZFS Health Check includes a timed restore.  
[klarasystems.com](https://klarasystems.com)