



barydata

Storage without a center

Technical Whitepaper

Datarecovery.com, Inc. · Edwardsville, Illinois
September 2026 · barydata.com

Abstract. barydata is a file sharing system in which no single machine, including the coordinating server, ever holds a readable copy of a user's data. Files are encrypted in the sender's browser, divided into fixed-size authenticated records, and placed in stripes across a network of storage nodes that hold only ciphertext. The decryption key travels exclusively in the URL fragment of the share link and is never transmitted to any server. Downloads reassemble a file from many nodes in parallel, so aggregate speed approaches the sum of the holders' uplinks rather than the speed of any one machine. The design assumes unreliable, untrusted, and slow nodes as the normal case: every stripe is anchored to an operator-class machine, redundancy is repaired within minutes of a node disappearing, and no node is trusted on its word about anything, including its own software version. This paper describes the architecture, the cryptographic format, the placement algorithm, and the measured behavior of the deployed system.

1 Design Goals

Five requirements shaped the system.

- **Instant shareability.** A share link must exist the moment sending begins, not after an upload completes. Recipients may begin downloading while the sender is still transmitting.
- **No readable data off its origin.** Plaintext exists only on the machine the file came from and in the browsers of link holders. Every other party, the coordinator included, handles ciphertext it cannot decrypt.
- **No storage at the center.** The coordinator maintains placement records, share locators, and sealed manifests. It relays data during transfers but retains none of it.

- **Commodity, untrusted machines.** Any computer may join the network and hold data without being trusted with anything, and machines are expected to disconnect frequently.
- **Speed from parallelism.** A download draws on every machine holding a piece of the file simultaneously, so slow contributors add to throughput instead of limiting it.

2 System Architecture

The network consists of one public coordinator (the hub) and any number of storage nodes. The hub is the only element with a public address. Nodes open no ports, require no VPN, and need no network configuration of any kind: each node dials out to the hub over HTTPS and holds several requests open, and the hub answers one of them whenever it has work for that machine, such as reading a range of a stored piece, storing a new record, or reporting status. The node performs the work and posts the result back. A machine behind NAT, a hotel network, or a corporate firewall therefore participates exactly like any other. Every request a node makes is authenticated with an HMAC over the method, target, timestamp, and body digest, keyed by a per-node secret issued at enrollment and valid within a five minute window.




Figure 1. The coordinator is the only public element. Content moves as ciphertext; keys move only inside links.

Enrollment is open by default. A machine that joins without an operator token is a *holder*: it stores and serves ciphertext and can do nothing else. Publishing files into the network is reserved for machines the operator has explicitly trusted. Capability is verified behaviorally rather than believed: before the coordinator entrusts a machine with a piece of work its software version implies it can do, it poses a request whose answer proves it, and a machine that fails the proof, or fails the same transfer repeatedly, is routed around.

3 Cryptographic Format

Files are processed as 4 MiB plaintext records, each encrypted with AES-256-GCM under a per-file key generated where the file lives. A record is stored as a 12 byte nonce, the ciphertext, and a 16 byte authentication tag.

Deterministic records. The nonce is synthetic: it is derived by keyed hash from the record index and the plaintext itself, using a subkey obtained from the file key. A nonce therefore never repeats for different content, and identical content always produces byte-identical ciphertext. Determinism is what makes multi-source assembly safe: every machine holding a record holds the same bytes, so a downloader can take any 256 KiB block of a record from any holder and the assembled record still authenticates.

Bound structure. The record index and total record count are bound into each record as authenticated associated data, so records cannot be reordered, substituted between positions, or truncated without detection. The receiving browser verifies every record before writing a byte of plaintext; a record that fails is discarded and fetched again from a different holder.

Keys never reach a server. The share link has the form `https://barydata.com/s/<locator>#<key>`. Browsers do not transmit the fragment after the # sign, so the coordinator resolves the locator and serves ciphertext without ever seeing the key. File names and sizes travel in a manifest sealed under the share key; the coordinator stores the sealed blob and cannot open it. Where the coordinator must recognize a secret later, such as a sender's session or the key that authorizes deleting an upload, it stores only a hash of it.

4 Placement: Stripes, Anchors, and Churn

Placement never handles whole files. A file is divided into *stripes* of about 256 MiB (runs of consecutive records), and each stripe is placed, copied, healed, and deleted independently. The stripe grid is fixed at the moment a file is created and never changes for that file. Striping removes the single-machine ceiling: a file may be larger than any machine in the network, since no machine ever needs room for more than a stripe, and capacity is bounded by the network's total free space divided by the number of copies kept.

Each stripe is maintained at a target number of copies on distinct machines, and two further placement invariants hold per stripe:

- **An anchor.** At least one copy of every stripe resides on an operator-class machine, selected steadiest-first by measured uptime. Unknown machines may leave at any moment without a stripe ever depending on them.
- **A fast holder.** At least one machine currently online and at or above the median measured speed holds every stripe, so no download is ever served exclusively by slow machines.

Churn handling is class-aware. An operator machine remains counted through a long outage; an unknown machine that has been silent for fifteen minutes stops counting and replacement begins immediately. Uptime itself is measured continuously by the coordinator rather than reported by the machine. Additional safeguards: the coordinator deducts space it has already promised a machine from that machine's reported free space, so room is never handed out twice while copies are in flight; each machine is given at most about ten minutes of transfer work at its own measured speed, so slow machines absorb work at their own pace; and the coordinator periodically asks each holder to confirm a piece it is recorded as holding, so a machine whose storage was wiped or replaced is detected and its stripes are re-copied without waiting for a download to hit the gap.

5 Data Transfer

Upload. The sending browser encrypts records locally and transmits them in order through the coordinator. The link is returned immediately, before the first record arrives. The first record of each stripe selects the machine that will receive that stripe, chosen among machines with room for a stripe rather than for the whole file, so a long upload spreads across the network as it proceeds. A machine that turns out to be full hands the stripe to another without the sender noticing. Several records are kept in flight so the sender's uplink never idles, and a receiving machine that restarts mid-upload causes an automatic resend from the last durable record.

Download. For each record, the coordinator identifies every machine whose stripe covers it and requests 256 KiB blocks from all of them in parallel, two outstanding blocks per holder. The next block goes to whichever holder frees up first, so fast machines naturally carry more of the transfer; a holder whose measured speed implies it would deliver late stops receiving new blocks, and a late block is re-issued to a faster machine. The receiving browser verifies each record cryptographically, decrypts in place, and streams the file to disk. In benchmark, four holders with uplinks of 2, 6, 6, and 0.4 MB/s delivered about 13 MB/s to a single browser, roughly ninety percent of their combined capacity, with the slowest machine contributing usefully rather than limiting the rate. A holder killed in the middle of a transfer costs nothing but its remaining blocks, which are reassigned automatically.

Local network path. Optionally, a recipient behind the same public address as a holder can fetch directly from that machine over HTTPS on the local network, authorized by per-file tokens the coordinator signs, with automatic fallback to the coordinator. Only ciphertext moves on this path as on every other.

6 Accountability

Anyone holding a link may open it; only sending is gated. Before creating a link, a sender verifies a mobile telephone number by texted code or an email address by emailed code, and that verified identity and the sender's network address are recorded with every link created. Verification codes are stored only as salted hashes with short lifetimes and bounded attempts, and issuance is rate-limited per destination, per address, and globally. Links are read-only and temporary: the sender chooses a lifetime between 1 and 30 days, and when the last link to an uploaded file ends, every piece of it is deleted from every machine that held one. The operator can revoke any link and bar any sender. barydata is a sharing mechanism, not an archive, and nodes perform no processing on the data they hold.

7 Assumed Failures

The system treats the following as ordinary events, exercised continuously in an automated test suite of several hundred checks: a node disconnecting mid-transfer; a node returning with its storage wiped; a node lying about its capabilities; a machine with almost no free space; every machine being individually too small for a file; the coordinator restarting under load; and sustained connection churn. In testing, a vanished unknown machine's stripes are re-placed in under half a minute, and a download in progress survives the loss of any holder.

8 Roadmap

Public contribution of storage is not yet open; the network currently runs on operator machines while the contribution experience is completed. Planned next: iPhone and Android applications, public node enrollment with the trust classes described above, and capability-bound release, under which a node would release ciphertext only upon proof derived from the link itself, so that even a compromised coordinator could not collect a file's pieces. The architecture described here is the subject of a pending United States patent application.

Datarecovery.com, Inc. has provided professional data recovery and digital forensics services for approximately three decades. barydata is built by its engineers. Contact: support@datarecovery.com.