

Backblaze Field Guide

Pitch, whiteboard, price and win B2 deals.
For Solutions Engineers and field sales.

Fred Nix · backblaze.nixfred.com

Edition 2026-10-17 · guide version 498215a6cb62

Unofficial and independent. Not affiliated with, endorsed by, or sponsored by Backblaze, Inc. Backblaze, B2 and related marks are trademarks of Backblaze, Inc., used only to identify its products. Prices and facts are from public sources on the dates shown; verify before use.

Contents

LEARN

1. **The pitch** 30-second and 2-minute pitches, by persona
2. **Company context** What Backblaze is in 2026: numbers, strategy
3. **ICPs and personas** Who buys, why, what they fear
4. **B2 Cloud Storage deep dive** Every feature and how it actually behaves
5. **Overdrive, B2 Neo, MCP** The AI and neocloud motion
6. **Computer Backup** Endpoint backup, Business, Enterprise Control
7. **Industry plays** Edu, healthcare, SLED, finserv, media, surveillance, gaming, MSP
8. **Feature matrix** What lives in which product; feature-to-pain map

WHITEBOARD

1. **Whiteboard fieldcraft** Ask, draw, validate
2. **Board: backup and ransomware** 3-2-1-1-0 with Object Lock
3. **Board: media workflows** Camera to archive to delivery
4. **Board: AI training data lake** GPU clouds, Overdrive, B2 Neo
5. **Board: CDN origin** Zero egress delivery through partners
6. **Board: hybrid and migration** NAS, tape, S3 to B2

EXECUTE

1. **TCO and economics** Formulas, worked examples, live calculator
2. **Competitive battlecards** S3, Wasabi, R2, Azure, GCS, on-prem, MinIO
3. **Partner plays** Veeam, Rubrik, Commvault, CDNs, MSPs
4. **Field lab** Discovery, qualification, objections, drills, flashcards, call plans
5. **POC and demo scripts** b2 CLI, aws CLI, boto3, rclone, Object Lock
6. **Migration playbook** Universal Data Migration, Fireball, rclone, cutover
7. **Deal story templates** How to write and tell a win
8. **Outreach templates** Recaps and follow-ups
9. **First 90 days** SE operating plan and metrics

REFERENCE

1. **Reference** CLI and API reference, technical deep Q&A
2. **Security and compliance FAQ** SOC 2, HIPAA BAA, GDPR, encryption, keys
3. **Customer FAQ quickfire** One-line answers
4. **Glossary** Terms an SE must use correctly
5. **Sources** Every URL, date checked

01. The pitch

Rule: the pitch is a hypothesis about their pain, not a product tour. Say it, then ask a question.

The 30-second pitch (default)

"Backblaze is S3-compatible object storage that stays hot and stays cheap. One storage class at \$6.95 a terabyte a month, no tiers, no retrieval fees, no minimum retention, and egress is free up to three times what you store. If you use Cloudflare, Fastly, CoreWeave, or Vultr, egress through them is free, period. Teams use it for immutable backup copies, media archives they actually touch, and AI datasets that need to move between GPU clouds. Most customers point their existing tools at a new endpoint and they are done. Where is storage cost or egress hurting you today?"

Why it works:

- Three numbers a CFO can repeat: \$6.95, 3x, \$0.
- "Point your existing tools at a new endpoint" kills migration fear early.
- Ends on their pain.

30-second variants by persona

CIO / VP Infrastructure "Your cloud storage bill has two lines that grow faster than the data: retrieval and egress. Backblaze removes both for most workloads. One always-hot tier at about a third of S3 Standard list, free egress up to 3x stored, SOC 2 Type 2, HIPAA BAA available, and Object Lock for ransomware-proof copies. It is a public company, 100,000+ B2 customers, and CoreWeave just signed a \$335 million multi-exabyte deal for it. What would you do with 60 to 70 percent of your object storage budget back?"

Storage / backup admin "If you run Veeam, Rubrik, Commvault, or a Synology, you can add an immutable off-site copy in about twenty minutes. B2 is an S3-compatible target with Object Lock, so your backup software sets the retention and nobody, including a compromised admin, can delete those restore points in compliance mode. Restores are not throttled or billed like Glacier. How do you test restores today, and what does a full restore cost you?"

ML platform lead "Your GPUs cost more per hour than your storage costs per month. Backblaze keeps the dataset in one S3 bucket and lets you pull it into any GPU cloud without paying egress tax: free to CoreWeave and Vultr through partnerships, and with B2 Overdrive you get dedicated throughput from 100 Gbps up to 1 Tbps with unlimited free egress anywhere. Decart went from 0 to 16 PB in 90 days on it. Where does your training data live today, and where do your GPUs live?"

MSP owner "You sell backup with margin. Backblaze gives you a storage cost you can quote a year out: \$6.95 a terabyte, no egress surprises when a client needs a full restore, Object Lock for the cyber insurance checklist, and a Partner API to provision and bill sub-accounts under your brand. Which of your backup stacks still lands on AWS or a NAS in a closet?"

Neocloud / GPU-as-a-Service operator "Your customers want object storage next to your GPUs. Building it on Ceph means hiring a storage team and owning hard drive operations for years. B2 Neo is white-label object storage you provision through the Partner API, with custom domains, SSO, and no egress fees. CoreWeave uses Backblaze for HDD tiers in its AI Object Storage. How many of your customers keep data in a hyperscaler and pay egress to reach you?"

The 2-minute pitch (whiteboard-ready)

Draw three boxes as you talk: **Data sources** on the left, **B2** in the middle, **Consumers** on the right.

1. **The problem (20s).** "Object storage pricing got complicated on purpose. Hyperscalers sell cheap cold tiers, then charge you to read, to retrieve, to delete early, and to move data out. The more useful your data is, the more you pay. That is backwards for backup restores, media reuse, and AI training, where the whole point is to read the data."
2. **What Backblaze is (25s).** "Backblaze B2 is S3-compatible object storage built on Backblaze's own hardware and software, run in data centers in Sacramento, Stockton, Phoenix, Reston, Amsterdam, and Toronto. One class: always hot. 11 nines designed durability using 17 data plus 3 parity shards across 20 storage pods. Over an exabyte under management."
3. **The economics (25s).** "List is \$6.95 per TB per month, billed per byte-hour. No minimum object size, no minimum storage duration. Free egress up to 3x your average stored data, then a penny a gig. Free egress, unlimited, through Cloudflare, Fastly, bunny.net, CacheFly, CoreWeave, Equinix Metal, Vultr, and phoenixNAP. API calls are free except Class D event notifications. Term deals of 1 to 5 years lock the rate."
4. **What it plugs into (20s).** Draw arrows from left box. "Backup software, NAS, MAM systems, data pipelines. Hundreds of integrations. If it speaks S3, change the endpoint and keys."
5. **Protection (15s).** "Object Lock in governance or compliance mode, legal hold, SSE-B2 encryption on by default for new uploads since September 2026, SSE-C if you hold the keys, scoped app keys, Cloud Replication across regions."
6. **The AI layer (15s).** Draw arrows right to GPU clusters. "For AI and HPC, B2 Overdrive gives dedicated vaults and 100 Gbps to 1 Tbps throughput with unlimited free egress starting at \$15/TB with a multi-PB commit. For GPU clouds, B2 Neo is white-label."
7. **Close (10s).** "Let's pick one workload and price it on your numbers, then prove it with a POC. Which workload costs you the most to read today?"

Pitch landmines (do not say these)

- "Backblaze is cheaper than Glacier." It is not, per GB stored. Win on retrieval, latency, and simplicity instead.
- "Unlimited free egress" without the qualifier. B2 is 3x free, then \$0.01/GB. Unlimited is partner egress and Overdrive.
- "Multi-region by default." A B2 account lives in one region. Multi-region needs Cloud Replication to a second account.
- "Same as S3." It is S3-compatible with documented gaps: no object-level ACLs, no IAM roles, no object tagging, no website configuration, no browser POST uploads to presigned URLs (docs checked 2026-10-14).
- "HIPAA certified." There is no such certification. Say "BAA available for covered entities".
- "FedRAMP." Backblaze lists GovRAMP Progressing Snapshot and TX-RAMP Provisional. Not FedRAMP.

02. Company context: Backblaze in 2026

Know the company well enough to answer “will you be around?” in one breath.

One-breath answer

“Backblaze is a public company, Nasdaq BLZE, founded 2007, profitable on an adjusted EBITDA basis at a 30 percent margin last quarter, with \$177 million ARR growing 21 percent, B2 growing 34 percent, and a \$335 million multi-exabyte contract with CoreWeave signed in June 2026.”

Snapshot (sources in 23-sources.md)

Fact	Value	Source / date
Ticker	Nasdaq: BLZE	ir.backblaze.com
Founded	2007, Palo Alto apartment	Backblaze blog author bio, CoreWeave post
CEO	Gleb Budman, co-founder, CEO since 2007, Chair since Jan 2021	backblaze.com/company/about
CFO	Marc Suidan (since 2024)	about page
CRO	Anuj Kumar (since Apr 13 2026), ex NetApp cloud, SUSE, HUMAN Security	IR press release
CTO	Dan Spraggins	Investor Day 2026 transcript
B2 customers	100,000+	backblaze.com/cloud-storage
Total customers	500,000 in 175+ countries	backblaze.com sales pages
Data under management	“over an exabyte”	CoreWeave announcement bio
Q2 2026 revenue	\$42.7M, +18% YoY	Q2 2026 results, Aug 3 2026
B2 revenue Q2 2026	\$26.6M, +34%, 62% of revenue	same

Fact	Value	Source / date
Computer Backup Q2 2026	\$16.1M, -2%	same
ARR	\$177.3M, +21%	same
Adj. EBITDA	\$12.8M, 30% margin	same
FY2026 guidance	\$172M to \$174M	same
Q3 2026 results	scheduled Nov 2, 2026	company.md notes [verify date on IR site]

Strategy in plain words

1. **From backup company to storage layer for AI.** The home page now reads "The object storage layer for AI and data-intensive workloads." Expect every enterprise conversation to ask "what is your AI story?" even when the workload is backup.
2. **Upmarket motion.** New CRO with a NetApp cloud background, a pilot of Solutions Engineers, and a plan for roughly 60 more GTM hires (recruiter note, [unverified publicly]). Guidance excludes deals over \$500K, which makes large deals pure upside. This is where SEs matter.
3. **Neocloud channel.** B2 Neo launched Feb 2026. Backblaze says 25 percent of its data transfer goes to neoclouds, not hyperscalers (B2 Neo page video title). CoreWeave deal: five and seven year order forms, ~\$335M TCV, multi-exabyte, B2 capacity in Backblaze DCs plus managed storage inside CoreWeave DCs, powering HDD tiers behind CoreWeave's LOTA cache with no code changes for CoreWeave customers.
4. **Flash plus capacity, not flash versus capacity.** Backblaze positions under VAST and WEKA (WEKA is a partner; Hume AI case study lists WEKA). Blog headline Sep 24 2026: "Why the Best AI Clouds Don't Run on Flash Alone."
5. **Agent-era tooling.** B2 MCP Server (open source, Sep 8 2026), a Claude skill, Pixeltable, CVAT, MLflow-style artifact store guides.

Recent news an SE should be able to mention (last 90 days)

Date (2026)	Item
Aug 3	Q2 beat, FY guidance raised to \$172M to \$174M
Aug 27	SSE-B2 default encryption announced for new uploads, live Sep 14

Date (2026)	Item
Sep 8	Backblaze B2 MCP Server released, open source, 40 tools
Sep 10	Suite Studios integration: mount B2 as a high-performance drive
Sep 24	Blog: "Why the Best AI Clouds Don't Run on Flash Alone"
Sep 29	Q2 2026 Drive Stats: quarterly AFR 1.73% (up from 1.24%), 350K+ drives, 20TB+ drives are 25%+ of fleet
Oct 1	New Jamf deployment options for Business Computer Backup

Drive Stats: the credibility asset

Backblaze has published hard drive failure data quarterly for over a decade. It is free, it is cited by the whole industry, and it proves one thing to a buyer: this company operates hundreds of thousands of drives and tells the truth about them, even when the number goes up (Q2 2026 AFR rose to 1.73 percent). Use it as a trust signal, not as a durability claim. Durability comes from erasure coding, not low AFR.

Why this matters at the whiteboard

- "Will you be around?" Public, adj. EBITDA positive, multi-year contracted demand.
- "Can you handle our scale?" Exabyte operator; CoreWeave multi-exabyte; Decart 0 to 16 PB in 90 days.
- "Are you just a backup company?" B2 is 62 percent of revenue and the growth engine.

03. ICPs and personas

Sell to the workload, then to the person. The same customer can be three ICPs.

Ideal customer profiles

ICP 1: Backup and cyber-resilience buyer

- **Who:** mid-market to enterprise IT running Veeam, Rubrik, Commvault, Cohesity, MSP360, Synology/QNAP, or Kasten.
- **Trigger events:** ransomware incident (theirs or a peer's), cyber insurance renewal questionnaire asking for immutable off-site copies, tape library end of life, AWS bill review, Wasabi price increase (Wasabi pay-as-you-go list moved to \$7.99/TB, per wasabi.com/pricing checked 2026-10-14).
- **Data shape:** 50 TB to multi-PB, write-heavy, read rarely except restores and restore tests. Retention-based churn.
- **Why B2 wins:** Object Lock at no extra cost, always-hot restores, no minimum duration so retention churn does not trigger early-delete fees, free egress for restore tests up to 3x.
- **Deal size signal:** 200 TB+ or a multi-year commit.

ICP 2: Media and entertainment

- **Who:** post houses, broadcasters, sports teams, agencies, houses of worship, public TV.
- **Triggers:** LTO generation migration, shuttle drives, remote collaboration, MAM rollout (iconik, Mimir, axle ai, CuttingRoom, Backlight), TPN audit.
- **Data shape:** 100 TB to 10 PB, large files (RED, BRAW, MXF), bursty reads, reuse of archive.
- **Why B2 wins:** archive stays online, free egress up to 3x, TPN Blue Shield status, Live Read for growing files, integrations with MAMs, Fireball and Universal Data Migration for tape.
- **Proof:** Philadelphia Eagles (800 TB, retired LTO-6), Austin City Limits (4,000+ hours of tape), Bell Bank studio, American Public Television, Jazz at Lincoln Center (1.3 PB).

ICP 3: AI/ML builders

- **Who:** model labs, AI startups, enterprise ML platform teams, research institutions.
- **Triggers:** egress bills moving data to GPU clouds, credits running out, multi-cloud GPU sourcing, dataset growth past 1 PB.
- **Data shape:** PB-scale, read-heavy in bursts, throughput-bound, multi-site.
- **Why B2 wins:** one S3 namespace that feeds any GPU cloud, partner free egress to CoreWeave and Vultr, Overdrive dedicated throughput to 1 Tbps with unlimited egress, no upload or delete fees.

- **Proof:** Decart (0 to 16 PB in 90 days), Hume AI (~10 GB/s uploads, consolidated multiple hyperscale stores), Mirage/Captions (Overdrive, 95% estimated savings), Segmed (750 TB medical imaging).

ICP 4: Application and SaaS storage

- **Who:** SaaS platforms storing user content, e-commerce, gaming, surveillance, IoT.
- **Triggers:** COGS review, CDN bill, scaling past free credits.
- **Why B2 wins:** drop-in S3 replacement (Tribute: "\$15K saved per month"), CDN partner egress at \$0, free transactions.
- **Watch out:** very high read-op, small-object, high-egress-to-internet apps without a partner CDN can exceed 3x egress. Model it.

ICP 5: MSPs and resellers

- **Who:** MSPs with 20 to 2,000 clients, VARs, distributors.
- **Why B2 wins:** predictable unit cost, Partner API and Groups for sub-account provisioning, Object Lock for insurance requirements, Computer Backup for endpoints.

ICP 6: Neoclouds

- **Who:** GPU-as-a-Service providers.
- **Why B2 Neo wins:** white-label storage in weeks, Partner API, custom domains, SSO, no egress fees, no storage team to hire.

Personas

CIO / VP Infrastructure

- **Owns:** budget, risk, vendor count, board narrative on ransomware and AI.
- **Fears:** surprise bills, a failed restore, vendor viability, migration risk.
- **Wins with:** a 3-year TCO with egress modeled, insurance and compliance evidence, a named reference in their industry, a migration plan that does not consume the team.
- **Language:** "predictable", "resilience", "exit cost", "board".
- **Do not:** go deep on erasure coding unless asked.

Storage / backup admin

- **Owns:** the backup jobs, the restore SLA, the 2 a.m. page.
- **Fears:** a new tool that breaks jobs, immutability that blocks legitimate deletes, throttled restores.
- **Wins with:** a 20-minute Veeam or Synology setup demo, Object Lock behavior proved live, restore speed test, clear docs.
- **Language:** "SOBR", "capacity tier", "immutable", "3-2-1-1-0", "RTO".

- **Do not:** oversell features they will find missing (no lifecycle transitions between classes, because there is one class).

ML platform lead

- **Owns:** GPU utilization, data pipelines, dataset versioning, cost per training run.
- **Fears:** idle GPUs, egress tax, throughput ceilings, rate limits, an S3 API gap that breaks a dataloader.
- **Wins with:** a throughput benchmark from their GPU cloud, partner egress confirmation, Overdrive sizing, dataloader compatibility (boto3, s3fs, PyTorch S3 connectors **[test before claiming]**).
- **Language:** "GB/s per node", "checkpoints", "shards", "time to first batch", "hot tier vs capacity tier".

MSP owner / partner manager

- **Owns:** margin, standardization across clients, insurance requirements.
- **Fears:** egress surprises during a client restore, support quality, billing complexity.
- **Wins with:** unit economics per client, Partner API demo, Object Lock in their backup stack, Computer Backup at \$99/year per computer.

CFO / procurement

- **Fears:** variable bills, commitment risk, prepayment.
- **Wins with:** pay-as-you-go vs 1 to 5 year commit comparison, rate lock, the formula sheet in 14-tco-economics.md.

CISO / security architect

- **Fears:** data exfiltration, insider deletion, compliance gaps.
- **Wins with:** Object Lock compliance mode, scoped app keys, SSE-B2 default, SSE-C, MFA, SOC 2 Type 2 report via Whistic, BAA, TPN, access logs.

06. B2 Cloud Storage deep dive

Be fluent enough to be honest. Every feature below has a “how it actually behaves” note, because the POC will find it if you do not.

What you will be able to do after this section

- Explain B2’s architecture from shard to region in 90 seconds.
- Configure buckets, keys, Object Lock, lifecycle, replication, and encryption from memory.
- Name every documented S3 compatibility gap before the customer finds it.

1. Architecture: from file to vault

Backblaze Vault. 20 storage pods, one per cabinet. Drives in the same slot position across the 20 pods form a **tome**. Each file is written to one tome.

Erasure coding. Each file is split into **20 shards: 17 data + 3 parity** using Reed-Solomon. Any 17 shards rebuild the file. The vault survives a pod failure, a full cabinet power loss, or a cabinet network outage. Writes still succeed with one pod down (leaving two parity shards). Reads succeed with three pods down.

Designed durability: 99.999999999 percent (11 nines) annually, per Backblaze docs “Resiliency, Durability, and Availability” (updated Aug 2023, checked 2026-10-14).

Overhead math for the whiteboard: $20/17 = 1.176$. Backblaze stores about 1.18 raw bytes per logical byte inside a vault, versus 3.0 for triple replication. That efficiency is a big part of how a \$6.95 price works.

How it actually behaves: durability is per region. A region is one or more data centers (US West uses Sacramento and Phoenix). If the requirement is “survive loss of a region”, the answer is Cloud Replication to a second account in another region, priced as a second copy.

2. Regions and data centers

Region	Data center(s)	Notes
US West	Sacramento CA and Phoenix AZ (Stockton CA also listed among DCs)	Pricing page compares against US West rates

Region	Data center(s)	Notes
US East	Reston VA (operated by Coresite per docs)	
EU Central	Amsterdam NL	GDPR residency conversations
CA East	Toronto ON	Newest region

Source: docs "Where are the Backblaze Cloud Storage Data Centers" (updated Feb 2025, checked 2026-10-14). Storage price is the same in every region.

How it actually behaves:

- Region is chosen at **account creation** and **cannot be changed**.
- One account = one region. Multi-region = multiple accounts, usually tied together under a Group.
- S3 endpoints look like `https://s3.<region-code>.backblazeb2.com`, for example `s3.us-west-004.backblazeb2.com`. The exact code appears on the bucket details page. Never guess it in a demo.

3. Pricing mechanics (checked 2026-10-14, backblaze.com/cloud-storage/pricing)

- **Storage:** \$6.95/TB/30 days, billed per byte-hour. First 10 GB free.
- **Egress:** free up to 3x average monthly stored (byte-hour average), then \$0.01/GB.
- **Partner egress:** unlimited free to Fastly, Cloudflare, bunny.net, CacheFly, CoreWeave, Equinix Metal, Vultr, phoenixNAP.
- **Transactions:** Class A, B, C free for pay-as-you-go. Class D (outbound Event Notifications) first 2,500/day free, then \$0.004 per 10,000.
- **No** minimum object size, **no** minimum storage duration, **no** upload fee, **no** delete fee.
- **Terms:** monthly PAYG (card or invoice) or annual capacity commitments of 1 to 5 years, billed in full at contract start, rate locked.
- **Add-ons:** implementation support, premium support, Universal Data Migration (with qualifying commitment).

Discrepancy to know: the transaction pricing page (checked 2026-10-14) still says storage is "\$0.005/GB/month", which is the old \$5/TB rate. The pricing page and docs say \$6.95. Quote \$6.95 and flag the stale page internally.

4. Buckets, files, versions

- Buckets are **private** or **public** (allPublic). No per-object ACLs.

- B2 keeps **file versions** by default. Uploading the same name creates a new version. "Deleting" through S3 with versioning creates a hide marker (delete marker).
- The web console Lifecycle Settings offers presets such as "Keep all versions" and "Keep only the last version". Veeam guides require **Keep all versions**, because Veeam manages retention itself.

5. Lifecycle rules

From docs (checked 2026-10-14):

- Up to **100 lifecycle rules per bucket**.
- Native API rule properties: `fileNamePrefix` (required, "" means all files), `daysFromUploadingToHiding`, `daysFromHidingToDeleting`, `daysFromStartingToCancelingUnfinishedLargeFiles`. Values must be positive integers or null.
- Example:

```
[{"fileNamePrefix":"backup/","daysFromUploadingToHiding":7,"daysFromHidingToDeleting":30,"daysFromStartingToCancelingUnfinishedLargeFiles":3}]
```

- **Lifecycle rules are native-only.** Set them with the native B2 API (`b2_update_bucket` / `b2_create_bucket` `lifecycleRules`), the b2 CLI, or the web UI. The S3-compatible API does **not** support lifecycle configuration (`PutBucketLifecycleConfiguration` is not available), so S3-only tools cannot manage B2 lifecycle.

How it actually behaves: there is no "transition to Glacier" because there is one class. Lifecycle is about hiding and deleting versions and cleaning up unfinished multipart uploads. A rule with prefix "" can delete everything. In a demo, always scope by prefix.

6. Object Lock (immutability)

From docs "Object Lock" (checked 2026-10-14):

- Enable at bucket creation, on an existing bucket, or with `b2_update_bucket`. **Once enabled, it cannot be disabled on that bucket.**
- Retention between **1 and 3,000 days**.
- **Governance mode:** keys with the right capability can shorten or remove retention.
- **Compliance mode:** nobody can remove retention; it can only be extended.
- **Legal hold:** independent of retention, no expiry, removed explicitly.
- **Default bucket retention** applies to new uploads; must be enabled at bucket creation; does not change existing files.
- Not allowed on restricted buckets (shared buckets, snapshots) or a bucket that is a replication **source**.

- **No extra cost.** Normal storage charges apply to locked data.
- Capabilities: `readBucketRetentions`, `writeBucketRetentions`, `writeFileRetentions`, `bypassGovernance`, `writeFileLegalHolds` and related.
[verify exact capability names in app-key table before a demo]
- Was formerly called "File Lock"; some API names still say file lock.
- Functionally matches AWS S3 Object Lock per Backblaze.

How it actually behaves: most customers never touch the API. Veeam, Rubrik, Commvault, MSP360, Synology Hyper Backup ask "how many days immutable" and set it per object. The SE job is to make sure the bucket has Object Lock enabled before the backup software connects.

Landmine: compliance mode means *you pay to store it until it expires*, even if you leave. Set realistic retention. Say this out loud; it builds trust.

7. Application keys

- **Master key:** all capabilities, all buckets, shown once, regenerating invalidates the old one. Needed for Partner API and account-level admin. **Does not work with the S3-compatible API:** S3 clients need a standard application key.
- **Standard keys:** scoped to one bucket, multiple buckets (Multi-Bucket Application Key), a file prefix, a capability set, and an expiry. Survive master key rotation.
- Limit: **100 million app keys per account** (contact sales for more).
- Use a standard key for the S3 API, integrations, and anything customer-facing.

How it actually behaves: least privilege is the security story. One key per app, per tenant, per backup server. Read-only keys for media distribution.

8. Server-side encryption

- **SSE-B2:** Backblaze-managed keys, AES-256, unique key per file, wrapped by a global key. **Default for new uploads since Sep 14 2026** (blog Aug 27 2026). Before that, encryption was opt-in per bucket and only applied to files uploaded after enabling.
- **SSE-C:** customer supplies an AES-256 key on each request; Backblaze does not store it. Lose the key, lose the data.
- No extra charge to encrypt. Nominal Class C calls to read or change bucket encryption settings.
- Works through both Native and S3-compatible APIs.

9. Cloud Replication

- Copies objects from a source bucket to one or more destination buckets, same region or across regions.
- Configured in the web console or Native API; **not** configurable through the S3 API, but replicates files written by either API.
- New uploads replicate as they arrive; existing files ("backfill") process daily at midnight UTC; non-fatal failures retry the next day.
- Cross-region requires a second account in that region.
- **Replicated data cannot be replicated again** (no chains).
- Metadata changes after initial replication are not replicated. Hidden files may not be hidden at the destination.
- First-time use needs a verified email and payment history. Over 50 TB: contact Sales.
- No service fee or egress charge for replication per Backblaze marketing; you pay storage for both copies.

10. Event Notifications

- Webhook HTTP POSTs on uploads, updates, deletions, filtered by bucket, event type, optional prefix; optional signing secret (HMAC).
- Managed via console or `b2_set_bucket_notification_rules` / `b2_get_bucket_notification_rules`.
- **Request access from Support** to enable (docs checked 2026-10-14).
- Billed as Class D: 2,500/day free, then \$0.004 per 10,000.
- Patterns: trigger transcoding, audit logs, metadata DB updates, Cloudflare Worker, forward to AWS SNS then Lambda (Backblaze guide exists).

11. Live Read

Read and edit large media files while they are still uploading, through B2 and connected tools (Telestream Vantage guide exists). Sell it to live sports and broadcast where time-to-air matters.

[verify availability terms: may require enablement]

12. S3-compatible API: the honest list

Supported: core object and bucket ops, multipart, presigned URLs (GET and PUT), SSE-B2/SSE-C, Object Lock, versioning, CORS, IPv6. Signature V4 only (Rubrik guide notes V2 fails until Rubrik support switches it).

Not supported (docs checked 2026-10-14): object ACLs, IAM roles, object tagging (Get returns empty, Put with x-amz-tagging is rejected), website configuration, browser POST uploads to presigned URLs. Bucket ACL is limited to private and public-read.

Naming rules differ slightly from AWS. Check bucket and file naming page before migrating.

13. Native B2 API, Partner API, Groups

- **Native API:** b2_* calls; required for some features (replication setup, event notification rules, some NAS sync per Object Lock FAQ).
- **Partner API:** create and manage accounts programmatically, usage reporting, billing workflows. Used by MSPs, Powered by Backblaze, and B2 Neo.
- **Groups / Organizations:** central admin and billing across member accounts. Group admins currently cannot manage event notification rules in member accounts.

14. Access and tooling

Web console (and an Enterprise Web Console), b2 CLI, AWS CLI, rclone, Cyberduck, SDKs, MCP server. Bucket access logs and CORS are listed security features.

Ask these, then listen

1. Which S3 features does your code use beyond Put/Get/List?
2. Do you need region-loss survival or just device and site loss?
3. Who can delete data today, and should anyone be able to?

Three moves, in this order

1. Map their API usage against section 12.
2. Pick the region by residency and GPU/user location.
3. Design keys and Object Lock before the first byte lands.

07. Overdrive, B2 Neo, B2 Reserve, MCP, and AI tooling

GPUs are the expensive part. Storage's job is to never be the reason they wait.

B2 Overdrive

What it is (backblaze.com/cloud-storage/b2-overdrive, checked 2026-10-14):

- Built on B2, with dedicated storage vaults and networking walled off from shared infrastructure.
- Throughput tiers from **100 Gbps up to 1 Tbps**.
- **Starts at \$15/TB/month. Multi-petabyte commitment required.**
- **Unlimited free egress to any destination**, including other clouds. Zero transaction fees.
- Private networking integration, white-glove onboarding, "PETA" support tier (highest offered).
- Enabled by Sales, including for testing.

When to lead with it:

- Read multiple well above 3x stored per month (repeated epochs from a remote GPU cloud).
- Aggregate throughput needs in the tens of GB/s.
- Training spread across multiple GPU providers.

Break-even vs standard B2 (list prices): Standard B2 monthly = $6.95 \cdot S + 10 \cdot \max(E - 3S, 0)$ where S = stored TB, E = egress TB, \$10/TB overage. Overdrive monthly = $15 \cdot S$. Set equal: $6.95S + 10(E - 3S) = 15S$ gives $E = 3.805 \cdot S$. **Above about 3.8x monthly read multiple to non-partner destinations, Overdrive is cheaper on list price alone**, before counting the dedicated throughput. Below that, standard B2 is cheaper unless throughput is the constraint.

Proof points:

- Mirage (formerly listed as Captions on the case-study URL): "0 GPUs waiting for data", "95% estimated cost savings", "\$0 egress", H100 clusters, after a competitive bakeoff.
- Decart: 0 to 16 PB in 90 days; "10x more efficient than competition"; "\$0 egress spend"; "the only technical thing we used in training these models that didn't crash the first time."

B2 Neo (B2 for Neoclouds)

What it is (backblaze.com/cloud-storage/b2-neo, checked 2026-10-14):

- White-label object storage for GPU-as-a-Service providers. Launched Feb 2026.
- **Partner API** to provision accounts, usage, reporting, billing from the neocloud's own portal.
- **Custom domains** so storage looks native.
- **SSO integration** with the operator's identity provider.
- Up to 1 Tbps throughput claim (meta description), SOC 2.
- "No egress fees" and "less operational overhead than building it yourself."
- Backblaze white paper: "Why Neoclouds Should Not Build Their Own Object Storage" covering the split-stack problem (customer data in a hyperscaler, compute on the neocloud, egress friction) and the cost of running Ceph or custom hardware.

The pitch to a neocloud: "Every customer who keeps data in a hyperscaler pays an egress tax to use your GPUs. That tax is your churn risk. B2 Neo puts S3 storage under your brand in weeks, and you never hire a storage ops team."

CoreWeave is the reference pattern, but be precise: the SEC 8-K does not label CoreWeave as a B2 Neo customer. Say "CoreWeave uses Backblaze to power HDD-based tiers in CoreWeave AI Object Storage", not "CoreWeave runs on B2 Neo."

Discovery for neoclouds:

1. What percent of your customers' training data lives outside your platform?
2. Do you offer object storage today? Built on what?
3. How many storage engineers would you need to run it at 10 PB? At 100 PB?
4. What do customers ask for: throughput, price, or just "S3 next to my GPUs"?
5. Billing: do you need per-customer usage feeds?

B2 Reserve and capacity pricing

B2 Reserve was Backblaze's prepaid annual capacity offering sold through channel partners (TB bundles, included egress and premium support). As of 2026-10-14 the old `/cloud-storage/b2-reserve` URL resolves to the main B2 page, and the pricing page describes "Pay as you go monthly or for capacity annually" with 1 to 5 year terms. The API still has `b2-reserve-trial-create-account` in docs. **Treat "B2 Reserve" as a legacy or channel SKU name and confirm the current name and terms with Backblaze deal desk before quoting. [unverified current status]**

B2 MCP Server (Sep 8 2026)

What it is: open source Model Context Protocol server that lets AI agents (Claude, Gemini CLI, Cursor-style clients, etc.) operate B2 through a defined tool set.

Facts from docs (checked 2026-10-14):

- **40 tools** in three groups: 17 B2 Native (buckets, notification rules, app keys, retention and legal holds, Partner API group/account ops), 19 S3-compatible (objects, metadata, multipart, presigned URLs). Note: lifecycle rules cannot be set through the S3-compatible API; they are native B2 API or web UI only, 4 analytics (storage growth, egress leaders, largest files, unfinished multipart uploads).
- Profiles: `full` (40 tools) and `read-only` (20 tools).
- Local mode runs as a subprocess over STDIO; the app key stays in the server environment and is **not** passed to the model.
- **Tools exposed are gated by the app key's capabilities.** A read-only, single-bucket key yields a smaller tool set. The server does not invent a new permission model.
- Read-only resources: `b2://server-config` , `b2://capabilities` , `b2://bucket/{bucketName}` (webhook secrets redacted).
- Optional prompts via `B2_ENABLE_MCP_PROMPTS=true` .
- Self-hosted HTTP deployments have configurable credential modes.

SE angle: this is the safest way to let agents touch storage, because the blast radius equals the key. Demo it with a read-only key first, then a bucket-scoped write key, and show the tool list shrink and grow.

Can an agent delete buckets or files through the B2 MCP Server?

Only if the application key you give it allows that. The server exposes tools based on the key's capabilities and does not add a permission model of its own. A key with `deleteBuckets` can delete buckets; a key with `deleteFiles` can delete files (Object Lock retention still applies). A key without those capabilities, or the `read-only` profile, cannot. To make deletion impossible, give the agent a read-only or bucket-scoped key without `deleteBuckets` or `deleteFiles` .

Say: "The agent can do exactly what its key can do. Want zero chance of deletes? Hand it a read-only key."

This guide's own MCP server (backblaze.nixfred.com/mcp)

Separate from Backblaze's B2 MCP Server. The field guide runs its own remote MCP endpoint at `https://backblaze.nixfred.com/mcp` with four tools: `search_guide` , `ask_guide` , `get_page` , `list_pages` . **It is read-only:** it reads this guide's text and nothing else. It has no B2 credentials and cannot create, change, or delete buckets, files, or keys. Install steps are on the MCP page (`/mcp/`).

Other AI integrations (backblaze.com/cloud-storage/b2-ai-integrations)

- Artifact store for model weights and data files (MLflow-style guide).
- **Pixeltable:** declarative multimodal dataset management.
- **CVAT:** source raw images/video and retain labeled datasets.
- Vibe Coding Starter Kit, a Claude skill for B2 (Agent Skills spec), OpenClaw state backups.
- Startup program **Flamethrower:** up to \$100,000 in storage credits ("Unlock 1PB free storage for startups").
- Caveat from the page: integrations are not covered by the SLA.

Positioning in the AI stack (draw this)

[Data sources] → [B2: capacity tier, system of record] ⇌ [Hot tier: VAST / WEKA / Lustre / local NVMe] → [GPUs]

↘ partner egress \$0 → CoreWeave, Vultr, Equinix

Metal

↘ Overdrive: dedicated 100G to 1T, \$0 egress

anywhere

Flash for the working set, Backblaze for everything else. Hume AI uses WEKA with B2 (case study integration tag). Never pitch B2 as a replacement for a parallel file system feeding GPUs at line rate; pitch it as the place the full corpus lives and the source the hot tier fills from.

08. Computer Backup (Personal, Business, Enterprise Control)

The original product. Smaller growth, but a clean land motion and a real ransomware story for endpoints.

Facts (backblaze.com/cloud-backup/pricing and [/business](https://backblaze.com/business), checked 2026-10-14)

	Personal	Business	Enterprise Control
Price	\$99/yr per computer (yearly); 2-year saves 8%	\$99/yr per computer	Custom quote
Data	Unlimited, all user data by default, external drives	Same	Same
Clients	Native Mac and Windows	Same plus centralized admin	Same
Deployment		MSI, Jamf, Munki, Kandji (now Iru), Addigy, Intune	Same
Version history	1 year included in Business; Forever retention for added fee		
SSO		Google and Microsoft	OIDC (Okta, Azure AD/Entra)
Extras		Group admin, centralized billing, pro-rated adds	Legal hold, backup lockdown, restricted restore access, granular permissions, advanced support

Restores: web zip restores up to 500 GB per request; free mailed encrypted USB drive (return within 30 days); dedicated restore app. **Security:** 2048-bit public/private keys protect a symmetric AES-128 key for data; optional private encryption key; 2FA (SMS and TOTP). **Revenue context:** Computer Backup was \$16.1M in Q2 2026, down 2 percent YoY. It is a cash engine and a door-opener, not the growth story.

When an SE should bring it up

- Customer has laptops with data that never touches a server (creative, legal, research, execs).
- Offboarding pain: "an employee left and their files went with the laptop."
- Ransomware on endpoints and no versioned copy.
- Mac-heavy fleets managed by Jamf.

Landmines

- It is not an M365 or SaaS backup. Point to Veeam M365 to B2 for that.
- It is not a server backup. Servers go to B2 via backup software.
- Personal and Business share the \$99 price; the difference is admin, billing, SSO.
- Business "Save files from Computer Backup to Cloud Storage" docs exist: restores can land in B2. Use it for legal hold or long-term retention stories.

Discovery

1. How many endpoints, Mac vs Windows?
2. Which MDM?
3. Who restores today, the user or IT?
4. Any legal hold or eDiscovery requests in the last year?
5. What happens to data when someone leaves?

27. Industry plays

Same product, different first sentence. Backblaze publishes industry pages for each of these (backblaze.com/cloud-storage/industries/*, checked 2026-10-14).

Higher education and research

- **Pain:** research data growth, grant-funded budgets, tape, ransomware at universities.
- **Hooks:** HECVAT completed, Internet2 Cloud Scorecard and I2PX connectivity, Whistic Education profile, no egress penalty for collaborators pulling data.
- **Proof:** Cal Poly Humboldt (eliminated surprise bills, shrank Google footprint), Pittsburg State (Object Lock off-site backups), UCSC Silicon Valley, Coast Community College District, Gladstone Institutes (900+ tapes eliminated, 450 researchers).
- **First question:** "Who pays when a collaborator downloads the dataset?"

Healthcare and life sciences

- **Pain:** imaging archives, genomics, ransomware targeting hospitals, HIPAA.
- **Hooks:** BAA, SSE-B2 default, compliance-mode Object Lock, Segmed 750 TB medical imaging migration for AI research.
- **Careful:** say "BAA available", never "HIPAA certified". Confirm PHI handling with their compliance team.
- **First question:** "Where is your immutable copy of PACS and EHR backups?"

State and local government (SLED)

- **Pain:** ransomware on municipalities, tight budgets, procurement vehicles.
- **Hooks:** GovRAMP Progressing Snapshot, TX-RAMP Provisional, predictable pricing for annual budgets, term commits.
- **Careful:** not FedRAMP. Check state contract vehicles [verify available cooperative contracts].
- **First question:** "What did your last cyber insurance or state audit ask for?"

Financial services

- **Pain:** retention mandates, audit, ransomware.
- **Hooks:** compliance mode, legal hold, SOC 2 Type 2, Boreal Capital Management and Bell Bank stories.
- **Careful:** SEC 17a-4 style WORM attestation is a specific requirement.

[verify whether Backblaze has a third-party 17a-4 assessment before claiming]

Media, sports, creative

- See whiteboard 10. TPN Blue Shield; Eagles, ACL, American Public Television, Falcons Creative Group, Steve McCurry Studios, Trailerfarm.

Surveillance and video security

- **Pain:** retention requirements for camera footage, NVR storage limits.
- **Hooks:** no minimum duration (footage rolls off on a schedule), lifecycle rules by prefix, Object Lock for evidentiary holds.
- **First question:** "How many days of footage must you keep, and how often do you pull it?"

Gaming

- **Pain:** asset delivery spikes, patch distribution, user-generated content.
- **Hooks:** CDN origin with \$0 partner egress; Overdrive for concurrent upload/download peaks (Overdrive page lists gaming).

E-commerce and SaaS

- **Pain:** COGS, image and media serving, backups of production DBs.
- **Hooks:** drop-in S3, presigned PUT uploads, CDN partner egress, Big Cartel and Tribute stories.

Houses of worship

- **Pain:** weekly service video, volunteer staff, small budgets.
- **Hooks:** simple setup with NAS sync, low cost per TB, streaming via CDN.

Kubernetes and Microsoft 365 backup

- **K8s:** Kasten K10 by Veeam with Object Lock; Velero-compatible S3 target **[test Velero config before claiming]**.
- **M365:** Veeam Backup for Microsoft 365 to B2. Pitch: keep the backup outside the Microsoft tenant's blast radius.

MSPs

- Partner API, Groups, Object Lock for insurer checklists, Computer Backup for endpoints, Fireball/MSP360 seeding.
- **First question:** "How many client backup stacks land on storage you do not control the price of?"

32. Feature matrix: what lives where

Teach these. Do not mislabel them.

Capability	B2 Cloud Storage	B2 Overdrive	B2 Neo	Computer Backup
Buyer	IT, dev, backup, media	AI/HPC, large media, gaming	Neocloud operators	IT for endpoints, individuals
Interface	S3 API, Native API, console, CLI	Same	Partner API, white-label, custom domain	Native Mac/Windows client, web admin
Price (list)	\$6.95/TB/mo	from \$15/TB/mo	Contact sales	\$99/yr per computer
Commitment	None (PAYG) or 1 to 5 yr term	Multi-PB	Contract	Monthly, 1 yr, 2 yr
Egress	3x free, then \$0.01/GB; partners free	Unlimited free	"No egress fees"	Restores free (web zip, USB by mail)
Throughput	Shared	Dedicated 100 Gbps to 1 Tbps	Up to 1 Tbps (claim)	Client-side
Object Lock	Yes	Yes	Yes (platform)	Enterprise Control: backup lockdown, legal hold
Encryption	SSE-B2 default, SSE-C	Same	Same	AES-128 data key, 2048-bit RSA, optional private key
Replication	Cloud Replication	Yes	[verify]	n/a
Event Notifications	Yes (access request, Class D)	Yes	[verify]	n/a
Private networking	No [verify]	Yes	[verify]	n/a
SSO	[verify]	[verify]	Yes	Google/Microsoft; OIDC on Enterprise Control
Support	Add-on premium	PETA tier included	[verify]	Standard; advanced on Enterprise Control

Feature-to-pain map

Pain	Feature	One-liner
Ransomware deletes backups	Object Lock compliance	"Nobody can delete it until the date."
Restore bills	3x free egress, hot class	"Restore tests are free up to 3x stored."
CDN origin cost	Partner egress	"\$0 from B2 to Cloudflare or Fastly."
GPU data movement	Partner egress, Overdrive	"One namespace, any GPU cloud."
Tape	UDM, Fireball	"We help fund and run the move."
Key sprawl	Scoped app keys	"One key per app, per bucket, per prefix."
Compliance evidence	SOC 2, BAA, TPN, HECVAT via Whistic	"The documents are already in Whistic."
Agents touching data	MCP server	"The agent can only do what the key allows."
Event-driven pipelines	Event Notifications	"Upload triggers your webhook."
Region loss	Cloud Replication	"Second region, second account, automatic."

25. Whiteboard fieldcraft

Anyone can demo. Learn to draw it.

The loop: Ask. Draw. Validate. Repeat.

1. **Ask** three questions before touching the marker.
2. **Draw** their current state first, in their words. Use their app names.
3. **Validate** each box: "Is this right?" Hand them the marker to fix it.
4. **Draw** the future state in a second color.
5. **Write numbers** on arrows: TB, GB/s, \$/month.
6. **Close** by circling the one thing that changes and asking what blocks it.

Board layout

- Top-right corner: the five numbers (stored TB, growth, read multiple, throughput, RTO).
- Left: sources. Center: B2. Right: consumers.
- Red: threats and costs. Green: restore and savings paths. Blue: data flow.
- Bottom: the one-line rule (3-2-1-1-0, "\$0 partner egress", "one namespace").

Five shapes only

Box (system), cylinder (storage), cloud (provider), stick figure (people), arrow (flow). Anything else slows you down.

Openers that work

- "Before I draw anything, where does the data start its life?"
- "Show me the last restore you ran."
- "Where are your GPUs this month?"

Mistakes

- Drawing Backblaze first. Draw them first.
- Logos instead of words. Words survive a photo.
- No numbers. A board without numbers is a slide.
- Arguing with the incumbent. Draw them in as a box and move on.

Virtual whiteboarding

Use a shared canvas (Excalidraw, Miro, FigJam). Pre-build the five shapes. Keep the same layout so screenshots look consistent across calls. Mermaid versions of every board are in sections 09 to 13.

The SE's 60-second architecture talk track

"Here is your data today [left]. Here is what it costs you to use it [red arrows]. Here is B2 [center], one hot class, S3 API. Your tools point here [blue arrows]. Restores and partners pull for free [green]. This copy cannot be deleted [lock]. That is the change. What would stop us testing it next week?"

09. Whiteboard: backup and ransomware immutability

The board is the deliverable. If the customer photographs it, you won.

Whiteboard outcome

By the end, the board shows: production, a local backup repository, an immutable off-site copy on B2 with Object Lock, an optional second-region replica, the attacker path that fails, and the restore path with time and cost written next to it.


Before you draw: four questions

1. Which backup software? (Veeam, Rubrik, Commvault, MSP360, Synology)
2. RPO/RTO for tier-1?
3. Retention schedule?
4. Who can delete backups today?

Draw it in 9 steps

1. **Left edge, draw three boxes stacked:** VMs, physical servers, M365/SaaS. Label the group "Production". Say: "This is what we protect."
2. **Middle-left, one box:** "Backup server" with their software name. Arrow from production. Say: "Your backup software stays. Nothing changes here."
3. **Below it, a cylinder:** "Local repo (fast restore)". Arrow from backup server. Write their RTO next to it.
4. **Right side, a large rounded rectangle:** "Backblaze B2 (region: US East)". Inside, a bucket icon labeled "Object Lock: ON, compliance, 30 days". Say: "This copy cannot be deleted or changed until the date passes. Not by you, not by an attacker with your admin password, not by Backblaze support."
5. **Arrow backup server to B2,** label it "S3 API, scoped app key, TLS". Write "SOBR capacity tier or direct-to-object" for Veeam.
6. **Draw the attacker:** a red figure top-center with a red arrow to the local repo (X: encrypted) and a red arrow to B2 that stops at a wall. Write "DeleteObject → AccessDenied (locked)". Say: "This is why insurers ask for immutable off-site copies."

7. **Optional, below B2:** a second rectangle "B2 EU Central (2nd account)" with a dashed arrow "Cloud Replication". Say: "If you need region-loss survival, we replicate to a second region. Note: Object Lock cannot be set on a replication source bucket, so design which side is locked."
8. **Restore path in green:** arrow from B2 back to a "Clean recovery environment". Write "Always hot. No retrieval tier. Restore tests free up to 3x stored per month."
9. **Bottom, write the rule:** "3-2-1-1-0: 3 copies, 2 media, 1 off-site, 1 immutable, 0 errors on restore test." Circle the "1 immutable".



Whiteboard diagram (rendered from the source mermaid)

The design decisions you must make with them

Decision	Options	Guidance
Lock mode	Governance vs compliance	Compliance for insurance and regulators; governance for teams that must be able to purge with a privileged key
Lock duration	7 to 90 days typical for backup chains	Match backup software retention plus block generation window; too long means paying for dead data
Bucket versioning	Keep all versions	Required by Veeam guide
Keys	One key per backup server, bucket-scoped	Never the master key
Region	Nearest DC with residency match	Region cannot be changed later

Decision	Options	Guidance
Second copy	Cloud Replication to second region account	Price both copies
Restore test cadence	Monthly or quarterly	Free egress up to 3x stored makes full tests affordable

Veeam specifics

- v12+ supports **direct-to-object** to B2 with immutability (Backblaze guide). v10/v11 or SOBR: capacity tier offload guide.
- Veeam Ready partner status per Backblaze integration page.
[verify current Veeam Ready category on veeam.com]
- Backblaze claims setup "takes less than 20 minutes".
- Also: Veeam Backup for Microsoft 365, Kasten K10 for Kubernetes with Object Lock.

Rubrik specifics

- Rubrik defaults to **SigV2** for S3-compatible targets; B2 requires **SigV4**. There is no UI switch; **open a Rubrik support case to enable V4 first** (Backblaze docs). Put this in the POC plan on day one.

Commvault, MSP360, Synology, TrueNAS

- Backblaze publishes integration guides for each (see 16-partner-plays.md). MSP360 and Synology Hyper Backup also have Fireball rapid-ingest guides for seeding.

What they say / what you say / how you prove it

They say	You say	Prove it
"We already have immutable local storage."	"Great. Is it off-site and on separate credentials? Insurers and 3-2-1-1-0 want one immutable copy that a site loss or admin compromise cannot touch."	Draw the attacker arrow to both.

They say	You say	Prove it
"Glacier is cheaper."	"Per GB stored, yes. Price a full restore: retrieval fee plus hours of latency plus egress. B2 restores are hot and egress is free up to 3x."	TCO sheet scenario with one full restore per year.
"What if Backblaze support deletes it?"	"Compliance mode blocks everyone until expiry. That is the point."	Live demo: delete attempt fails (17).
"Restores over the internet are too slow."	"Local repo handles fast restores. B2 is the copy that survives when local does not. For bulk, size bandwidth and test it in the POC."	Restore throughput test.

From the SE chair

Lead with the restore, not the backup. Everyone can write data somewhere. Ask "when did you last restore everything from off-site?" and let the silence do the work.

10. Whiteboard: media workflows

Archive that you can actually use is not an archive. It is a library. Sell the library.

Whiteboard outcome

Camera to edit to archive to delivery on one board, with B2 as the always-online library in the middle, the MAM as the brain, and the CDN as the free exit.


Before you draw

1. Cameras, codecs, TB per week?
2. MAM/PAM and NLE?
3. LTO generation and tape count?
4. Who pulls from archive and how often?

Draw it in 10 steps

1. **Far left:** camera icon "Ingest: RED / BRAW / MXF". Write their TB/week.
2. **Next:** "On-prem NAS / SAN (edit storage)" box. Arrow from camera. *Say: "Your edit bays keep working on fast local storage."*
3. **Above NAS:** "MAM" box (iconik, Mimir, axle ai, CuttingRoom, Backlight). Dotted line to NAS: "indexes, proxies".
4. **Center:** big "Backblaze B2" rectangle labeled "Always-online library". Arrow NAS → B2: "sync / archive policy (Synology Cloud Sync, Hyper Backup, Archiware P5, Tiger Bridge, etc.)".
5. **Arrow MAM → B2:** "MAM writes proxies and originals; keeps search". *Say: "Search in the MAM, click, and the original comes back in minutes, not a tape robot queue."*
6. **Below B2:** tape stack with an X: "LTO: retire". Arrow tape → B2 labeled "Universal Data Migration / Fireball". *Say: "Backblaze helps fund and run the tape migration on a qualifying commitment."*
7. **Right of B2:** "Remote editors" stick figures. Arrow labeled "presigned URLs / Suite Studios mount / LucidLink".
8. **Further right:** "CDN (Cloudflare / Fastly / bunny.net)" cloud. Arrow B2 → CDN labeled "\$0 egress (partner)". Then CDN → "Viewers / OTT / social".
9. **Small box under B2:** "Event Notifications → transcode / QC". *Say: "Upload triggers a webhook. Your transcoder picks it up. No polling."*

10. **Write the money line at bottom:** "Archive reads: free up to 3x stored/month. Delivery via partner CDN: \$0 egress." And for live: "Live Read: start editing while the file is still uploading."



Whiteboard diagram (rendered from the source mermaid)

Proof points to use

- **Philadelphia Eagles:** 800 TB under management, retired LTO-6, 100% geographic redundancy, Mimir MAM, Quantum QXS, Cloud Replication, UDM. Quote: "We went from frequent LTO crashes and long restore times to near-instant access for every stored clip."
- **Austin City Limits:** 40+ years of content, 4,000+ hours of tape migrated, Fireball, UDM.
- **Jazz at Lincoln Center:** 1.3 PB archived, 3x cost savings, 400+ hard drives migrated.
- **Bell Bank (Bell Studios):** 100% automated backups, 400+ hours saved per year, replaced LTO.
- **TPN Blue Shield** status aligned with MPA Content Security Best Practices: answer for studio work.

Sizing notes

- 1 hour of 8K RED RAW can be hundreds of GB depending on compression **[codec-dependent; get their numbers]**.
- Ask for proxy-to-original ratio. Proxies may live on the NAS and in B2.
- Estimate monthly archive pulls as percent of stored. Most libraries read well under 3x.

Landmines

- No object tagging in S3 API: if a MAM relies on S3 tags for metadata, test it.
- Many MAMs and NAS sync tools use the Native B2 API or S3. Check which and whether Object Lock is needed.
- Live Read may require enablement and supported tools. Do not promise without checking.

11. Whiteboard: AI training data lake with GPU clouds

Start with the lifecycle, not the product. Ingest, prepare, train, checkpoint, serve, retain.

Whiteboard outcome

One S3 namespace on B2 as the system of record, feeding GPU clusters in two or more clouds through \$0-egress paths, with a hot tier where it belongs, checkpoints flowing back, and the throughput number written on every arrow.

Before you draw: the five numbers


1. Corpus size today and in 12 months (PB).
2. Full-dataset reads per month (the read multiple).
3. Target aggregate throughput (GB/s) and per-node.
4. Where the GPUs are (CoreWeave, Vultr, Lambda, Crusoe, on-prem, hyperscaler).
5. Hot tier in place? (VAST, WEKA, Lustre, local NVMe, none)

Write all five in the top-right corner before drawing anything else.

Draw it in 11 steps

1. **Left:** "Raw sources": web crawl, licensed media, sensors, app logs, partners. Arrow right labeled "ingest (free upload, free PUT)".
2. **Center-left:** big "Backblaze B2: system of record" rectangle with three bucket icons: `raw/`, `curated/`, `checkpoints/`. Say: "One namespace. Every stage writes back here."
3. **Above B2:** "Data prep / curation" box (Spark, Ray, Pixeltable, CVAT labeling). Two-way arrows to `raw/` and `curated/`. Say: "Upload and delete are free, so you can iterate on curation without a delete penalty."
4. **Right, top:** "GPU cloud A: CoreWeave" box. Arrow from `curated/` labeled "\$0 egress (partner)".
5. **Right, middle:** "GPU cloud B: Vultr or Equinix Metal". Arrow labeled "\$0 egress (partner)".
6. **Right, bottom:** "GPU cloud C: other / on-prem". Arrow labeled "3x free, then \$0.01/GB, or Overdrive: \$0 anywhere".

7. **Between B2 and a GPU box:** small box "Hot tier: VAST / WEKA / local NVMe". Say: "If you need line-rate to every GPU, flash caches the working set. B2 holds everything and refills the cache."
8. **Arrows from GPU boxes back to checkpoints/** : label "checkpoint writes, versioned, Object Lock optional for model provenance".
9. **Far right:** "Inference / serving" with arrow from checkpoints/ labeled "load weights". Add "RAG docs" bucket feeding it.
10. **Overdrive box around B2 (dashed):** "B2 Overdrive: dedicated vaults, 100 Gbps to 1 Tbps, unlimited egress, from \$15/TB, multi-PB commit". Say: "If your read multiple is above about 3.8x or your throughput target is tens of GB/s, this is the tier."
11. **Agents corner:** small robot icon "Agents via B2 MCP Server, scoped key". Say: "Agents can list, upload, presign, and report egress, but only what the key allows."



B2 Neo variant (selling to a neocloud)

Replace "GPU cloud A" with the neocloud's logo and draw B2 *inside* its boundary with a label "B2 Neo: white-label, your domain, your SSO, Partner API billing". Draw the hyperscaler outside with a red "egress tax" arrow to show what B2 Neo removes. Reference pattern: CoreWeave uses Backblaze for HDD tiers under its LOTA cache, customers get new tiers with no code changes.

The throughput math

- 1 Tbps $\approx$ 125 GB/s. 100 Gbps $\approx$ 12.5 GB/s.
- Time to read a 1 PB corpus once at 12.5 GB/s: $1,000,000 \text{ GB} / 12.5 = 80,000 \text{ s} \approx 22.2 \text{ hours}$. At 125 GB/s: $\approx 2.2 \text{ hours}$.
- GPU idle cost framing: $\text{idle_cost} = \text{GPUs} \times \$/\text{GPU-hr} \times \text{hours_waiting}$. $512 \text{ GPUs} \times \$2.50 \times 10 \text{ hours} = \$12,800 \text{ per stall}$ **[use their GPU rate]**.
- Hume AI reported close to 10 GB/s uploads from its clusters to B2 (case study).

Discovery traps

- "Our dataloader uses S3 tagging for splits." Not supported. Use prefixes or a manifest.
- "We need IAM roles from our GPU nodes." Not supported in S3 API docs; use scoped app keys distributed by a secrets manager.
- "We read everything every epoch from another cloud." Model the read multiple. Above $\sim 3.8x$ to non-partner destinations, Overdrive wins on list price.

From the SE chair

Never argue with the flash vendor in the room. Draw them in. "VAST/WEKA for the working set, Backblaze for the full corpus and checkpoints" makes you the architect, not the competitor.

12. Whiteboard: B2 as CDN origin


Origin egress is the bill nobody budgets. Remove it.

Whiteboard outcome

Users → CDN edge → B2 origin, with "\$0" written on the origin-to-edge arrow, signed URLs for private content, and the cache-miss math in the corner.

Draw it in 7 steps

1. **Right side:** world map scribble with users. "Users / players / viewers".
2. **Middle:** cloud "CDN edge: Cloudflare / Fastly / bunny.net / CacheFly". Arrow to users.
3. **Left:** "Backblaze B2 bucket (origin)". Arrow B2 → CDN labeled "**\$0 egress, Bandwidth Alliance / partner**".
4. **Under CDN:** "Worker / Compute@Edge: auth, rewrite, image resize". Say: "Private buckets work too: the edge signs requests to B2 so the bucket stays private." (Backblaze has guides for private and public content through Cloudflare and for Fastly Compute.)
5. **Arrow from app servers → B2:** "uploads via presigned PUT URLs". Say: "Users upload straight to B2. Your servers never touch the bytes."
6. **Event Notifications box:** B2 → webhook → "Worker: purge cache / generate thumbnails".
7. **Corner math:** $\text{origin_egress} = \text{total_delivered} \times (1 - \text{cache_hit_ratio})$. Write their numbers. Then: "With partner CDN, origin egress cost = \$0 regardless."



Whiteboard diagram (rendered from the source mermaid)

Facts

- Backblaze is a founding member of the Bandwidth Alliance (with Cloudflare, DigitalOcean and others). Downloads from B2 to Cloudflare: zero transfer cost (docs).
- Partner free egress list on the pricing page: Fastly, Cloudflare, bunny.net, CacheFly, CoreWeave, Equinix Metal, Vultr, phoenixNAP.
- Backblaze is also a member of the CDN Alliance (nonprofit industry association).
- No static website configuration in the S3 API. Use the CDN for index documents and pretty URLs.

Worked example

SaaS delivering 200 TB/month from a 50 TB library at 90 percent cache hit:

- Origin pulls = $200 \times 0.10 = 20$ TB.
- On S3 origin to a non-AWS CDN: 20 TB egress $\approx \$1,742$ ($10,000 \text{ GB} \times \$0.09 + 9,900 \text{ GB} \times \$0.085$) using AWS tiers (first 100 GB free, \$0.09 next 10 TB, \$0.085 next 40 TB) plus storage $50,000 \text{ GB} \times \$0.023 = \$1,150$. Total $\approx \$2,892$ /month. (Using CloudFront instead changes the math: S3 to CloudFront transfer is free, but CloudFront delivery is billed.)
- On B2 behind Cloudflare: storage $50 \times \$6.95 = \$347.50$; origin egress \$0. Total $\approx \$347.50$ /month.

Landmines

- Cloudflare's own terms govern caching large video on certain plans. Do not promise CDN pricing; B2 only controls the origin side.
- If the customer is on CloudFront, the comparison is CloudFront delivery vs the new CDN's delivery, not just origin egress.

13. Whiteboard: hybrid cloud and migration


Nobody migrates for fun. Draw the end state first, then the path, then the risk controls.

Whiteboard outcome

Today-state on the left (NAS, tape, S3/Azure/GCS/Wasabi), end-state on the right (B2 primary plus on-prem cache), and three migration lanes in the middle with TB and days written on each.

Draw it in 9 steps

1. **Left column, "Today":** boxes for each source with TB: "AWS S3: 600 TB", "NAS: 200 TB", "LTO: 900 tapes".
2. **Right column, "End state":** "Backblaze B2" big box; small "on-prem NAS cache" box; arrows from apps.
3. **Lane 1 (cloud to cloud):** arrow S3 → B2 labeled "Universal Data Migration: Backblaze helps cover egress and transfer with qualifying commit; MD5 verified".
4. **Lane 2 (NAS):** arrow NAS → B2 labeled "rclone / NAS Cloud Sync / backup software; bandwidth = X Gbps → Y days".
5. **Lane 3 (offline):** truck icon NAS/tape → "Fireball (96 TB Synology, 10GbE)" → B2. Say: "When the wire is too slow, ship it."
6. **Under lanes:** "Dual-write window" bracket. Say: "New writes go to both for N days, then we cut reads over."
7. **Verification box:** "Inventory diff: object count, bytes, MD5/SHA1 sample".
8. **Cutover box:** "Flip endpoint + keys in app config. Keep source read-only 30 days."
9. **Write the transfer formula:** $\text{days} = \text{TB} \times 8,000 / (\text{Gbps} \times 86,400 \times \text{efficiency})$. With 1 Gbps at 80%: 100 TB ≈ 11.6 days.



Whiteboard diagram (rendered from the source mermaid)

Transfer time table (80 percent link efficiency)

Data	1 Gbps	10 Gbps	40 Gbps
100 TB	11.6 days	1.2 days	7 hours
1 PB	116 days	11.6 days	2.9 days
10 PB	3.2 years	116 days	29 days

Formula: $\text{seconds} = \text{bytes} \times 8 / (\text{bps} \times 0.8)$. 1 PB = 8×10^{15} bits; at 10 Gbps $\times 0.8 = 8 \times 10^9$ bps $\rightarrow 10^6$ s $\approx$ 11.6 days.

Full playbook in 19-migration-playbook.md.

14. TCO and economics

Every number comes from them. Formulas on the board, inputs from their bill, sources in the footer. Never a universal savings claim.

Price inputs (list, US, checked 2026-10-14)

Provider	Storage	Egress to internet	Requests	Minimums	Source
Backblaze B2	\$6.95/TB/mo	Free to 3x avg stored, then \$0.01/GB; unlimited free via partners	A/B/C free; Class D \$0.004/10K after 2,500/day	None	backblaze.com/storage/pricing
B2 Overdrive	from \$15/TB/mo	Unlimited free, any destination	Free	Multi-PB commit	backblaze.com/overdrive
AWS S3 Standard (us-east-1)	\$0.023/GB first 50 TB, \$0.022 next 450 TB, \$0.021 over 500 TB	First 100 GB free; \$0.09/GB next 10 TB; \$0.085 next 40 TB; \$0.07 next 100 TB; \$0.05 over 150 TB	PUT/COPY/POST/LIST \$0.005/1K; GET \$0.0004/1K	None for Standard; IA/Glacier have minimum durations and sizes	aws.amazon.com JS-rendered; fig via upstash.com 2026 summaries AWS calculator use]
Wasabi PAYG	\$7.99/TB/mo	"No fees for egress", subject to fair-use policy	No API fees	90-day minimum storage duration (deleted early, still billed for 90 days)	wasabi.com/pricing
Cloudflare R2 Standard	\$0.015/GB-mo (\$15/TB)	Free	Class A \$4.50/M; Class B \$0.36/M	IA class has minimum duration and \$0.01/GB retrieval	developers.cloudflare.com/r2/ (last updated Oct)

Provider	Storage	Egress to internet	Requests	Minimums	Source
Azure Blob Hot LRS (East US)	~\$0.0208/GB first 50 TB; \$0.020 next 450 TB	~\$0.087/GB first 10 TB [unverified]	Write \$0.055/10K; read \$0.0044/10K	Cool/Cold/Archive minimum durations	learn.microsoft.com page sample takes verify on Azure
GCS Standard regional (US)	~\$0.020/GB (\$0.000027397/GiB-hr)	~\$0.12/GB first TB tiers [unverified]	Class A/B per 1K	Nearline 30d, Coldline 90d, Archive 365d minimums	cloud.google.com

Units: this guide uses 1 TB = 1,000 GB for all providers so the math is comparable. AWS and others bill in GB (some in GiB). Differences are under 10 percent and do not change conclusions.

The formulas

B2 standard

```
B2_month = S × 6.95 + max(E_internet - 3S, 0) × 1,000 × 0.01 + ClassD
S = average stored TB
E_internet = egress TB to non-partner destinations
partner egress (Cloudflare, Fastly, CoreWeave, Vultr, ...) = $0
```

Note: Backblaze says partner egress is free and “egress above 3x ... is also free through many CDN and compute partners.” Model partner egress as \$0.

[verify with Backblaze whether partner egress counts toward the 3x pool]

B2 Overdrive

```
OD_month = S × 15 (list floor; real price is negotiated; multi-PB commit)
```

AWS S3 Standard

```
S3_month = tiered_storage(S) + tiered_egress(E) + PUTs/1000×0.005 +
GETs/1000×0.0004
```

Wasabi

$W_{\text{month}} = S \times 7.99$ (valid only while egress $\leq$ stored, per fair-use; above that Wasabi may require a different plan)

Cloudflare R2

$R2_{\text{month}} = S \times 15 + \text{ClassA_millions} \times 4.50 + \text{ClassB_millions} \times 0.36$

Overdrive break-even: $6.95S + 10(E - 3S) = 15S \rightarrow E = 3.805S$. **R2 break-even vs B2 (ignoring R2 ops):** $6.95S + 10(E - 3S) = 15S \rightarrow$ same 3.8x. With R2 ops added, the crossover shifts higher.

Worked examples (monthly, list prices, storage + egress; requests excluded except R2 reads)

Assumptions: all egress to non-partner internet destinations; AWS us-east-1 tiers; R2 reads computed at 8 MB average object (reads = egress_bytes / 8 MB); Wasabi flat. Script: `research/tco.py`.

100 TB stored

Egress / month	B2	AWS S3	Wasabi	R2
10% (10 TB)	\$695	\$3,141	\$799	\$1,500
50% (50 TB)	\$695	\$6,542	\$799	\$1,502
100% (100 TB)	\$695	\$10,043	\$799	\$1,504
300% (300 TB)	\$695	\$21,045	\$799 (beyond fair use)	\$1,514
500% (500 TB)	\$2,695	\$31,045	n/a (beyond fair use)	\$1,522

AWS breakdown at 100 TB / 10% egress: storage $50,000 \times 0.023 + 50,000 \times 0.022 = \$2,250$; egress $(10,000 - 100) \text{ GB} \times 0.09 = \$891$.

1 PB stored

Egress / month	B2	AWS S3	Wasabi	R2
10% (100 TB)	\$6,950	\$29,343	\$7,990	\$15,004
50% (500 TB)	\$6,950	\$50,345	\$7,990	\$15,022

Egress / month	B2	AWS S3	Wasabi	R2
100% (1 PB)	\$6,950	\$75,345	\$7,990	\$15,045
300% (3 PB)	\$6,950	\$175,345	n/a	\$15,135
500% (5 PB)	\$26,950	\$275,345	n/a	\$15,225

AWS storage at 1 PB: $50 \text{ TB} \times 23 + 450 \text{ TB} \times 22 + 500 \text{ TB} \times 21 = \$1,150 + \$9,900 + \$10,500 = \$21,550$.

10 PB stored

Egress / month	B2	B2 Overdrive (list floor)	AWS S3	Wasabi	R2
10% (1 PB)	\$69,500	\$150,000	\$264,345	\$79,900	\$150,045
100% (10 PB)	\$69,500	\$150,000	\$714,345	\$79,900	\$150,450
300% (30 PB)	\$69,500	\$150,000	\$1,714,345	n/a	\$151,350
500% (50 PB)	\$269,500	\$150,000	\$2,714,345	n/a	\$152,250

3-year view at 10 PB, 10% egress: B2 \$2.50M; AWS \$9.52M; difference $\approx$ \$7.0M before discounts.

Hyperscaler EDP discounts commonly reduce list prices [discount levels vary; ask the customer], so always restate with their effective rates.

What the tables really say (be honest in the room)

1. **B2 wins on storage plus moderate egress.** Up to 3x read multiple, B2 is the lowest list price of the five.
2. **Wasabi is close on storage** (\$7.99 vs \$6.95, about 15 percent more) but Wasabi charges no egress fees only within its fair-use policy (monthly egress at or below stored volume; above that Wasabi may flag the account or move it to a different plan, it does not bill per-GB egress) and it has a 90-day minimum storage duration. Position B2 on egress headroom (3x vs 1x) and no minimum duration.
3. **R2 wins at very high internet egress** (above roughly 4x per month without a partner CDN) because egress is always free. But R2 storage is more than 2x B2, and Class A ops at \$4.50/M punish write-heavy and small-object workloads. Counter with Overdrive (unlimited egress) at multi-PB, or partner CDN egress at \$0.
4. **AWS is dominated by egress.** At 1 PB with 100% monthly egress, egress is 71 percent of the bill.
5. **Overdrive** makes sense above $\sim 3.8x$ read multiple or when dedicated throughput is the requirement.

Scenario cards (copy to the whiteboard)

Backup with annual full restore test, 300 TB:

- B2: $300 \times 6.95 \times 12 = \$25,020/\text{yr}$. Restore test of 300 TB = 1x stored, inside 3x free. \$0 egress.
- AWS S3 Standard: storage $50,000 \times 0.023 + 250,000 \times 0.022 = \$6,650/\text{mo} \rightarrow \$79,800/\text{yr}$; plus one 300 TB restore egress $\approx \$18,795$ (100 GB free, 10 TB @ \$0.09 = \$900, 40 TB @ \$0.085 = \$3,400, 100 TB @ \$0.07 = \$7,000, 149.9 TB @ \$0.05 = \$7,495). Total $\approx \$98,595/\text{yr}$.
- AWS Glacier Deep Archive would be far cheaper to store ($\sim \$0.00099/\text{GB}$) but has 180-day minimum, 12+ hour restores, retrieval fees, and the same egress. Say so.

Media library, 1 PB, 20% monthly reuse, delivery via Cloudflare:

- B2: \$6,950/mo storage; archive pulls 200 TB inside 3x; CDN origin egress \$0.
- S3 Standard: \$21,550 storage + 200 TB egress $\approx \$13,795 = \$35,345/\text{mo}$.

AI corpus, 5 PB, 4 full reads/month from a non-partner GPU cloud:

- B2 standard: $5,000 \times 6.95 + (20,000 - 15,000) \times 10 = \$34,750 + \$50,000 = \$84,750/\text{mo}$.
- B2 Overdrive: $5,000 \times 15 = \$75,000/\text{mo}$, plus dedicated throughput. Overdrive wins.
- If the GPU cloud is CoreWeave or Vultr (partner): B2 standard egress \$0 $\rightarrow \$34,750/\text{mo}$.

TCO beyond the invoice (ask, do not assume)

Line	On-prem	Hyperscaler	B2
Hardware refresh (3 to 5 yr)	Yes	No	No
Support contracts	Yes	Included / support plan	Premium support add-on
Power, space, cooling	Yes	No	No
Admin FTE time	High	Medium (tiering, lifecycle tuning)	Low (single class)
Data migration in	n/a	Egress from source	UDM can subsidize with commit
Exit cost	Hardware disposal	Egress	Egress (3x free/month helps)

Commit structures

- PAYG monthly: card or invoice, no commitment.
- Capacity term 1 to 5 years: billed in full at contract start, rate locked, volume and term discounts. Ask Backblaze deal desk for discount bands; none are public.
- Overdrive: multi-PB commitment.

Landmines

- Do not compare B2 to S3 Glacier on storage alone. Compare total cost of a real restore.
- Do not claim “1/5 the cost of AWS” without context. Backblaze uses that line on its Veeam page. It holds for storage vs S3 Standard list ($6.95/23 \approx 30$ percent, i.e. about 1/3.3) only with egress included. Use your own math.
- Prices change. Wasabi moved from \$6.99 to \$7.99 (Jul 1 2026 per company notes); B2 moved from \$6.00 to \$6.95. Re-check before every proposal and write the check date on the slide.

15. Competitive battlecards

Position on fit, never on their engineering. Every competitor here is a good product for someone. Find the workload where they are the wrong answer.

Card format: **When they win / When we win / Landmines to plant / Landmines to avoid / Objections / Proof.**

AWS S3

When they win: deep AWS-native apps using IAM roles, S3 event to Lambda, Athena, object tagging, S3 Express One Zone latency; customers burning down an EDP commit; data that is genuinely cold (Glacier Deep Archive at ~\$0.99/TB/mo).

When we win: egress-heavy workloads (restores, media reuse, multi-cloud AI, CDN origin to non-AWS CDN); teams tired of tiering and lifecycle tuning; COGS-sensitive SaaS.

Landmines to plant (questions):

- "What percent of last month's S3 bill was data transfer?"
- "Have you paid early-delete fees on IA or Glacier?"
- "How long would a full Glacier restore take, and what would retrieval cost?"
- "How many storage classes do you manage, and who tunes the lifecycle rules?"

Landmines to avoid:

- Do not claim feature parity. Gaps: IAM roles, object tagging, object ACLs, website hosting, browser POST.
- Do not compare to Glacier on storage price.
- Do not forget EDP discounts; ask for effective rates.

Objections:

- "We are all-in on AWS." → "Keep compute there. Put the copy that must survive an AWS account compromise somewhere else. That is also the cheapest copy to restore from."
- "S3 has 11 nines too." → "Yes. Durability is table stakes. Differences are egress, tiers, and who can delete your backups."

Proof: Tribute (\$15K/month saved, "drop-in API replacement"); Hume AI consolidated from multiple hyperscale object stores; Cal Poly Humboldt shrank Google footprint.

Wasabi

When they win: flat-rate buyers with egress reliably below stored volume; existing Wasabi reseller relationships; customers who want the “no egress” headline.

When we win: reuse-heavy data (egress 1x to 3x stored), short-retention data (no minimum duration on B2), AI workloads and GPU-cloud partnerships, CDN partner egress, Overdrive needs; public-company transparency; Drive Stats credibility.

Egress, precisely: Wasabi charges **no egress fees**. Its fair-use policy expects monthly egress to stay at or below the volume you store; if egress regularly exceeds stored volume, Wasabi may flag the account, ask you to reduce it, or move you to a different plan. There is no per-GB egress charge to quote. The real risk is the policy, not a fee: a 1 PB customer that serves 2 PB a month does not fit Wasabi’s fair-use model, while B2 gives 3x stored free and a published \$0.01/GB beyond that.

Plant:

- “What does Wasabi’s fair-use policy say if your monthly egress exceeds your stored volume?”
- “Do you delete data before 90 days? Wasabi has a 90-day minimum storage duration, so deleted data is still billed for the full 90 days.”
- “Your price went from \$6.99 to \$7.99 this year. How do you budget the next increase?”

Avoid: calling Wasabi unreliable or making claims about their finances or outages without a cited source.

Objections: “Wasabi has no egress fees at all.” → “As long as egress stays within their fair-use terms. B2 gives 3x stored free and unlimited to partners, with a published \$0.01/GB if you go over. Predictable both ways.”

Math: 1 PB: B2 \$6,950 vs Wasabi \$7,990 per month = \$12,480/yr difference at list.

Cloudflare R2

When they win: extremely egress-heavy internet delivery (above ~4x stored per month) without a partner CDN; teams already deep in Workers; small datasets where storage cost is trivial.

When we win: storage-dominant workloads (R2 is \$15/TB, B2 \$6.95); write-heavy or small-object workloads (R2 Class A \$4.50/M); backup (B2 Object Lock is mature and certified by the major backup ISVs; R2 has had bucket locks since March 2025, so win on certifications and economics, not on “R2 has no immutability”); you can use Cloudflare CDN *with* B2 and get \$0 egress anyway.

Plant:

- “How many PUTs and LISTs a month? R2 bills them.”
- “R2 has had bucket locks since March 2025. Does your backup vendor certify R2 as an immutable target, and does R2 lock match the S3 Object Lock API your vendor calls?”

- "If you are on Cloudflare already, why not keep the CDN and put a cheaper origin behind it?"

Avoid: saying R2 is not durable or not enterprise. Stay on economics and fit.

Math: 1 PB, 10% egress: B2 \$6,950 vs R2 \$15,004/month. Writes: 100M PUTs/month on R2 = \$450.

Best line: "Cloudflare is our partner. B2 behind Cloudflare gives you R2's egress story at less than half R2's storage price."

Azure Blob Storage

When they win: Microsoft-centric shops with MACC commitments, Entra ID integration, Azure-native analytics, Archive tier for truly cold data.

When we win: off-Azure backup copy (Veeam, M365 backup to a non-Microsoft target for separation), egress-heavy use, simplicity (Hot/Cool/Cold/Archive tiers with minimum durations and retrieval fees vs one class).

Plant: "If your Microsoft tenant is compromised, is your backup in the same blast radius?" "What do Cool and Archive early-deletion and retrieval fees cost you per year?"

Numbers: Hot LRS first 50 TB ~\$20.80/TB/mo (sample table) vs B2 \$6.95. [\[verify current Azure list\]](#)

Google Cloud Storage

When they win: BigQuery/Vertex-centric data platforms, GCP credits, dual-region and turbo replication needs.

When we win: training on GPUs outside GCP (egress), media archives, backup copies outside Google.

Plant: "Where are your GPUs actually coming from this quarter?" "What does pulling the dataset to that provider cost?"

Numbers: Standard regional ~\$20/TB/mo vs B2 \$6.95; Nearline 30-day, Coldline 90-day, Archive 365-day minimums (GCS pricing page).

On-prem: Pure Storage (FlashBlade), NetApp (StorageGRID / ONTAP S3), Dell ECS/ObjectScale

When they win: latency-critical primary storage, air-gapped or sovereign sites, existing large installed base, capex-friendly budgets, customers who need sub-ms or file protocols.

When we win: off-site copy (on-prem cannot be its own off-site), capacity growth without refresh cycles, archive tier behind flash, budget moving from capex to opex, staff constraints.

Plant: "When is the next refresh and what is the quote?" "Who manages drive failures and firmware?" "Where is your off-site immutable copy?"

Avoid: pitching B2 as primary storage for databases or VMs. Pitch it as the tier behind them. NetApp, Pure, and Dell all support tiering or replication to S3 targets; validate the specific product's support for B2 before claiming. [verify per-product B2 certification]

Line: "Keep the flash. Stop buying disk shelves for data you touch twice a year."

MinIO / Ceph (self-built object)

When they win: strict on-prem or edge requirements, Kubernetes-native S3 inside a data center, teams with storage engineers who want control.

When we win: when the team counts the real cost: hardware, 3x or erasure-coded overhead, drive replacement, upgrades, on-call, and a second site for off-site copies. For neoclouds, B2 Neo's whole pitch is "do not build this."

Plant: "How many people are on call for the cluster?" "What happens when a 20 TB drive fails during a rebuild?" "Where is the off-site copy?"

Math prompt: $\text{self_built_}\$/\text{TB}/\text{mo} = (\text{hardware}/36 + \text{power} + \text{space} + \text{support} + \text{FTE} \times \text{salary}/12) / \text{usable_TB}$. Fill with their numbers. Include usable vs raw overhead.

Quick comparison matrix (list, checked 2026-10-14)

	B2	Wasabi	R2	S3 Std	Azure Hot	GCS Std
\$/TB/mo	6.95	7.99	15	~23	~20.8	~20

	B2	Wasabi	R2	S3 Std	Azure Hot	GCS Std
Egress	3x free, then \$0.01/GB; partners \$0	Free within fair use	Free	\$0.05 to 0.09/GB	Paid [verify]	Paid [verify]
Request fees	Free (Class D only)	None	A \$4.50/M, B \$0.36/M	Yes	Yes	Yes
Min duration	None	90 days PAYG [verify]	None (Std)	None (Std)	None (Hot)	None (Std)
Single hot class	Yes	Yes	Std + IA	No (many)	No	No
Object Lock	Yes	Yes	[verify]	Yes	Immutability policies	Bucket Lock / retention
Public company	Yes (BLZE)	No	Yes (NET)	Yes	Yes	Yes

16. Partner ecosystem plays

Partners carry the deal into the account. Know who sources, who influences, and who fulfills.

Partner types (backblaze.com/partners, checked 2026-10-14)

- **Technology / integration partners:** backup, NAS, MAM, CDN, compute.
- **Channel:** resellers, distributors, VARs.
- **MSPs:** dedicated program page.
- **Powered by Backblaze:** embed B2 in your own product (Partner API).
- **Alliance partners:** CoreWeave, Cloudflare, Fastly, Vultr, etc. (free egress).

Play 1: Veeam (backup, enterprise and mid-market)

- **Fit:** VBR v12+ direct-to-object or SOBR capacity tier; Veeam Backup for Microsoft 365; Kasten K10.
- **Hook:** immutable off-site copy for insurance; “set up in under 20 minutes” (Backblaze claim).
- **SE checklist:** bucket with Object Lock enabled, lifecycle “keep all versions”, bucket-scoped key, correct S3 endpoint, Veeam immutability days set, test restore.
- **Co-sell motion:** find Veeam customers whose capacity tier is S3 Standard or Azure Hot. Run 14-tco math with a full restore scenario.

Play 2: Rubrik

- **Fit:** archival location / cloud archive.
- **Gotcha:** SigV4 must be enabled by Rubrik support first (default SigV2 fails). File the case in week one.

Play 3: Commvault

- Backblaze publishes a Commvault integration guide. Position B2 as cloud library for secondary copies with immutability. [verify WORM support specifics in Commvault guide before claiming]

Play 4: Synology and QNAP (SMB, media, branch)

- Synology Hyper Backup and Cloud Sync to B2; Fireball rapid ingest guide for Hyper Backup seeding; Fireball itself is a Synology NAS.
- QNAP Hybrid Backup Sync to B2.
- **Motion:** NAS resellers and MSPs. "Your NAS is not your off-site copy."

Play 5: TrueNAS

- Cloud Sync tasks to B2 (Backblaze guide). Strong with research labs and homelab-to-SMB crossovers.

Play 6: MAM / media (iconik, Mimir, axle ai, CuttingRoom, Backlight, Archiware P5, Tiger Bridge, Suite Studios, LucidLink, Telestream Vantage)

- iconik has a Backblaze integration page and a Fireball with iconik gateway guide.
- Telestream Vantage with Live Read guide.
- Suite Studios (Sep 10 2026): S3-native streaming that turns B2 into a high-performance drive.
- **Motion:** sell into post houses and sports with the MAM vendor; B2 is the library under their UI.

Play 7: MSP360 and MSP platforms

- MSP360 integration and Fireball rapid ingest guide.
- **Motion:** MSPs standardize backup storage across clients; Object Lock satisfies insurer checklists; Partner API for sub-accounts.

Play 8: CDN partners (Cloudflare, Fastly, bunny.net, CacheFly)

- \$0 egress from B2 to these CDNs.
- Guides: deliver public and private B2 content through Cloudflare; Fastly CDN and Fastly Compute integrations.
- **Motion:** origin swap. Find customers on S3 origin behind Cloudflare/Fastly.

Play 9: Compute and GPU partners (CoreWeave, Vultr, Equinix Metal, phoenixNAP)

- \$0 egress from B2.
- CoreWeave: strategic agreement; Backblaze powers HDD tiers in CoreWeave AI Object Storage.
- **Motion:** AI teams renting GPUs from these providers keep data in B2 with no egress tax.

Play 10: AI tooling (WEKA, Pixeltable, CVAT, MLflow-style artifact stores, MCP clients)

- WEKA as hot tier above B2 (Hume AI tags WEKA).
- **Motion:** architecture partner; joint whiteboard where both win.

Partner qualification questions

1. Which of our joint customers are on hyperscaler storage today?
2. Who owns the storage decision in your deals?
3. Do you resell storage or refer it?
4. What breaks in the field when you integrate with S3-compatible targets?
5. Will you co-present a POC?

Integration landmines

- Some tools use the Native B2 API, some S3. Object Lock with NAS sync uses the Native API per Backblaze FAQ.
- SigV2-only tools fail (Rubrik default).
- Tools that rely on object tagging or per-object ACLs need testing.

04. Discovery question bank

Ask, then shut up. Write the numbers on the board. Every number in your TCO must come from their mouth.

Format: question, then *why you ask* in italics. Codes let you reference questions in call notes (Q-GEN-03).

General (Q-GEN)

1. Q-GEN-01: What prompted you to look at storage now? *Trigger event. No trigger, no deal.*
2. Q-GEN-02: What happens if you do nothing for twelve months? *Cost of inaction.*
3. Q-GEN-03: How much data do you store in object storage today, and how fast is it growing? *Sizing; growth rate drives TCO.*
4. Q-GEN-04: Where does it live: which clouds, regions, on-prem systems? *Map the estate.*
5. Q-GEN-05: Who signs, who influences, who can block? *Decision process.*
6. Q-GEN-06: Is there a budget line, and when does the fiscal year turn? *Timing.*
7. Q-GEN-07: What does success look like in 90 days? *Defines POC criteria.*
8. Q-GEN-08: Which vendors are you also evaluating? *Picks the battlecard.*
9. Q-GEN-09: Have you tried an alternative cloud before? What happened? *Uncovers scar tissue.*
10. Q-GEN-10: Any contractual commitments elsewhere (AWS EDP, Azure MACC, GCP commit)? *Commit drawdown is a real blocker.*

Economics (Q-ECON)

11. Q-ECON-01: What was last month's storage bill, split into storage, requests, retrieval, egress? *Egress share is the wedge.*
12. Q-ECON-02: How much data leaves the cloud per month, and to where? *Egress destination decides free vs paid.*
13. Q-ECON-03: Which storage classes do you use, and do you ever pay early-deletion or retrieval fees? *Tiering pain.*
14. Q-ECON-04: Average object size and object count? *Request fees and small-object penalties elsewhere.*
15. Q-ECON-05: Does finance prefer opex monthly or a prepaid term? *Pay-as-you-go vs 1 to 5 year commit.*
16. Q-ECON-06: What did the last restore or data pull cost you? *Makes egress pain concrete.*
17. Q-ECON-07: What does a TB cost fully loaded on-prem: hardware, support, power, space, people? *For on-prem displacement.*

- 18. Q-ECON-08: When is your next hardware refresh or support renewal? *Timing anchor.*
 - 19. Q-ECON-09: Do you have credits from a hyperscaler, and when do they run out? *Startup trigger.*
 - 20. Q-ECON-10: Would a migration subsidy change your timeline? *Universal Data Migration lever.*
-

Backup and ransomware (Q-BKP)

- 21. Q-BKP-01: Which backup software and versions? *Integration fit (Veeam, Rubrik, Commvault, etc.).*
 - 22. Q-BKP-02: Where do copies live today: primary, secondary, off-site? *Map 3-2-1.*
 - 23. Q-BKP-03: Is any copy immutable? How is that enforced? *Object Lock opening.*
 - 24. Q-BKP-04: What does your cyber insurer require? *Insurance checklist is a forcing function.*
 - 25. Q-BKP-05: What are your RPO and RTO for tier-1 apps? *Always-hot vs Glacier.*
 - 26. Q-BKP-06: When did you last test a full restore from off-site? How long did it take? *Usually never. That is the pain.*
 - 27. Q-BKP-07: Retention policy: daily, weekly, monthly, yearly, how long? *Sizing plus churn plus lock duration.*
 - 28. Q-BKP-08: Who can delete backups today? Could a compromised admin account delete them? *Compliance mode value.*
 - 29. Q-BKP-09: Do you back up Microsoft 365 or Kubernetes? *Veeam M365, Kasten.*
 - 30. Q-BKP-10: Is tape still in the picture? Who manages it? *LTO replacement.*
 - 31. Q-BKP-11: Do you use a scale-out repository or direct-to-object? *Veeam SOBR capacity tier vs direct.*
 - 32. Q-BKP-12: What dedupe ratio do you see? *Translates logical to stored TB.*
 - 33. Q-BKP-13: Do you need copies in two geographies? *Cloud Replication across regions.*
 - 34. Q-BKP-14: How many sites and endpoints? *Computer Backup cross-sell.*
 - 35. Q-BKP-15: Have you been hit, and what did recovery look like? *Listen. Do not pitch for a full minute.*
-

Media (Q-MED)

- 36. Q-MED-01: What cameras and codecs? Average file size? *RED/BRAW/MXF large-object fit.*
- 37. Q-MED-02: How many hours or TB per week of new footage? *Ingest sizing.*
- 38. Q-MED-03: Which MAM/PAM/NLE do you use? *iconik, Mimir, axle ai, Premiere, Resolve, Avid.*
- 39. Q-MED-04: How does archive get pulled back for reuse? How often? *Egress pattern.*
- 40. Q-MED-05: LTO generation and number of tapes? *Migration scope.*
- 41. Q-MED-06: Do remote editors or freelancers need access? *Presigned URLs, Suite Studios, LucidLink.*
- 42. Q-MED-07: Do you distribute content to viewers? Through which CDN? *Partner egress.*
- 43. Q-MED-08: Any TPN or studio security requirements? *TPN Blue Shield.*
- 44. Q-MED-09: Do you need to start editing while uploads are in progress? *Live Read.*
- 45. Q-MED-10: Who owns the archive budget: production, IT, or library? *Buyer.*

AI and ML (Q-AI)

- 46. Q-AI-01: Where are your GPUs: which clouds, regions, on-prem? *Partner egress and latency.*
- 47. Q-AI-02: How big is the training corpus, and how fast is it growing? *Capacity tier sizing.*
- 48. Q-AI-03: How many times per month do you read the full dataset? *Egress multiple; 3x threshold; Overdrive trigger.*
- 49. Q-AI-04: Target throughput per node and aggregate? *Overdrive tier (100 Gbps to 1 Tbps).*
- 50. Q-AI-05: Do you have a hot tier (VAST, WEKA, Lustre, local NVMe)? *Position B2 as capacity tier beneath it.*
- 51. Q-AI-06: How do dataloaders access data: S3 API, FUSE mount, prefetch to local? *Compatibility test.*
- 52. Q-AI-07: Checkpoint size and frequency? *Write bursts and retention.*
- 53. Q-AI-08: Do you version datasets or models? With what? *Object versioning, Pixeltable, artifact stores.*
- 54. Q-AI-09: Do you move data between providers to chase GPU availability? *The one-namespace story.*
- 55. Q-AI-10: Any data residency limits on training data? *US East/West, EU Central, CA East.*
- 56. Q-AI-11: Are agents writing outputs to storage? *MCP server.*
- 57. Q-AI-12: What does an idle GPU-hour cost you? *Turns throughput into dollars.*
- 58. Q-AI-13: Inference: where do model weights load from, and how often? *Cold start reads.*
- 59. Q-AI-14: RAG: where do source documents and embeddings live? *Object as source of truth.*
- 60. Q-AI-15: Who benchmarked storage last time, and what was the test? *Reuse their methodology.*

CDN and application (Q-APP)

- 61. Q-APP-01: What CDN do you use, and what is the origin today? *Bandwidth Alliance fit.*
- 62. Q-APP-02: Cache hit ratio? *Origin egress = misses.*
- 63. Q-APP-03: Requests per second to origin, read and write? *Transaction model.*
- 64. Q-APP-04: Public or private content? Signed URLs? *Cloudflare private bucket pattern.*
- 65. Q-APP-05: Do you use S3 features like IAM roles, object tagging, website hosting, ACLs per object? *Known gaps. Find them now, not in the POC.*
- 66. Q-APP-06: SDK and language? *boto3, aws-sdk-js, Go, etc.*
- 67. Q-APP-07: Do you need event triggers on upload? *Event Notifications (webhooks; request access).*
- 68. Q-APP-08: Multi-tenant? One key per customer? *App keys, up to 100M per account.*
- 69. Q-APP-09: Latency targets for first byte? *Region choice and CDN.*
- 70. Q-APP-10: Do you serve users outside North America and EU? *Region plus CDN design.*

Migration (Q-MIG)

- 71. Q-MIG-01: Total TB and object count to move? *Method choice.*
- 72. Q-MIG-02: Available bandwidth and windows? *Network vs Fireball (96 TB Synology).*
- 73. Q-MIG-03: Source egress cost if moved over the wire? *Universal Data Migration subsidy.*
- 74. Q-MIG-04: Can apps tolerate dual-write for a period? *Cutover strategy.*

75. Q-MIG-05: Integrity requirements? *MD5 verification per UDM FAQ.*
76. Q-MIG-06: Who owns the cutover runbook? *Resource reality.*
77. Q-MIG-07: Any objects in Glacier/Archive tiers needing restore first? *Hidden time and cost.*
78. Q-MIG-08: Do bucket and object names follow S3 rules that B2 also accepts? *Naming restrictions differ slightly.*

Security and compliance (Q-SEC)

79. Q-SEC-01: Which frameworks apply: HIPAA, SOC 2, PCI, GDPR, CJIS, state RAMP? *Map to the compliance page.*
80. Q-SEC-02: Do you need a BAA or DPA? *Available.*
81. Q-SEC-03: Who manages keys? Need customer-managed keys? *SSE-C.*
82. Q-SEC-04: SSO for the console? *Enterprise web console and OIDC for Computer Backup Enterprise Control* [verify B2 console SSO options].
83. Q-SEC-05: Do you require private connectivity? *Overdrive private networking.*
84. Q-SEC-06: Log retention and audit needs? *Bucket access logs.*
85. Q-SEC-07: Data residency? *Region selected at account creation, cannot change.*
86. Q-SEC-08: Vendor questionnaire format: SIG, CAIQ, HECVAT? *Whistic profiles.*

Persona closers

- **CIO:** "If we prove the TCO and the restore time in 30 days, what stops you from signing?"
- **Admin:** "If I set this up with you in 20 minutes today, will you run your next restore test against it?"
- **ML lead:** "If we hit your GB/s target from your GPU cloud, is storage solved for this training run?"
- **MSP:** "If unit cost and restore egress are fixed, which client do you move first?"

05. Qualification: the STORE framework

A POC is a decision instrument, not a science project. Qualify before you build.

Use MEDDPIC if your team does. STORE is the storage-specific overlay that tells you whether Backblaze is the right answer.

S: Size and shape

- Stored TB today and in 36 months.
- Object count and average size.
- Read multiple per month (egress / stored).
- **Green:** 50 TB+, objects over 1 MB average, egress under 3x stored or flowing to partners.
- **Yellow:** tiny objects with very high request rates, or egress 3x to 5x stored to the internet.
- **Red:** under 5 TB with no growth, or a hard requirement for sub-millisecond latency.

T: Trigger and timeline

- Named event: renewal, incident, insurance, credits ending, refresh, migration.
- Date that forces a decision.
- **Red:** "just benchmarking prices."

O: Owner and organization

- Economic buyer identified and met.
- Technical champion who will run the POC.
- Who blocks: security, procurement, an AWS commit owner.

R: Requirements fit

Check every hard requirement against documented capability:

Requirement	Fit	Notes
S3 API core (Put/Get/List/Multipart/Presign)	Yes	
Object Lock (governance, compliance, legal hold)	Yes	1 to 3,000 days per docs
Versioning	Yes	B2 buckets keep versions; lifecycle rules manage them
SSE - B2 / SSE - C	Yes	SSE - B2 default for new uploads since Sep 14 2026

Requirement	Fit	Notes
Customer KMS integration (bring-your-own KMS)	Not documented	SSE-C means the client supplies the key per request [verify roadmap]
Object-level ACLs	No	Bucket-level private/public only
IAM roles / STS	Partial	Docs list an "IAM-STS API" in nav; S3-compatible doc lists IAM roles unsupported [verify current]
Object tagging	No	Returns empty set
Static website config	No	Use a CDN in front
Lifecycle transitions to colder class	N/A	Single hot class; lifecycle hides/deletes
Cross-region replication	Yes	Cloud Replication, separate account per region
Event notifications	Yes	Webhooks, request access, Class D charged
Private connectivity	Overdrive	Private networking integration
Regions	US West, US East, EU Central, CA East	
FedRAMP	No	GovRAMP Progressing, TX-RAMP Provisional

E: Economics proven

- TCO built with their bill, not list estimates.
- Commit structure agreed in principle (PAYG, 1 to 5 year term, Overdrive multi-PB).
- Migration funding discussed (Universal Data Migration with qualifying commitment).

Red flags

- They need a feature listed as unsupported and will not change the app.
- Their egress goes to the open internet at 5x+ stored with no CDN (R2 may genuinely be cheaper; see 15).
- A large hyperscaler commit they must burn down this year.
- No champion willing to run a test.

Deal sizing quick math

Annual value $\approx$ stored_TB $\times$ \$/TB/month $\times$ 12

- 500 TB at list = $500 \times 6.95 \times 12 = \$41,700/\text{yr.}$
- 2 PB at list = $2,000 \times 6.95 \times 12 = \$166,800/\text{yr.}$
- 5 PB Overdrive at \$15 list = $5,000 \times 15 \times 12 = \$900,000/\text{yr.}$ Note: deals over \$500K sit outside public guidance, which is why they matter.

Exit criteria for each stage

Stage	Exit when
Discovery	STORE S, T, O filled; TCO inputs collected
Technical validation	Hard requirements checked against table above; no red
POC	Written success criteria met (see 17)
Commercial	Term and price agreed; migration plan signed off
Close	Paper signed; onboarding date set

20. Objection bank (60)

Format: **They say / You say / Prove it.** Acknowledge first, ask second, answer third. Never argue with the customer's experience.

Price and value

1. **"Wasabi is basically the same price."** You say: "On storage, close: \$7.99 vs \$6.95. The difference is what happens when you read. B2 includes 3x your stored data in free egress every month, unlimited through partners, and no minimum storage duration." Prove: 14-tco table at their read multiple and deletion pattern.

2. **"Glacier is way cheaper."** "It is, for data you never touch. What is your restore SLA, and what did the last retrieval cost?" Prove: restore scenario including retrieval fee, hours, egress.
 3. **"Your price just went up from \$6 to \$6.95."** "It did. Wasabi went up too. That is why term commits of 1 to 5 years lock the rate." Prove: term option on quote.
 4. **"R2 has zero egress."** "True. R2 storage is \$15/TB, more than double, and writes cost \$4.50 per million. If you use Cloudflare, put B2 behind it and get \$0 egress at \$6.95." Prove: 1 PB comparison.
 5. **"We get big discounts from AWS."** "Great, let's use your effective rate. What is it on storage and egress?" Prove: rerun TCO with their numbers.
 6. **"\$0.01/GB overage could surprise us."** "Overage only starts above 3x stored. At 1 PB you get 3 PB a month free. If you will exceed that, Overdrive gives unlimited egress." Prove: break-even 3.8x.
 7. **"We cannot prepay."** "Pay-as-you-go monthly by card or invoice. No commitment." Prove: pricing page.
 8. **"Savings are not worth the migration effort."** "Universal Data Migration offsets migration cost on a qualifying commitment and Backblaze runs it. What is your team's cost per month of delay?" Prove: UDM page.
 9. **"Our storage bill is small."** "Then let's find the line that hurts: restores, data movement, or the next refresh." Disqualify politely if nothing hurts.
 10. **"Free egress is a gimmick; you will change it."** "It is published on the pricing page with a \$0.01 overage. Term contracts lock pricing." Prove: contract terms.
-

Technical fit

11. **"Is it really S3 compatible?"** "Core APIs yes, with documented gaps: no object ACLs, IAM roles, object tagging, website config, or browser POST. Which do you use?" Prove: run their app in POC.
12. **"We use IAM roles."** "Not in the S3-compatible API. Use scoped app keys distributed through your secrets manager." Prove: key scoping demo.
13. **"We rely on object tags."** "Tags are not supported. Use prefixes or a metadata index. Many apps only need tags for lifecycle, which B2 handles by prefix." Prove: lifecycle rule demo.
14. **"Is it fast enough?"** "For which workload? Hume AI hit about 10 GB/s uploads. Overdrive goes 100 Gbps to 1 Tbps dedicated. Let's measure from your network." Prove: throughput test.
15. **"Latency for small files?"** "B2 is optimized for throughput and capacity. For sub-millisecond small I/O, keep a hot tier." Prove: honest architecture.
16. **"Single region per account is a limitation."** "It is a design choice that keeps residency explicit. Cloud Replication covers multi-region with a second account." Prove: replication setup.

17. **"No storage classes? We want tiering."** "One hot class at a price below most cool tiers. No retrieval fees, no minimum durations. Tiering is complexity you no longer need." Prove: price compare vs S3 IA (\$12.50/TB) and Azure Cool.
18. **"Does it work with our backup software?"** "Veeam, Rubrik, Commvault, MSP360, Synology, TrueNAS and hundreds more have guides. Rubrik needs SigV4 enabled by their support." Prove: integration guide.
19. **"What about Kubernetes?"** "Kasten K10 and Velero-style S3 targets work; Object Lock with Kasten is documented on the Veeam page." Prove: test.
20. **"Can we mount it like a drive?"** "Via tools like Suite Studios, LucidLink, rclone mount, Mountain Duck. Object storage is not a POSIX filesystem; know the trade-offs." Prove: Suite Studios integration.
21. **"Event triggers?"** "Event Notifications send signed webhooks on create, update, delete. Access is enabled by Support." Prove: demo 6.
22. **"We need versioning."** "B2 keeps versions by default. Lifecycle rules control how long." Prove: console.
23. **"Can we host a static site?"** "Not via S3 website config. Put Cloudflare or Fastly in front, with \$0 egress." Prove: CDN guide.
24. **"Our app uses SigV2."** "B2 requires SigV4. Most SDKs default to V4; update the client." Prove: SDK config.
25. **"Max object size?"** "Large files via multipart.
[verify current max object and part size in docs before quoting]"

Trust, risk, and company

26. **"Backblaze is a consumer backup brand."** "B2 is 62 percent of revenue, 100,000+ customers, and CoreWeave signed a \$335M multi-exabyte deal." Prove: Q2 2026 results, 8-K.
27. **"Will you be around?"** "Public company, adj. EBITDA margin 30 percent in Q2 2026, ARR \$177M growing 21 percent." Prove: IR.
28. **"Your drive failure rate went up."** "Yes, Q2 2026 AFR was 1.73 percent, and Backblaze published it. Durability comes from 17+3 erasure coding across 20 pods, not from drives never failing." Prove: Drive Stats, durability doc.
29. **"You are too small for our scale."** "Over an exabyte under management; Decart went 0 to 16 PB in 90 days." Prove: case study.
30. **"Support quality?"** "Premium support add-on; Overdrive includes the highest tier; customers like Mirage cite 'real human support'." Prove: references.

31. **"What if you get acquired?"** "Contract terms and S3 compatibility protect you. Your data is portable with 3x free monthly egress." Prove: egress math for exit.
32. **"Hard drives are old tech."** "HDD is how the industry stores exabytes economically. Even CoreWeave uses Backblaze for its HDD tiers under a flash cache." Prove: CoreWeave announcement.
33. **"We had a bad experience with an alt-cloud."** "What broke? Let's put that exact test in the POC." Prove: POC criterion.
34. **"Your SLA?"** "[pull current SLA]. Durability designed at 11 nines."

Security and compliance

35. **"Are you HIPAA compliant?"** "We sign BAAs for covered entities. HIPAA compliance is shared." Prove: compliance page.
36. **"FedRAMP?"** "No. GovRAMP Progressing and TX-RAMP Provisional." Prove: compliance page. Qualify out if hard requirement.
37. **"Is data encrypted?"** "SSE-B2 AES-256 by default on new uploads since Sep 14 2026; SSE-C if you hold keys." Prove: blog.
38. **"Can your staff delete our backups?"** "Not in compliance mode, not before expiry." Prove: demo 3.
39. **"We need customer-managed KMS."** "SSE-C today means the client supplies the key per request. Native KMS integration is not documented." Prove: SSE docs. Escalate roadmap ask.
40. **"Do you have SOC 2?"** "SOC 2 Type 2; report via Whistic." Prove: Whistic.
41. **"GDPR residency?"** "EU Central in Amsterdam; DPA available; data stays in the selected region." Prove: regions doc.
42. **"Is it air-gapped?"** "Logically: off-site, separate credentials, immutable. Not a physical gap." Prove: architecture.
43. **"What about insider threats?"** "Scoped keys, compliance-mode lock, MFA, access logs." Prove: key demo.

Organization and process

44. **"We have an AWS commit to burn."** "Then start with the copy that must sit outside AWS: the immutable backup. It does not touch your commit's core workloads." Prove: backup design.
45. **"Procurement takes six months."** "Start PAYG on a card for the POC; paper the term deal in parallel." Prove: PAYG.

46. **"Security review will take forever."** "Whistic profiles have SOC 2, HECVAT, and questionnaires ready." Prove: Whistic.
47. **"We do not have staff for a migration."** "UDM is white-glove." Prove: UDM.
48. **"Our MSP decides."** "Let's bring them in; Backblaze has an MSP program and Partner API." Prove: partner page.
49. **"We are standardizing on one cloud."** "For compute, sensible. For the backup copy, standardizing on the same provider puts the copy in the same blast radius." Prove: 3-2-1-0.
50. **"Not a priority this quarter."** "What changes next quarter? Insurance renewal, refresh, credits ending?" Prove: trigger.
-

AI specific

51. **"We need flash for training."** "For the working set, yes. B2 holds the full corpus and checkpoints and refills the flash tier." Prove: whiteboard 11.
52. **"Our GPUs are in a hyperscaler."** "Then model egress from B2 at 3x free or Overdrive unlimited, versus hyperscaler storage price." Prove: TCO scenario.
53. **"Neoclouds can build storage with Ceph."** "They can. Then they run a storage company. B2 Neo launches in weeks." Prove: white paper.
54. **"Is CoreWeave a B2 Neo customer?"** "CoreWeave signed a multi-exabyte storage agreement; Backblaze powers HDD tiers in CoreWeave AI Object Storage. The SEC filing does not label it B2 Neo." Prove: 8-K wording.
55. **"Can agents safely use storage?"** "B2 MCP Server exposes only tools the key permits." Prove: demo 7.
56. **"Throughput claims are marketing."** "Agreed. Test from your GPU nodes with your shards." Prove: demo 8.
-

Media specific

57. **"Our MAM uses tags."** "Test it. Most MAMs keep metadata in their own DB." Prove: POC.
58. **"Tape is cheaper per TB."** "Media is. Add library maintenance, LTO migrations, off-site handling, restore delays. Eagles retired LTO-6." Prove: case study.
59. **"Uploads of 8K are too slow."** "Size the link; use Fireball for backlog; Live Read starts editing while uploading." Prove: transfer table.

60. "Studios require TPN." "Backblaze has TPN Blue Shield status." Prove: compliance page.

24. Field lab: drills, flashcards, quiz

Fifteen minutes to get sharp. Do one scenario drill out loud every morning of your first month.

Scenario drills

Each drill: read the setup, set a 10-minute timer, draw the board, say the first three questions, name the risk, name the proof.

Drill 1: The insurer letter

Setup: 400-person manufacturer, Veeam v12, 180 TB of backups on a local linux repo, cyber insurer renewal in 60 days demands an immutable off-site copy. **Board:** whiteboard 09. **First questions:** retention schedule? RTO for ERP? who has admin on the Veeam server? **Risk:** Object Lock not enabled before connecting Veeam. **Proof:** 20-minute setup and failed delete demo. **Numbers:** $180 \times 6.95 = \$1,251/\text{month}$ list.

Drill 2: The egress shock

Setup: SaaS video platform, 600 TB in S3 Standard, 400 TB/month egress to Fastly. Ask for the bill, then check it with the math below. **Math check:** storage $50,000 \times 0.023 + 450,000 \times 0.022 + 100,000 \times 0.021 = 1,150 + 9,900 + 2,100 = \$13,150$. Egress $399,900 \text{ GB} : 900 + 3,400 + 7,000 + 249,900 \times 0.05 = 12,495 \rightarrow \$23,795$. Total $\approx \$36,945$. **B2:** $600 \times 6.95 = \$4,170$; egress to Fastly \$0. **Risk:** app uses object tagging for moderation flags. **Proof:** POC with their SDK; tag dependency replaced by DB column.

Drill 3: The GPU nomad

Setup: AI lab, 3 PB corpus in GCS, renting GPUs at CoreWeave this quarter and maybe Lambda next. Reads full corpus $\sim 2x/\text{month}$. **Board:** whiteboard 11. **Answer:** B2 standard: $3,000 \times 6.95 = \$20,850/\text{month}$; $2x$ read to CoreWeave = \$0 partner egress. To non-partner at $2x$: still inside $3x$ free. **Risk:** dataloader uses GCS-specific client. **Proof:** throughput test from CoreWeave nodes.

Drill 4: The tape library

Setup: regional broadcaster, 1,200 LTO-5 tapes (~1.5 TB each native ≈ 1.8 PB), library end of support, archive pulled weekly for retrospectives. **Answer:** UDM tape migration with commit; B2 library + MAM; delivery through Cloudflare. **Risk:** tape read failures; budget for unreadable media. **Proof:** Eagles and ACL references.

Drill 5: The Wasabi renewal

Setup: MSP with 900 TB on Wasabi, renewal at \$7.99, clients restore ~1.2x stored per month in aggregate during incidents. **Answer:** 3x free egress vs fair-use; no minimum duration for short-retention clients; Partner API. **Math:** \$7,191 vs \$6,255/month at list. **Risk:** migration of backup chains; plan re-seed vs copy.

Drill 6: The neocloud founder

Setup: 8,000-GPU neocloud, customers keep data in AWS, founder considering Ceph. **Board:** whiteboard 11, B2 Neo variant. **First questions:** customers' data location? storage headcount? billing integration? **Proof:** Partner API demo, custom domain, CoreWeave reference pattern (stated precisely).

Drill 7: The Azure shop

Setup: hospital system, all Microsoft, M365 backups to Azure Blob in the same tenant. **Answer:** separate-blast-radius copy on B2 via Veeam M365, BAA, compliance mode. **Risk:** MACC commitment politics. **Proof:** BAA, SOC 2, failed delete demo.

Drill 8: The R2 developer

Setup: startup storing 40 TB user images on R2, 300M Class A ops/month (thumbnails), Cloudflare CDN. **Math:** R2: $40 \times 15 = \$600 + 300 \times 4.50 = \$1,350 \rightarrow \$1,950+$. B2 behind Cloudflare: $40 \times 6.95 = \$278$, egress \$0, ops free. **Risk:** Workers bindings to R2 need rework to S3 calls.

Flashcards (front / back)

1. B2 list price? / \$6.95 per TB per month, billed per byte-hour.
2. Free egress allowance? / 3x average monthly stored; then \$0.01/GB.
3. Partner free-egress list? / Fastly, Cloudflare, bunny.net, CacheFly, CoreWeave, Equinix Metal, Vultr, phoenixNAP.
4. Overdrive starting price? / \$15/TB/month, multi-PB commit.
5. Overdrive throughput? / 100 Gbps to 1 Tbps.
6. Overdrive egress? / Unlimited free, any destination.
7. Erasure coding scheme? / 17 data + 3 parity across 20 pods.
8. Designed durability? / 11 nines annually.
9. Regions? / US West, US East, EU Central, CA East.

10. Can you change an account's region? / No.
11. Object Lock retention range? / 1 to 3,000 days.
12. Can Object Lock be disabled on a bucket? / No.
13. Compliance vs governance? / Compliance cannot be shortened or removed; governance can with privileged key.
14. Legal hold expiry? / None until removed.
15. Lifecycle rules per bucket? / Up to 100.
16. Lifecycle rule fields? / fileNamePrefix, daysFromUploadingToHiding, daysFromHidingToDeleting, daysFromStartingToCancelingUnfinishedLargeFiles.
17. App keys per account? / Up to 100 million.
18. Default encryption? / SSE-B2 AES-256 on new uploads since Sep 14 2026.
19. SSE-C? / Customer supplies AES-256 key per request.
20. S3 API gaps? / Object ACLs, IAM roles, tagging, website config, browser POST.
21. Signature version? / SigV4 only.
22. Rubrik gotcha? / Defaults to SigV2; support must enable V4.
23. Veeam lifecycle setting? / Keep all versions.
24. Class D cost? / 2,500/day free then \$0.004 per 10,000.
25. Cross-region replication requirement? / Second account in target region.
26. Can replicated data be replicated again? / No.
27. Existing-file replication timing? / Daily at midnight UTC.
28. Fireball capacity? / 96 TB Synology NAS, 10GbE RJ-45.
29. UDM integrity method? / MD5 byte-for-byte verification.
30. MCP server tool count? / 40 (17 native, 19 S3, 4 analytics); read-only profile 20.
31. Computer Backup price? / \$99/year per computer.
32. HIPAA answer? / BAA available for covered entities.
33. Gov programs? / GovRAMP Progressing Snapshot, TX-RAMP Provisional.
34. CoreWeave deal size? / ~\$335M, multi-exabyte, five-year (5 and 7 year order forms).
35. Q2 2026 B2 growth? / +34% to \$26.6M.
36. ARR Q2 2026? / \$177.3M, +21%.
37. Overdrive break-even read multiple? / ~3.8x stored per month.
38. Q2 2026 Drive Stats AFR? / 1.73%.
39. Free storage? / First 10 GB.
40. Minimum storage duration? / None.

Quiz (answers at bottom)

1. A customer stores 200 TB and egresses 700 TB/month to the internet. B2 monthly list?
2. Which mode would you use for an insurer-mandated backup copy?
3. A customer needs survival of a full region loss. What do you design?
4. Which S3 call will be rejected if it includes x-amz-tagging?
5. Name two things Fireball pricing includes.
6. Why should you never say "unlimited free egress" for standard B2?
7. What is the first thing you do in a Rubrik POC?

8. At what read multiple does Overdrive beat standard B2 on list price?
9. What happens to existing files when SSE-B2 default is enabled?
10. What does the MCP server do with a read-only key?

Answers: 1) $200 \times 6.95 = \$1,390 + (700 - 600) \text{ TB} \times \$10 = \$1,000 \rightarrow \$2,390$. 2) Compliance mode. 3) Cloud Replication to a bucket in a second account in another region. 4) PutObject or CopyObject. 5) Security deposit, non-refundable shipping fee, per-30-day rental. 6) It is 3x free then \$0.01/GB; unlimited applies to partners and Overdrive. 7) Open a Rubrik support case to enable SigV4. 8) ~3.8x. 9) Not encrypted retroactively; only new uploads. 10) Exposes only read tools (read-only profile, 20 tools max).

31. Call plans

Every meeting has one job. Name it at the start, check it at the end.

A. First discovery call (30 minutes)

Job: decide whether there is a qualified storage problem worth a whiteboard.

Minute	Move
0 to 3	Agenda and permission: "I'll ask a lot of questions, then tell you honestly if we fit."
3 to 18	Q-GEN-01 to 05, then Q-ECON-01 and 02. Get stored TB, growth, egress, trigger.
18 to 23	30-second pitch matched to persona; one proof point from their industry.
23 to 27	Confirm STORE gaps; ask who else should be in the whiteboard.
27 to 30	Book the whiteboard; send the recap template the same day.
Leave with: five numbers, a trigger date, a second meeting.	

B. Technical whiteboard (60 minutes)

Job: agree on a target architecture and POC criteria.

Minute	Move
0 to 5	Restate what you heard. Ask what changed.
5 to 20	Draw current state with them. Numbers on arrows.
20 to 40	Draw future state using the matching board (09 to 13).
40 to 50	Walk the requirement fit table (05). Surface gaps honestly.
50 to 60	Write POC criteria on the board. Assign owners. Photograph.
Leave with: a photographed board and signed-off POC criteria.	

C. Executive briefing (30 minutes, CIO/CFO)

Job: get sponsorship for the POC and commercial path.

1. Their situation in three bullets (2 min).
2. The economics: one table, their numbers, sources dated (8 min).
3. Risk controls: immutability, compliance, migration support, company stability (8 min).
4. Proof: one reference in their industry (4 min).
5. Ask: approve POC and identify the commercial path (PAYG vs term) (8 min).

D. POC kickoff (45 minutes)

Job: zero ambiguity on scope.

- Confirm environment, data set, keys, endpoint, region.
- Walk the scorecard (17).
- Schedule check-ins (day 3, day 7, mid-point, final).
- Name the decision meeting date.

E. POC readout (30 minutes)

Job: convert results into a decision.

- Scorecard criterion by criterion.
- TCO refreshed with measured numbers.
- Proposed commercial structure.
- Ask for the decision or the specific blocker.

F. Partner QBR (45 minutes)

Job: pipeline and enablement.

- Joint wins and losses; what broke in integration.
- Top 10 overlapping accounts.
- One whiteboard enablement (09 or 12).
- Next quarter targets.

G. Neocloud architecture session (90 minutes)

Job: decide build vs buy.

1. Their GPU footprint and customer data locations.
2. Draw the split-stack problem: customer data in hyperscaler, compute on neocloud, egress in between.
3. Build-it-yourself cost model: hardware, Ceph/MinIO ops, on-call, drive replacement, upgrades.
4. B2 Neo: Partner API flow, custom domain, SSO, billing feed.
5. Time to launch and commercial model.

Call hygiene

- Notes in the CRM within one hour, with Q-codes.
- Every number tagged "customer said" or "we estimated".
- Next step has a date, or it is not a next step.

17. POC and demo scripts

A POC is a decision instrument. Write the success criteria before you write the first command.

POC charter template

Field	Example
Workload	Veeam SOBR capacity tier, 40 TB subset
Success criteria	1) Backup job completes nightly within window. 2) Object Lock prevents deletion by admin key. 3) Restore of 2 TB VM set completes in under X hours. 4) Projected 3-year TCO under \$Y.
Duration	14 to 30 days
Owners	Customer: backup admin. Backblaze: SE.
Data	Non-production or copy
Exit	Signed decision meeting on day N

Backblaze advertises a free POC for 50 TB+ (cloud migration and CDN pages: “Want to run a free proof of concept with more than 50TB?”).

Demo 0: setup (2 minutes)

```
# b2 CLI (python)
pip install --upgrade b2
b2 account authorize          # prompts for keyID and applicationKey
b2 bucket create --help      # check flags for your CLI version
```

CLI syntax changed in b2 CLI v4 (subcommands like `b2 bucket create`, `b2 file upload`). Older docs show `b2 create-bucket`. Run `b2 --help` before the demo.

[verify flags against installed version]

Create a bucket-scoped key in the web console for the demo. Never demo with the master key.

Demo 1: S3 drop-in with aws CLI (5 minutes)

```
export AWS_ACCESS_KEY_ID=<b2 keyID>
export AWS_SECRET_ACCESS_KEY=<b2 applicationKey>
export EP=https://s3.us-west-004.backblazeb2.com # copy exact endpoint from
bucket page

aws s3 mb s3://nix-demo-bucket --endpoint-url $EP
aws s3 cp ./sample.mp4 s3://nix-demo-bucket/media/sample.mp4 --endpoint-url $EP
aws s3 ls s3://nix-demo-bucket --recursive --endpoint-url $EP
aws s3 presign s3://nix-demo-bucket/media/sample.mp4 --expires-in 3600 --endpoint-
url $EP
```

Say: "Same CLI. I changed two things: the endpoint and the keys."

Demo 2: boto3 (5 minutes)

```
import boto3
s3 = boto3.client(
    "s3",
    endpoint_url="https://s3.us-west-004.backblazeb2.com",
    aws_access_key_id=<keyID>,
    aws_secret_access_key=<applicationKey>,
)
s3.upload_file("sample.mp4", "nix-demo-bucket", "media/sample.mp4")
for o in s3.list_objects_v2(Bucket="nix-demo-bucket",
Prefix="media/").get("Contents", []):
    print(o["Key"], o["Size"])
url = s3.generate_presigned_url(
    "get_object", Params={"Bucket": "nix-demo-bucket", "Key": "media/sample.mp4"},
    ExpiresIn=900
)
print(url)
# presigned PUT for direct browser/device upload
put_url = s3.generate_presigned_url(
    "put_object", Params={"Bucket": "nix-demo-bucket", "Key": "uploads/user1.jpg"},
    ExpiresIn=900
)
```

Demo 3: Object Lock, the money demo (7 minutes)

```
import boto3, datetime
s3 = boto3.client("s3", endpoint_url=EP, aws_access_key_id=KID,
aws_secret_access_key=KEY)

s3.create_bucket(Bucket="nix-immutable-demo", ObjectLockEnabledForBucket=True)
until = datetime.datetime.utcnow() + datetime.timedelta(days=1)
r = s3.put_object(
    Bucket="nix-immutable-demo", Key="backups/restorepoint.vbk", Body=b"pretend
backup",
    ObjectLockMode="GOVERNANCE",          # use COMPLIANCE only if you accept
paying until expiry
    ObjectLockRetainUntilDate=until,
)
vid = r["VersionId"]
try:
    s3.delete_object(Bucket="nix-immutable-demo", Key="backups/restorepoint.vbk",
VersionId=vid)
except Exception as e:
    print("Delete blocked:", e)          # AccessDenied: object is locked
print(s3.get_object_retention(Bucket="nix-immutable-demo",
Key="backups/restorepoint.vbk", VersionId=vid))
```

Say: "That was my key. Imagine it is the attacker's. The restore point is still there."

Demo hygiene: use GOVERNANCE and 1 day in demos. Compliance mode on a demo bucket means paying for it until expiry. Object Lock cannot be turned off on the bucket afterwards.

Demo 4: rclone migration (8 minutes)

```
rclone config create b2s3 s3 provider=Other \
    access_key_id=<keyID> secret_access_key=<applicationKey> \
    endpoint=s3.us-west-004.backblazeb2.com

# or native backend
rclone config create b2native b2 account=<keyID> key=<applicationKey>

# copy from AWS (configured as remote 'aws') to B2 with parallelism and checksum
compare
rclone copy aws:source-bucket b2s3:nix-demo-bucket \
    --transfers 64 --checkers 32 --s3-chunk-size 64M --checksum --progress
```

```
rcclone check aws:source-bucket b2s3:nix-demo-bucket --one-way
```

Tune `--transfers` and chunk size to bandwidth. Show `rcclone check` output as integrity proof.

Demo 5: lifecycle and versions (3 minutes)

Web console: Bucket Settings → Lifecycle Settings. Show "Keep only the last version" vs custom rule with prefix. Explain hide vs delete. Or native:

```
b2 bucket update nix-demo-bucket allPrivate \  
  --lifecycle-rules  
'[{"fileNamePrefix":"tmp/","daysFromUploadingToHiding":7,"daysFromHidingToDeleting":1}]'
```

[verify flag name for your CLI version]

Demo 6: Event Notifications (5 minutes, needs access enabled by Support)

1. Spin up a webhook receiver (e.g., a Cloudflare Worker or `python -m http.server` behind a tunnel).
2. Console: Bucket → Event Notifications → add rule: event `b2:ObjectCreated:*` [verify event type names], prefix `media/`, endpoint URL, signing secret.
3. Upload a file; show the POST payload; verify the HMAC.

Demo 7: MCP server with a scoped key (5 minutes)

1. Create a read-only key on one bucket.
2. Configure an MCP client (Claude Desktop, Cursor, etc.) with the B2 MCP Server per the docs repo.
3. Ask: "List the 10 largest files in nix-demo-bucket." Then ask it to delete one. Show the tool is not available.
4. Swap to a write key scoped to `outputs/` prefix; show the tool list grow.

Demo 8: throughput test (AI POCs)

```
# parallel download test from a GPU node
rclone copy b2s3:dataset/shards/ /local/nvme/ --transfers 128 --multi-thread-streams 16 --progress
# or s5cmd for high-parallel S3
s5cmd --endpoint-url $EP --numworkers 256 cp 's3://dataset/shards/*' /local/nvme/
```

Record GB/s, node count, object size, region pair. Compare to target from discovery. If short, discuss Overdrive.

POC scorecard

Criterion	Target	Result	Pass
Ingest rate	X TB/day		
Restore / read rate	Y GB/s		
Immutability enforced	Delete fails		
Integration works unmodified	Yes		
Projected 3-yr TCO	≤ \$Z		
Security review	Passed		

Failure modes to pre-empt

- Wrong endpoint region → ‘InvalidAccessKeyId’ or auth errors.
- Using the master key with the S3 API → fails. The master application key does not work with the S3-compatible API; create a standard (bucket-scoped) application key.
- SigV2 client → fails.
- Object Lock not enabled before Veeam connects → immutability option unavailable.
- Tiny objects with huge parallelism → measure ops, not GB/s.

19. Migration playbook

Draw the end state, then the path. Migration risk is the number one silent objection.

Phase 0: Assess (week 1)

- Inventory per source: TB, object count, average size, largest object, storage class (Glacier needs restore first), versioning, Object Lock, tags/ACLs in use.
- Network: available bandwidth, windows, egress costs at source.
- Apps: which read/write the data; how endpoints and keys are configured.
- Compliance: residency, encryption, retention locks at source (locked source objects cannot be deleted after copy).

Output: migration sheet with TB per lane and target dates.

Phase 1: Choose lanes

Lane	When	Tool	Notes
Universal Data Migration (UDM)	Cloud, NAS, SAN, tape, drives, cloud drives (Dropbox, Box, Google Drive, M365), MAM/DAM	Backblaze-managed	Requires qualifying commitment; Backblaze helps offset egress and transfer costs; MD5 byte-for-byte verification; white-glove
Self-serve online	Under ~100 TB or good bandwidth	rclone, aws s3 sync, s5cmd, backup software re-seed	Parallelism is the main lever
Fireball	Slow links, 10s of TB per batch	96 TB Synology NAS, 1GbE/10GbE RJ-45	Security deposit + non-refundable shipping + per-30-day rental, not prorated; ships from US; Backblaze uploads and verifies, then wipes
Backup re-seed	Backup chains	Veeam/Commvault/etc.	Often better to start new chains on B2 and age out old

Phase 2: Pilot (week 2)

- Move 1 to 5 percent including the largest and smallest objects.
- Validate: count, bytes, checksums, app read/write against B2 endpoint.
- Measure throughput; extrapolate.

Phase 3: Bulk move

- Run lanes in parallel. Track daily TB moved vs plan.
- For S3 sources, consider S3 Batch Operations or rclone from compute in the same region as the source to maximize throughput. **[egress still applies unless UDM covers it]**

Phase 4: Delta and dual-write

- Option A: app writes to both for N days.
- Option B: final delta sync during a short freeze.
- Option C: read-through: app reads B2 first, falls back to source.

Phase 5: Cutover

- Change endpoint and keys in config. Use DNS or config management, not hand edits.
- Keep source read-only 30 days. Then delete per policy.

Phase 6: Verify and close

- `rclone check --one-way` or manifest diff.
- Sign-off document: counts, bytes, checksum sample, app tests.

Fireball mechanics (docs, checked 2026-10-14)

1. Order through Support.
2. Receive Synology NAS plus drives; install drives.
3. Connect via ethernet; find with Synology Assistant.
4. Copy data (some backup ISVs require loading with their tools).
5. Remove drives, repack, ship back.

6. Backblaze uploads to B2, verifies, clears the NAS. Email notices at each step. International: contact Support.

Risk register

Risk	Mitigation
Source egress bill	UDM subsidy with commit; move from in-region compute
Glacier objects	Restore in advance; budget retrieval
Naming incompatibility	Check B2 naming rules; rename map
Tag/ACL dependency	Replace with prefixes/manifests; bucket-level public/private
App downtime	Dual-write / read-through
Integrity	MD5 verification (UDM), rclone check
Locked source data	Copy but cannot delete until expiry; budget overlap

Proof

Segmed: 750 TB for AI medical imaging. Eagles: LTO to cloud. ACL: 4,000+ hours of tape. Jazz at Lincoln Center: 1.3 PB, 400+ hard drives.

22. Deal story templates

Stories beat feature lists. Tell three workflow stories, not thirty features.

The template (use for every win, loss, and public case study)

1. **Who:** industry, size, persona who championed.
2. **Before:** where data lived, what hurt, with a number.
3. **Trigger:** why now.
4. **Alternatives:** who else they evaluated.
5. **What we drew:** the architecture in one sentence.
6. **Proof:** the POC criterion that decided it.
7. **After:** outcome with a number and a quote.
8. **Lesson:** what the next SE should copy.

Keep it to 150 words spoken, 90 seconds.

Public stories rewritten in the template (from backblaze.com case studies, checked 2026-10-14)

Decart (AI)

- **Who:** AI startup focused on real-time inference; CEO Dean Leitersdorf.
- **Before:** burning free credits at a traditional cloud; data growing from hundreds of thousands to hundreds of millions of video hours.
- **Trigger:** credits and scale.
- **What we drew:** B2 as the single store feeding GPU clusters around the world; Overdrive.
- **After:** 0 to 16 PB in 90 days; "10x more efficient than competition"; \$0 egress spend.
- **Quote:** "B2 Cloud Storage was literally the only technical thing we used in training these models that didn't crash the first time we tried it."
- **Lesson:** lead with reliability under scale, not just price.

Hume AI (AI)

- **Before:** audio data across multiple hyperscale object stores past PB scale.
- **What we drew:** one unified object store with WEKA in the stack.
- **After:** close to 10 GB/s upload from clusters; \$0 egress fees; one store.

- **Quote:** "Our growth doesn't follow a curve, it follows a trigger."
- **Lesson:** consolidation is a story on its own.

Mirage (AI video; URL slug "captions")

- **Before:** high egress and inflexible regions at prior provider while moving to H100 clusters.
- **Proof:** competitive bakeoff.
- **After:** 0 GPUs waiting for data, 95% estimated cost savings, \$0 egress (Overdrive).
- **Lesson:** run the bakeoff; Overdrive wins on dedicated throughput plus egress.

Philadelphia Eagles (sports media)

- **Before:** failure-prone LTO-6, slow access, no geographic redundancy.
- **What we drew:** Mimir MAM + Quantum QXS + B2 with Cloud Replication; UDM.
- **After:** 800 TB under management, LTO retired, 100% geographic redundancy.
- **Quote:** "We went from frequent LTO crashes and long restore times to near-instant access for every stored clip."

Austin City Limits (public TV)

- **Before:** 4,000+ hours of deteriorating videotape.
- **What we drew:** Fireball + UDM into B2.
- **After:** 40+ years preserved; 40 TB transferred without delay.

AcenTek (telecom IT)

- **Trigger:** cyber insurance requirement.
- **After:** proved immutable off-site backups to insurer.
- **Quote:** "Adding Backblaze to our infrastructure allowed us to satisfy our insurance carrier's requirements."
- **Lesson:** insurance is the forcing function for backup deals.

Tribute (SaaS)

- **After:** \$15K saved per month. "Backblaze was a drop-in API replacement with a much better cost structure."

Pittsburg State University (education)

- **After:** off-site, ransomware-protected backups with Object Lock; "4x data protection", "5 decades of data".

Discovery-to-story worksheet (fill after each call)

Field	Notes
Pain number	
Trigger date	
Competitor	
Board drawn	
POC criterion	
Champion quote	

Telling it

- Open with the “before” number.
- Pause after the trigger.
- Draw the architecture in five strokes.
- Close with the “after” number and ask: “Does that look like your situation?”

30. Outreach and follow-up templates

Templates for SEs and AEs. Personalize the first line or do not send it.

Post-discovery recap

Subject: Recap: [Company] storage discovery and next step Body: Thanks for the time today. What I heard:

- ☒ TB in [provider], growing [Y]% a year.
- Pain: [restore cost / egress / tape / GPU data movement], about \$[N] last quarter.
- Trigger: [insurance renewal / refresh / credits] on [date]. Proposed next step: a [14/30]-day POC on [workload] with these success criteria: [1, 2, 3]. Can we lock [date] for kickoff?

Post-whiteboard

Subject: The board from today Body: Photo attached. The one change we circled: [immutable off-site copy / B2 as origin / one namespace for GPU clouds]. TCO draft with your numbers is in the sheet, every price has a source and date. What would you change?

POC close-out

Subject: POC results: [pass/fail per criterion] Body: Results against the criteria we agreed: [table].
Recommendation: [proceed / adjust]. Decision meeting proposed for [date].

Partner intro

Subject: Joint customers on hyperscaler storage Body: Your customers running [Veeam/iconik/Synology] often land backups or archives on S3 or Azure Hot. B2 is a documented target with Object Lock and 3x free egress. Worth 20 minutes to map overlapping accounts?

26. SE operating plan: first 90 days

A field guide is only useful if it changes what you do on Monday.

Days 1 to 30: learn and prove

- Open a B2 account; run every demo in 17 on your own dime (first 10 GB free).
- Pass internal certification; shadow 10 calls; record your own pitch and fix it.
- Build the TCO calculator from 14 as a shareable sheet with source dates.
- Learn three integrations cold: Veeam, Synology, rclone.

Days 31 to 60: run the motion

- Own discovery and whiteboard on live opportunities.
- Run two POCs with written success criteria.
- Write two deal stories with the template in 22.
- Build a partner list: top 10 Veeam, MAM, and GPU-cloud partners in territory.

Days 61 to 90: scale it

- Turn repeated questions into one-page answers.
- Run a partner enablement session (whiteboard 09 and 12).
- Bring one \$500K+ opportunity to technical win (upside to public guidance).
- Publish an internal “what broke in POCs” list and feed it to product.

Metrics an SE should track

Metric	Why
Technical win rate	Core SE output
POC cycle time	Speed to decision
TB under POC → TB contracted	Conversion

Metric	Why
Partner-sourced pipeline	Leverage
Average deal size	Upmarket motion

28. CLI and API quick reference

Working tools, not diagrams. Verify flags against your installed versions before a live demo.

Endpoints

- S3: `https://s3.<region>.backblazeb2.com` (copy from bucket details page, e.g. `s3.us-west-004.backblazeb2.com`, `s3.eu-central-003.backblazeb2.com` [\[verify code on bucket page\]](#)).
- Native API: authorize with `b2_authorize_account`, which returns the account's API URL and download URL.
- Friendly download URL pattern for public buckets:
`https://<downloadUrl>/file/<bucket>/<key>`.

b2 CLI (v4 style)

```
pip install --upgrade b2
b2 account authorize
b2 bucket list
b2 bucket create my-bucket allPrivate
b2 file upload my-bucket ./local.bin path/remote.bin
b2 ls b2://my-bucket/path/
b2 sync ./dir b2://my-bucket/dir
b2 key create --bucket my-bucket demo-key listFiles,readFiles,writeFiles
```

[\[Commands and capability names vary by CLI version; run `b2 --help` and `b2 key create --help`\]](#)

aws CLI against B2

```
aws configure --profile b2                                # keyID / applicationKey
aws --profile b2 --endpoint-url $EP s3 ls
aws --profile b2 --endpoint-url $EP s3 sync ./dir s3://my-bucket/dir
aws --profile b2 --endpoint-url $EP s3api put-object-lock-configuration \
  --bucket locked-bucket \
```

```
--object-lock-configuration '{"ObjectLockEnabled":"Enabled","Rule":
{"DefaultRetention":{"Mode":"GOVERNANCE","Days":7}}}'
aws --profile b2 --endpoint-url $EP s3api get-object-retention --bucket locked-
bucket --key k
aws --profile b2 --endpoint-url $EP s3api put-object-legal-hold --bucket locked-
bucket --key k --legal-hold Status=ON
```

Note: default bucket retention must be enabled at bucket creation per Backblaze docs.

boto3 config tips

```
from botocore.config import Config
cfg = Config(signature_version="s3v4", retries={"max_attempts": 10, "mode":
"adaptive"},
              max_pool_connections=128)
s3 = boto3.client("s3", endpoint_url=EP, config=cfg, aws_access_key_id=KID,
aws_secret_access_key=KEY)
```

Use `boto3.s3.transfer.TransferConfig(multipart_chunksize=64*1024*1024, max_concurrency=32)` for big files.

rclone

```
rclone config create b2 b2 account=$KID key=$KEY # native backend
rclone copy src b2:bucket/path --transfers 32 --b2-chunk-size 96M --fast-list --
progress
rclone mount b2:bucket /mnt/b2 --vfs-cache-mode full # convenience, not a POSIX
filesystem
rclone config create b2crypt crypt remote=b2:bucket/enc password=$(rclone obscure
'pass')
```

s5cmd (high parallel)

```
s5cmd --endpoint-url $EP --numworkers 256 cp 's3://bucket/prefix/*' ./local/
```

Native API calls worth knowing

Call	Use
b2_authorize_account	Get auth token and URLs
b2_create_bucket / b2_update_bucket	Bucket type, lifecycle rules, Object Lock, encryption, CORS
b2_create_key / b2_delete_key	Scoped keys
b2_update_file_retention / b2_update_file_legal_hold	Object Lock per file
b2_set_bucket_notification_rules	Event Notifications
b2_list_file_versions	Version inventory
b2_hide_file	Hide marker
b2_get_download_authorization	Token for private downloads by prefix
b2_start_large_file / b2_upload_part / b2_finish_large_file	Multipart (large files)

Transaction classes (checked 2026-10-14)

- **A (free):** uploads, deletes, retention/legal hold updates, multipart ops, PutObjectLockConfiguration.
- **B (free):** downloads, GetObject, HeadObject, get retention/legal hold.
- **C (free):** list, copy, bucket ops, key ops, notification rule get/set, CORS, encryption config.
- **D (paid after 2,500/day):** outbound Event Notification calls, \$0.004 per 10,000.

Error decoder

Symptom	Likely cause
InvalidAccessKeyId / SignatureDoesNotMatch	Wrong region endpoint, SigV2 client, wrong key
AccessDenied on delete	Object Lock retention or legal hold (working as designed) or key lacks capability
Object Lock options greyed out in backup tool	Bucket created without Object Lock

Symptom	Likely cause
503 / slow down	Back off and retry; increase parallelism carefully
Tags missing	Tagging unsupported

29. Technical deep Q&A for skeptics

The storage architect in the room will test you. Answer precisely, admit limits, offer a test.

Q: How is a file actually placed? A file goes to one tome in one vault. It is split into 17 data shards plus 3 parity shards (Reed-Solomon), one shard per pod across 20 pods in 20 cabinets. Any 17 rebuild it. Source: durability doc.

Q: What is the storage efficiency? $20/17 \approx 1.18x$ raw-to-logical inside a vault, versus 3x for triple replication.

Q: What does "11 nines" assume? It is a calculated annual design durability for the vault architecture. Ask for the methodology under NDA if they need it. It is not a guarantee of availability.

Q: Is data spread across data centers within a region? US West stores data in Sacramento and Phoenix per regions doc. A single vault is described within cabinets; how vaults map across DCs is not documented publicly. [unverified: whether a single object spans DCs] For site-loss survival requirements, design Cloud Replication.

Q: Consistency model?

[unverified: Backblaze public docs do not state strong read-after-write for S3 API in pages checked]
Test list-after-write in the POC if the app depends on it.

Q: How are uploads handled for big files? Multipart (S3) or large file (Native) with parts uploaded in parallel. Unfinished large files can be auto-cancelled by lifecycle rule `daysFromStartingToCancelingUnfinishedLargeFiles`.

Q: Throughput scaling? Scale with parallel connections and multiple objects. Standard B2 is shared infrastructure; Overdrive is dedicated vaults and networking with 100 Gbps to 1 Tbps tiers. Hume AI reported ~10 GB/s uploads.

Q: Rate limits? Backblaze returns 503 under load; clients should retry with backoff.

[verify published limits]

Q: How do versions and hide markers interact with Object Lock? Hiding a locked file creates a hide marker; the locked version persists until retention expires. You pay for retained versions.

Q: What breaks with Object Lock and replication? Object Lock cannot be enabled on a bucket with a source replication configuration. Destination buckets can apply default retention, computed from the source file's upload timestamp.

Q: Encryption key hierarchy for SSE-B2? Per-file AES-256 key, wrapped by a global key. Rotation details: [\[ask Backblaze\]](#).

Q: Can we bring our own KMS? Not documented. SSE-C (per-request customer key) or client-side encryption.

Q: IPv6? S3-compatible API supports IPv6 (docs).

Q: CORS? Supported (PutBucketCors etc. in Class C list).

Q: Presigned POST for browser forms? Not supported. Use presigned PUT. Presigned URLs must be generated with SigV4 (SigV2 presigning fails).

Q: How does billing measure storage? Byte-hours averaged over the month. Hidden versions and locked versions count.

Q: Egress metering? 3x allowance against byte-hour average stored. Partner egress free. Overage \$0.01/GB.

Q: Multi-tenant isolation? Per-tenant buckets or prefixes plus per-tenant keys (100M key limit), or separate accounts via Partner API/Groups.

Q: Observability? Usage reports, bucket access logs, Event Notifications, MCP analytics tools (storage growth, egress leaders, largest files, unfinished multipart).

Q: Disaster scenario walk-through? Ransomware encrypts primary and local backup. Attacker has backup admin creds; tries deletes on B2; compliance-locked versions refuse. Team stands up clean environment; restores from B2 at full available bandwidth with no retrieval fee; egress inside 3x allowance.

18. Security and compliance FAQ

Narrow language always. Say "supports", "provides", "available on request". Never "certified for HIPAA".

Source for this section unless noted: backblaze.com/cloud-storage/compliance and [/cloud-storage/security](https://backblaze.com/cloud-storage/security), docs on SSE, Object Lock, app keys (checked 2026-10-14).

Q1. What certifications does Backblaze hold? SOC 2 Type 2 (independent third-party audit). Data centers are also SOC 2 compliant. Listed status: GovRAMP Progressing Snapshot, TX-RAMP Provisional, TPN Blue Shield, HECVAT lite and full v3.06 completed, VPAT (Section 508), Internet2 Cloud Scorecard completed. ISO 27001 applies to the data centers Backblaze primarily operates in, not to Backblaze itself; say it that way.

Q2. HIPAA? Backblaze provides a Business Associate Agreement (BAA) on request for business customers who are covered entities. Compliance remains shared responsibility.

Q3. GDPR? Backblaze adheres to GDPR and UK GDPR; Data Processing Addendums for EEA/EU and UK are available. EU Central region (Amsterdam) for residency.

Q4. PCI DSS? Backblaze uses Stripe for card data and states its controls contribute to PCI DSS adherence. That is about Backblaze's billing, not a PCI attestation for storing cardholder data in B2. Customers storing PAN must apply their own controls (encrypt with SSE-C or client-side).

Q5. FedRAMP? Not listed. GovRAMP Progressing and TX-RAMP Provisional are the government statuses.

Q6. How is data encrypted at rest? SSE-B2 (AES-256, unique per-file key wrapped by a global key) is applied by default to new uploads since Sep 14 2026. SSE-C lets the customer supply the key per request. Files uploaded before default encryption remain as they were unless rewritten. **[ask Backblaze whether existing unencrypted objects are being backfilled]**

Q7. In transit? HTTPS/TLS endpoints for S3 and Native APIs. [verify minimum TLS version in docs]

Q8. Can Backblaze employees read my data? With SSE-C, Backblaze does not retain the key. With client-side encryption (e.g., rclone crypt, backup software encryption), Backblaze only stores ciphertext. Recommend client-side or backup-software encryption for maximum separation.

Q9. How does immutability work? Object Lock: governance or compliance retention 1 to 3,000 days, plus legal hold. Compliance mode cannot be shortened or removed by anyone. Enabled per bucket, irreversible on that bucket.

Q10. How do we enforce least privilege? Standard application keys scoped to bucket(s), prefix, capabilities, and expiry. Up to 100M keys per account. Master key only for account-level and Partner API work, and it does not work with the S3-compatible API at all. No IAM roles in the S3-compatible API.

Q11. MFA and SSO? MFA available for accounts. Computer Backup Business supports Google/Microsoft SSO; Enterprise Control supports OIDC (Okta, Entra ID). B2 Neo advertises SSO integration for operators.

[verify B2 console SSO options for enterprise B2 accounts]

Q12. Audit logging? Bucket Access Logs are listed as a security feature. Event Notifications can feed a SIEM via webhook.

Q13. Physical security? Data centers staffed 24/7/365 with "best-in-class security features" per Backblaze. Specific controls are in the SOC 2 report via Whistic.

Q14. How do we get the SOC 2 report and questionnaires? Whistic profiles: Education Industry, EU Customers, All Other Customers.

Q15. Where is my data? Only in the region chosen at account creation (US West, US East, EU Central, CA East) unless you replicate it. Region cannot be changed.

Q16. Durability vs availability? Durability: designed 11 nines via 17+3 erasure coding across 20 pods. Availability: governed by SLA. [pull current SLA document and uptime commitment before quoting]

Q17. Ransomware insurance questionnaires ask for "air-gapped" backups. Is B2 air-gapped? Say "logically air-gapped / virtually air-gapped": off-site, separate credentials, and immutable via Object Lock. Not a physical air gap. Backblaze marketing uses "virtually air-gapped".

Q18. Can we restrict access by IP or private network? Overdrive includes private networking integration.

[verify IP allow-listing options for standard B2]

Q19. Data deletion on exit? Customer deletes objects; locked objects persist until retention expires.

[verify media sanitization statement in SOC 2 / policies]

Q20. TPN for studios? TPN Blue Shield aligned with MPA Content Security Best Practices.

Q21. Education? HECVAT completed; Internet2 Peer Exchange connectivity; multiple university case studies.

Q22. Penetration testing / bug bounty? [unverified; check Trust Center]

33. Customer FAQ quickfire

One-sentence answers for the questions that come up in every call. If the answer needs more, point to the section.

1. **Is there a free tier?** First 10 GB of storage is always free; no credit card required to start.
2. **How am I billed?** Monthly per byte-hour of stored data, or prepaid annual capacity on a 1 to 5 year term (06).
3. **Do you charge for uploads?** No. Uploads and Class A calls are free.
4. **Do you charge for deletes?** No, and there is no minimum storage duration.
5. **Is there a minimum object size charge?** No.
6. **What counts as egress?** Data downloaded from B2; free up to 3x your monthly average stored, then \$0.01/GB.
7. **Which destinations are always free?** Fastly, Cloudflare, bunny.net, CacheFly, CoreWeave, Equinix Metal, Vultr, phoenixNAP.
8. **Is there a cold tier?** No. Everything is hot, priced below most providers' cool tiers.
9. **Can I pick my region?** Yes, at account creation: US West, US East, EU Central, CA East. It cannot be changed later.
10. **Can I store in two regions?** Yes, with a second account and Cloud Replication.
11. **Is my data encrypted?** New uploads are encrypted with SSE-B2 (AES-256) by default since Sep 14 2026; SSE-C is available.
12. **Can I hold my own keys?** Yes, with SSE-C or client-side encryption.
13. **How do I make backups immutable?** Enable Object Lock on the bucket and set retention in your backup software.
14. **Can I undo Object Lock?** Not on the bucket. Governance-mode retention can be changed by privileged keys; compliance mode cannot be shortened.
15. **How long can I lock data?** 1 to 3,000 days, or legal hold indefinitely.
16. **Does Object Lock cost extra?** No, you pay normal storage.
17. **Is it S3 compatible?** Yes, with documented gaps: object ACLs, IAM roles, tagging, website config, browser POST.
18. **Which SDKs work?** Standard AWS SDKs with a custom endpoint, plus Backblaze native SDKs and community SDKs.
19. **Does rclone work?** Yes, with the native b2 backend or the S3 backend.
20. **Does Veeam work?** Yes, direct-to-object (v12+) or SOBR capacity tier, with immutability.
21. **Does Rubrik work?** Yes, after Rubrik support enables SigV4.
22. **Can I trigger functions on upload?** Yes, Event Notifications (webhooks), enabled on request.
23. **Can I serve a website from a bucket?** Not with S3 website config; put a CDN in front.
24. **How durable is it?** Designed for 11 nines annual durability with 17+3 erasure coding across 20 pods.
25. **How fast is it?** Depends on parallelism and network; Overdrive offers dedicated 100 Gbps to 1 Tbps.
26. **How do I move 500 TB in?** Universal Data Migration (with commitment), rclone at high parallelism, or Fireball (96 TB per unit).

27. **Will you help pay for migration?** Universal Data Migration can offset egress and transfer costs with a qualifying commitment.
28. **How is migration verified?** MD5 byte-for-byte verification per Backblaze.
29. **Do you sign BAAs?** Yes, for covered entities on request.
30. **Do you have SOC 2?** SOC 2 Type 2; documents via Whistic.
31. **Are you FedRAMP?** No; GovRAMP Progressing Snapshot and TX-RAMP Provisional.
32. **Are you TPN?** TPN Blue Shield status.
33. **Do you support GDPR?** Yes; EU region and DPAs.
34. **Can agents use B2?** Yes, through the open source B2 MCP Server, scoped by app key.
35. **Is there a startup program?** Flamethrower: up to \$100,000 in storage credits.
36. **What is Overdrive?** Dedicated high-throughput B2 with unlimited free egress, from \$15/TB, multi-PB commit.
37. **What is B2 Neo?** White-label B2 for neoclouds with Partner API, custom domains, SSO.
38. **What does Computer Backup cost?** \$99 per computer per year, unlimited data.
39. **Can I get a hard drive of my data?** Computer Backup: yes, encrypted USB shipped free if returned within 30 days.
40. **Who are your big customers?** Public stories include CoreWeave, Decart, Hume AI, Mirage, Philadelphia Eagles, Austin City Limits, Jazz at Lincoln Center, and 100,000+ B2 customers.
41. **Is Backblaze profitable?** Adjusted EBITDA \$12.8M (30 percent margin) in Q2 2026; GAAP net loss \$5.1M.
42. **What if I want to leave?** Your data is S3-accessible and 3x monthly free egress applies to exit as well; compliance-locked data stays until expiry.
43. **Can an AI agent delete my buckets through the B2 MCP Server?** Only if its application key has `deleteBuckets` (or `deleteFiles` for files); give it a read-only or bucket-scoped key and it cannot (07).
44. **Can this guide's MCP server touch my B2 account?** No. The field guide's MCP server at backblaze.nixfred.com/mcp is read-only and only searches this guide; it has no B2 credentials (07).
45. **Does Wasabi charge for egress?** No. Wasabi charges no egress fees, but its fair-use policy expects monthly egress to stay at or below stored volume; if it regularly exceeds that, Wasabi may flag the account or move it to a different plan (15).
46. **Governance or compliance mode for a regulated customer (hospital, bank, public sector)?** It depends on who must be unable to delete. If a regulator or cyber-insurer requires that nobody, including admins, can shorten or remove retention, use compliance mode for the backup copy. If the team must be able to purge data early with a privileged key, use governance mode. Either way, test with governance and a 1-day retention first, because compliance mode cannot be undone and is billed until it expires (09, 17, 18).

21. Glossary

3-2-1-1-0. Three copies, two media types, one off-site, one immutable or air-gapped, zero errors on restore verification. **3x free egress.** B2 allowance: monthly egress up to 3x average stored data (byte-hour average) is free; beyond is \$0.01/GB. **AFR.** Annualized failure rate of hard drives. Backblaze publishes it quarterly in Drive Stats (Q2 2026: 1.73 percent). **App key (application key).** B2 credential: keyID + applicationKey. Master or standard. Standard keys scope bucket(s), prefix, capabilities, expiry. **Bandwidth Alliance.** Cloudflare-led group of providers that reduce or waive transfer fees between members. Backblaze is a founding member. **B2 Cloud Storage.** Backblaze's S3-compatible object storage service. **B2 Neo.** White-label B2 for neocloud operators, with Partner API, custom domains, SSO. **B2 Overdrive.** Dedicated high-throughput B2 tier: 100 Gbps to 1 Tbps, unlimited free egress, from \$15/TB, multi-PB commit. **B2 Reserve.** Legacy/channel prepaid capacity SKU name. Confirm current status. **Byte-hour.** Billing unit: storage measured continuously and averaged over the month. **Class A/B/C/D transactions.** B2 API call classes. A, B, C free; D (outbound event notifications) charged after 2,500/day. **Cloud Replication.** B2 feature copying objects between buckets, within or across regions (cross-region needs a second account). **Compliance mode.** Object Lock mode where retention cannot be shortened or removed by anyone. **CoreWeave AI Object Storage / LOTA.** CoreWeave's object storage with a distributed cache (LOTA); Backblaze powers HDD tiers beneath it. **Drive Stats.** Backblaze's public quarterly hard drive reliability reports. **Egress.** Data transferred out of a storage provider. **Erasure coding (Reed-Solomon).** Splitting data into data and parity shards; B2 uses 17+3 across 20 pods. **Event Notifications.** B2 webhooks on object create/update/delete, optional HMAC signing. **Fireball.** Backblaze's 96 TB Synology NAS shipped to customers for offline data ingest. **Governance mode.** Object Lock mode where privileged keys can modify or remove retention. **Groups / Organizations.** Multi-account administration and billing. **Hide marker.** B2 version marker that hides a file (S3 delete marker equivalent). **Hot tier.** Fast storage (flash, parallel FS) for the active working set. **Legal hold.** Object Lock protection with no expiry until removed. **Lifecycle rules.** B2 rules to hide and delete versions and cancel unfinished large files by prefix; up to 100 per bucket. Set via the native B2 API or web UI only, not the S3-compatible API. **Live Read.** Read/edit large media files while still uploading. **MCP (Model Context Protocol).** Standard for AI agents to discover and call tools. Backblaze B2 MCP Server exposes 40 tools gated by key capabilities. **Multipart upload / large file.** Uploading an object in parts; B2 native calls it "large file". **Native API.** B2's own API (b2_* calls), alongside the S3-compatible API. **Neocloud.** GPU-as-a-Service provider outside the hyperscalers. **Object Lock.** WORM immutability for objects; 1 to 3,000 days; formerly "File Lock". **Partner API.** B2 API to create and manage accounts, usage, billing; used by MSPs, B2 Neo, Powered by Backblaze. **Partner egress.** Unlimited free egress from B2 to listed CDN and compute partners. **Powered by Backblaze.** Program to embed B2 inside a partner's product. **Presigned URL.** Time-limited URL granting GET or PUT on one object without sharing keys. On B2 they must be signed with SigV4. **Read multiple.** Monthly egress divided by stored data. Key TCO input. **Region.** B2 data location chosen at account creation: US West, US East, EU Central, CA East. **SigV4.** AWS Signature Version 4, required by B2's S3 API. **SOBR.** Veeam Scale-Out Backup Repository; B2 can be the capacity tier. **SSE-B2 / SSE-C.** Server-side encryption with Backblaze-managed or customer-provided keys, AES-256. **Storage pod.** Backblaze's storage server design; 20 pods form a vault. **Tome.** The set of drives in the same slot across a vault's 20 pods; a file lives on one tome. **TPN Blue Shield.** Trusted Partner Network status aligned with MPA content security practices. **UDM (Universal Data**

Migration). Backblaze-run migration service with cost relief for qualifying commitments. **Vault (Backblaze Vault).** 20 storage pods across 20 cabinets storing erasure-coded data. **WORM.** Write once, read many.

23. Sources

Official first. Narrow language always. All checked 2026-10-14 (ET) unless noted.
Raw captures in `/workspace/backblaze/site/research/` .

Backblaze primary

- Pricing: <https://www.backblaze.com/cloud-storage/pricing>
- Transaction pricing: <https://www.backblaze.com/cloud-storage/transaction-pricing> (note: still shows \$0.005/GB storage)
- B2 overview: <https://www.backblaze.com/cloud-storage>
- B2 Overdrive: <https://www.backblaze.com/cloud-storage/b2-overdrive>
- B2 Neo: <https://www.backblaze.com/cloud-storage/b2-neo>
- Neocloud white paper form: <https://www.backblaze.com/contact-sales/why-neoclouds-should-not-build-their-own-storage>
- AI/ML industry: <https://www.backblaze.com/cloud-storage/industries/ai-ml>
- AI integrations: <https://www.backblaze.com/cloud-storage/b2-ai-integrations>
- Security: <https://www.backblaze.com/cloud-storage/security>
- Compliance: <https://www.backblaze.com/cloud-storage/compliance>
- Universal Data Migration: <https://www.backblaze.com/cloud-storage/solutions/cloud-migration>
- Object Lock solution: <https://www.backblaze.com/cloud-storage/solutions/object-lock>
- Ransomware: <https://www.backblaze.com/cloud-storage/solutions/ransomware-readiness-and-recovery>
- Media: <https://www.backblaze.com/cloud-storage/solutions/media-workflow>
- CDN: <https://www.backblaze.com/cloud-storage/solutions/cdn>
- LTO replacement: <https://www.backblaze.com/cloud-storage/solutions/lto-replacement>
- Integrations: <https://www.backblaze.com/cloud-storage/integrations>
- Veeam: <https://www.backblaze.com/cloud-storage/integrations/veeam>
- Synology: <https://www.backblaze.com/cloud-storage/integrations/synology>
- iconik: <https://www.backblaze.com/cloud-storage/integrations/iconik>
- Partners: <https://www.backblaze.com/partners> ; MSP: <https://www.backblaze.com/partners/roles/msp>
- Computer Backup pricing: <https://www.backblaze.com/cloud-backup/pricing> ; Business: <https://www.backblaze.com/cloud-backup/business>
- Case studies index: <https://www.backblaze.com/case-studies> ; Decart, Hume AI, Captions (Mirage), Philadelphia Eagles, Austin City Limits, Bell Bank under <https://www.backblaze.com/cloud-storage/case-studies/>

Backblaze docs

- Data regions: <https://www.backblaze.com/docs/cloud-storage-data-regions> (updated 27 Feb 2025)
- Durability: <https://www.backblaze.com/docs/cloud-storage-resiliency-durability-and-availability> (updated 21 Aug 2023)
- S3-compatible API: <https://www.backblaze.com/docs/cloud-storage-s3-compatible-api>
- Object Lock: <https://www.backblaze.com/docs/cloud-storage-object-lock>
- Lifecycle rules: <https://www.backblaze.com/docs/cloud-storage-lifecycle-rules>
- Event Notifications: <https://www.backblaze.com/docs/cloud-storage-event-notifications>
- Cloud Replication: <https://www.backblaze.com/docs/cloud-storage-cloud-replication>
- Application keys: <https://www.backblaze.com/docs/cloud-storage-application-keys>
- SSE: <https://www.backblaze.com/docs/cloud-storage-server-side-encryption>
- Fireball: <https://www.backblaze.com/docs/cloud-storage-backblaze-fireball-program>
- MCP server: <https://www.backblaze.com/docs/cloud-storage-connect-ai-clients-to-cloud-storage-operations-using-b2-mcp-server>
- Cloudflare: <https://www.backblaze.com/docs/cloud-storage-cloudflare-integrations>
- Fastly: <https://www.backblaze.com/docs/cloud-storage-integrate-fastly-cdn-with-backblaze-b2>
- Veeam direct-to-object: <https://www.backblaze.com/docs/cloud-storage-integrate-veeam-direct-to-cloud-backups-to-backblaze-b2-with-immutability>
- Rubrik: <https://www.backblaze.com/docs/cloud-storage-configure-rubrik-with-backblaze-b2>
- Commvault: <https://www.backblaze.com/docs/cloud-storage-integrate-commvault-with-backblaze-b2>
- TrueNAS: <https://www.backblaze.com/docs/cloud-storage-integrate-truenas-with-backblaze-b2>

Backblaze blog and IR

- MCP server launch: <https://www.backblaze.com/blog/introducing-the-backblaze-b2-mcp-server/>
- CoreWeave: <https://www.backblaze.com/blog/coreweave-announcement-2026/>
- Default encryption: <https://www.backblaze.com/blog/backblaze-b2-to-encrypt-new-uploads-by-default/>
- Q2 2026 results: <https://ir.backblaze.com/news/news-details/2026/Backblaze-Announces-Second-Quarter-2026-Financial-Results/default.aspx>
- CRO appointment: <https://ir.backblaze.com/news/news-details/2026/Backblaze-Appoints-Anuj-Kumar-as-Chief-Revenue-Officer/default.aspx>
- Q2 2026 Drive Stats: <https://ir.backblaze.com/news/news-details/2026/Backblaze-Publishes-Q2-2026-Drive-Stats-Quarterly-Failure-Rate-Rises-to-1-73/default.aspx>
- CoreWeave 8-K: <https://www.sec.gov/Archives/edgar/data/1462056/000162828026044804/blze-20260616.htm>

Competitors

- Wasabi: <https://wasabi.com/pricing>
- Cloudflare R2: <https://developers.cloudflare.com/r2/pricing/> (page last updated Oct 1 2026)
- AWS S3: <https://aws.amazon.com/s3/pricing/> (JS-rendered; cross-checked with <https://upstash.com/blog/aws-s3-pricing-explained-2026> and <https://cloudburn.io/blog/amazon-s3-pricing>)
- Azure Blob: <https://azure.microsoft.com/en-us/pricing/details/storage/blobs/> and <https://learn.microsoft.com/en-us/azure/storage/blobs/blob-storage-estimate-costs> (sample prices)
- GCS: <https://cloud.google.com/storage/pricing>

Secondary (context only)

- Blocks & Files on CoreWeave deal: <https://www.blocksandfiles.com/public-cloud/2026/06/23/backblaze-gets-335-million-coreweave-bonanza/5260288>
- Q2 2026 call transcript: <https://www.fool.com/earnings/call-transcripts/2026/08/10/backblaze-blze-q2-2026-earnings-call-transcript/>

Open verification list

SLA terms; max object size; Azure/GCS egress list; B2 console SSO; IP allow-listing for standard B2; current B2 Reserve status; app key capability names; event type names; Q3 earnings date; 17a-4 assessment.