ANALYSIS · INGRESS ECONOMICS · CO-BRANDING · TERANODE

The cost of being reachable: transaction-intake economics for Teranode miners

A miner's transaction flow lands wherever submitters can reach with the lowest latency, so global intake is not optional. This paper prices the ways to get it: building your own anycast footprint, steering with DNS, renting reach from an edge provider by the gigabyte or by the megabit (Cloudflare as the worked example; the shapes generalise), and the 1BSV model, where **admission is free and delivery is the paid product**. The comparison is honest in both directions: rented reach is priced from real market bands, and the 1BSV side states plainly that scale still costs money and where that money goes. Companion analysis: *The Asset service at scale*, which prices the download half of the same problem.

Author: Jeff Harris · Lightweb Inc. · jeff@lightweb.net

Audience: Teranode operators and engineers. **We** = the 1BSV fabric; **you** = a miner. All vendor pricing is a modelled band, not a quote; see §8.

The short version

- **Being reachable is competitive, which makes it mandatory.** Submit latency decides where transaction flow lands. Every miner needs intake near every sender, and every rung of the market prices that differently.
- **Renting reach puts a meter on the raw stream.** Per-gigabyte edge proxying runs roughly $84 k to $1.7 M per month at 1 M TPS for a single copy of the transaction stream, and climbs into nine figures at 1B TPS. Per-megabit prefix protection is an order cheaper but leaves the footprint problem with you.
- **Rent both halves and the same vendor meters you coming and going.** Per-gigabyte intake plus the CDN-fronted download bill from the companion paper combine to roughly $0.2 M to $3.2 M per month at 1 M TPS, with one third party terminating TLS on both sides of the consensus path.
- **Rented reach terminates globally; it does not ingest globally.** A proxy hands you your own firehose at one origin. The shard mapping, the regional capacity, and the abuse control remain your problem.
- **1BSV admission is free because its amplification is bounded.** An admitted transaction is protocol work plus one emission into the fabric. There is no per-copy cost to meter at the front door, so we do not meter it.
- **Delivery is the paid product, and that is the honest place for the meter.** Consumers pay for the stream they elect to receive, once, over one tunnel. Nobody pays per gigabyte to be reachable.
- **Identity survives every tier.** From a shared hostname to your own domain to your own address block announced on your behalf, attribution follows the brand at admission while delivery stays uniform on the fabric.
- **Scale still costs.** At 1B TPS someone must terminate 2.6 Tbit/s of submissions and carry them. Those costs are real; the structural win is that they are paid for one shared footprint and one emitted copy, not duplicated across every miner.

1. The competitive geometry of intake

A wallet submits where it can reach. Round-trip time to the submission endpoint decides which miner sees a transaction first, and at the margin, which miner earns it. Being close to senders wherever they are is therefore not an optimisation a miner may defer; it is competitive, which makes it mandatory.

That geometry has a price. Whoever wants the flow must terminate submissions in every region the senders occupy, absorb abuse on an endpoint that must stay open to strangers, and do it continuously. The rest of this paper is about who pays that price, under which meter, and what each option leaves on the miner's plate.

2. What the stream weighs

The planning basis from the companion paper: a 90 %/10 % mix of 250 B and 1 000 B transactions gives a mean of `S = 325 B`. Intake is **one copy** of the stream; there is no fan-out on the upload half.

Aggregate TPS	Sustained rate	Volume per month
10 k	26 Mbit/s	8.4 TB
100 k	260 Mbit/s	84 TB
1 M	2.6 Gbit/s	842 TB
10 M	26 Gbit/s	8.4 PB
100 M	260 Gbit/s	84 PB
1B	2.6 Tbit/s	842 PB

Two readings of this table matter. First, intake is far lighter than the download half: one copy, not `N`. Second, it is still a firehose at the top of the ladder, and any pricing model that meters it per byte will scale with it.

3. The rented options, priced honestly

Transaction submission is raw TCP/UDP toward a propagation ingest, not cacheable HTTP, so the classic CDN tier does not apply. What edge providers actually sell for this shape of traffic falls into three categories. We price them through Cloudflare's catalogue because it is the most familiar and the best documented, not because it is unusual: every major edge vendor sells the same three shapes under different names, and the analysis turns on the meters, not the logos. Every band below is modelled from list pricing and public reporting, because enterprise pricing is negotiated and confidential.

Free tunnels (Cloudflare Tunnel class). HTTP-oriented, fair-use bounded, and terminating at your origin. Free does not survive contact with sustained multi-gigabit financial ingest; the upgrade path leads to the two products below.

Per-gigabyte anycast proxying (Spectrum class). The genuine fit for raw TCP/UDP: global anycast termination, proxied to your origin. Metered on bytes through the edge, both legs, so roughly twice the payload. Modelled band $0.05 to $1.00 per GB:

Aggregate TPS	Billed volume (×2 legs)	Cost band per month
10 k	17 TB	$0.8 k – $17 k
1 M	1.7 PB	**$84 k – $1.7 M**
10 M	17 PB	$0.8 M – $17 M
1B	1.7 EB	**$84 M – $1.7 B**

For scale: at 1 M TPS this band overlaps the entire CDN download bill priced in the companion paper, for one copy of the stream, before a single byte is fanned out to peers. That is what renting a per-gigabyte meter on the generation rate looks like.

Per-megabit prefix protection (Magic Transit class). BGP anycast of a prefix you own, delivered to tunnel endpoints you operate. Priced on clean bandwidth at the 95th percentile, with contract minimums that make small deployments moot. Modelled band $4 to $10 per Mbps per month:

Aggregate TPS	95th-percentile demand	Cost band per month
10 k	26 Mbps	below minimums
1 M	2 600 Mbps	**$10 k – $26 k**
10 M	26 000 Mbps	$104 k – $260 k
1B	2 600 000 Mbps	**$10 M – $26 M**

This is the honest budget option among the rented tiers, and it concedes the architectural point: it anycasts an address block, which is the right mechanism, but it hands raw packets to endpoints you still provision in every region where you want low-latency intake. You have bought routing and DDoS absorption. The footprint is still yours.

The double meter: paying the same vendor coming and going. The companion paper priced the download half: CDN-fronting the Asset service still bills the network for `N` deliveries of every object. A miner who also rents intake has now placed a per-byte meter on **both halves of the same bytes**: a transaction is billed inbound at submission, and its subtree and block forms are billed outbound `N` times at propagation. Network-wide, at the modelled bands:

Aggregate TPS	Intake meter (per-GB proxy)	Delivery meter (CDN-fronted pull[1])	Combined
1 M	$84 k – $1.7 M /mo	$88 k – $1.49 M /mo	**$0.2 M – $3.2 M /mo**
1B	$84 M – $1.7 B /mo	$88 M – $1.49 B /mo	**$0.2 B – $3.2 B /mo**

1. Full-object-stream band at `N = 19`, from the companion paper.

The money is only half of the whammy. The same third party now terminates TLS on the submission path and on the propagation path, observes both sides of the miner graph, and is a correlated failure domain for intake and delivery at once. Two meters, one landlord, both halves of consensus.

What every rented option leaves with you. Three things, regardless of price:

- **Termination is not ingestion.** The provider's edge is global; your origin is not. A submitter far from your ingest gets a fast local handshake, then the transaction rides a tunnel across the world to the one region your propagation cluster lives in. Fixing that means regional origins, which is the self-built footprint again with a provider invoice on top.
- **The edge cannot do admission.** Real ingress computes the shard mapping (`groupIdx = f(TxID)`), deduplicates, attributes, and injects into the distribution fabric. A proxy hauls bytes. The ingest brain, and the capacity behind it, remain yours at every rung.
- **Attribution degrades.** Proxied flows arrive with rewritten source addresses, exactly where per-submitter attribution and abuse control need to live.

Self-built, for reference. The ceiling option from the companion paper: your own sites, hardware, transit, routing and DDoS expertise in every region, at competitive quality tens of thousands of dollars per month and up, per miner, with every site sized for its peak. Highest control, highest cost, and the entire burden duplicated by every miner that makes the same choice. DNS-based geographic steering is the floor option: cheap, but slower failover, coarser proximity, and resolver-cache exposure; the flow you lose to those gaps lands elsewhere.

4. The 1BSV model: free admission, paid delivery

1BSV splits the economics at the only clean joint: **submission is free; delivery is the paid product.**

4.1 Why admission can be free

Free is not a promotion; it is a property of the architecture.

- **Bounded amplification at the front door.** The open ingress classes are election-scoped: a submitted transaction targets exactly one shard group. There is no broadcast-amplified class on the open path, so an admission cannot be made expensive by construction. The amplified classes (subtree and block push) are not public; they ride provisioned consumer tunnels.
- **Emit once.** An admitted transaction is protocol work (stamp, shard, deduplicate, attribute) plus one emission into the multicast fabric. The fabric replicates at branch points; the admission itself never pays a per-copy cost. With no per-copy cost at the front door, there is nothing honest to meter there, so we do not meter it.
- **One shared footprint.** One anycast ingress fleet serves every miner and every co-brand. The fixed cost of being globally reachable is paid once and amortised across the ecosystem, instead of once per miner.

4.2 The tiers: identity at every level

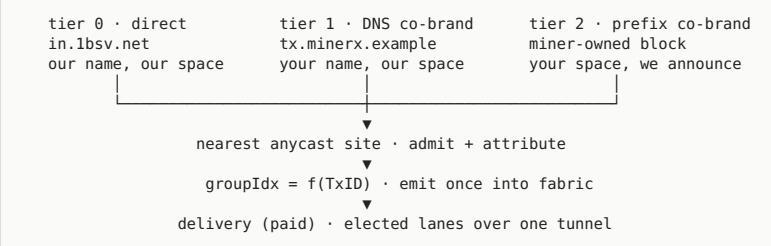

Figure 1 - three intake tiers, one admission path. Attribution happens at the front door; everything behind it is identical.

- **Tier 0, direct.** Submit to the shared ingress hostname. No branding, no setup, free.
- **Tier 1, DNS co-brand.** Your domain points at the 1BSV anycast ingress. Submitters resolve and connect to **your hostname**; admission attributes every transaction to your brand for metrics and billing. No routing infrastructure, no capital outlay.
- **Tier 2, prefix co-brand.** We announce **your address block** as anycast from our ingress edges, with cryptographically authorised routing (RPKI). Packets are addressed into your space, from locations you do not operate. Identity at the routing layer, not merely in DNS.

The design property that keeps every tier cheap: a co-brand is **never a source on the fabric and never rides in the frame**. Attribution is a per-domain record at admission; delivery to every subscriber stays uniform. The hundredth co-brand costs what the first did: a DNS record or a prefix announcement, and a label on a counter.

4.3 Delivery: where the meter honestly belongs

The paid product is delivery: a consumer (a miner's cluster, a service) subscribes to the lanes it wants (transactions by shard, subtrees, blocks) and receives its **elected stream, once, over one provisioned tunnel**, with sequence numbering, gap detection, and bounded repair included.

That placement of the meter has three properties the rented options cannot offer:

- **You pay for what you elect, not for the network's size.** A delivery bill scales with the consumer's chosen lanes and the network's TPS, never with the number of other participants. There is no `N` in it.
- **You pay once per copy.** The fabric's emit-once economics mean your stream is not re-billed per receiver, per PoP, or per leg.
- **Your footprint is one tunnel.** The global intake problem, the DDoS surface, and the anycast operations are on our side of the demarcation. Your side is the cluster you already run.

5. Side by side

Option	Pricing model	@ 1 M TPS (modelled)	@ 1B TPS (modelled)	You still operate	Identity
Self-built anycast	capex + transit + ops per region	tens of $k/mo and up	$M+/mo class	everything, everywhere, sized for peak	yours, at full cost
DNS geo-steering	cheap, quality-limited	$10s–100s/mo	not viable	all origins; failover and proximity gaps	yours
Per-GB proxy (Spectrum class)	$0.05–$1.00/GB, both legs	$84 k – $1.7 M/mo	$84 M – $1.7 B/mo	origins, ingest brain, regional capacity	yours behind provider TLS
Per-Mbps protection (Magic Transit class)	$4–10/Mbps 95th pct + minimums	$10 k – $26 k/mo	$10 M – $26 M/mo	tunnel endpoints, footprint, all capacity	your prefix, your footprint
1BSV intake (any tier)	**free admission**	**$0**	**$0**	your cluster, one delivery tunnel	your hostname and/or prefix
1BSV delivery	**metered on your elected stream**	scales with your election, flat in `N`	same, flat in `N`	reception and validation	uniform frames, attributed intake

6. What scale still costs

Free admission does not mean scale is free, and it would be dishonest to imply it. At 1B TPS, someone terminates 2.6 Tbit/s of submissions across a global fleet, carries 842 PB per month of admitted stream into the fabric, holds retransmission headroom above it, and absorbs the abuse that an open endpoint attracts. As a reference point, one carried copy of the 1B TPS stream is roughly $0.3 M to $1.3 M per month in blended wholesale transit alone (modelled at $0.10 to $0.50 per Mbps), before hardware, sites, and operations.

Those costs are real, and they are funded by the delivery side of the ledger. The claims we make are narrower than "free," and structural:

- The cost of the stream is paid on **one emitted copy and each consumer's elected delivery**, not on `N` redundant hauls of the same bytes.
- The cost of reachability is paid on **one shared anycast footprint**, not once per miner per metro.
- The meter sits on the party receiving value (the consumer of the stream), not on the open door the ecosystem needs to stay open.

Every alternative in §3 fails at least one of those three. Per-gigabyte proxying meters the open door. Per-megabit protection shares nothing. Self- building duplicates the footprint `N` times. The 1BSV split is not cheaper by magic; it is cheaper by not paying for the same bytes and the same buildings more than once.

7. Choosing a tier

- Testing, tooling, or low volume: **tier 0**, submit to the shared ingress, free, today.
- A brand in front of your flow with zero infrastructure: **tier 1**, one DNS record, attribution and per-domain metrics included.
- Identity at the routing layer, your own address block visible to the world: **tier 2**, we announce your prefix with RPKI authorisation on your behalf.
- Receiving the network's stream into your cluster: a **delivery subscription**, priced on the lanes you elect, delivered once over one tunnel.

Tiers are not lock-in rungs; a tier-1 co-brand becomes tier-2 by allocating a prefix, and the admission path behind every tier is byte-identical.

8. Assumptions and limits

Quantity	Value	Basis
Transaction size	325 B weighted mean	**stated**: 90 % × 250 B + 10 % × 1 000 B
Months	30 days	**stated**
Per-GB proxy band	$0.05 – $1.00 / GB, ×2 legs	**modelled**: list and reported enterprise bands
Per-Mbps protection band	$4 – $10 / Mbps / mo, 95th pct	**modelled**: reported enterprise bands, minimums apply
Wholesale transit band	$0.10 – $0.50 / Mbps / mo	**modelled**: blended committed transit at scale
Self-built intake cost	tens of $k/mo and up	**modelled**: order of magnitude at competitive quality
95th percentile ≈ sustained	firehose traffic shape	**stated**: continuous stream, low burst-to-mean ratio

Limits, stated plainly:

1. **Every vendor number is a modelled band, and a category band.** Enterprise pricing is negotiated and confidential, so the bands are wide on purpose, and they describe the pricing shape, not one vendor: substitute any major edge provider and the same three meters reappear under different names. The argument does not depend on the constants; per-byte meters scale with the stream at any realistic discount.
2. **1BSV delivery pricing is not published in this paper.** The claim made is structural: the meter is on elected delivery, once per copy, flat in the number of participants. Commercial terms are a conversation, not a table.
3. **Free admission is sustainable because of what is open.** Only bounded-amplification classes ride the open path. If that boundary moved, the economics would move with it; the design intent is that it does not.

References

The Asset service at scale: why fronting DataHubURL with a CDN extends the gossip problem instead of solving it (Lightweb Inc.): https://1bsv.net/papers/asset-pull-and-the-cdn-reflex.pdf

Pushed delivery into a Teranode cluster over unicast TCP (Lightweb Inc.): https://1bsv.net/papers/pushed-delivery-teranode.pdf

---

1BSV - BSV Layered Multicast - cash moves fast on 1BSV. · 1bsv.net

---

Lightweb Inc. · Jeff Harris · jeff@lightweb.net · github.com/lightwebinc