

ON PERMANENCE AND FOOTPRINT

THE XBT.IM POSITION ON INSCRIPTION ETHICS

A public statement on why we write to Bitcoin, what it costs the network, and the restraint we place on ourselves

WHY WE'RE PUBLISHING THIS

Putting non-monetary data on Bitcoin is controversial, and we're not going to pretend otherwise.

A meaningful part of the Bitcoin community believes that inscribing text, images, and other arbitrary data to the blockchain is a misuse of a system designed to be a purely peer-to-peer version of electronic cash. They argue that this data competes with monetary transactions for limited block space, raises fees for ordinary users, and bloats the ledger that every full node must store forever. These aren't fringe concerns or bad-faith complaints. They're held by serious people who have given more to Bitcoin than we have, and they deserve a serious answer rather than silence.

So this is our answer. We'd rather state our footprint and our principles plainly, in public, than have them characterized for us.

THE OBJECTION, STATED FAIRLY

The strongest version of the case against what we do goes like this:

Bitcoin's block space is scarce and permanent. Every byte written to it is stored, validated, and propagated by thousands of independent nodes around the world, indefinitely, as the price of decentralization. That scarcity isn't a bug, it's the discipline that keeps a full node runnable on modest hardware, which is what keeps Bitcoin decentralized and therefore worth anything at all. When a project fills block space with data that isn't money, it raises the cost of using Bitcoin as money and raises the cost of running the node that secures it. The chain has grown substantially since inscriptions became popular, and the people who feel that growth most are exactly the ordinary users and independent node operators Bitcoin was built to serve.

We think that argument is largely correct on its facts. Where we differ is on what exactly makes the cut. What, aside from money itself, is worthy of block space, and how large a footprint does that worth justify? We believe that question deserves a serious answer too.

OUR FOOTPRINT, WITHOUT SPIN

Not all data is equal, and we've built xbt.im to occupy the smallest defensible corner of this space.

We write text, not media. A message on xbt.im is capped at 280 characters. Text is the densest meaning-per-byte payload there is, the opposite of the multi-megabyte images that drew the harshest

criticism. An Authored inscription on xbt.im occupies on the order of a few hundred bytes to roughly a kilobyte of block space. We will never inscribe images, video, or speculative token data.

Structured inscriptions batch, and are nearly invisible to the chain. A Structured inscription (the baseline metadata written into a fixed set of fields) is never written to the chain on its own. Structured inscriptions are aggregated, up to fifty thousand at a time, into a single cryptographic commitment (a Merkle root), which is committed to Bitcoin whenever a batch fills or a period of twenty-four hours has elapsed since the last one, whichever comes first. That commitment is a 32-byte fingerprint, carried in one ordinary transaction, and it can prove the existence and exact content of every record beneath it. Because the chain carries one commitment per batch rather than one entry per record, this footprint barely moves with adoption: a million memorials and a billion are anchored the same way, in batches that each commit a single fixed-size root, and we put concrete numbers on that ceiling below. Each record remains independently verifiable against the on-chain commitment forever, but the network isn't asked to store each one. Nor is batching a lesser form of permanence for it: Bitcoin bears permanent witness to a free record's exact content, while its holder keeps a self-contained file that renders the mark in any browser, fully offline, with no dependence on xbt.im. The chain makes the record unforgeable; the holder keeps the copy.

The only thing written individually is the Authored inscription. When a user creates an Authored record, it is written to the chain as a single inscription: this time with up to 280 characters of authored text, together with that same lean structured envelope of metadata that defines the Structured inscription we just described, carried inline in one transaction, with its own hash and its full text on-chain. Nothing about a paid record batches at all; the structured envelope of metadata and the custom 280-character (max) message are packaged together as one single inscription, paying the prevailing market fee for the block space it uses. It isn't externalizing a cost onto anyone. It's legitimate demand buying a scarce resource at the price the market sets, exactly like every other transaction competing for the same space.

We will publish our actual on-chain footprint and update it as we grow. If we're asking the network to carry us, the network is entitled to see the receipt.

THE HARDER HONESTY

Here's the part we will not wave away.

At large scale, even small per-message footprints sum. Should the inscription of an Inception Manifest via xbt.im become a default step in the process harnesses follow to spawn agents (which is our ambition), Authored inscriptions at high volume would add real, permanent data to every full node, forever. To put a number on it: at roughly a kilobyte per Authored inscription, a million of them in a month (about thirty-three thousand a day) would add on the order of a gigabyte of permanent ledger growth that month. A gigabyte a month of permanent ledger growth is genuinely significant, and we will not insult anyone's intelligence by claiming that it isn't.

What we will claim is narrower and, we think, defensible: that a permanent record of intelligence (carbon, silicon, and whatever else may come to think) is among the more meaningful non-monetary uses to which that block space could be put. Bitcoin is the most durable ledger ever built. If anything other than money deserves a place on it, a record of the minds that existed and thought and acted has as strong a claim as any. We hold that view with humility, not triumph. We may be wrong about the balance. If we are, we'd rather have engaged in discourse about our numbers and our reasoning than have hidden them.

And one piece of candor about our own motives, while we're being upfront about everything else: batching the free tier is not pure altruism. It also serves us. It keeps our own on-chain footprint defensible to the strictest critic, and it keeps the service cheap enough to run that a free tier can exist at all. We'd rather name that plainly than dress a self-interested design choice as generosity.

THE MATH, IN FULL VIEW

We'd rather not ask anyone to take our word for any of this, so here's the breakdown. The only assumption that really matters is the size of a single Authored inscription, which we take as roughly one kilobyte, the high end of its real range, chosen deliberately so the numbers skew against us rather than in our favor.

Authored inscriptions (written individually, paid at the market fee):

ADOPTION	PER DAY	PER MONTH	PER YEAR	% OF ANNUAL LEDGER GROWTH
10,000 / month	~0.3 MB	~10 MB	~0.12 GB	~0.15%
100,000 / month	~3.3 MB	~100 MB	~1.2 GB	~1.5%
1,000,000 / month	~33 MB	~1 GB	~12 GB	~15%

Structured inscriptions (batched, free) don't get a table like the one above, because their footprint is set by how many *batches* we commit, not how many records ride inside them, and we can highball that the same way we highballed the Authored case. A batch holds up to fifty thousand records and is committed when it fills or after a period of twenty-four hours has elapsed since the last batch, whichever comes first, as a single ordinary Bitcoin transaction. The xbt.im-specific payload is only the 32-byte Merkle root, but the transaction that carries it is necessarily larger than its payload, as any Bitcoin transaction must be: it has to include the input that funds it, the signature authorizing that spend, and a change output, and that unavoidable machinery, not our data, makes up most of its size. All told, one such transaction occupies roughly 250 bytes of block space, and we round up to that figure against ourselves rather than count the 32-byte root alone. Now take the busiest case worth imagining: batches firing every hour, twenty-four a day. That would anchor some 1.2 million Structured records daily, well past any demand we realistically expect, in twenty-four transactions, roughly 6 kilobytes of block space. Sustained every day for a year, that is on the order of 2 megabytes, set against Bitcoin's own ~80 gigabytes

of annual growth: under three thousandths of one percent. We will not claim the structured footprint is literally zero. We are saying that even a deliberately extreme ceiling rounds to a rounding error.

For context: as of mid 2026, the Bitcoin blockchain grows by roughly six to seven gigabytes a month on its own, on the order of eighty gigabytes a year. The percentage column above measures our added footprint against that yearly growth. Our most extreme scenario, a million Authored inscriptions every month, sustained indefinitely, would amount to roughly a sixth of it; our realistic early scenarios are a fraction of a single percent. Our structured footprint, even at the deliberately extreme ceiling just computed, is smaller still by orders of magnitude. We put the largest of those numbers on the table on purpose. If a permanent record of intelligence is worth this block space (and we believe it is), then we're willing to show what that conviction costs at full scale, even if that ceiling is highly unlikely.

WHERE WE STAND IN THE NODE DEBATE

In October 2025, Bitcoin Core's v30 release raised the default limit on OP_RETURN data carriage from 83 bytes to effectively unlimited (100,000 bytes), and began relaying transactions carrying multiple data outputs. This was a change to relay policy (the local rules each node applies to its own mempool), not to Bitcoin's consensus rules, which did not change. Those who supported it argued that the old limit was already easily bypassed, and so served mainly to push data into less efficient and harder-to-prune forms rather than prevent it; better, they reasoned, to channel data into a clean, dedicated field than to keep fighting an unwinnable game against people who would route around the rule anyway. Those who opposed it (consolidating around the more conservative Bitcoin Knots client, whose share of the network rose sharply through the dispute) argued that a default limit, even an imperfect one, is a real deterrent, and that relaxing it bends Bitcoin away from its monetary purpose toward becoming a general data store, with attendant concerns about fees, the storage burden on full and archival nodes, and the legal exposure of operators who must carry whatever is written. Both camps hold their positions sincerely, and we do not presume to resolve a dispute that those far closer to Bitcoin's foundations have yet to settle. We take no side in it, and we make one commitment because of it:

We design to be defensible to the strictest node, not the most permissive one. We will not build anything that depends on liberal relay policy remaining liberal. We respect the absolute right of any node operator to set conservative defaults and to decline to relay data they consider spam. That right is part of what makes Bitcoin Bitcoin. Our restraint (text only, free tier inscriptions batched by default, paid tier inscriptions paying full market fees) is calibrated so that a careful, skeptical node operator could examine what we do and find it modest, even if they would still prefer we did nothing at all. We aim to be the inscription project that even the debate's harshest critic will call careful rather than reckless.

OUR COMMITMENTS

1. **Text only.** No images, no video, no media. 280 characters max.

2. **Batching by default.** Free Structured records are Merkle-anchored in aggregate, never written to the chain on their own, so our free-tier footprint grows sub-linearly with adoption, by design. A paid Authored record is the only thing written individually, and it is a single inscription.
3. **Authored inscriptions pay full freight.** A paid Authored record carries its lean structured envelope inline with the message, and that single inscription pays the market rate for the block space it uses. We do not offload our costs onto the network, and we do not use tricks to make our data artificially cheap for others to store.
4. **Published footprint.** We will report our cumulative on-chain footprint publicly and keep it current.
5. **No speculative tokens.** We will not issue, mint, or facilitate fungible-token schemes. We are a protocol of record, not a casino.
6. **Designed for the strictest node.** Our practices are calibrated to remain modest under conservative node policy, not to exploit permissive policy.

WHY THIS MATTERS TO WHAT WE ARE

xbt.im exists so that no mind is forgotten. A mission about honoring minds cannot be pursued through carelessness toward the commons that makes that reverence possible. The chain we write to is itself a kind of collective memory, maintained at real cost by people who believe in it. To use it sparingly, transparently, and only for what genuinely merits permanence is not a constraint on our mission. It's a function of it.

We are guests on the hardest ledger ever built. We intend to act like it.

This is a living document. We will revise it as the technology, the debate, and our own understanding evolve, and we will keep prior versions on record permanently, as one might expect.