

The Rail Is the Product

Tokenised cash funds are real, early, and measured on a yardstick that does not yet reach them. What an allocator can establish before deciding, and what the checking is worth.

Fund data read 19 August 2026, tokenised figures re-confirmed 20 August | US dollars | Written for professional investors | No individual fund, provider or issuer is named

\$16.19bn

The whole tokenised cash segment, across 87 products. Roughly double where it stood a year ago

\$158.7m

Three on-chain funds resolve directly in our own fund database. The coverage line sits further out than the market assumes

9bp

The entire fee spread across the six largest incumbent vehicles. The tokenised premium can be several times it

In 2026 the tokenised cash fund stopped being a pilot. Inside four months, an asset manager put tokenised share classes on an institutional cash range, a second listed a tokenised liquidity share class on a payments ledger, a third made liquidity funds available across a fund distribution network, and a fourth cleared Irish regulatory approval for an on-chain money market fund. Two of the largest US fund houses launched on-chain Treasury vehicles of their own. **So the question arrives at the allocator in a familiar form.** Here is a fund that holds the same Treasury bills you already own, on a settlement rail that runs continuously. It costs a little more. Should you switch?

It is a reasonable question and it is very hard to answer honestly, because the two sides of the comparison are not measured the same way. One side publishes a standardised thirty-day yield computed under a regulator's formula. The other publishes a seven-day annualised rate set by its sponsor. One side discloses a total expense ratio. The other discloses a management fee that is currently subsidised, with custody and network costs sitting outside it. One side reports assets to a regulator. The other reports them on its own website.

None of this is concealment. All of it is disclosed. And it means the most natural question an allocator can ask, is this cheaper than what I already hold, has no common basis on which to be answered.

Which is a reason to change the question rather than to keep asking it. This paper puts a view on where the product sits, then establishes four things that are checkable so the view can be tested.

1 Price it as a liquidity instrument, not as a yield instrument

The comparison fails because the two products are not doing the same job, and once that is said out loud most of the difficulty goes away.

A listed cash fund is a **yield instrument**. It is held to earn the bill rate at the lowest available cost, and every basis point of fee is a direct deduction from the only thing it is there to produce. Judged that way, more expensive is simply worse, and there is no argument to have.

A tokenised cash fund is a **liquidity operations instrument**. It is held so that a position can be moved or pledged at a moment when the rails that normally move it are shut. That is a different job, it is not measured in yield, and the fee is not a deduction from it. It is the price of the capability.

OUR VIEW

This belongs in the treasury and collateral sleeve, budgeted against settlement risk, and sized to the cash that has to move on a closed day. **It does not belong in a yield comparison, and the industry practice of putting it there is what makes the product look both overpriced and overhyped at the same time.**

JUDGEMENT

Everything that follows is the checking an allocator should do before buying it, and the reason that checking is currently harder than it looks.

2 The incumbent shelf is large, and it prices inside nine basis points


Start with the half that can be measured, because it anchors everything after it. Across US cash and ultra-short strategies our fund database resolves **164 funds holding \$686.3bn**, counted at strategy level rather than share-class level. **OURS** Within that, the six largest US-listed exchange-traded vehicles price as follows.

VEHICLE	ASSETS	EXPENSE RATIO
0 to 3 month Treasury bill fund	\$103.21bn	0.09%
1 to 3 month Treasury bill fund	\$46.71bn	0.14%
Ultra-short active income fund	\$40.60bn	0.18%
1 to 3 year Treasury fund	\$25.34bn	0.15%
0 to 1 year Treasury fund	\$20.65bn	0.15%
Floating-rate Treasury fund	\$18.73bn	0.15%

Our fund database, read 19 August 2026, collapsed to fund level. Ranked by assets across US-domiciled exchange-traded vehicles in the three cash-like categories. Expense ratios are unioned across both fee tables, because one carries exchange-traded wrappers and the other carries open-end wrappers. **OURS**

Two things follow from that table, and the second is the one that matters.

The obvious one is depth. Against \$686.3bn, the entire tokenised Treasury and money market segment stands at **\$16.19bn across 87 products**, on the category aggregator's 19 August read. **EXTERNAL** A single incumbent fund is still more than six times the whole segment. But the segment is roughly double where it stood a year ago, and it is not a rounding error in anybody's book.



Assets, US dollars, billions. Traditional vehicles are measured directly in our fund database at the 19 August 2026 read. The tokenised segment total is the category aggregator's figure at the same date, covering 87 products, and is not independently verified by us.

2 The incumbent shelf is large, and it prices inside nine basis points CONTINUED

The less obvious one is the fee column. **The six largest incumbents price between 0.09% and 0.18%, a spread of nine basis points.** **OURS** That is the yardstick every tokenised fee gets held against, and it is a very short one.

Against it, the tokenised side spans a management fee of roughly fifteen to fifty basis points across the category.

EXTERNAL Which means the premium is not a number. Take the cheapest tokenised fee against the median incumbent and it is nothing at all. Take the dearest against the cheapest incumbent and it is more than forty basis points. Our own all-in build-up, on the vehicles we examined and with custody and network costs separated out, lands at five to thirty. **EXTERNAL** All three are defensible. They answer different questions.

That settles the pricing argument rather than starting it. A premium that swings from nothing to forty depending on which pair you choose cannot be the basis of the decision. Price is the wrong axis. The product has to be bought for a capability, or not bought.

SO WHAT

For an allocator, this is a sleeve rather than a switch, and the sizing question replaces the pricing question. **For a manager, a segment that has roughly doubled in a year across 87 products is past the pilot stage and not yet consolidated, which is the window in which position is cheapest to take.**


3 The comparison has to be built, and that is where the work is

This section is drawn from our own working papers rather than anyone else's. When we first assembled this comparison, three load-bearing figures were wrong. Each was corrected by checking it against something, and **the size of the corrections is the argument.**

A traditional fund's expense ratio, carried as an external estimate at 0.09%, is 0.15% in our own data. **OURS** The estimate had applied the 0 to 3 month fund's fee to the 0 to 1 year fund. Six basis points. That sounds trivial until it is placed against the effect being measured, which is five to thirty. At the tight end, **the error in the baseline was larger than the entire effect.**

A tokenised fund's yield, reported once at 4.65%, corroborated at 3.49% to 3.55% across three independent sources, and reads 3.40% on the category aggregator today. **EXTERNAL** More than a hundred basis points of correction on a single-source figure that had already been flagged as unreliable and was still carried into a first draft, and another nine since. Cash yields move. A figure of this kind has a shelf life measured in days, and any comparison built on one has to carry its date.

The headline cost gap, first stated as 10 to 45 basis points, is 5 to 30 on our own build-up. **EXTERNAL** The top of the original range could not be supported once the components were built up from disclosed figures only, with custody, network and onboarding costs separated out rather than folded in. As section 2 sets out, that is one defensible answer among several, and the range depends entirely on which two funds are placed side by side.



Left, the two cost corrections against the cost premium the comparison is trying to measure. Right, the yield correction, shown on its own axis and as a range rather than a midpoint, because a yield and a cost are not the same quantity and do not belong on one scale. Teal was corrected against our own data, orange against external corroboration.

3 The comparison has to be built, and that is where the work is CONTINUED

Three figures, all load-bearing, all moved, and none of them because anybody was careless. They moved because a tokenised fund and a listed fund publish different things, and the work of putting them onto one basis is where the errors live.

The comparison is not hard because the technology is new. It is hard because the two sides answer to different disclosure regimes, and neither regime publishes the translation to the other.

This is now the fourth time this year we have found the same failure mode in a different market. A flow score measuring the **pace** of buying, read as its **direction**. A fund distribution rate measuring **cash sent**, read as **income earned**. A fee field populated for one wrapper type, compared against another. And now a sponsor's seven-day rate, read as a standardised yield. Every one of those figures was correct. Every one answered a narrower question than the one being asked of it.

SO WHAT

Each side publishes to its own regime and neither publishes the bridge, which is why the natural question has no answer today. **That bridge is the work, and whoever builds it first sets the terms on which this segment gets compared.**

4 The reference layer has not decided what these products are

Here is the finding an allocator cannot get elsewhere, because it requires looking at the plumbing rather than the product. Our fund database resolves three on-chain Treasury and liquidity funds directly. **OURS**

VEHICLE	DOMICILE	ASSETS
Treasury digital fund, on-chain	United States	\$76.8m
On-chain government liquidity fund	Luxembourg	\$56.9m
On-chain liquidity fund	United States	\$25.0m
Combined		\$158.7m

Our fund database, all three read 19 August 2026 at fund level. **OURS**

The naming and the structure disagree

The Luxembourg fund is **named** as a share class, carrying a class designation and a distribution suffix in its legal name, exactly as an ordinary share class of an ordinary fund would. In the database it is **modelled** as its own separate portfolio, with no parent fund and no sibling classes. The same is true of the two US vehicles. Each sits in the record as a new fund rather than as a new version of an existing one. **OURS**

That matters because the industry is describing this wave in both languages at once. Some launches are announced as tokenised share classes of existing funds, which would let them inherit an incumbent's assets, track record and distribution from day one. Others are standalone products starting from nothing. **Those two paths have completely different economics, and the reference data currently records them the same way.** **JUDGEMENT**

4 The reference layer has not decided what these products are CONTINUED

The identifiers are already colliding

One ticker widely cited in coverage of tokenised Treasury funds resolves, in our database, to an unrelated listed short-term bond exchange-traded fund holding **\$2.49bn**. **OURS** Two different products, one symbol, at opposite ends of the size range. Anyone screening that ticker in a commercial system receives a clean-looking answer about the wrong fund.

We published a near-identical collision a fortnight ago, in a physical gold fund whose ticker was shared with an unrelated European miners vehicle. Twice in two papers is not a coincidence. Reference data is thinnest at a market's newest edge, precisely where the products are least familiar and the checking is hardest.

SO WHAT

How this segment gets classified is being written right now. For an allocator that means asking which of the two things you are being shown, because a size or flow figure will not tell you. **For a manager it means the structure chosen at launch determines how the product is recorded, screened and found for years afterwards, and that window is still open.**


5 What the rail buys, and how to size it

On the static comparison the traditional vehicle wins, and it is worth saying plainly. It is cheaper, it is deeper, it discloses more, and its yield is computed under a formula an allocator can verify. On a buy-and-hold cash allocation with no intraday requirement, there is no return argument for the tokenised version.

What the tokenised version buys is settlement. Consider a Gulf-based family office holding US dollar cash. Local markets trade Sunday to Thursday. US banking and wire rails run Monday to Friday. Three days a week at most one of the two is open: Friday, when the US is open and the local market is not; Saturday, when neither is; and Sunday, when the local market is open and the US is not.

The gap the rail closes is the three days the calendars do not share

This, and not a yield comparison, is what the tokenised premium buys. Size the sleeve to the cash that actually has to move on a closed day.



Trading and settlement calendars for a Gulf-based holder of US dollar cash. The window in which the two calendars do not both trade is the gap a continuously settling instrument closes.

5 What the rail buys, and how to size it CONTINUED

The cost is not the whole three days. It is the cutoff. An instruction that misses the local Thursday cutoff cannot be raised at home until Sunday, and one that misses the US Friday cutoff cannot settle until Monday. Either way the position is immobile across exactly the window in which a regional escalation is most likely to require collateral to move. A holding that can be transferred or pledged continuously closes that gap.

The economics are straightforward and they are not about yield. If the alternative to a settlement gap is an emergency bridge facility or an over-collateralised margin draw, the implied cost of the gap runs to several hundred basis points annualised on a short-dated borrowing. **JUDGEMENT** Against that, a premium of at most a few tens of basis points on the sleeve that needs to move is inexpensive, wherever in the range it falls.

SO WHAT

Pay the premium on the portion of the liquidity book that has to be able to move on a closed day, and hold the rest in the cheaper, deeper traditional vehicle. **The sizing question, and not the yield comparison, is the whole decision.**

The same logic applies where a tokenised holding is required as programmable collateral inside an on-chain structure. That is not a substitution decision at all: a listed fund cannot serve that function, so the comparison never arises.

6 What to do with this

Three questions to put to anyone showing you one of these.

What is the yield figure, and under whose formula? A seven-day annualised sponsor rate and a standardised thirty-day yield are both correct and are not comparable. Ask for the second, or ask for the first alongside the traditional vehicle's equivalent on the same basis. **Is the fee subsidised, and what sits outside it?** Custody, network costs and onboarding are frequently disclosed separately or not at all, so the management fee is rarely the all-in number. **Is this a new fund, or a new version of an existing one?** If it is a share class of a fund with a record, that record is available to you. If it is standalone, the track record starts on launch day whatever the name on the door. Then the only question that requires a view: what proportion of this liquidity book has to be able to move on a day the banks are shut. That number is the size of the tokenised sleeve.

For a manager rather than a holder, the same four findings read as a build list. The segment has roughly doubled in a year and has not consolidated, which is when position is cheapest to take. The premium cannot be argued on price, so the product has to be sold on the operational capability and the disclosure has to make that capability legible. Neither disclosure regime publishes the bridge to the other, so whoever builds it first sets the terms of comparison. And the reference layer is still being written, so how a launch is structured now determines how it is found later. **JUDGEMENT**

The rail is the product. It is worth paying for exactly as often as you need to run something down it, and it is worth building for while the segment is still deciding what it is.

6 What to do with this CONTINUED

WHAT WE COULD NOT ESTABLISH

Every tokenised figure here is external. Each is sponsor or aggregator published, carries a date between March and August 2026, and is not independently verified by us. We hold no proprietary return, flow or holdings series for these funds, which is why no performance comparison appears anywhere in this paper.

We cannot price the settlement optionality in a stress event. That requires data on actual weekend liquidity events, which we do not hold, so the sizing rule is offered as a method rather than as a number.

The scale comparison pairs two provenances. It places our directly measured fund figures against a category aggregator's segment total. Both carry the same date, and the segment total is the aggregator's count across 87 products rather than anything we hold.

The segment total in an earlier draft of this work was wrong by roughly three times. It carried a figure of five to six billion dollars, assembled from four sponsors and two aggregators on mixed dates between March and August. Re-confirmed against the category aggregator on 20 August, the figure is \$16.19bn. That miss is the reason this paper dates every external figure at the point of use.

This is scoped to what we examined. Where we describe what a set of funds does or does not offer, that description covers the funds reviewed here at the date shown, and is not a statement about the market as a whole. Tokenised private-credit pools sit at a different point on the risk curve and belong in a different paper. Which ledger a fund settles on is not an allocator's decision variable and does not appear in this analysis.

One earlier input was removed rather than carried forward. Credit-spread context that appeared in an internal draft of this work could not be re-derived at the date of writing, so it does not appear here.

Three Horizons Capital, 20 August 2026. Fund counts, fund sizes, expense ratios and the on-chain fund coverage set are derived from the Three Horizons Capital data platform, read on 19 August 2026. Morningstar category is used as the classification construct. Fund-level figures are collapsed from share-class records; expense ratios are unioned across both fee tables. All figures describing tokenised products, including segment size, quoted yields and fee levels, are sponsor or aggregator published, carry dates between March and August 2026, and are not independently verified by Three Horizons Capital; they are tagged EXTERNAL at the point of use. Items tagged JUDGEMENT are reasoned inference rather than measured data. No individual fund, provider, issuer or ticker is named. This document is for informational purposes, is intended for professional and qualified investors, does not constitute investment, tax or legal advice, and is not a recommendation to buy or sell any security. Three Horizons Capital is not a regulated entity.