

Vaults and Covenants

James O’Beirne
vaults@au92.org

January 9, 2023

Abstract

Vaults are a technique for substantially reducing the risk of Bitcoin theft. In this paper, we examine how they might be implemented using different covenant designs. We survey the vault implementations that are currently usable, as well as those that would be possible only with proposed consensus changes.

A new approach, OP_VAULT, is presented which avoids the pitfalls of a general covenant proposal while still enabling the behavior necessary for a featureful vault implementation. The design assumes the deployment of package relay and ephemeral anchors for dynamic fee management, but allows for future fee management approaches, e.g. transaction sponsors, should they come. OP_VAULT uniquely allows batching operations, partial unvaultings, dynamic withdrawal targets, and recursive deposits.

Last updated

This article was last updated on February 21, 2023. The technical details for the OP_VAULT proposal have been superseded by a Bitcoin Improvement Proposal [OBe23b] which incorporates a number of changes based on feedback from the ensuing mailing list discussion [OBe23a] after this paper’s initial release.

Introduction

Custodying bitcoin is notoriously hazardous. It is roughly equivalent to another task famous for its difficulty: keeping sensitive data accessible but out of unauthorized reach. Luckily, due to the programmability of Bitcoin script, custodians do not necessarily have to rely solely on the failure-prone task of key management.

In Bitcoin, a *covenant* is a constraint put on how a coin can be spent on the basis of its spending transaction, above and beyond the traditional requirements of satisfying a one-time unlocking script. Covenants can be written in such a way that they express recursive constraints on outputs, so that the effective period of covenant enforcement can span an arbitrary number of transactions.

Currently there is no way to enforce covenants “on-chain,” or endogenously within Bitcoin’s validation rules, but there are many pending proposals for how to modify Bitcoin’s validation rules to allow for this functionality. Covenant proposals can be divided into two kinds:

- **precomputed:** where the state machine of possible transactions within the lifecycle of a covenant is predetermined and therefore bound (e.g. OP_CTV [Rub20b], ANYPREVOUT [Chr18], OP_TX [Rus22]), and
- **general:** where the state machine is expressed within Bitcoin script, and can reapply itself indefinitely, making it potentially unbounded (e.g. transaction introspection opcodes and OP_CHECKSIGFROMSTACK [Kan21]).

Vaults are an especially useful kind of covenant that give Bitcoin users operational simplicity during expected use, but heightened security in the event of unauthorized access to private keys. When they were originally presented in [MS16], a vault was defined as a simple covenant that ensured that the spend of a coin was only allowed after broadcasting an intent to unvault and waiting some period. During this delay, the funds could be “clawed back” into a prespecified recovery path in case the proposed spend was unexpected.

In order to make use of this covenant, a user would need to spend their coins into a vault, then configure a watchtower process to monitor the chain. If the watchtower detected an unexpected unvaulting, it could alert the owner or automatically broadcast a transaction that sweeps the funds to safety.

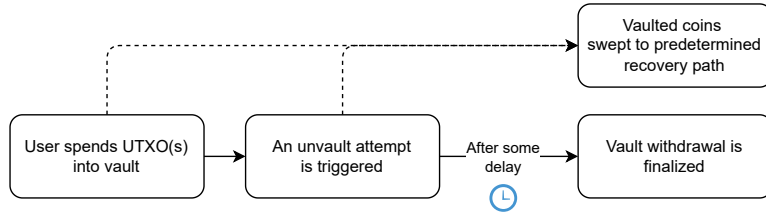


Figure 1: The basic lifecycle of a Bitcoin vault.

An idealized vault

The most useful incarnation of a vault design would allow Bitcoin custodians to benefit from a predetermined contestation period and recovery method with minimal overhead.

Some features that are crucial to this include

- **efficient reuse of an existing vault configuration.** A single vault scriptPubKey should be able to “receive” multiple deposits and still facilitate efficient administration of the vaulted coins. Batched operations for recovery and unvault should be supported in order to allow managing such a vault efficiently.
- **partial unvaultings**, which would allow users to withdrawal partial vault balances without having to perform the setup ceremony for a new vault.
- **dynamic unvault targets**, or allowing the proposed withdrawal target for a vault to be specified at unvault time rather than when the vault is first created. This would remove the need for a statically-specified, intermediate live wallet during withdrawal that only exists to route unvaulted funds to their desired destination.
- **dynamic fee management** that, like dynamic targets, defers the specification of fee rates and source of fees to unvault time rather than vault creation time.

As we will see, no existing vault design meets this set of features.

Precomputed vaults

Given Bitcoin’s existing validation rules, the only current, viable means of implementing a vault that is enforced directly by the validation rules¹ is by presigning a limited graph of transactions, as proposed in [Bis19] and implemented in [Bis20]. A user must generate an ephemeral key,

¹Note that [Rev] is a method of emulating vaults using large multi-sig quorums, but this comes with significant operational complexity.

sign a transaction sending coins to an address controlled by that key, presign a tree of possible transactions using that key, and then delete the key.

This locks the coins into a predetermined flow, and technically satisfies the definition of a vault. A robust implementation of this approach is described in detail in [Swa+20], as well as many related considerations like running a watchtower.

Presigning transactions to construct vaults in this way does add the safety of a limited recovery window, but it has a number of drawbacks:

- Key deletion cannot be proved, so ensuring that the vault isn't backdoored is more difficult.
- The custodian must not lose the presigned transactions, since there is no other way of spending the bitcoin. The sensitive data that is necessary to store indefinitely grows linearly with the number of vaults created.
- The spend target for the vault is static, and presumably must correspond to some kind of hot (or "warm") wallet. Loss of control of that hot wallet necessitates sweeping all presigned transactions to the likely difficult-to-access recovery wallet.
- Arbitrary vault withdrawal amounts are not possible after the structure of the vault is locked in by presigning.
- Vault operations, namely recoveries and unvaultings, cannot be batched together. This is especially unfortunate because in the case of a key leak, it may be critical to sweep all vaults to the recovery path as soon as possible; otherwise the custodian may end up racing the attacker to spend out of the vault.
- Once initiated, other deposits cannot (or *should not*) be sent to after the initial deposit. This prevents e.g. exchanges from initiating vaults and then letting customers deposit directly into vaults at-will.

Precomputed vaults with covenants

Having some kind of general on-chain covenant mechanism improves the situation somewhat. Script functionality like `OP_CHECKTEMPLATEVERIFY` (CTV) [Rub20b] allows us to use a vault scheme very similar to presigned transaction vaults, but with the considerable benefit that we do not need to engineer and operate an ephemeral key signing and deletion ceremony, nor do we have to persist critical presigned transactions indefinitely.

Similar to the presigned transaction approach, the entire allowable transaction state machine needs to be generated ahead of time and committed to with an `OP_CHECKTEMPLATEVERIFY` hash. This is demonstrated in [OBe22].

This approach has the benefits of not requiring storage of anything aside from the vault parameters used to generate the CTV transaction graph, which is not particularly sensitive. It also doesn't require ephemeral keys due to the nature of CTV, which uses an on-chain commitment to enforce a precomputed covenant.

Unfortunately, this approach still has major limitations.

- By nature of being precomputed, the number of vault operations is limited and predefined, ruling out arbitrary partial unvaults and recursive re-vaults.
- The destinations are still fixed to a set of keys. Unvaulting must be done through a single predetermined path. Fee management keys must be precommitted; if the corresponding wallet is lost, ability to finalize vault transactions may be impaired (more on fees later).

- Vault operations cannot be batched together; this turns out to be a big limitation when responding to attackers.

These caveats also apply to other, similar vault implementations using different precomputed-covenant approaches in which the allowed graph of transactions must be precomputed, e.g. ANYPREVOUT ([Poi22]) and the (proposed but not implemented) OP_TX ([Rus22]).

Recursive vaults with general covenants

If we introduce sufficiently powerful script functionality to allow arbitrary introspection of the spending transaction, we would have the ability to write any kind of covenant – including a fully-featured, recursive vault – without having to resort to presigned transactions, covenant emulation via multi-sig, or CTV/APO-style precomputed transaction graphs.

Such a vault would be free from the limitations of the precomputed approaches described above, but at the cost of significant script complexity. Writing the necessary script for such a vault based on primitives described in e.g. [Kan21] would not only be very complex, but it would be very verbose on-chain. The `scriptPubKey` sizes would be quite large for what might be a very common operation.

There is also a significant chance that the Bitcoin community might not reach agreement on how to proceed towards a general covenant solution. Vaults have considerable utility for a great many users in the immediate term, since all users are concerned with custody. Gating this highly practical feature on establishing broad consensus around a general covenant mechanism (which isn't even guaranteed to happen) seems suboptimal.

OP_VAULT

Instead of resigning ourselves to the limitations of precomputed vaults, or waiting for a general-covenant mechanism to be deployed to script (a daunting prospect), we propose to evaluate the the creation of opcodes `OP_VAULT` and `OP_UNVAULT`, which have covenant-like characteristics but do not attempt to address the general problem of covenants.

This approach has a complete set of desirable features for safer custodial operations, none of the limitations of precomputed vaults, and is more concise and usable than a vault implemented with more general covenant scripting primitives.

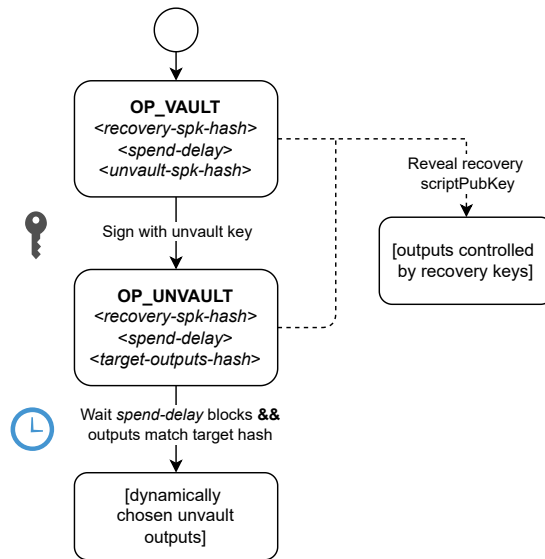


Figure 2: A high-level description of the OP_VAULT state machine.

Components of the vault

Each OP_VAULT-style vault makes use of a few pieces of essential data. These include

- **a recovery path:** the destination that vault funds can be swept to at any point prior to the finalization of withdrawal to the *unvault target*.

This recovery address would usually correspond to a spending script that is inconvenient to exercise but is highly secure, e.g. a key generated offline or a geographically distributed multisig. It could be a Taproot `scriptPubKey` that incorporates a number of different spending conditions.

Vaults which share the same recovery path can be swept in batch operations, which is an important practical aspect of managing large numbers of vaults.

- **an unvault key:** used to authorize beginning an unvault process, i.e. the spending of an OP_VAULT into a suitable OP_UNVAULT, which “announces” the intent to unvault and begins the withdrawal lock-in period. If an attacker obtains access to this key, the outcome is not catastrophic because any unvault can be interrupted and swept to the recovery address. This can be an arbitrary `scriptPubKey` so long as it is a valid witness program.

Vaults which have this in common can be unvaulted in batch.

- **an unvault target:** an arbitrary target or destination that is specified as a parameter to OP_UNVAULT, and dictates where unvaulted funds go. This target consists of a list of destination outputs (including amounts) which the vault will be spent into after the delay period. These destinations can of course include recursive OP_VAULT outputs to facilitate partial-amount unvaults.

Initial vaulting

To create a vault, a user would spend to an output with the following `scriptPubKey` (or equivalent taproot structure):

```
<recovery-spk-hash> <delay-period> <unvault-spk-hash> OP_VAULT
```

where

- `<recovery-spk-hash>` is the tagged SHA256 hash of the recovery `scriptPubKey`. This could e.g. correspond to a P2TR address with many spending conditions.
- `<delay-period>` indicates the delay between when the `OP_UNVAULT` confirms and when the funds can be spent to the `unvault` target outputs. It is specified in the same way that the lower 24 bits of `nSequence` timelocks are, allowing a relative lock in terms of time or block height.
- `<unvault-spk-hash>` is the tagged SHA256 hash of the `scriptPubKey` whose witness is required to start the `unvault` process.

An output encumbered by an `OP_VAULT` `scriptPubKey` can only be spent one of two ways:

1. it can be swept to the recovery address at any point with no witness data, or
2. it can be sent to an `OP_UNVAULT` output matching the specification described below.

Triggering an `unvault`

To withdraw bitcoin from an `OP_VAULT`-encumbered output to some arbitrary destination, it must first be spent into an `OP_UNVAULT` output.

Witness requirements

A witness must be provided for the `OP_VAULT` input being spent which contains

1. a witness program `scriptPubKey` that, when hashed, yields the `<unvault-spk-hash>`, and
2. a witness stack satisfying that witness program.

This witness is crucial for preventing unauthorized `unvault` attempts.

Spending output requirements (covenant)

The `OP_UNVAULT` `scriptPubKey` must have the form

```
<recovery-spk-hash> <delay-period> <unvault-target-hash> OP_UNVAULT
```

where the first two parameters are carried over from the input `OP_VAULT` (described above), and `unvault-target-hash` is defined as

```
sha256d(serialized(d.scriptPubKey) || serialized(d.nValue)  
for d in target_outputs)
```

In other words, the target hash commits to the set of outputs that are proposing to be withdrawn to.

Withdrawal to the unvault target

The unvaulting party must now wait for the timelock specified by the [Fri+15]-compatible *delay-period* to mature before being able to broadcast a transaction spending the OP_UNVAULT output into a set of outputs satisfying *unvault-target-hash*. No witness is needed for this finalizing “withdrawal” transaction, which represents the completion of the vault lifecycle for all amounts which are not revaulted.

If the owner of the vault does not recognize this proposed withdrawal, he can sweep the vault into the recovery path before the delay period ends.

Sweeping to recovery

In order to sweep an OP_VAULT or OP_UNVAULT output to the recovery address, a transaction is broadcast spending that output into a new output corresponding to the recovery address, consuming the full amount of the output(s) being swept. This can happen at any point before withdrawal to the unvault target is finalized.

The witness for the swept output is empty, since the only authorization necessary is simply that the `scriptPubKey` being spent to matches the `<recovery-spk-hash>` committed to in the parent OP_VAULT output.

Denial-of-service protection

During the creation of OP_VAULT and OP_UNVAULT outputs, the recovery address remains hidden behind a hash. This avoids denial-of-service (DoS) attacks that involve a third party attempting to blindly sweep vault transactions to their recovery path, which could result in a temporary halt in liquidity while the vault owner goes through the potentially lengthy process of activating the recovery keys for funds retrieval.

If it becomes necessary to make use of the recovery path, the recovery `scriptPubKey` will be revealed, which means that any other vaults with that recovery path may be swept there by an unauthenticated party.

This seems like a reasonable trade-off, since presumably if the vault’s owner had to make use of the recovery path for one vault, any other vault sharing the recovery path may be at risk due to a compromise in the owner’s infrastructure.

Batching

If multiple vaults share a common *recovery-spk-hash*, they can be swept in batch. This is an especially important feature for users that maintain many vaults, since a key compromise might necessitate rapidly sweeping many vaults at once before a delay period has elapsed. Batching the sweep reduces the chainspace required significantly, because we’re able to sweep to a total of two outputs instead of creating $2v^2$ outputs for v vaults.

Note that this batching option is not possible in vault schemes that rely on precomputed covenant techniques.

²An anchor output for fee control is needed for each sweep transaction

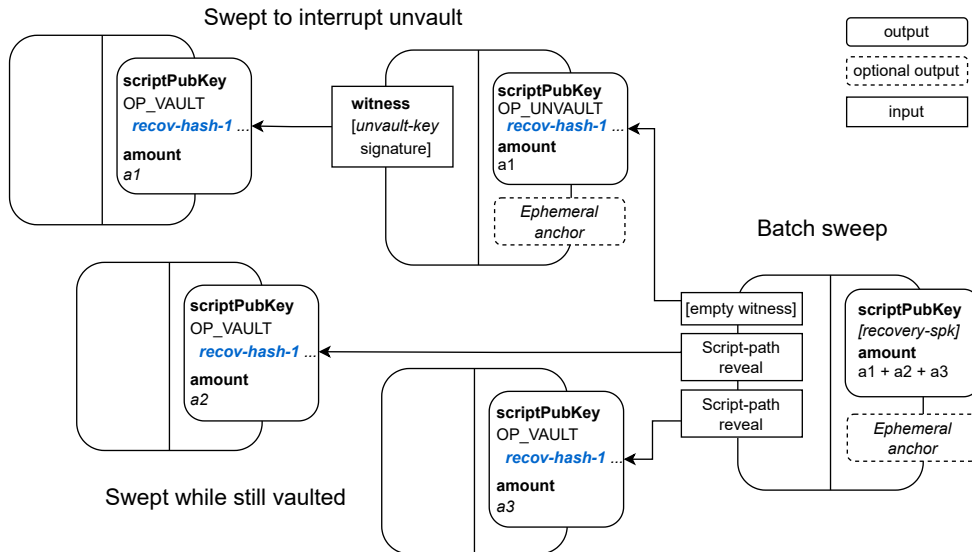


Figure 3: Demonstration of multiple vaults being batch-swept simultaneously. Two of the vaults (bottom) are still vaulted while one (top) is interrupting an attempted unvault. This pattern is demonstrated with functional tests in `test_batch_sweep()` of `[OBe23c]`.

Managing fees safely

One difficulty when designing this scheme was determining how to adjust the fee rate of transactions that happen sometime after vault creation. The fee environment may differ considerably from when the vault was initially created to when it will be accessed; fees may be much higher.

The most common method of increasing fee rate is to lower output amounts, which leaves more bitcoin available to miners, adding to the incentive to include a given transaction. During an OP_VAULT lifecycle, allowing variable amounts would unfortunately present some problems.

If an attacker discovers the recovery preimage, they may trigger undesired recovery sweeps. If the OP_VAULT validation rules allowed output amounts to vary for the recovery transaction, an attacker could simply set that value to 0 and burn the vaulted coins. A similar consideration must be applied to OP_UNVAULT transactions.

In order to avoid such a failure, there are two options:

1. script validation rules could require some allowable “range” of amount discounts during unvault/recovery to facilitate fee payment, or
2. script validation rules could require that the unvault/recovery outputs preserve the full value of any vaults being acted on.

1. seems like a bad design. 2. seems like the preferable approach, but it introduces a complication: how are we to pay fees for sweep and unvault transactions if the full value of the vaulted coins must be preserved?

Using package relay and ephemeral anchors

Fee management is a thorny part of designing any contracting protocol built atop Bitcoin script. Transactions which are presigned (or covenant structures that fix amounts) cannot vary their amounts to adjust fees dynamically. A common workaround is to attach an “anchor output” [Vara], which is an output that is intended to be spent by some child transaction. Spending the child can raise the effective feerate of its parent, the precommitted contract transaction. This is referred to as child-pays-for-parent (CPFP).

This technique alone is insufficient for reliable fee control. If fees have risen considerably since the initial contract negotiation and the feerate of a precommitted transaction is beneath the minimum mempool feerate, it may be unable to broadcast in the first place and CFP cannot take place.

Luckily, it is looking increasingly likely that package relay policies ([Varb]) that remedy this will become part of bitcoin. The OP_VAULT proposal detailed here not only assumes the presence of package relay, but also makes use of *ephemeral anchors* ([San22]), zero-value outputs which must be spent within the package they're relayed with, in order to avoid unnecessarily wasting vaulted value for fee control.

Our rationale for the permissibility of this hard dependency is that many – perhaps all – contracting schemes built on Bitcoin become unworkable without a robust solution for dynamic fee management, which itself appears to depend on package relay. We fully expect that if OP_VAULT were found to be a desirable soft-fork, its deployment would naturally come after the deployment of package relay and some kind of fee control mechanism like ephemeral anchors.

Given the deploy of these mempool policies, the OP_VAULT strategy is able to both provide efficient fee management and rule out coin-burning DoS attacks.

Future-proofing fee control

To avoid wasting value and to allow dynamic fee control, each OP_UNVAULT transaction, recovery sweep transaction, and unvault target transaction *may* have an ephemeral anchor output that facilitates robust fee management with CFP and package relay.

But there are other possible designs for dynamic fee management. One is transaction sponsors [Rub20a]. Sponsors allow fee rate improvement of any unconfirmed transaction by any party, and they do not require any structural “awareness” from the transaction whose fee is being bumped. In the proposed design, sponsors would be required by consensus to be mined alongside other transactions which they are sponsoring. This allows CFP-like adjustment of fee rate, but without a wasted output or specific planning.

In order to allow for the possible adoption of such a proposal, any ephemeral anchor outputs in the OP_VAULT scheme are strictly optional. To avoid pinning by an adversary who discovers the recovery path, any transaction for which the ephemeral anchor is optional must be marked for replacement by fee (RBF) (by policy, as a part of this proposal). This ensures that the vault owner can override spurious sweep transactions that exclude the anchor with higher-fee packages making use of the anchor.

The importance of dynamic unvault targets

In vault schemes which are usable today, the unvault target is always fixed at vault creation time. Even in schemes that assume the use of a precomputed covenant mechanism like CHECKTEMPLATEVERIFY or ANYPREVOUT, the set of withdrawal targets is similarly fixed.

The design presented here allows for the withdrawal target to be decided at unvault time, before the delay period, rather than needing to be specified during the creation of the vault. This presents a significant benefit, because if the withdrawal target is fixed, an intermediary hot or warm wallet must be used to unvault and then send the coins to their ultimate destination.

If a withdrawal target can be set at unvault time, no such intermediary wallet is needed. This saves on operational complexity, avoids a potential avenue of compromise, and also removes the need for an additional transaction – which may be onerous when considering possible future demand for chainspace.

Batching

Dynamic unvault targets also allow multiple vaults to be unvaulted (or recovered) to a single set of output targets instead of requiring separate transactions for each vault deposit. In the case of a withdrawal, unvaulted funds can skip the “warm” wallet step that precomputed vault funds must pass through on their way to destinations only known at unvault time.

This isn’t only important as blockspace becomes scarcer; the ability to batch may prevent a situation in which, given a large number of vaults and a low timeout period, a custodian otherwise may not be able to get all the confirmations needed to thwart an attacker.

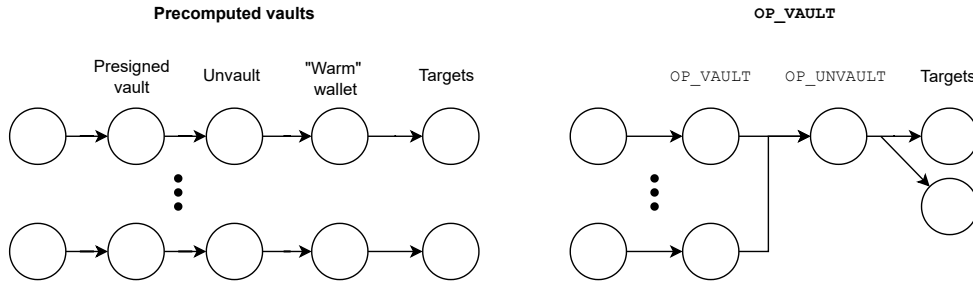


Figure 4: Comparison of the “expected” withdrawal flow for precomputed vaults vs. an OP_VAULT usage. Note that due to (i) the ability to consolidate the unvault and (ii) skip the warm wallet, the OP_VAULT design uses something on the order of $2v$ fewer outputs (for v vaults) than the precomputed vault does.

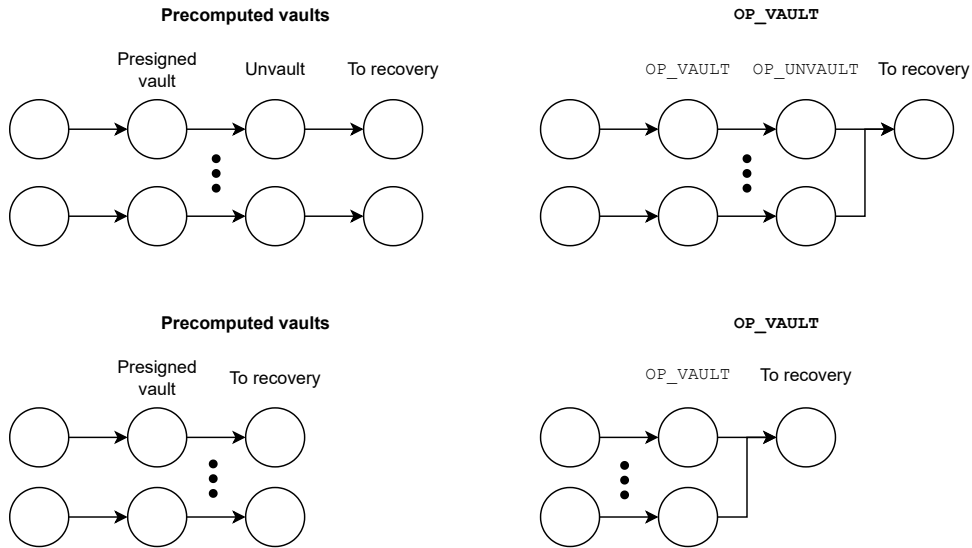


Figure 5: Comparison of the recovery process. The top row shows an interrupted unvault, the bottom shows recovery prior to an attempted unvault. In each case, OP_VAULT reduces output creation on the order of v .

Attacks on batching

Possible DoS attacks on batching have been hypothesized. One possible attack can happen if a custodian of a large number of vaults is attempting to batch sweep; an attacker who has mapped the mempool may “unbundle” the batch, choose a single vault output to sweep, and attempt to front-run and pin the batch sweep with a single conflicting transaction at mempools near miners. This would have the effect of preventing the propagation of a batch sweep.

A way around this might be for the custodian to monitor the confirmation of the sweep; if a batch sweep isn't confirmed despite having an appropriate fee rate, the batch could be split in two and rebroadcast. If a confirmation doesn't happen in a subsequent block, the batch could be split into four, and so on and so forth until the attacker is forced to permit the confirmation of the sweeps.

This attack may not be specific to vaults per se; more robust anti-pinning techniques are necessary for many applications in Bitcoin. This discussion does however point to the necessity for any vault user to take caution that the number of vaults which share *recovery-spk-hash* and *unvault-spk-hash* parameters can be swept in the number of blocks that their *spend-delay* provides.

Recursive deposits (revaults)

Another benefit of supporting dynamic unvault targets is that an arbitrary number of partial unvaults can happen by withdrawing to some output outside the vault, but then redepositing the remaining balance back into the same `OP_VAULT [...]` construction.

This could enable, for example, an exchange to offer a vault address to a customer, allow them to make multiple deposits, and support partial withdrawals without having to redo a vault initialization.

Downsides relative to other designs

We have already written about the important benefits that this scheme provides relative to alternatives. While we think the set of trade-offs for `OP_VAULT` are appealing, it is important to acknowledge downsides.

A scheme that does not codify vault semantics at the script interpreter layer, e.g. one that uses only a general covenant primitive like transaction introspection within script, allows for more flexibility in the particular control flow of the vault, since the rules of the vault are specified by the end-user instead of the consensus engine. It more easily facilitates user experimentation.

For example, perhaps a user prefers to safeguard their recovery path with a signature rather than revealing the preimage to a witness (program) hash, so that discovery of the recovery path doesn't allow an unauthenticated user to sweep open vaults to recovery.

The downside, of course, is that this results in very large on-chain script sizes. In order to make vault constructions concise, they ultimately must be codified in validation rules, barring some breakthrough in a zero-knowledge proof system that could be supported on-chain.

Implementation

A draft implementation is provided at [OBe23d]. The draft lacks a proposed activation mechanism and some minor policy and wallet changes, but the major functionality described above is implemented and comprehensively tested in [OBe23c].

If there is sufficient interest, an accompanying BIP will be written.

Conclusion

We have presented a new set of opcodes, `OP_VAULT` and `OP_UNVAULT`, which enable featureful vaults in Bitcoin. These opcodes allow encumbering a set of coins in such a way that their spending requires passing through a delay period, during which the coins can be recovered to a set address. Enabling on-chain enforcement of such a control flow presents significant benefits to

custodians of bitcoin, whether large or small, because this scheme in essence offers multisig-like security for an expected operational burden roughly on par with single sig use.

Thanks to use of package relay and ephemeral anchors, the OP_VAULT proposal is unique in enabling robust dynamic fee management, previously a major challenge for vaults. It also enables batch management, a critical feature when considering future scarcity of chainspace and the need to respond to key-compromise efficiently. Dynamic unvault-target specification removes the need for intermediate routing wallets, making the use of vaults simpler and safer.

Acknowledgements

Thanks to BL, Jeremy Rubin for review and discussion. Thanks to Antoine Riard for finding attacks. Thanks to Greg Sanders for lots of discussion and ideas on script execution.

References

- [Bis19] Bryan Bishop. *Bitcoin vaults with anti-theft recovery/clawback mechanisms*. 2019. URL: <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2019-August/017229.html>.
- [Bis20] Bryan Bishop. *python-vaults*. 2020. URL: <https://github.com/kanzure/python-vaults>.
- [Chr18] Anthony Towns Christian Decker. *BIP-118: SIGHASHANYPREVOUT for Taproot Scripts*. 2018. URL: <https://github.com/bitcoin/bips/blob/master/bip-0118.mediawiki>.
- [Fri+15] Friedenbach et al. *BIP-68: Relative time-lock*. 2015. URL: <https://github.com/bitcoin/bips/blob/master/bip-0068.mediawiki>.
- [Kan21] Sanket Kanjalkar. *Elements Taproot Introspection Opcodes*. 2021. URL: https://github.com/ElementsProject/elements/blob/master/doc/tapscript_opcodes.md.
- [MS16] Ittay Eyal Malte Möser and Emin Gün Sirer. *Bitcoin Covenants*. 2016.
- [OBe22] James O’Beirne. *simple-ctv-vault*. 2022. URL: <https://github.com/jamesob/simple-ctv-vault>.
- [OBe23a] James O’Beirne. *[bitcoin-dev] OP_VAULT: a new vault proposal*. 2023. URL: <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2023-January/021318.html>.
- [OBe23b] James O’Beirne. *OP_VAULT BIP*. 2023. URL: <https://github.com/jamesob/bips/blob/jamesob-23-02-opvault/bip-vaults.mediawiki>.
- [OBe23c] James O’Beirne. *OP_VAULT draft functional tests*. 2023. URL: https://github.com/jamesob/bitcoin/blob/2023-01-opvault/test/functional/feature_vaults.py.
- [OBe23d] James O’Beirne. *OP_VAULT draft implementation*. 2023. URL: <https://github.com/jamesob/bitcoin/tree/2023-01-opvault>.
- [Poi22] Antoine Poinot. *simple-anyprevout-vault*. 2022. URL: <https://github.com/darosior/simple-anyprevout-vault>.
- [Rev] Revault. *Revault*. URL: <https://revault.dev>.
- [Rub20a] Jeremy Rubin. *A Replacement for RBF and CFPF: Non-Destructive TXID Dependencies for Fee Sponsoring*. 2020. URL: <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2020-September/018168.html>.

- [Rub20b] Jeremy Rubin. *BIP 119: CHECKTEMPLATEVERIFY*. 2020. URL: <https://github.com/bitcoin/bips/blob/master/bip-0119.mediawiki>.
- [Rus22] Rusty Russell. *OP_TX: generalized covenants reduced to OPCHECKTEMPLATEVERIFY*. 2022. URL: <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2022-May/020450.html>.
- [San22] Gregory Sanders. *Ephemeral Anchors: Fixing V3 Package RBF against package limit pinning*. 2022. URL: <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2022-October/021036.html>.
- [Swa+20] Jacob Swambo et al. *Custody Protocols Using Bitcoin Vaults*. 2020. URL: <https://arxiv.org/pdf/2005.11776>.
- [Vara] Various. *Bitcoin Optech: Anchor Outputs*. URL: <https://bitcoinops.org/en/topics/anchor-outputs/>.
- [Varb] Various. *Bitcoin Optech: Package Relay*. URL: <https://bitcoinops.org/en/topics/package-relay/>.