THORChain

Research & Analysis

Institutional

December 2019
85 Broad Street
New York, NY, 10004
www.delphidigital.io

Table of Contents

Analysts

Yan Liberman, CFA, CAIA

yan@delphidigital.io

Medio Demarco

medio@delphidigital.io

Executive Summary

THORChain is an interoperable, application-specific blockchain that facilitates cross-chain liquidity pools. BEPSwap, its first mainnet iteration, will go live in Q1 2020 and run for 6 weeks, enabling BEP2 assets to be pooled on Binance Chain. Shortly after, ASGARDEX, the full cross-chain implementation, will launch. THORChain has raised $2.1m from seed-funding in 2018, a Binance IDO in July 2019 and private sales.

Details	
Ticker	RUNE
Price (USD)	$0.07
Market Cap	$8,585,718
Current Supply	122,000,000
Maximum Supply	500,000,000

Returns in USD	7D	1M	3M	1Y
RUNE	-0.78%	-23.99%	+674.07%	-
BTC	+0.18%	-1.15%	-12.36%	+95.23%
ETH	+5.53%	-11.28%	-25.05%	-1.23%

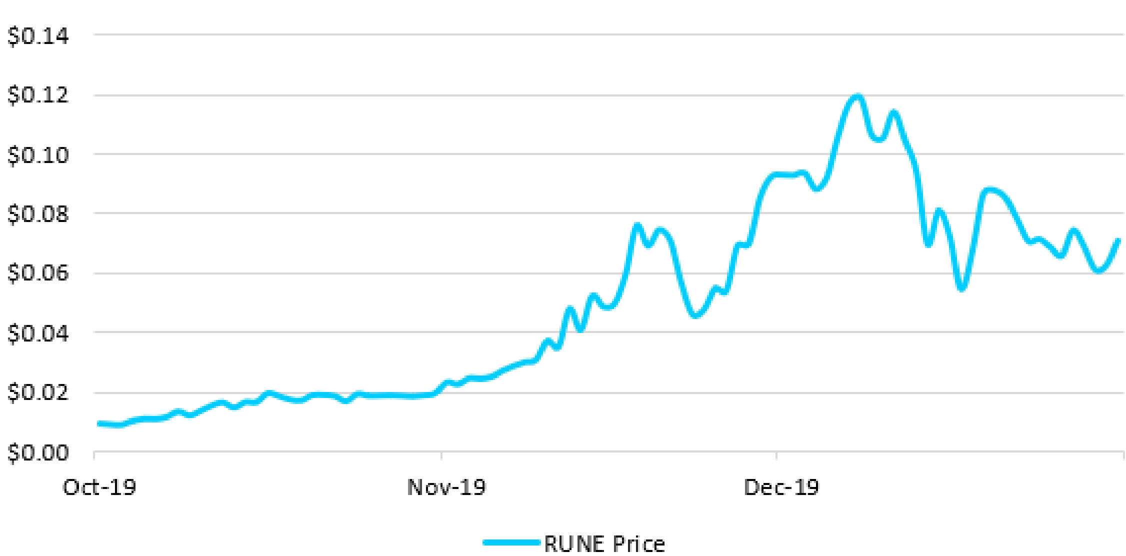

Data as of: December 31st, 2019

Key Takeaways - THORChain

- The team leveraged codebases from other projects and combined them in a novel way to construct THORChain, helping mitigate execution risk. Key components of its architecture were built using Cosmos SDK, Tendermint and Binance Chain's TSS code library.
- As an application-specific blockchain, THORChain is optimized to best serve its specific use case (i.e. liquidity pools). The economic design has incentives in place to ensure the financial risk of the application is always proportional to the security of its network. If successful, we may see similar cross-chain designs take DeFi by storm.
- Trade mechanics aside, THORChain operates much like a centralized exchange but with one key difference - it decentralizes control over the exchange wallets. THORChain's validator nodes control the on-chain liquidity pools and much reach consensus to move funds. The network will start accepting liquidity once 12 nodes are live, gradually increasing to 99+.
- Staked liquidity is pooled on its native chain into "vaults". Pools maintain a single balance for each asset, with liquidity divided among, and regularly moved between, different vaults.
- TSS is used to unlock vaults and move funds (i.e. execute a swap) and has several advantages over MultiSig. It's chain agnostic and already compatible with BTC & ETH.
- The project has a large focus on privacy to improve the exchange's resiliency.
- With its unique design, THORChain has a head start on its closest competition.

Key Takeaways - CLPs & RUNE Economics

- THORChain's Automated Market Maker (AMM) called the Continuous Liquidity Pool (CLP) reduces a Liquidity Provider's (LP) under-performance relative to a buy and hold strategy. It also incorporates a Slip Based Liquidity Fee, which increases the fee that LP's receive as transactions increase in size relative to the pool. It's an effective way to incentivize liquidity providers (supply) to actively search for pools where transaction size and volume (demand) suggests there's a need for more liquidity.
- This model also enables the less cost conscious traders, who are just looking to execute, to pay a premium while allowing the transaction costs of smaller sized trades to come down
- When providing liquidity, each pool is a combination of RUNE and another asset. Based on the incentive mechanisms built into the system, as assets are staked for liquidity provision, the price of RUNE will appreciate. More specifically, the value of Bonded+Staked RUNE must be 3x the value of non-RUNE staked assets, creating a growing floor value as more assets are staked.
- The network will bootstrap initial participation from Liquidity Providers and Node Operators by economically incentivizing them with RUNE issuance. The yields early on will likely get liquidity pools and security to an operational size where economically rational actors will be organically incentivized to participate.

Value Proposition & Token Economics

Use Cases

1. Cross-Chain Exchange: Traders can swap assets directly between chains without needing to use a centralized exchange. Major blockchains such as Bitcoin, Ethereum, Binance Chain and others can be supported today, due to THORChain's technical design.

2. Stake Unproductive Assets for Yield: Holders of unproductive assets (e.g. BTC) can stake them in liquidity pools to earn a yield from trading fees and new RUNE issuance. THORChain's liquidity pool design insulates stakers from losses better than other models (Uniswap), further incentivizing staking.

3. Run a Node for Yield: Anyone can bond the minimum amount of RUNE required to operate a node (currently 1m RUNE) and run the network. In return, they can earn a yield from trading fees and new RUNE issuance.

4. Arbitrage Opportunities: Arbitrageurs can earn a profit, while bringing the ratio of assets in liquidity pools back to their equilibrium.

Key Benefits

1. Permissionless & Private: Anyone in the world can directly swap assets cross-chain without needing permission or disclosing their identity. In addition, node operators are anonymous and no one can identify which of them signed off on a specific transaction.

2. Staked Liquidity is Non-Custodial: Assets are pooled on-chain, rather than entrusted to a third-party custodian, and secured by a decentralized network.

3. Security Linked to Liquidity: The value of bonded RUNE securing the network must be 2x the value of staked external assets (e.g. BTC) provided as liquidity. This means that as the value of staked liquidity grows, the value backing its security grows proportionally.

4. Unique Fee Model: The Slip Based Liquidity fee used by Continuous Liquidity Pools provides a unique way to incentivize liquidity providers (supply) to actively search for pools where transaction size and volume (demand) suggests there's a need for more liquidity.

5. Bootstrapping: RUNE is used to bootstrap liquidity and security through preset rewards early on. This reward declines, but is sufficient enough to also offset any risk of under-performance relative to a buy and hold strategy for many years.

RUNE

- RUNE has two primary uses - 1) it needs to be staked in the liquidity pools, where it acts as the base pair and 2) it needs to be bonded as collateral by the network's nodes for security.
- Each liquidity pool must have a 1:1 ratio of external assets to RUNE. For every $1 of BTC pooled there must be $1 of RUNE also pooled. If this ratio breaks down, an arbitrage opportunity exists that brings it back to equilibrium. Liquidity providers earn 1/3rd of the 'System Income', which is comprised of trading fees and new RUNE issuance.
- Each node must bond a 2:1 ratio of RUNE to staked external assets. For every $1 of BTC pooled, there must be $2 of RUNE bonded by the nodes to secure it. Node operators earn 2/3rds of the System Income for running the network. If the ratio breaks down, and the value of staked assets becomes greater than the value of bonded assets securing it, then node operators receive a larger share of System Income until it is fixed. This temporarily removes the incentive for staking liquidity while boosting the incentive for nodes to bond more RUNE, bringing the ratio back to equilibrium.

THORChain Overview

If successfully launched, THORChain's permissionless, cross-chain liquidity pools would have strong product-market fit (e.g. BTC & USDT). In fact, it's hard to think of another use case that could generate as much interest in the near-term. The first project to deliver on this will usher in a paradigm shift for how digital assets are traded, away from centralized exchanges and towards a truly decentralized future. This would improve the resiliency of the entire space. Due to the inherent network effects of pooled liquidity, this project could find themselves in a winner-take-most situation. Our team typically focuses our research on applications and networks that are already live, since there have been too many failed promises in crypto. While the full implementation is expected to go live by Q1 2020, it's safest to assume it won't. As a result, we'll focus the first half of this report trying to figure out the only thing that really matters right now - will THORChain actually work?

Key Takeaways

- The development team leveraged codebases from other projects and combined them in a novel way to construct THORChain. While ambitious, this helps to mitigate execution risk. Key components of its architecture were built using Cosmos SDK, Tendermint and Binance Chain's TSS code library.
- As an application-specific blockchain, THORChain is optimized to best serve its specific use case (i.e. liquidity pools). The economic design has incentives in place to ensure the financial risk of the application is always proportional to the security of its network. If successful, we may see similar cross-chain designs take DeFi by storm.
- While using Tendermint provides THORChain with fork resistance and near-instant finality, the PoW chains it connects to will not have these characteristics. As a result, measures have been put in place to mitigate the negative impact that forks and double-spends could have on the liquidity pools.
- Trade mechanics aside, THORChain operates much like a centralized exchange but with one key difference - it decentralizes control over the exchange wallets. THORChain's validator nodes control the on-chain liquidity pools and much reach consensus to move funds. Anyone can operate a node as long as they bond the minimum amount of RUNE required (currently 1 million RUNE). THORNodes must run nodes for each external chain they "observe" in order to properly monitor and manage the pooled liquidity on those chains. Fast sync can be used to make this easier. The network will start accepting liquidity once 12 nodes are live, gradually increasing to 99+ nodes after.
- Staked liquidity is pooled on its native chain into addresses ("vaults"). There are 2 types of vaults - Asgard and Yggdrasil. Asgard receives inbound flow (new liquidity) and Yggdrasil handles outbound flow (executing trades). Pools maintain a single balance for each asset, with staked liquidity divided among, and regularly moved between, different vaults.
- A Threshold Signature Scheme (TSS) is used to unlock vaults and send outgoing transactions (i.e. execute a swap). This offers key advantages over MultiSig such as increased privacy and cheaper transactions. The nodes never hold their own private keys, instead they come together in an off-chain, trustless computation to generate a private key and sign transactions. TSS is chain-agnostic (adding support for new chains is easier) and already compatible with most major blockchains, including Bitcoin and Ethereum.
- The project has a large focus on privacy and utilizes an off-chain, minimalistic approach to governance. To summarize - 1) vaults are difficult to identify on-chain and funds are constantly moved, 2) nodes are anonymous and their coordination is purposely limited, 3) transactions have no identifying traits and 4) the team is pseudo-anonymous.
- Relative to other interoperability focused projects, THORChain differentiates itself by not using "pegged tokens" to move assets across chains. Real BTC can be swapped for real ETH, permissionlessly and with on-chain settlement. It also has a superior security model by ensuring that the value of RUNE bonded by nodes is always 2x the value of non-RUNE assets staked in liquidity pools. This is an improvement over Ren's approach where, hypothetically, $100m of pegged assets could be secured by only $20m of REN. Copying THORChain's approach would require changing their design and overall strategy, giving THORChain a head start on its closest competition.
- BEPSwap was initially set to happen on January 3rd, 2020 but was delayed pending the completion of 3 external audits on the code and econ.

THORCHAIN

An Interoperable, Application-Specific Chain

THORChain has undergone several design iterations since inception, leading to some confusion in the broader community regarding how its current approach will actually work. For example, the whitepapers from 2018, available on Github, are stale with outdated information. Due to this, we've been in communication with the development team to verify the accuracy of the information cited within this report.

Blockchain interoperability has been a goal many, higher-profile projects have spent years pursuing. With this in mind, how did THORChain's relatively unknown team develop a new solution to a problem that has proven technically challenging for so long? The short answer is that they didn't. Instead, they leveraged codebases developed by other projects and combined them in a novel way. This is a testament to open source development.

Let's begin with how the core of THORChain's infrastructure was developed. The team leveraged Cosmos SDK, for its intended purpose, and built an application-specific blockchain. The application, in this situation, being a liquidity pool. Up until recently, applications were usually deployed on top of blockchains with general-purpose VMs (e.g. Ethereum). This makes sense for a few reasons. It would have been impractical for a developer to build a new, custom blockchain from scratch each time they wanted to launch a new application. Even if they did, those blockchains would probably have been siloed off and not interoperable with other networks. Cosmos SDK changed this dynamic by providing developers with a framework they can use to easily create custom, interoperable blockchains. Sorry ATOM holders, as evidenced here, the resulting value doesn't have to accrue to you.

As an application-specific blockchain, THORChain can be optimized to best serve its use case. It's not constrained by limitations that it would otherwise face being built on a generalized platform. While security can be a challenge with this type of approach, THORChain has a viable path for attaining it. The incentives are designed in such a way that the financial risk of the application is always proportional to the security of its network. Would building on Ethereum be more secure? Yes, but does it actually need to be more secure? Time will tell but we may see similar cross-chain designs take DeFi by storm, if successful.

By using Cosmos SDK, THORChain is built out of composable modules (depicted by hexagons to the right), letting the development team plug-and-play different functionality that others have already built. One such module is the consensus engine (the green hexagon).

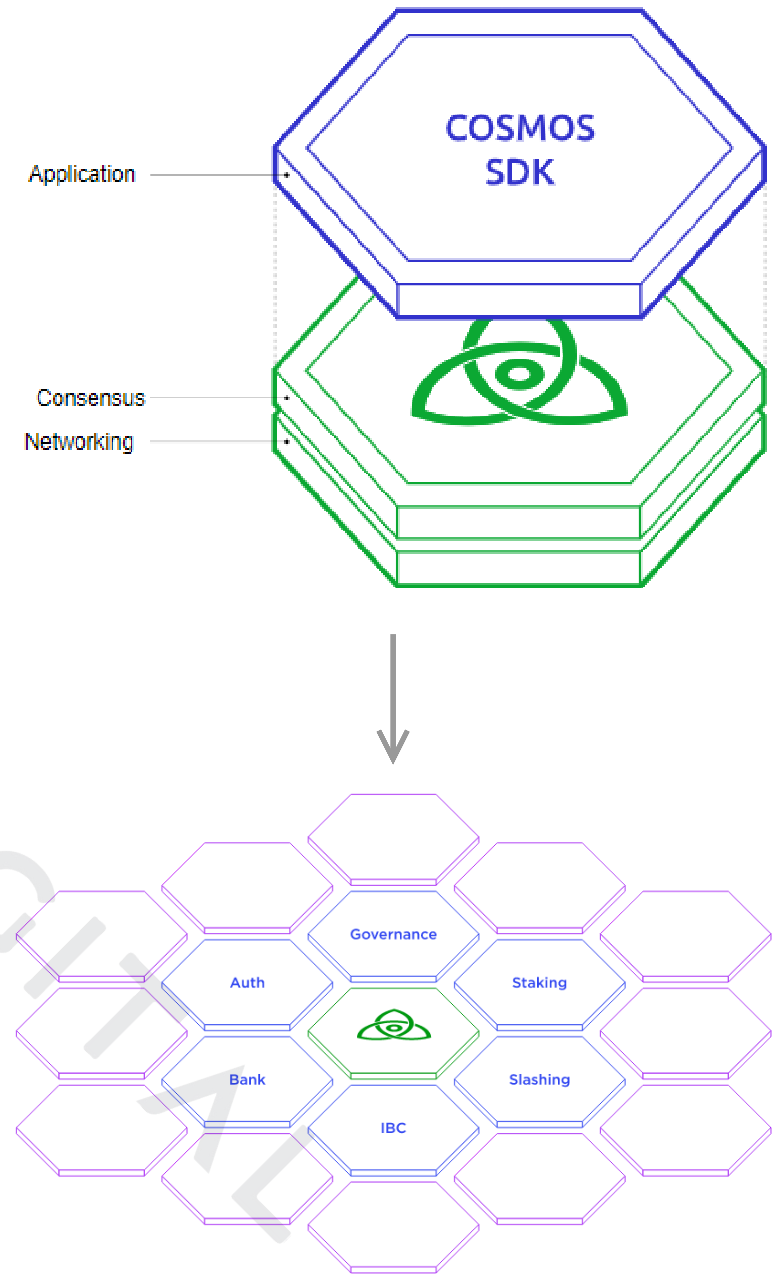

Consensus, Finality & Forks

THORChain uses Tendermint, the default consensus engine of Cosmos SDK. Initially created in 2014 by Jae Kwon. Tendermint has since become a leading Proof-of-Stake ("PoS") consensus mechanism in the space. As some bonus content, here's an old Reddit thread asking why Ethereum doesn't use Tendermint, with Vitalik and Jae both commenting. Tendermint provides certain key advantages over Proof-of-Work ("PoW") consensus, including fork resistance and near instant finality. While THORChain will benefit from these traits, it's important to emphasize that the PoW chains it connects to will not, resulting in additional attack vectors that need to be considered. There's a lot to unpack here so let's walk through this slowly with some examples to make sure you understand the concepts at play. We'll start by explaining how THORChain adds support for new external chains. Then, we'll elaborate on the challenges we just alluded to regarding PoW chains, and the steps that are taken to mitigate the relevant risks.

Let's use an example where THORChain's developers enable support for Bitcoin. In order for this new external chain to go live, more than 2/3rds of THORChain's nodes need to start observing it. After they do, consensus has been reached and the chain is now activated. Liquidity providers can start staking their BTC, and traders can start swapping for it. Staked assets, in this case BTC, are always pooled on their own native chains in "vaults" (addresses), which the nodes monitor and control. For a while there is peace in the Bitcoin world and everything runs smoothly. Then, out of nowhere, a big blocker appears convincing member's of the Bitcoin community that change has to happen, and it has to happen now. A contentious hard fork ensues, splitting the Bitcoin blockchain. For the THORChain nodes, this results in a confusing situation. They're supposed to monitor the BTC that's staked in the liquidity pools but now they don't know which chain is the correct one to track. If this occurs, the safest path forward is for the nodes to simply stop observing either, or both, chains. When more than 2/3rds of the validating nodes stop observing an external chain, "Ragnarok" is automatically initiated and the assets staked in the liquidity pools are safely returned to their owners (we explain Ragnarok in greater detail on page 13). In the future, support can be re-enabled for the chain if more than 2/3rds of the nodes start observing it again. This is a simple yet effective process that gives control to the community while at the same time protecting staked assets.

Earlier, we stated that Tendermint provides "near instant finality". What does that mean? Finality is the assurance that once a transaction has been completed, it can not be undone. Unfortunately, with PoW chains even though a transaction might have been confirmed in a block, it can still be reversed. Although difficult to pull off, a 51% attack could occur resulting in a double-spend. As more blocks get added to the chain though, it becomes increasingly hard for an attacker to revert an old transaction. This is why PoW chains are described as having probabilistic finality. That old transaction is "probably" final. Mitigating the risk of a double-spend impacts both stakers and traders on THORChain. Stakers withdrawing their liquidity from a PoW chain will need to wait 24 hours for their funds to unlock. For traders, however, its impractical to assume they would wait that long and for small transactions its completely unnecessary. As a result, the developers created an excellent framework that calculates the number of confirmations a trader needs to wait. At a high level, this is based on the size of a trade relative to that network's security. For example, using this method a $25k BTC transaction would only need to wait 1 confirmation due to the strength of Bitcoin's security. In addition, because THORChain's liquidity pools use an automated market maker model ("AMM"), smaller trades usually fair better than larger ones in terms of slippage. This is also good for security purposes since no rational actor would ever attempt such an expensive and difficult attack just so they could double-spend a small trade size.

Bonded Validator Nodes

DELPHI DIGITAL

While we briefly mentioned nodes on the previous slide, it's a topic worth discussing more thoroughly. Before we dive in though, let's take a step back and consider how THORChain's approach compares to what centralized exchanges ("CEX") already do. With a CEX, traders send an on-chain transaction to an on-chain exchange wallet. The CEX manages this on-chain exchange wallet, internally accounts for trading and then, when a trader wants a withdrawal, the CEX pays them with an out-going transaction from an exchange wallet. THORChain's design is similar but with a few key differences, particularly who controls the "exchange wallet". To start, traders always maintain custody of their assets and trades settle on-chain (this is a DEX after all). Only the liquidity providers ("LPs") need to move and stake their assets. However, instead of trusting a centralized third-party with these assets, the LPs trust a distributed network of collateralized nodes, economically incentivized to protect their staked liquidity. At its core the difference is really that simple. THORChain just decentralizes who has control over the liquidity. In the future, THORChain may also have specialized validating nodes that enable people to stake liquidity in a non-pooled, self-sovereign way. We believe this would have significant appeal to the Bitcoin community if successfully launched.

Anyone can run a THORNode as long as they bond the minimum amount of RUNE required. This is currently set at 1 million RUNE. THORNodes must run nodes for each external chain they "observe" in order to directly monitor the pooled liquidity on those chains. These nodes are managed by the THORNode. When you consider the storage and bandwidth required to run nodes for Bitcoin, Ethereum, etc. all at once, it's clear this can increase the operational cost of running a THORNode, beyond just the bonding requirement. However, this can be mitigated by fast syncing external nodes that have old, unnecessary in this case, data pruned from them. As we'll see on the following slides, there are other limitations imposed by consensus (Tendermint) and the signature scheme (TSS) used that will probably limit the total number of nodes to less than 100 over the near-term. Based on the RUNE currently staked on RuneVault, 25 of the 781 stakers have enough RUNE to run a node if the network launched today.

The launch of BEPSwap (ChaosNet), the first Mainnet deployment, will start accepting liquidity once 12 public nodes are live (we discuss this more on page 15). Pending a successful ChaosNet, the full implementation will go live, with the expectation that the amount of nodes will gradually increase from 12 to 33 over the course of a month. From there, the goal is to reach 99 nodes in the near-term. Are these enough nodes to be sufficiently "decentralized"? Regardless, it is ultimately more decentralized and secure than trusting a single exchange. That's what really matters.

Staked Liquidity Vaults

Anyone can stake liquidity without permission, although the network must reach consensus to sign and approve liquidity withdrawals. The latter point is no different than Ethereum needing to reach consensus before releasing liquidity pooled in Uniswap. Assets staked as liquidity are pooled on their native chains into addresses referred to as "vaults". For reasons we'll discuss on the next slide, these vault addresses are indistinguishable from regular addresses. It's important to understand that every external chain has 2 types of vaults which serve different purposes - Asgard and Yggdrasil. At a high level, Asgard receives all of the inbound flow (new liquidity) and Yggdrasil handles all of the outbound flow (executing trades). We'll now explain each type more in depth.

Each chain has at least 1 Asgard vault, which receives the staked liquidity for that respective chain. An Asgard vault is managed by a set of validating nodes that have the power to move funds out of the vault when they reach consensus. Ideally, in the future the network's entire set of active validating nodes (99+) will be part of each chain's sole Asgard. However, due to current limitations related to the signature scheme (TSS) and consensus mechanism (Tendermint), the network will likely have multiple Asgard vaults ("Multi-realm Asgard") on each chain in the near-term. Instead of a single Asgard with 66 of 99 nodes necessary to reach consensus, there will be three Asgards each requiring 22 of 33 nodes for consensus.

After Asgard receives new liquidity, the assets are then split up and moved into smaller Yggdrasil vaults that have distinct on-chain addresses. To calculate the number of Yggdrasil vaults, you take the size of the Asgard validator set and divide it by 3. For example, an Asgard vault with 33 validators would have 11 Yggdrasil vaults associated with it. Each of those Yggdrasil vaults would have 3 validators, requiring 2 of the 3 to reach consensus and sign a transaction. The smaller validator sets allow outgoing transactions (trades) to be executed much faster by the network. For example, signing speeds for a 22/33 Asgard take several minutes, whereas a 2/3 Yggdrasil vault can sign in less than a second. Despite the smaller validator set, attacking this would be difficult based on how trades are delegated and the incentive structure in place. Conceptually, Yggdrasil can be thought of as an insured hot wallet given that its less secure but more performant. Similarly, Asgard is like a cold wallet, being more secure but less performant. THORChain maintains a single balance for each asset, with its staked liquidity divided among, and regularly moved between, different vaults. A validating node is part of 1 Asgard and 1 Yggdrasil for each external chain. While we're on this topic, it's important to highlight that RUNE will be pooled on Binance Chain (that's where the token is issued). Due to Binance Chain's centralization, there may be risk of censorship and other attacks since a successful THORChain could be an existential threat to Binance itself. There's no precedent for this, however, the team is monitoring it and is prepared to move RUNE to another chain, possibly even THORChain itself, in the future.

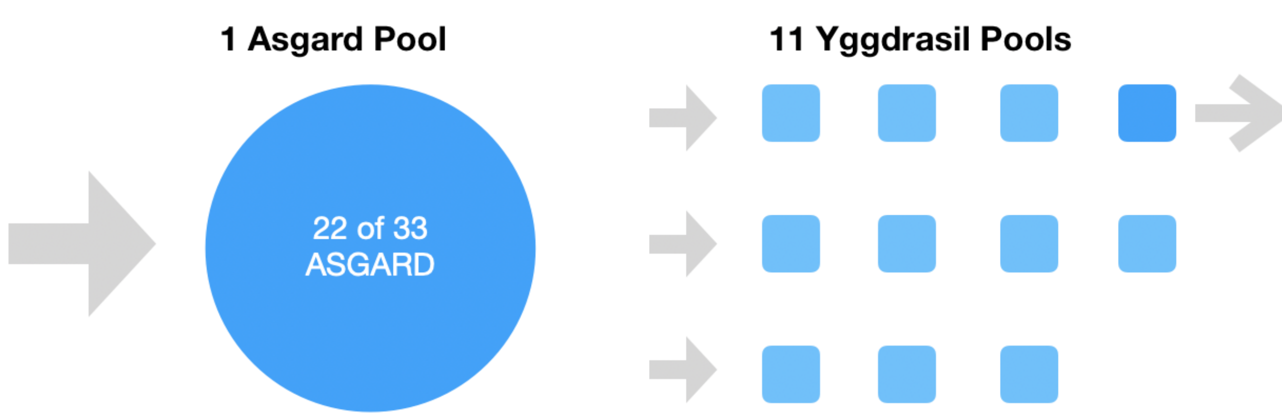

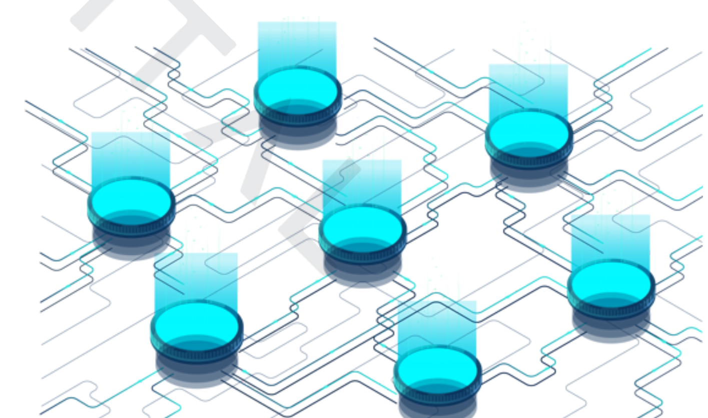

Threshold Signature Scheme vs MultiSig

Now that we've provided an overview of THORChain's general architecture, we can dive into a technical component that plays a critical role in how it all comes together. Let's start by doing a quick refresher on some blockchain basics. In order to move assets out of an address (or spend a UTXO), a person needs to have the private key associated with that address (or UTXO) to sign a valid transaction. At a base level, all THORChain really does when it executes a swap is - 1) comes to consensus off-chain about the swap and 2) signs a transaction to move assets out of the vault. That transaction pays the trader in the asset they just bought, on its native chain.

The concept of multiple parties coming together to sign a transaction might remind you of how MultiSig addresses work. That intuitively makes sense, however, THORChain is actually using a new type of Threshold Signature Scheme ("TSS"), which has some key advantages. While TSS and MultiSig both require multiple parties (M-of-N) to a sign a transaction, that's where the similarities end. With MultiSig, each of the parties has their own private key that corresponds to a different lock. For example, if there were 5 parties (N=5), there would also be 5 locks for each key. It may require at least 3 of the parties (M=3) to use their private key to sign an outgoing transaction. These private keys are static, stored locally by each party and the signing process happens entirely on-chain. It's also very easy to identify a MultiSig address on-chain and observe the signers. Using Bitcoin MultiSigs to illustrate this point, in the table to the right you can clearly see how much BTC is held in each type of MultiSig*. MultiSig transactions also have higher fees (more data per signature).

With TSS, multiple parties (THORNodes in this case) are still required to come together and sign a transaction (unlocking the vault) but they do so using a single private key. How? To start, each node has a piece of information they store locally and keep private, never revealing it to the others. In a trustless, interactive and off-chain computation, the nodes combine their secret pieces of information to generate a single private key and sign a transaction. The private key is never revealed throughout this process. Not only is the vault address difficult to identify, because it looks like any other address with a single lock, but the signers (nodes) also remain private. External watchers can't identify who took part in the key generation and signing process. Compared to MultiSig, this method offers improved privacy and lower transaction fees (less data with 1 signature). In addition, TSS is chain-agnostic, allowing THORChain to support new external chains with less development effort.

MultiSig Vault

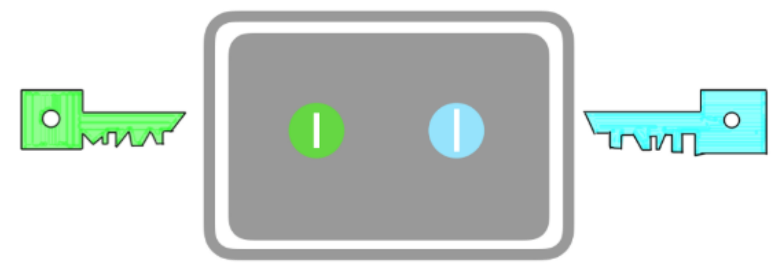

Threshold Signature Vault

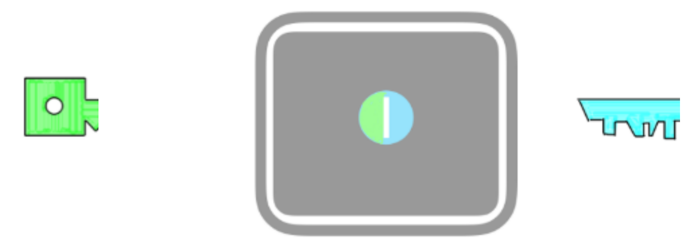

Image modified; Source: ZenGo

BTC in P2SH Addresses by Type

Type of P2SH Address	Amount of BTC	Value of BTC
UNSPENT	3,527,000	$26047 M
P2WPKH	1,018,000	$7518 M
2 OF 3	314,000	$2319 M
OTHER MULTISIGS	290,000	$2142 M
3 OF 4	264,000	$1950 M
3 OF 6	77,000	$569 M
3 OF 5	72,000	$532 M
2 OF 6	69,000	$510 M
sw 2 OF 3	63,000	$465 M
2 OF 2	58,000	$428 M
2 OF 4	16,000	$118 M
OTHER NON-MULTI	10,000	$74 M
3 OF 3	852	$6 M
TOTAL	**5,778,852**	**$42677 M**

Data as of: December 29th, 2019
Source: txstats.com,

**The UNSPENT category represents MultiSigs that haven't spent an output yet, which would reveal what type (M-of-N) of lock it is*

Key Generation & Validator Churn

DELPHI DIGITAL

Basically, THORChain uses TSS to communicate cross-chain instead of relying on IBC (although it may use it in the future). THORChain is implementing Binance Chain's TSS code library, which is based on a paper published in 2018 by Rosario Gennaro and Steven Goldfeder. Their approach improved upon previous threshold ECDSA designs by being faster and by finally removing the need for a trusted dealer to distribute the pieces of secret information used in the key generation process. Hence the name of their paper - "Fast Multiparty Threshold ECDSA with Fast Trustless Setup". Binance recently open sourced their TSS library in November 2019.

Even if TSS was a game changer in certain regards, if it's not compatible with important external networks then THORChain wouldn't stand to benefit from implementing it. Fortunately, because it generates ECDSA signatures, it works with most cryptographic curves used by prominent blockchains, particularly secp256k1. This is important because Bitcoin and Ethereum use secp256k1, meaning that TSS already works for both chains right now. As a result, THORChain is not dependent on yet-to-be-invented technology or waiting on an external network to make an update outside of its control. However, we should emphasize that using new cryptographic protocols are higher risk due to the fact the code is less battled tested in a live, adversarial environment. To help mitigate this risk, Binance Chain's TSS code library underwent a third-party audit in October 2019 by Kudelski Security. 10 security issues, all medium to low severity, were identified and patched. In addition, the auditors stated "*None of the issues found in the frame of this audit could be exploited per se to completely break the security of the scheme, or recover secret data*". The THORChain team also has its own third-party TSS audit ongoing.

Earlier in this report, we discussed THORNodes but didn't explain how they're cycled in and out of the active validator set. Every 3 days (50,000 blocks), THORChain automatically replaces the 2 oldest, most unreliable validators with 2 new ones. The new validators are selected based on their amount of bonded RUNE. This regularly scheduled rotation of validators is done in an effort to help resist network capture and stagnation, but there's more to this process than just that. Effectively, the churn event represents a "reset" every time it takes place. The new set of validators come together to generate new keys, for new vault addresses. Once completed, all of the staked assets that were held in the old vaults are moved to the new vaults. However, there are situations where this may not run smoothly, negatively impacting liquidity providers. If, during this event, 1/3rd of the validators leave before being churned out, and they delete necessary information they're storing forever, it could lead to a situation where assets in a vault become locked up for good. While clearly concerning, one comforting fact is that the validators would also lose their own collateral. As a result, this type of "attack" would probably not be done on purpose by rational, profit-seeking nodes. Irrational actors, such as nation states, are of course a risk but the same could be said for any other crypto-economic system as well.

As we alluded to on page 9, the more active validators there are, the more difficult it becomes to have a single Asgard for each chain. As more validators get added to the key generation, the time it takes to complete the process grows quadratically. If this takes longer than 10 minutes, the validator set automatically shards into smaller subsets, resulting in "Multi-realm Asgard". This speeds up the process significantly. While we've grouped these together so far, it's important to specify that TSS sets the upper limit for how big an Asgard can be (22/33) and Tendermint sets the upper limit for how many active nodes there can be (currently ~100).

Anonymity & Governance

It's important to understand that THORChain's primary value proposition is permissionless, private, cross-chain exchange. Despite the recent emergence of DEXs, trading assets cross-chain has always been, and still is, conducted through centralized exchanges (CEXs) due to a lack of interoperability between blockchains. It should come as no surprise that, based on the global regulatory environment, CEXs are ramping up KYC procedures, delisting assets and geo-blocking users from certain countries with increased frequency. This may reduce their profitability but it will likely increase their longevity. THORChain could be well positioned to fill the void left by CEXs. Our team, of course, **does not** support using THORChain to circumvent laws, just as Bitcoin shouldn't be used to that end. Individuals should follow all laws that are applicable to them and potential node operators in particular should take caution.

Earlier, we explained why vaults are difficult to identify on-chain by external watchers and how they're frequently moved. Even if you could isolate one, there's no way of knowing which validators signed which transactions. Remember, a TSS transaction looks like any other transaction signed by a single key. There's no way to determine which nodes participated. Even the nodes themselves do not keep track of this. The focus on privacy goes much further though. Nodes are anonymous and their ability to coordinate is purposely limited by the network's governance model. Even the development team has focused on remaining pseudo-anonymous, which makes sense in the broader context (Satoshi, anyone?). In summation, a pseudo-anonymous team is launching a decentralized network of anonymous nodes, which facilitate permissionless, non-custodial cross-chain trading. If you could distill the ethos of this space, that would be it.

THORChain will utilize an off-chain, minimalistic approach to governance. Initially, the team planned to implement on-chain governance through a mechanism they were developing called AEsir Protocol. However, they later became disillusioned with this idea due to low participation on other chains, the possibility of voters social signalling, fear of manipulation and regulatory risk. These are all valid concerns. Nodes can only generate valid blocks if the software they run is fully-compliant with that of the super-majority. On the previous slide, we mentioned how the churn event acts as a reset for the network. This also pertains to when nodes can change the software they run. Once a node joins the current validator set, there is no way for it to make any changes to its software, without being churned back out first. This means that the churn event is also when new software upgrades can take place. In order to change the system, an off-chain improvement proposal ("TIP") needs to be submitted for testing, validation and discussion among the developer community. Updates must be backwards compatible. If the update proves beneficial, it can be merged into the node software, allowing new nodes to run it and signal their support. Once the super-majority are running the same software, the system auto-updates.

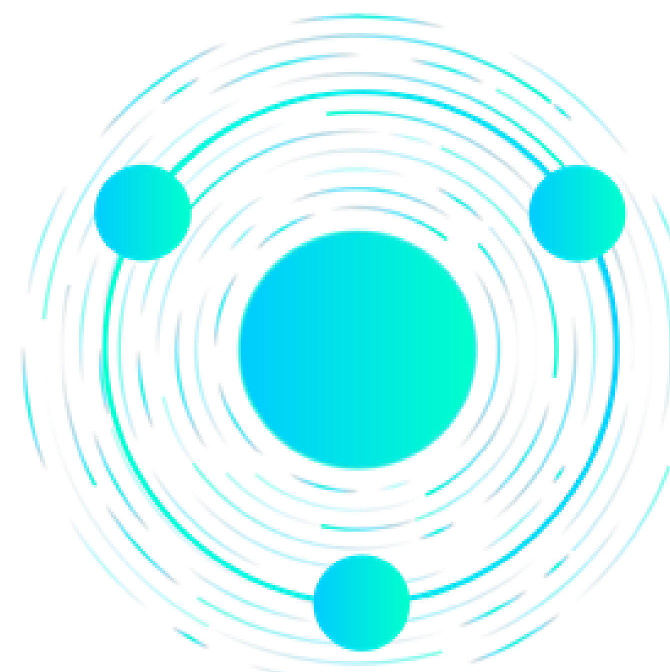

Security & Capital at Risk

What happens if there's an emergency and an important update can't be coordinated fast enough? Or what happens if most of the nodes fall offline while there is still liquidity staked in the system? The answer to both these questions is simple - invoke Ragnarok. This Norse term for the end and rebirth of the world is a fitting name for THORChain's shutdown procedure. Once a minimum threshold of active nodes is hit (the current number is 4), the system automatically shuts down and all of the staked assets are refunded back to the liquidity providers ("LPs"). The network can be restarted later on. THORChain prioritizes the security of staked capital since liquidity providers are ultimately the ones driving value to the network. If THORChain isn't safe for stakers it doesn't matter what the yield is. No one will provide liquidity if they think their assets will be lost or stolen. This is where THORChain's economic design really excels. The value of bonded capital locked up by nodes must always be equal to, or greater than, the value of staked capital provided by LPs. If this ratio breaks down, and the value of staked assets becomes greater than the value of bonded assets securing it, then node operators receive all of the 'System Income' until it is fixed. This temporarily removes the incentive for staking liquidity while boosting the incentivize for nodes to bond more RUNE. This self stabilizing mechanism should bring the ratio back to equilibrium. In effect, the value of the assets securing the network should always grow in tandem with its liquidity. As we'll discuss on the next slide, this is an important feature others lack.

Based on the rules of its consensus mechanism (Tendermint), if more than 1/3rd of the nodes act maliciously the network grinds to a halt but staked assets remain safe. In fact, the capital bonded by the malicious nodes remains locked and is now at risk of being slashed as a penalty. However, if more than 2/3rds of the nodes act maliciously they can steal staked assets from the vaults and their bonded capital probably wouldn't be slashed, since they're now the network's "source of truth". As seen below though, THORChain's economic design acts as a strong deterrent for rational, profit-driven actors. While we show two different scenarios, the column farthest to the right is the only realistic outcome. Sure, the malicious nodes won't have their bonded RUNE slashed, but it doesn't have to be. Such an attack would probably drive the price of RUNE to $0 on its own. As seen in the first line below, for every $66 of capital the nodes have at risk, at most they can only steal $50 of pooled, non-RUNE assets. Assuming their bonded RUNE is now worthless, due to their attack, they walk away with a solid $16 loss. However, the caveat is that this assumes the internal (pool) price of staked assets matches their external price. There may be edge cases where external prices are different and this dynamic breaks down.

Bonded Capital

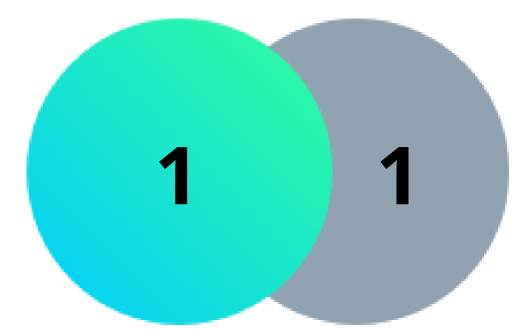

Staked Capital

CaR = Capital at Risk

					RUNE Price Unaffected	RUNE Price Goes to $0
% of Malicious Nodes	Value of Bonded CaR (RUNE)	Value of Staking CaR (Non-RUNE)	Value of Staking CaR (RUNE)	Value of Total Staked CaR	Total Capital Post-Attack	Potential Profit (Non-RUNE)
66%	$66	$50	$50	$100	$166	-$16
75%	$75	$50	$50	$100	$175	-$25
90%	$90	$50	$50	$100	$190	-$40
100%	$100	$50	$50	$100	$200	-$50

Competitive Advantage

DELPHI DIGITAL

Being able to provide permissionless, cross-chain liquidity is clearly an important use case with clear market fit today. With that being said, it's only logical that there would be several competing projects working to deliver it in some form or another. A few worth noting are Ren, Cosmos and Bancor. Let's focus on Ren, a project our team first wrote about in February 2019, although its focus has changed since. Ren and THORChain share some aspects of their design in common. For example, they both utilize Tendermint for consensus and require that their token be bonded by validators to secure their respective networks. Oddly enough, they both also named their tests "Chaosnet". Speaking of Ren's Chaosnet, they currently have a live DEX deployed on it that supports swapping between DAI, BTC, BCH and ZEC. Personally, we've never been able to execute a trade through it but we'll give that a pass for now since its really just a proof of concept and not a product the team is actually focused on building out. Ren doesn't want to launch its own cross-chain exchange. Instead, it wants to be a piece of infrastructure (RenVM) that existing DeFi applications can integrate to bring cross-chain assets to their apps. Rather than competing with Uniswap, they want to be the tool it uses to bring BTC to its liquidity pools (yes, we're aware there is already a wBTC pool). This is where the nuance of their design differences come into play. THORChain describes itself as a "1-way state peg", while Ren, Cosmos and Bancor are "2-way state pegs".

We'll use BTC and ETH in our example to explain the differences between the two. Ren's 2-way peg involves locking BTC on its native chain and minting an equivalent amount of pegged tokens on Ethereum. Ren calls this pegged token 'zBTC' and it's an ERC-20 representation of BTC. Afterward, these pegged tokens must be burned in order to unlock the BTC backing them. In comparison, THORChain doesn't use pegged tokens, it just determines how to move assets on their own chains. As a trader, I can swap real BTC for real ETH without either asset ever "leaving" its chain. This isn't to discredit the potential usefulness of pegged tokens but they do have unique risks. For example, despite zBTC being backed by BTC 1:1, its price can diverge from it. Why? Simply put, holding zBTC has more risk than holding BTC. If I'm a trader who swaps for zBTC to use in DeFi, as long as I hold that asset I'm at risk of Ren's cross-chain bridge becoming insecure and the peg breaking. In that situation, the zBTC I own could become worthless.

This begs the question, how will Ren's security compare to THORChain's once its fully live? To put it bluntly, it has a few disadvantages. To start, it doesn't have the same economic model we outlined on the previous page. Hypothetically, it could be responsible for securing $100m of pegged assets but only require $20m of REN to attack. THORChain, fundamentally, can not have this type of imbalance. Copying this mechanism would require Ren to rework parts of its design (i.e. on-chain oracle) and rethink their overall strategy. THORChain's high issuance rate in the near-term is also a strong incentive that can help bootstrap its liquidity, leading to a positive feedback loop for growth. Ren doesn't have this at its disposal. In addition, because Ren is infrastructure for other projects to use, it's adoption is reliant on external parties. It may have a difficult time generating enough cross-chain transaction fees to keep its network secure. Ren's team is already considering a proposal that would add an annual fee of 1% for pegged assets, on top of transaction fees. While pegged tokens are always used in some capacity, based on Ren's commentary from November 2019, it appears that interacting with them can be abstracted away for a trader. This is good, since we believe people are more interested in permissionlessly trading assets directly across chains, than they are using pegged tokens, at least for now. Liquidity has strong network effects and the first project to reach a critical mass could find themselves in a winner-take-most situation.

Ren　　COSMOS　　Bancor

RuneVault, BEPSwap & ASGARDEX

For THORChain to successfully launch it needs to be technically sound, have the right incentives in place to bootstrap liquidity and attract interest from the broader community. Hopefully, by this point in the report you have a good understanding for how THORChain works at a technical level. Before we analyze the economic incentives in the next section, let's discuss how THORChain leveraged its token to generate interest and how it all feeds into their launch strategy.

Most people currently familiar with the project have probably heard of RuneVault, which allows holders to stake their RUNE and earn a weekly yield of 1%, derived from new issuance. This relatively high yield has attracted 781 stakers who in turn have locked up ~81% of the circulating supply. But if the network hasn't launched yet, why exactly does it need stakers right now? The answer is that it doesn't. While staking on RuneVault freezes a user's RUNE, they do not incur any additional risk. RuneVault is simply an incentivized way for the community to learn THORChain's staking mechanics, while staying interested through a frequent, weekly reward. It also helps disperse RUNE's supply. Ideally, their plan is that once the network goes live, RuneVault stakers will stay around to provide liquidity and run nodes. This strategy is far superior to what is commonly seen, where a project issues a large amount of tokens from the start, which then get dumped on the market rather than being used to spur adoption.

RuneVault will come to an end once BEPSwap launches. This was set to take place on January 3rd, 2020 but it was delayed pending the completion of 3 separate external audits, which the team expects will take 6 weeks. BEPSwap's testnet can be viewed here. Once live, BEPSwap will be THORChain's first mainnet exchange, allowing people to trade and pool liquidity for BEP2 assets on Binance Chain. Why start with Binance Chain? After all, doesn't that chain already have a native DEX built in? Well, as we mentioned on page 7, PoW blockchains necessitate taking additional risk factors into account given their lack of finality. Binance Chain, just like THORChain, was built using Cosmos SDK and leverages Tendermint for consensus, providing it with near instant finality. This reduction in complexity makes connecting to Binance Chain an ideal first step in their roll-out. BEPSwap will mirror the full implementation as much as possible and has a few key objectives - 1) capture real trading data from the liquidity pools, 2) test public nodes, 3) rehearse churn/Ragnarok and 4) involve the community. BEPSwap has a hard-coded Ragnarok scheduled to take place after 6 weeks. Once BEPSwap winds down, and assuming it ran smoothly, ASGARDEX, the fully implemented cross-chain exchange, is expected to launch shortly after in the first half of 2020.

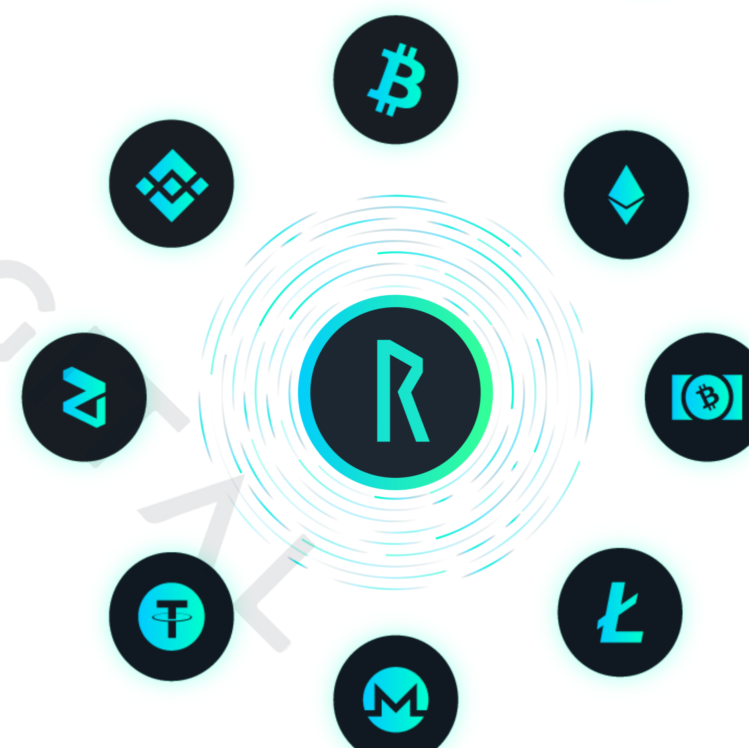

CLP & RUNE Economic Overview

This section will take a deep dive into how the RUNE token can accrue value by breaking down the utility it provides to various users. We will also analyze the mechanics of its unique liquidity pool design and how it matches up against competitors. From a high level, it addresses the needs of two of the largest existing and future audiences/use cases, non-productive assets and trading, by making them provide value to each other. Non-productive assets refers to those that are used as a store of value or medium of exchange. It's not intended to be a slight, but is used to categorize their mostly idle functionality within their ecosystem. Creating utility for otherwise idle assets by enabling them to provide liquidity in return for a fee can be useful for an incredibly large audience, if done correctly. On the other side of the table, you're allowing users to tap that liquidity to trade anonymously between any assets that exist in any pool, in a non-custodial manner.

Key Takeaways

- THORChain uses a unique Automated Market Maker (AMM) called a Continuous Liquidity Pool (CLP) to facilitate how assets move in and out of these pools. It shares major characteristics with other AMM's, like slippage increasing as transaction size increases relative to the size of the pool and the potential for under-performance relative to a buy and hold portfolio when price begin to deviate from their starting point.
- A unique aspect that the CLP model enables is reducing that under-performance, particularly compared to Uniswap's XYK Constant Product model, which we break down in further detail.
- Another unique aspect is the Slip Based Liquidity Fee. This mechanism increases the fee that Liquidity Providers receive as transactions increase in size relative to the pool. It's an effective way to incentivize liquidity providers (supply) to actively search for pools where transaction size and volume (demand) suggests there's a need for more liquidity. While increasing the cost for users who are less cost conscious, it also enables smaller transactions to take place at costs lower than most centralized exchanges. It's a feature that a fixed rate AMM is unable to offer.
- On the liquidity side, each pool is made up of RUNE and another asset. For security, nodes bond RUNE in order to validate the network. The total value of nodes is economically incentivized (we expand on this) to be at least as large as the combined value of all the pools. This means the value of bonded and staked RUNE will be 3x the value of the non-RUNE assets being staked to provide liquidity. The value of RUNE grows directly with the increase in staked assets.
- The token is uniquely used to bootstrap liquidity by appropriately incentivizing both Liquidity Providers and Node Operators early on with issuance. This reward declines, but is sufficient enough to offset any risk of under-performance relative to a buy and hold strategy for many years. Total Supply is fixed, there's a pool of tokens set aside to support this.

Continuous Liquidity Pools

THORChain employs an Automated Market Making (AMM) mechanism called a Continuous Liquidity Pool (CLP) to facilitate trading. An Automated Market Maker is a rules-based trading program that responds to changes in different variables based on those rules. The individual rules that govern AMM's are what distinguish them from each other. Although the rules can be very complex with certain mechanisms, the fact that all actions follow a set of rules creates predictability and stability.

CLP's share some characteristics with the XYK Constant Product AMM used by Uniswap, however, there are some notable differences as well. It's worth explaining the mechanics of a CLP before addressing the differences between THORChain's CLP and Uniswap's XYK. Continuous Liquidity Pools are made up of two assets, with one of them always being the RUNE Token. We'll use BTC as the other asset for the purpose of this example. When a trader interacts with the pool, they can deposit (sell) RUNE and withdraw (buy) BTC or, conversely, sell BTC and withdraw RUNE. In reality, there will be many pools, each with a different asset staked alongside RUNE. This enables cross-pool transactions where you'll be able to leverage an ETH / RUNE pool to buy and sell BTC for ETH. As new pools get added, the underlying utility of existing ones will also increase as it will enable an increasing number of cross pair combinations. This will also make RUNE a highly liquid asset.

Relevant Formulas

(We'll Walk Through These Over The Next Few Pages)

Unit	Definition
x	Input
y	Output
X	Balance of TKN in the input side of the pool
Y	Balance of TKN in the output side of the pool

$$x = \frac{\left(\frac{XY}{y} - 2X\right) \pm \sqrt{\left(\frac{XY}{y} - 2X\right)^2 - 4X^2}}{2}$$

$$y = \frac{xYX}{(x+X)^2} \qquad Liquidity\ Fee = \frac{x^2 Y}{(x+X)^2}$$

Price Slippage

As with just about all liquidity pools, slippage is a function of transaction size relative to liquidity pool depth. The larger your transaction, the more slippage you'll experience. Total slippage is measured as the decline in what you received relative to what you should've received assuming the full transaction occurred at the prevailing price. Transaction output, which dictates overall slippage, is derived by the formula below.

$$y = \frac{xYX}{(x+X)^2}$$

It's clear that slippage increases significantly when the transaction represents a meaningful portion of the liquidity pool. This shows why it'll be important for these pools to grow in size in order for the system to gain traction with traders. It also shows how incentives are naturally aligned. Large pools enable larger transaction capacity and more overall trading volume. This leads to growth in fees and yield to liquidity providers, incentivizing people to further add liquidity to these pools. The chart on the bottom right will be initially relevant as pools first begin to grow, but realistically, the chart on the top will be more applicable when trying to gauge overall transaction costs. For perspective, average 2019 transaction sizes for BTC USD on Coinbase Pro and BTC USDT on Binance were ~$1,350 and $900, respectively. Using the midpoint of maker and taker fees on each exchange, a THORChain pool would allow you to trade the average transaction size more cheaply than the aforementioned exchanges before one side of the pool passes $2 million.

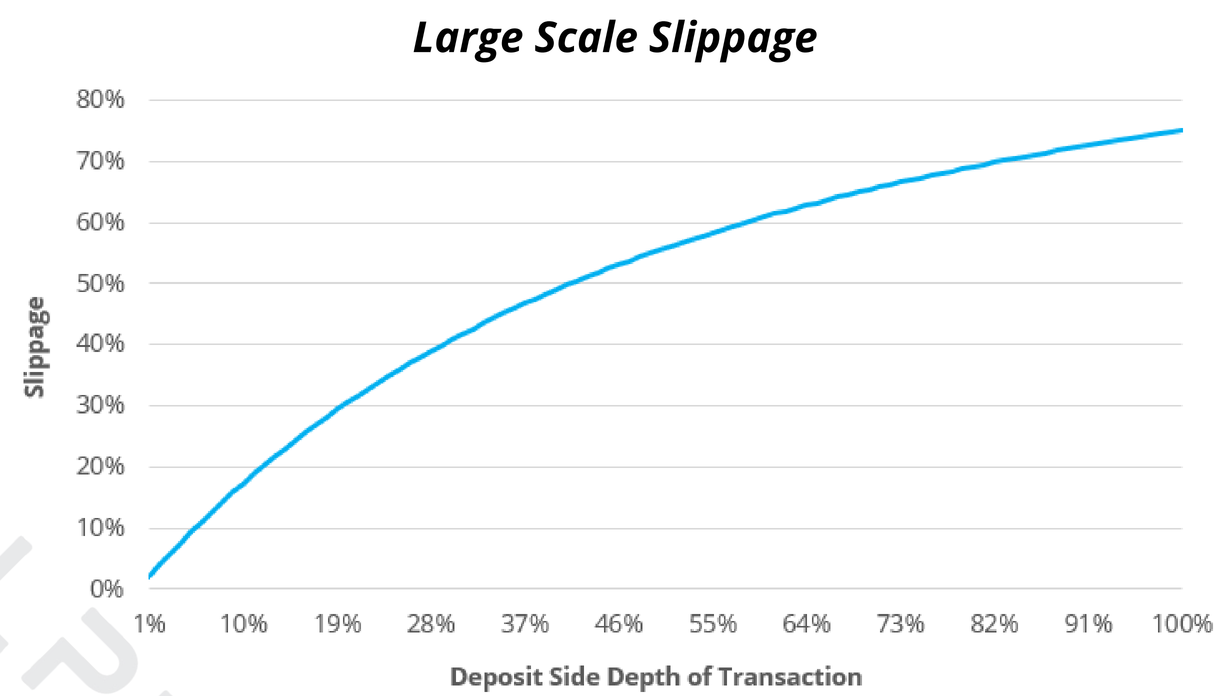

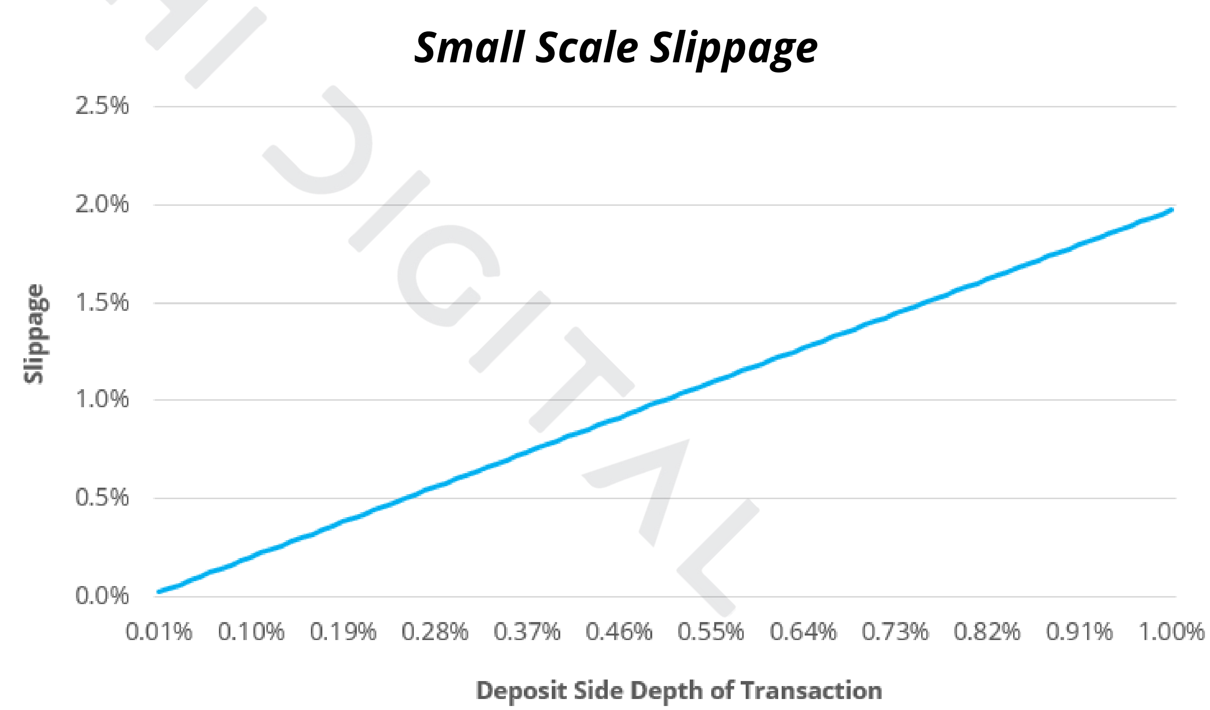

Trading Premium

It was logical to calculate slippage in the previous slide because the scenario was one where the trader knew how much he wanted to sell, but was still looking to understand how much less he was going to receive. It's also useful to think of it in terms of premium paid, particularly when you know how much you'd like to receive, but need to figure out how much extra it'll cost on top of the going rate. In order to do that, we just rearranged the equation in the previous example to instead solve for x [input] based on a known [output] y *(as seen below)*.

Much like with slippage, excessively large premiums should only be a concern in the very beginning when pools are small. The premiums quickly shrink as pool sizes grow and the size of individual transactions decline relative to the pools. The CLP model is also structured to have a limit on any individual transaction. If you're thinking about it in terms of the previous example, the single purchase maximum size is met when your input to the pool is the same size as the entire depth of the input asset in the pool. This would set the output you receive to 25% of the output side of the pool, or containing 75% of slippage. Applied to this example, you'd have to pay a 300% premium to purchase 25% of the pool. If you were to increase your input beyond 100% of what's currently in the input side of the pool, the output you receive would actually begin to decline. This is a function of the original CLP formula where, when calculating the output, the denominator begins to decline at a slower rate than the numerator as your bid size grows beyond 100%, leading to output decline. More simply, if you're buying BTC with Rune, the most BTC you can buy in one transaction is 25% of what's in the pool, and it would require you to effectively double the amount of Rune in the pool.

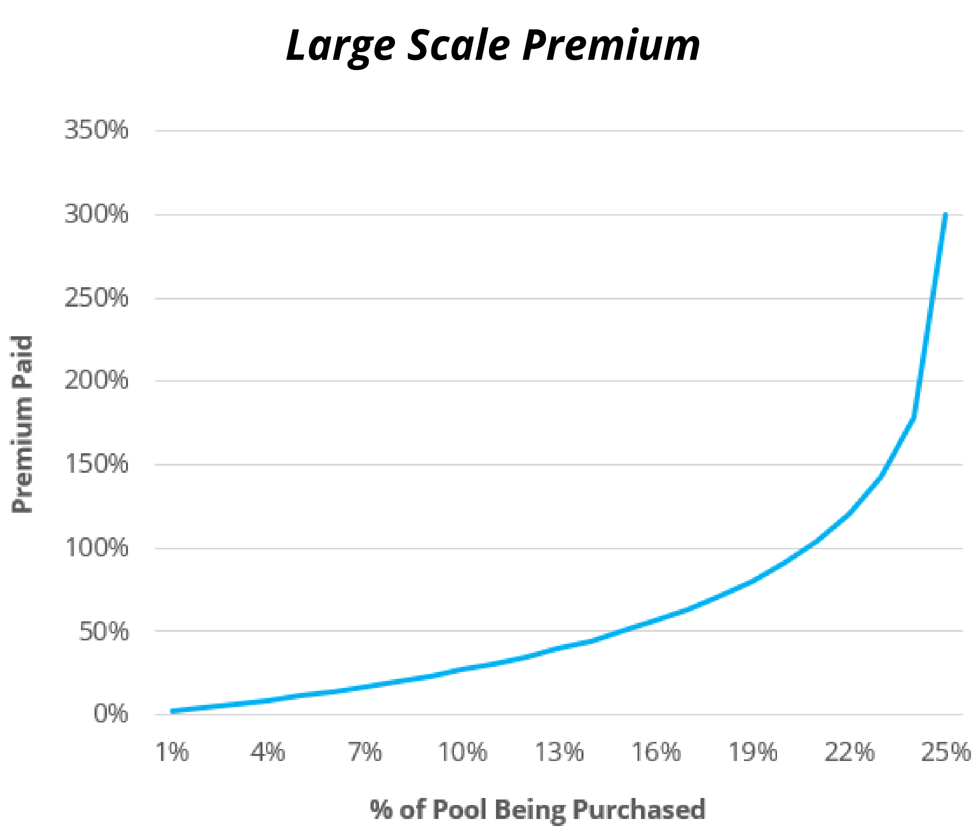

$$x = \frac{\left(\frac{XY}{y} - 2X\right) \pm \sqrt{\left(\frac{XY}{y} - 2X\right)^2 - 4X^2}}{2}$$

Unit	Definition
x	Input
y	Output
X	Balance of TKN in the input side of the pool
Y	Balance of TKN in the output side of the pool

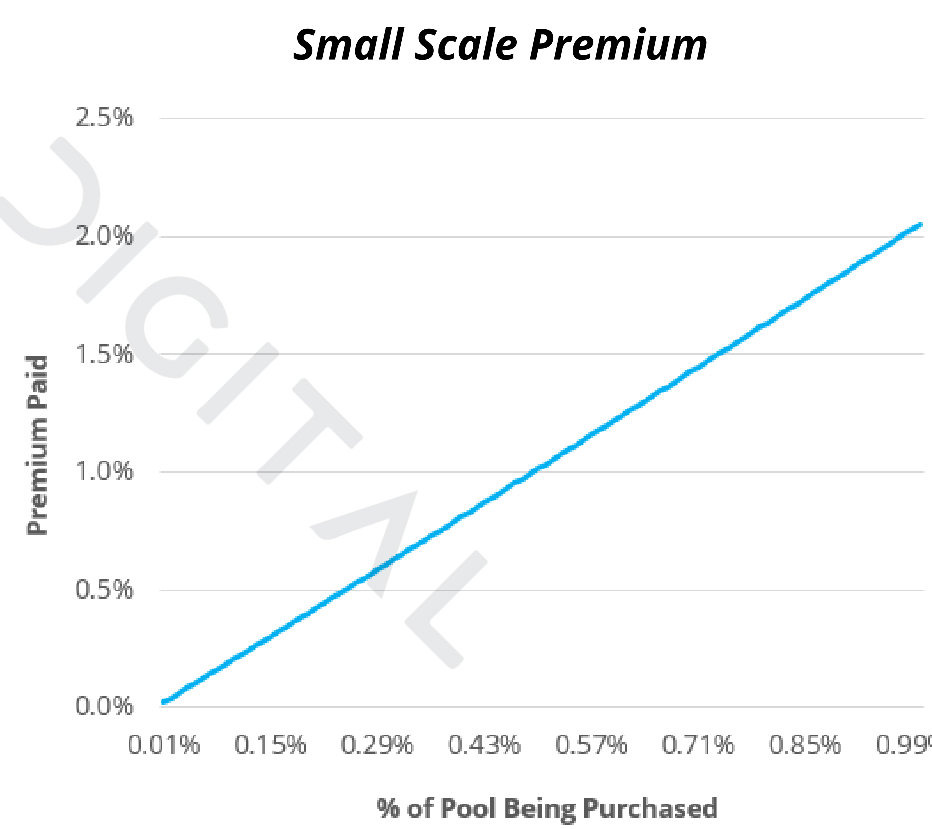

Liquidity Incentive Mechanisms

The relationship between transaction size and slippage is directionally consistent across liquidity pools for a trader, however, the underlying mechanics of the slippage can vary, and that's where creative incentive mechanisms can be introduced. THORChain incorporates a slip-based liquidity fee, which accrues to liquidity providers (LPs). As the transaction size grows relative to the overall pool, the liquidity fee to LPs increases. From the perspective of a trader, the underlying fee distribution isn't as relevant. The liquidity fee and general slippage are grouped together to reduce the output, implying a trader's interest would be focused on total slippage. This is, however, important for LPs as it incentivizes them to seek out liquidity demand. You can now have both passive and active providers, creating increased efficiency in liquidity provision.

In the chart on the right we show what composes total slippage. The area below the grey line represents the portion that goes to LPs, and the remainder of it can be thought of the portion that can be captured by arbitrageurs. Although liquidity fees to LP's does increase with transaction size, it grows at a slower rate than slippage, causing the share of slippage that goes to LP's to begin at 50% and end at 33%. Realistically, this decline isn't important as transactions will rarely be that large relative to the overall pool. Even if the transaction were to match 20% of the deposit side depth of the pool, LPs would still take home over 45% of total slippage as fees.

This structure helps make providing liquidity highly appealing by shifting yield toward LPs rather than mostly enabling arbitrage traders to take advantage. In general, liquidity pool providers will achieve the best yield when they maximize the ratio of transaction volume to pool depth because it helps them maximize their share of fees.

THORChain's CLP model builds on that by allocating a considerably larger portion to LPs. This creates an environment where liquidity providers are encouraged to actively seek out the largest supply (in the form of liquidity) and demand (in the form of desire to trade) mismatches. This helps create a bootstrapping effect for new pools that get added, while also encouraging existing capital to be distributed more efficiently. We'll address what else is being done to help bootstrap liquidity, along with how CLP mechanics compare to those of an XYK pool later on.

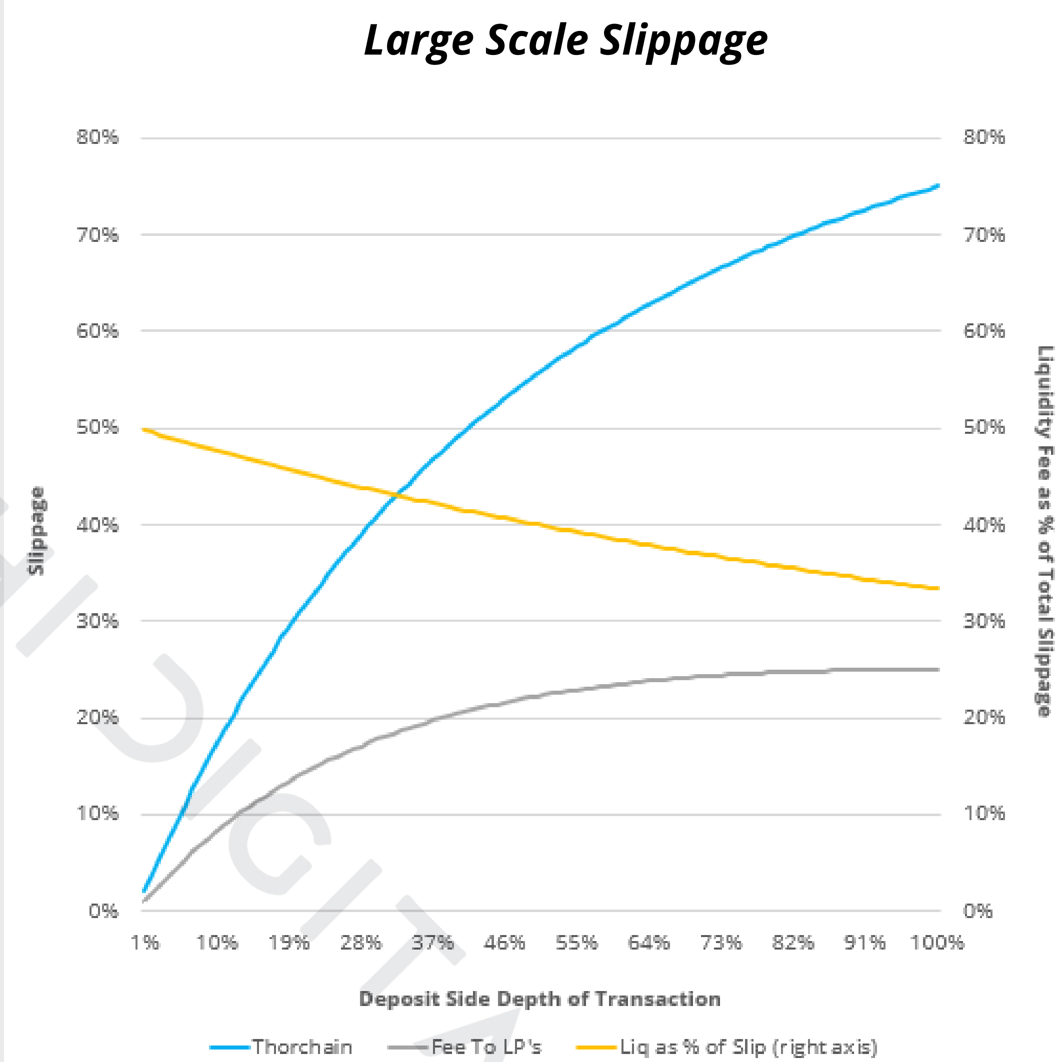

Arbitrage

To best grasp this section, it helps to take a step back and understand how these pools operate, and how different individuals interact with them. These pools are just a pair of assets, with the pool price of each asset dictated by their ratio in the pool. Regular traders use these pools purely with the intention of moving in and out of assets, while arbitrageurs are rewarded for maintaining balance between the ratio and external prices. Arbitrage opportunities are born by two situations:

(1) Someone trades with the pool by depositing one asset and withdrawing the other, causing the ratio to lose parity with external prices.
(2) External prices have moved and the prevailing pool ratio no longer coincides with external prices.

This is why price volatility often results in the largest volume days for AMM's. Arbitrage traders are constantly rebalancing the pool to ensure parity.

Arbitrage, to an extent, comes at the expense of LP's. The CLP model makes maximizing arbitrage slightly more difficult because it's impossible to capture the arbitrage in one trade. Instead, multiple trades must be made to gradually close the price gap between the pool and the external price. For context, even a 25 basis point price difference will net 0 gain for the arbitrage trader when the gap is closed in one trade. When the price difference increases, the profit maximizing trade implies that you would close less of the gap. Conversely, when the price gap decreases, a larger portion of the gap is closed in the ideal trade. In the chart, you can see the relationship between the most profitable arbitrage trade and the price spread between the pool and external markets.

The need for traders to be patient and gradually capture arbitrage over time has positive knock-on effects. Since multiple trades need to take place, this creates opportunity for additional traders to participate, potentially leading to increased transaction volume. Simultaneously, increasing the probability of inefficient arbitrage translates to less under performance for LPs.

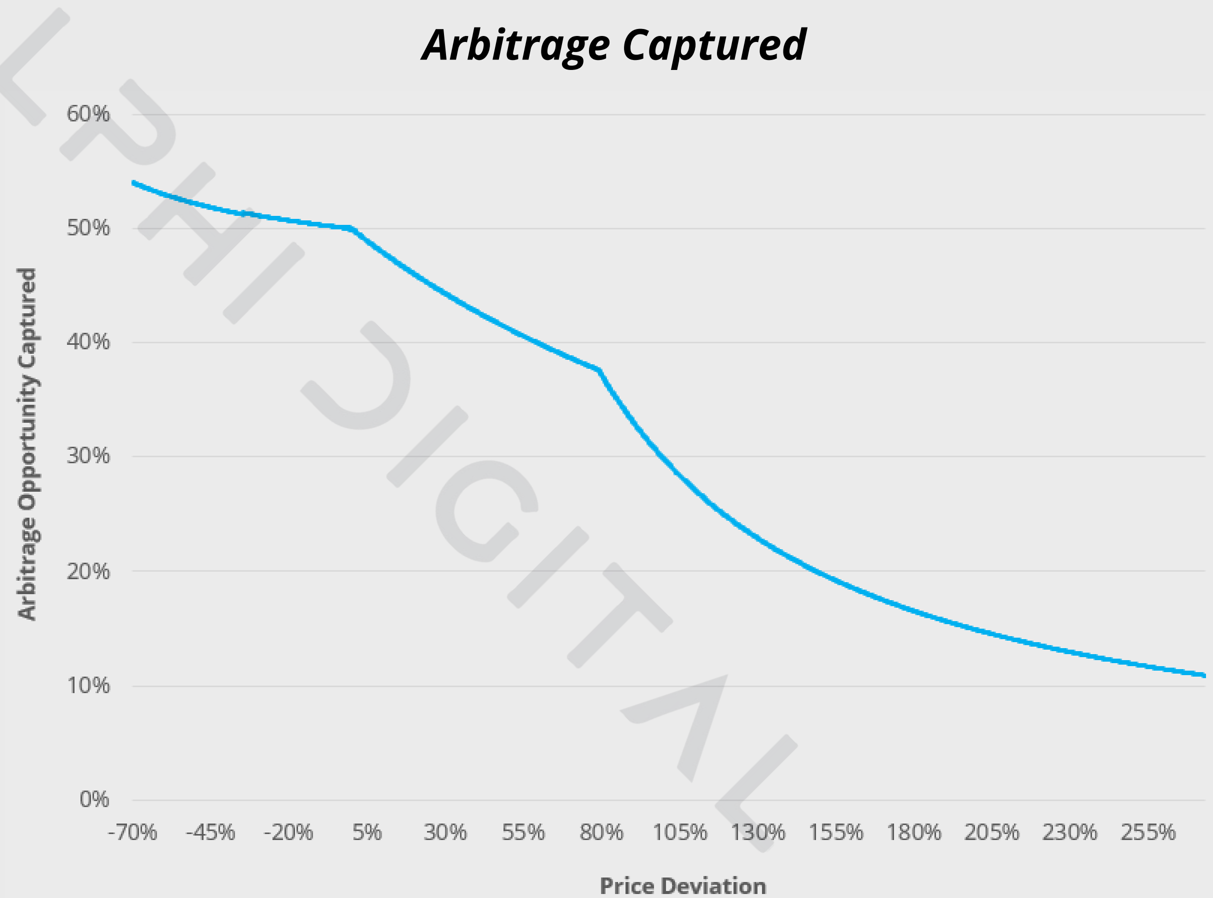

Arbitrage

An important question any LP is concerned with is "How much will my portfolio under-perform a buy and hold strategy by providing liquidity when prices deviate from my entry point?" We'll compare this with Uniswap's XYK model later on, but first it's worth understanding exactly how this AMM works. As discussed in the previous slide, peak losses for LP's occur when arbitrage is executed to perfection. It's very difficult to derive a direct formula to asses the losses associated with a particular price move because multiple specific trades are necessary to close the price gap, and they vary according to the prevailing price spread. We have, however, built a model which accomplishes exactly that, so we were able to derive what those potential losses would be given any price move.

It's important to remember that this under-performance would take place assuming each trade was executed to perfection, all the way until price parity. Since that's nearly impossible, this can be more so thought of as a worst case scenario. Initial trades are the ones that capture the most arbitrage, so if those are poorly executed, losses for LPs can be substantially mitigated. The poor execution more directly refers to an individual that tries to capture an arbitrage opportunity impatiently, and closes too much of the gap between the internal and external price in one trade. When a trader executes an arbitrage trade that's well below the optimal level, it reduces the potential future losses caused by subsequent arbitrage since potential losses to LP's increase with the price spread. Although this will be explained later on, it's also important to understand that these aren't impermanent losses like with Uniswap, but are unrealized losses that can be realized if taken advantage of correctly. Since these scenarios won't be happening in a vacuum, other non-arbitrage motivated individuals will also be trading during these opportunities, reducing the likelihood of achieving max losses. Some trades can actually result in a loss for the trader, particularly when their goal is to transact rather than arbitrage, which is value that accumulates back to the pool and LPs.

In order to decide if staking is worthwhile, the next logical thought exercise is to establish a potential yield, and understand the the likelihood of it offsetting potential losses relative to a pure buy and hold strategy. Before we do that, however, it's worth taking a look at the overall token economics to understand the source and distribution of that yield.

Price Shift Relative to Start of 100%

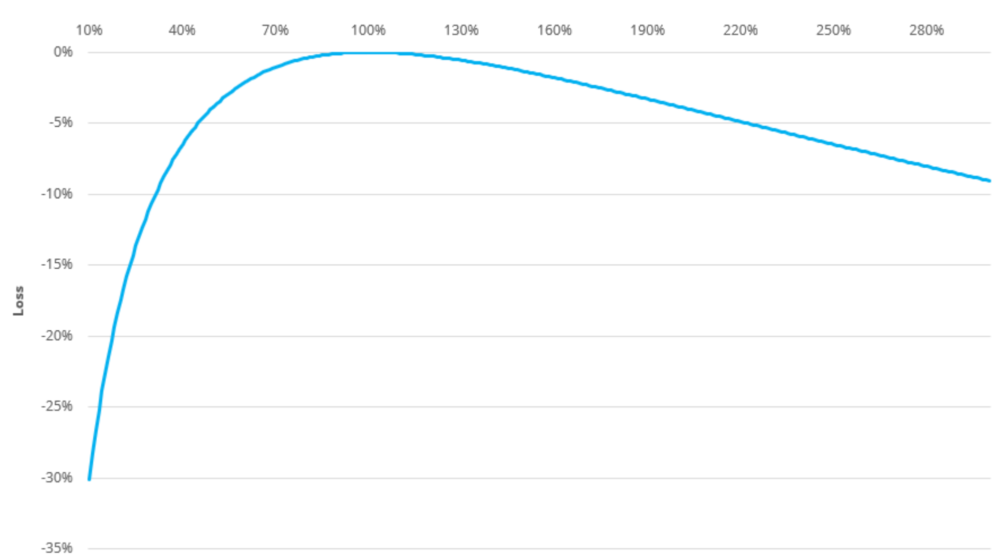

% Move	-90%	-50%	-25%	-10%	10%	25%	50%	90%
% Underperform	-30.08%	-3.83%	-0.68%	-0.09%	-0.08%	-0.41%	-1.35%	-3.31%

Portfolio Value Comparison

The Incentive Pendulum

The RUNE token has been integrated in a uniquely functional way that makes its value directly tied to the growth of the platform and integral to its security. For every $1 of non-RUNE assets staked as liquidity, there must also be $1 worth of RUNE staked in the pool alongside it. On the security side you have Bonded Capital, which represents all of the RUNE bonded by nodes, and is structured to match total Staked Capital (comprised of both RUNE and external assets) 1:1. This means the total RUNE value locked up will be ~3x the value of external, non-RUNE assets staked as liquidity.

Security is fortified through their incentive pendulum, which helps ensure that the value bonded by nodes is larger than the external capital being secured. As we demonstrated on page 13, the non-RUNE assets staked in the liquidity pools are the real bounty here because an attack by malicious nodes would likely cause the price of RUNE to fall off a cliff. In the table to the right, the "Distribution" portion is based solely on RUNE. "Bond" represents the total value of RUNE bonded in nodes, and "Staked" represents RUNE being staked for liquidity. The other half of staked capital that's not being represented here is the cumulative value of all other assets being used to provide liquidity (i.e the bounty).

In theory, the more that's bonded, particularly relative to the potential bounty, the more secure the network becomes. Conversely, as bonded capital declines relative to staked capital, the network becomes less secure as the attack vector becomes more exposed and the relative bounty increases. This structure ensures that if bonded capital begins to fall relative to stake, there are immense incentives on the bond side in the form of a larger portion of total payout, in tandem with strong disincentives on the stake side in the form of a reduced payout, to help bolster the over security of the network.

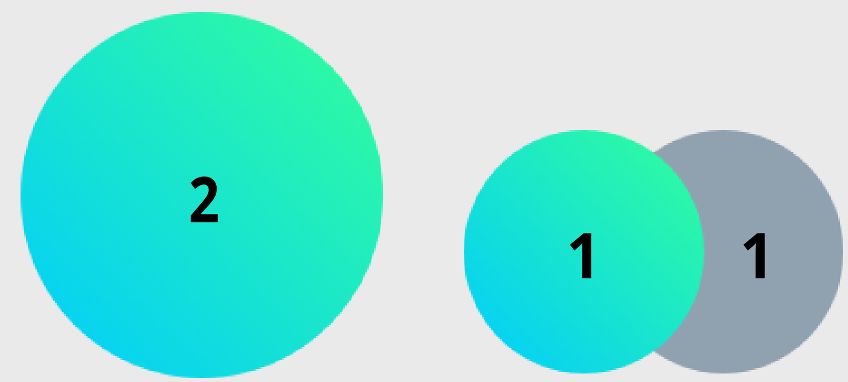

Bonded Capital **Staked Capital**

Value Distribution		Share of System Income	
Bond	***Staked***	***Bond***	***Staked***
100	0	0%	100%
90	10	20%	80%
80	20	40%	60%
70	30	60%	40%
66.7	33.3	67%	33%
60	40	80%	20%
50	50	100%	0%

This table implies that nodes will certainly command a larger yield, effectively double, when the network is at parity. That's reasonable considering the costs that come with running a node, the general operational frictions of running a node relative to providing liquidity, and the overall incentive of all honest participants to have as secure a network as possible. The payout that's being portioned between the two groups is total system income.

System Income - Issuance

Total System Income, comprised of block rewards and liquidity fees, is the fuel that helps keep all actors financially incentivized to participate and participate honestly. It's also a great example of how a project like this can benefit from having a token in the first place. Using the token, particularly in the early years, to help bootstrap liquidity by providing additional payments to liquidity providers, is an effective way to incentivize LPs and grow liquidity pools. Goes without saying that this a potential advantage over an AMM that doesn't have a token. To be clear, these tokens have already been issued, but are set aside in a separate pool and are distributed at a preset schedule.

Yield will be determined by multiple factors, so it's worth analyzing them separately. First, we'll concentrate on the portion that's sourced from issuance. There's 220M tokens that are set aside for issuance yield to node operators and liquidity providers. Each month, the system distributes 1/72nd of the remaining emission, eventually asymptoting to the final 500M supply hard cap. This issuance is already predetermined and is issued agnostic of the percent of supply staked, meaning that as the portion of supply staked for nodes and liquidity provision goes down, the yield goes up. This creates an underlying tailwind for individuals to participate in the network using the token to capture that yield, reducing potential risks surrounding attack vectors, since it becomes more difficult to actually get the amount of tokens necessary to maliciously attack the network. In terms of sustainability, this yield will certainly decline because of a combination of declining issuance and increased circulating supply from various unlocks. An overwhelming majority of the tokens will be unlocked by March 2021, at which point the decline in yield will become very gradual. However, at that point, only 16.6% (36.6m) of the tokens used to supplement yield will be used, ensuring that this bootstrapping mechanism will survive for an extended period. As of March 2021, total network issuance will be 11%, where it will continue declining gradually over time, going below 4% as soon as March 2025.

Issuance Based Yields To Liquidity Providers

% Staked	50%	60%	70%	80%	90%	100%
Mar-20	24.28%	20.23%	17.34%	15.18%	13.49%	12.14%
Apr-20	22.70%	18.92%	16.22%	14.19%	12.61%	11.35%
May-20	21.29%	17.74%	15.21%	13.31%	11.83%	10.65%
Jun-20	18.43%	15.36%	13.16%	11.52%	10.24%	9.21%
Jul-20	17.43%	14.52%	12.45%	10.89%	9.68%	8.71%
Aug-20	16.51%	13.76%	11.80%	10.32%	9.17%	8.26%
Sep-20	14.64%	12.20%	10.46%	9.15%	8.14%	7.32%
Oct-20	13.95%	11.63%	9.97%	8.72%	7.75%	6.98%
Nov-20	13.32%	11.10%	9.51%	8.32%	7.40%	6.66%
Dec-20	12.00%	10.00%	8.57%	7.50%	6.67%	6.00%
Jan-21	11.50%	9.58%	8.21%	7.19%	6.39%	5.75%
Feb-21	11.15%	9.30%	7.97%	6.97%	6.20%	5.58%
Mar-21	10.28%	8.57%	7.34%	6.42%	5.71%	5.14%
Apr-21	9.99%	8.32%	7.13%	6.24%	5.55%	4.99%
May-21	9.70%	8.09%	6.93%	6.07%	5.39%	4.85%
Jun-21	9.43%	7.86%	6.74%	5.90%	5.24%	4.72%
Jul-21	9.17%	7.64%	6.55%	5.73%	5.10%	4.59%
Aug-21	8.92%	7.44%	6.37%	5.58%	4.96%	4.46%
Sep-21	8.68%	7.24%	6.20%	5.43%	4.82%	4.34%
Oct-21	8.45%	7.04%	6.04%	5.28%	4.69%	4.22%
Nov-21	8.28%	6.90%	5.91%	5.17%	4.60%	4.14%
Dec-21	8.11%	6.75%	5.79%	5.07%	4.50%	4.05%
Jan-22	7.94%	6.62%	5.67%	4.96%	4.41%	3.97%
Feb-22	7.78%	6.48%	5.56%	4.86%	4.32%	3.89%

Year	Supply Unlocked
2020	130.07%
2021	25.22%
2022	7.41%
2023	5.83%
2024	4.66%
2025	3.76%
2026	3.07%
2027	2.52%
2028	2.08%
2029	1.72%
2030	1.43%

To Clarify:
Supply Unlocked Column Represents The Rate of Supply Being Unlocked. All Tokens are Already Issued.

Issuance Based Yields To Nodes

% Staked	50%	60%	70%	80%	90%	100%
Mar-20	48.56%	40.47%	34.69%	30.35%	26.98%	24.28%
Apr-20	45.41%	37.84%	32.43%	28.38%	25.23%	22.70%
May-20	42.58%	35.48%	30.41%	26.61%	23.66%	21.29%
Jun-20	36.85%	30.71%	26.32%	23.03%	20.47%	18.43%
Jul-20	34.86%	29.05%	24.90%	21.79%	19.36%	17.43%
Aug-20	33.03%	27.52%	23.59%	20.64%	18.35%	16.51%
Sep-20	29.29%	24.41%	20.92%	18.30%	16.27%	14.64%
Oct-20	27.91%	23.26%	19.93%	17.44%	15.50%	13.95%
Nov-20	26.63%	22.19%	19.02%	16.64%	14.79%	13.32%
Dec-20	24.01%	20.01%	17.15%	15.00%	13.34%	12.00%
Jan-21	23.00%	19.17%	16.43%	14.38%	12.78%	11.50%
Feb-21	22.31%	18.59%	15.94%	13.94%	12.39%	11.15%
Mar-21	20.56%	17.13%	14.68%	12.85%	11.42%	10.28%
Apr-21	19.97%	16.64%	14.27%	12.48%	11.10%	9.99%
May-21	19.41%	16.17%	13.86%	12.13%	10.78%	9.70%
Jun-21	18.87%	15.72%	13.48%	11.79%	10.48%	9.43%
Jul-21	18.35%	15.29%	13.11%	11.47%	10.19%	9.17%
Aug-21	17.85%	14.87%	12.75%	11.15%	9.91%	8.92%
Sep-21	17.36%	14.47%	12.40%	10.85%	9.65%	8.68%
Oct-21	16.90%	14.08%	12.07%	10.56%	9.39%	8.45%
Nov-21	16.55%	13.79%	11.82%	10.34%	9.19%	8.28%
Dec-21	16.21%	13.51%	11.58%	10.13%	9.01%	8.11%
Jan-22	15.88%	13.23%	11.34%	9.92%	8.82%	7.94%
Feb-22	15.56%	12.96%	11.11%	9.72%	8.64%	7.78%

System Income - Liquidity Fees

Issuance will be the mechanism that bootstraps liquidity early on, but liquidity fees will keep LP's incentivized in the long run. As depicted by the formula and visualization, liquidity fees scale sub linearly with deposit depth. This relationship doesn't take effect until transactions get significantly large, meaning that an overwhelming majority of transactions will exist in the deposit depth window where they actually scale mostly linearly.

$$Liquidity\ Fee = \frac{x^2 Y}{(x + X)^2}$$

In a previous section we discussed how the liquidity fee mechanism incentivizes liquidity pool growth by paying a higher fee to LP's in situations where transactions constituted a larger portion of the pool. This is a great way to match transaction demand with liquidity supply, and a feature that a fixed rate AMM is unable to offer, because it's fixed rate by nature. With a fixed fee AMM, the transaction fee will always be the same portion of your overall transaction. As a liquidity provider, that doesn't create any additional incentives for you to seek out liquidity and demand mismatches. This means the existing capital being used to provide liquidity isn't as motivated to provide capital efficiently, which actually comes at the direct cost of the trader, as their transaction will now incur additional slippage from a lack of adequate pool depth. This slippage is then captured by arbitrageurs who bring the pool back in balance, resulting in less gain for liquidity providers and additional cost for traders.

From the perspective of a trader, slippage is slippage. Their main concern is what they get out of the pool relative to what they put into it. The underlying make up of that slippage is irrelevant, they mostly look to minimize it. In reality, total slippage consists of multiple factors, with the major driver typically being the mechanism by which the pool transacts, and secondary factors being fees and gas. You can safely say all traders would generally prefer to minimize slippage, but there are many instances where the preference to transact quickly can supersede it (we'll expand on non-monetary motivations later on). From the perspective of an LP, the important factors are yield and the under-performance of their provided liquidity relative to a buy hold portfolio. Most AMM's don't have an effective way of allowing liquidity providers to take advantage of the temporary semi-agnostic attitude towards fees, while also allowing for the flexibility of addressing the needs of regular traders. We'll tangibly show how the CLP's slip-based liquidity fee structure is able to take advantage of this dynamic in a value accretive way.

Small Scale Liquidity Fee

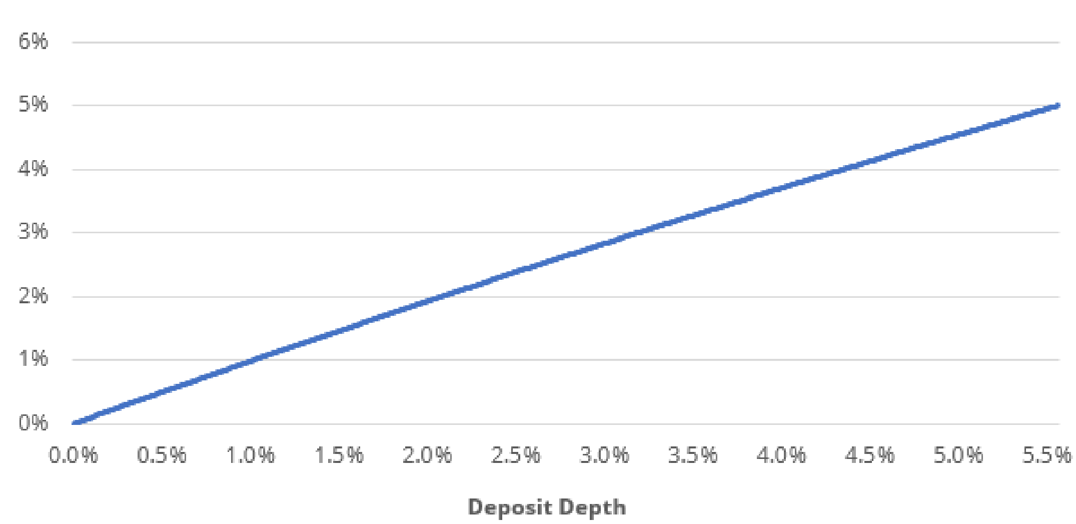

Large Scale Liquidity Fee

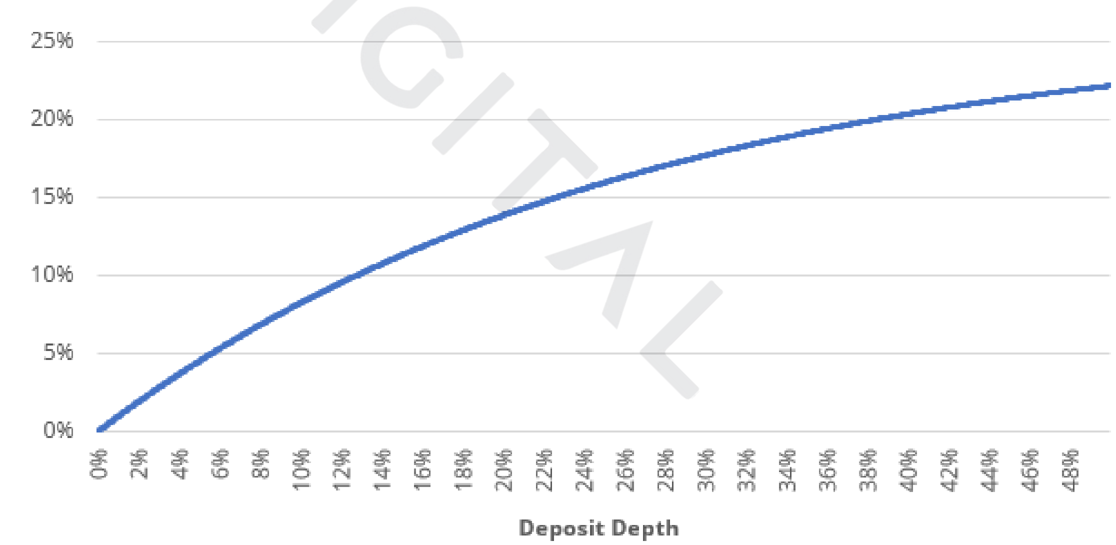

Uniswap vs THORChain - Intro

DELPHI DIGITAL

In order to run our analysis and comparison, we took every transaction in December from one of Uniswap's largest and most active pools, ETH/MKR. Due to complications from internal transactions and transactions between various DeFi platforms, getting a completely exhaustive list that included the necessary information for each transaction is difficult. We had to settle for a list that encompassed over 90% of all transactions (11k+), which we deemed to be sufficient. Another worthwhile preface for this analysis is that it's impossible to assume that every individual would've reacted the exact same way for every trade, but after going through the results, it'll make sense to assume that an overwhelming majority wouldn't have done much differently. A major point we're trying to illustrate is that the slip based liquidity fee allows market participants with different underlying motivations to continue acting in their own best interest while simultaneously benefiting others who are doing the same.

At the highest level, this analysis is assessing costs for Traders and revenue for Liquidity Providers (LPs), oriented around Liquidity Fees and Total Slippage. Liquidity Fees, a part of Total Slippage, are the fees that accrue to LPs. Total slippage represents the difference between what the trader received and what she would've received in a 0 fee vacuum. While THORChain's CLP model allowed it to accrue almost 2.5x as much in liquidity fees as Uniswaps XYK constant product model, the more important part is to understand at whose expense. The dollar value average Liquidity Fee is much higher for THORChain, as expected, but the median is considerably lower, indicating it's likely to be top heavy. Average % Liquidity Fee is considerably lower, which might come as a surprise considering the dollar value relationship, however, it's important to remember that liquidity fees as a % scale with transaction size, so the dollar value growth of liquidity fees can have a compounding effect. Uniswap's fixed fee model logically has mean and median fee at the same 30 bps level. Average dollar value liquidity fee per transaction for Uniswap is much lower, as cumulative fees traders are paying liquidity providers is much lower, however, the median liquidity fee is considerably higher. This is exactly where the opportunity exists for THORChain, with their median fee implying most transactions are actually considerably cheaper.

	Thorchain	Uniswap
Total Liquidity Fees	$161,861	$67,451
Average Liquidity Fee	$14.15	$5.90
Median Liquidity Fee	$0.32	$2.09
Average Liquidity Fee %	0.129%	0.30%
Median Liquidity Fee %	0.047%	0.30%
Average Total Slip	0.260%	0.427%
Median Total Slip	0.094%	0.347%
Ann. Total Yield	65.0%	27.1%
Ann. Node Yield	43.3%	
Ann. LP Yield	**21.7%**	**27.1%**
Ann. Node Yield + Issuance	73.7%	
Ann. LP Yield + Issuance	**36.8%**	**27.1%**

	Thorchain			Uniswap			Applies To Both	
	% of Fees	Avg Liquidity Fee	Avg Total Slip	% of Fees	Avg Liquidity Fee	Avg Total Slip	% of Volume	% of Trades
Relating to transactions larger than 1% Pool Depth	60.9%	1.60%	3.24%	15.1%	0.30%	1.93%	15.1%	1.1%
Relating to transactions between .5% and 1% of Pool Depth	17.6%	0.71%	1.43%	17.2%	0.30%	1.01%	17.2%	3.1%
Relating to transactions between .25% and .5% of Pool Depth	15.1%	0.34%	0.68%	30.9%	0.30%	0.64%	30.9%	12%
Relating to transactions below .25% of Pool Depth	6.3%	0.06%	0.12%	36.9%	0.30%	0.36%	36.9%	83.8%

Uniswap vs THORChain - Liquidity Fees

DELPHI DIGITAL

Slip based liquidity fee forces costs to meet demand. Individuals transacting in small amounts likely aren't in need of significant liquidity on a particularly short time horizon. This implies they're likely to be more cost sensitive than someone who is more motivated by immediate access to liquidity. Again, all these traders are likely to be cost sensitive, but certain traders have other more pressing motivations. By allowing the cost sensitive to transact more cheaply than a fixed rate fee mechanism would allow, you're able to address their primary interests and benefit from their transaction volume. On the other hand, you have the large scale traders. We know that total slippage increases with transaction size, so when you see the sizable transactions come through, you can already assume there's less price elasticity shown by these individuals, and they're more interested in getting the asset quickly. The CLP model effectively allows all traders and LP's to benefit from the ability to charge these individuals more, something a fixed rate model doesn't have the ability to take advantage of.

We provide additional color on how these dynamics play out in the tables on the bottom by separating transactions into different pockets, and showing how the liquidity fee accumulation and overall trading expenses vary between the two models. Uniswap's fixed liquidity fee causes total slip to be elevated on smaller transactions while remaining relatively muted on larger ones. As liquidity fee stays fixed and total slip goes up, that margin bypasses the liquidity providers and eventually gets consumed by arbitrage traders. This makes liquidity fee sources relatively evenly distributed for Uniswap, but still skewing towards the lower end as that's where a great deal of volume comes from. THORChain's model caters to the price sensitive by allowing average total slip on the lower end to be .12%, by only extracting an average liquidity fee of .06% from this group. THORChain's CLP model certainly causes more total slip on the higher end, outpacing Uniswap total slip as transactions go beyond .50% of pool depth. The important aspect is that liquidity fees grow along with it. Looking at the .50% to 1% bucket, CLP average total slip is .40% higher than that of Uniswap, however, all of that slip basically accrues to the liquidity fee as that's also roughly .40% higher. This occurs to an even higher degree in the largest bucket, where non-liquidity slip is even between the two, but average liquidity fee for THORChain is over 5x the size.

	Thorchain	Uniswap
Total Liquidity Fees	$161,861	$67,451
Average Liquidity Fee	$14.15	$5.90
Median Liquidity Fee	$0.32	$2.09
Average Liquidity Fee %	0.129%	0.30%
Median Liquidity Fee %	0.047%	0.30%
Average Total Slip	0.260%	0.427%
Median Total Slip	0.094%	0.347%
Ann. Total Yield	65.0%	27.1%
Ann. Node Yield	43.3%	
Ann. LP Yield	**21.7%**	**27.1%**
Ann. Node Yield + Issuance	73.7%	
Ann. LP Yield + Issuance	**36.8%**	**27.1%**

	Thorchain			Uniswap			Applies To Both	
	% of Fees	Avg Liquidity Fee	Avg Total Slip	% of Fees	Avg Liquidity Fee	Avg Total Slip	% of Volume	% of Trades
Relating to transactions larger than 1% Pool Depth	60.9%	1.60%	3.24%	15.1%	0.30%	1.93%	15.1%	1.1%
Relating to transactions between .5% and 1% of Pool Depth	17.6%	0.71%	1.43%	17.2%	0.30%	1.01%	17.2%	3.1%
Relating to transactions between .25% and .5% of Pool Depth	15.1%	0.34%	0.68%	30.9%	0.30%	0.64%	30.9%	12%
Relating to transactions below .25% of Pool Depth	6.3%	0.06%	0.12%	36.9%	0.30%	0.36%	36.9%	83.8%

Uniswap vs THORChain - Yields

This is an efficient way to take advantage of those trader's price inelasticity, greatly to the benefit of Liquidity Providers. They are able to extract maximum value from that group to source 61% of their total fees despite it only making up 15.1% of total volume and 1.1% of trades. This extraction subsidizes the smaller traders, enabling an impressively low median total slip for all transactions of .094%. That's cheaper than most centralized exchanges, while still being anonymous and from the convenience of your wallet. The variance in the distribution of fee revenue sources between the two models really illustrates how they function differently.

While THORChain does command a much higher liquidity fee total, it's important to remember that this still gets pooled with issuance based rewards to form system income, which then gets distributed to both, Validators and Liquidity Providers. Despite THORChain commanding an annual total yield of 65% on the capital provided for liquidity in this example, only 21.7% would actually flow to LP's while the remaining 2/3's flows to Nodes. This 21.7% represents the liquidity fee portion, but it gets further supplemented by issuance yield. As seen in the system income issuance slide, yield can vary depending on the % of supply that's locked up in Nodes + Liquidity Pools. The combined yields in the table on the right assume that 80% is staked. We think that's a reasonable expectation for a few reasons. On one end you'll have pretty sizable guaranteed rewards in the early months/years that'll incentivize individuals to participate in the network to capture that yield, creating upward pressure on total supply staked. On the other end, there are a few sources of excess float that'll exist in perpetuity. One of them being node candidates that are waiting in line to be selected as a validator, and another being pure speculators that are taking directional exposure with the token, but don't plan on participating in the network. Either way, this supplemental issuance is a strong liquidity bootstrapping mechanism because it comfortably offsets any potential under-performance a liquidity provider might be exposed to, removing price risk as a deterrent for potential Liquidity Providers.

	Thorchain	Uniswap
Total Liquidity Fees	$161,861	$67,451
Average Liquidity Fee	$14.15	$5.90
Median Liquidity Fee	$0.32	$2.09
Average Liquidity Fee %	0.129%	0.30%
Median Liquidity Fee %	0.047%	0.30%
Average Total Slip	0.260%	0.427%
Median Total Slip	0.094%	0.347%
Ann. Total Yield	65.0%	27.1%
Ann. Node Yield	43.3%	
Ann. LP Yield	**21.7%**	**27.1%**
Ann. Node Yield + Issuance	73.7%	
Ann. LP Yield + Issuance	**36.8%**	**27.1%**

	Thorchain			Uniswap			Applies To Both	
	% of Fees	Avg Liquidity Fee	Avg Total Slip	% of Fees	Avg Liquidity Fee	Avg Total Slip	% of Volume	% of Trades
Relating to transactions larger than 1% Pool Depth	60.9%	1.60%	3.24%	15.1%	0.30%	1.93%	15.1%	1.1%
Relating to transactions between .5% and 1% of Pool Depth	17.6%	0.71%	1.43%	17.2%	0.30%	1.01%	17.2%	3.1%
Relating to transactions between .25% and .5% of Pool Depth	15.1%	0.34%	0.68%	30.9%	0.30%	0.64%	30.9%	12%
Relating to transactions below .25% of Pool Depth	6.3%	0.06%	0.12%	36.9%	0.30%	0.36%	36.9%	83.8%

Slippage & Liquidity Fees Comparison

The portion of total slippage that doesn't contribute to liquidity fees effectively transfers value to arbitrage traders, because it causes the pool to be further out of balance. Although that's not necessarily a negative, since those traders are often the primary drivers of volume, it does reduce the potential earnings of Liquidity Providers. Uniswap's fixed rate liquidity fee of .30% creates a Total Slip floor, preventing small transactions from scaling down in cost. This causes liquidity fees, as a portion of total slip, to peak at the smallest possible transaction size for Uniswap, as illustrated in the Small Scale Slippage chart. THORChain's slip based fee model allows cost conscious traders to continue scaling down transaction fees with their transaction size. This dynamic explains why Uniswap sources so much more of its total fees from small traders than THORChain, which can certainly be thought of as a positive for Uniswap, particularly as these pools grow. Over time you expect more and more of these transactions to fall into the lowest bucket if these pools continue to grow, favoring Uniswap to earn more liquidity fees. This yield concern is partially mitigated by the fact that those individuals who want to trade one non-RUNE asset for another non-RUNE asset will have to pay two layers of fees since each pool will have RUNE in it. The transactions will happen behind the scenes to make the process easier for the trader, but the fees still apply. This has to be weighed against the potential volume another potential mitigant is the fact that low fees might attract additional volume, so lower fees would need to be weighed against increased volume.

Once transaction size surpasses 30bps, THORChain becomes more expensive to use, and that difference expands over time. While this can certainly act as a deterrent for traders, the difference in price between the two systems isn't overly significant at first. THORChain, however, maintains an elevated level of liquidity fees as a share of Total Slip, enabling LP’s to capture an outsized portion of fees relative to their Uniswap counterparts at comparable trade sizes.

Small Scale Slippage

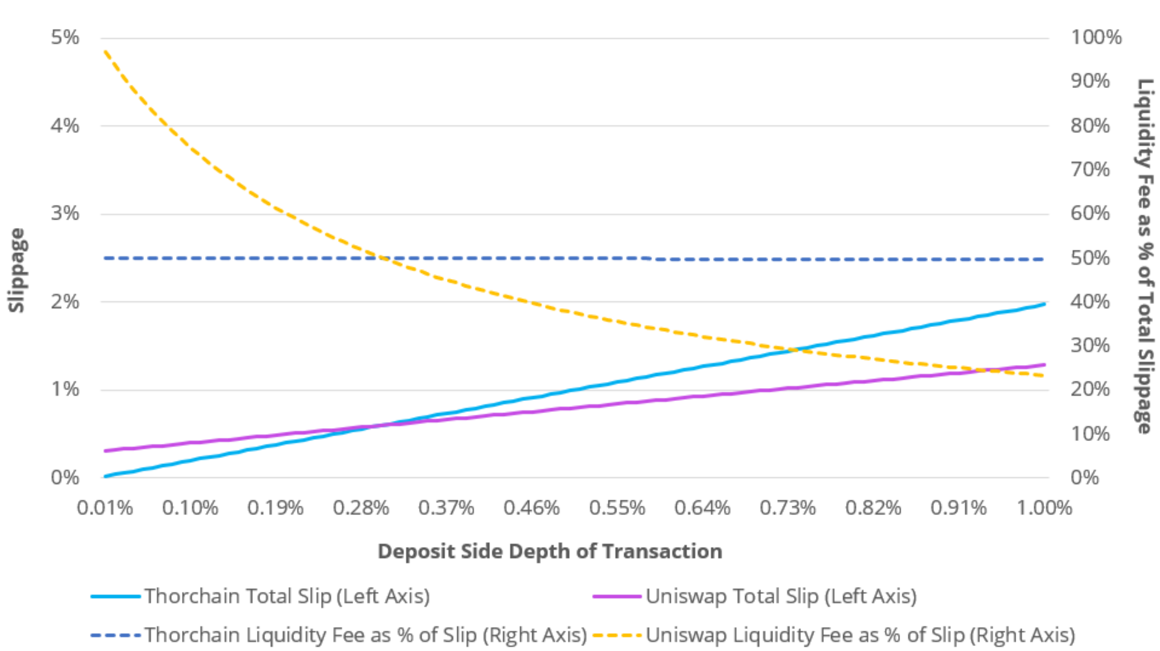

Large Scale Slippage

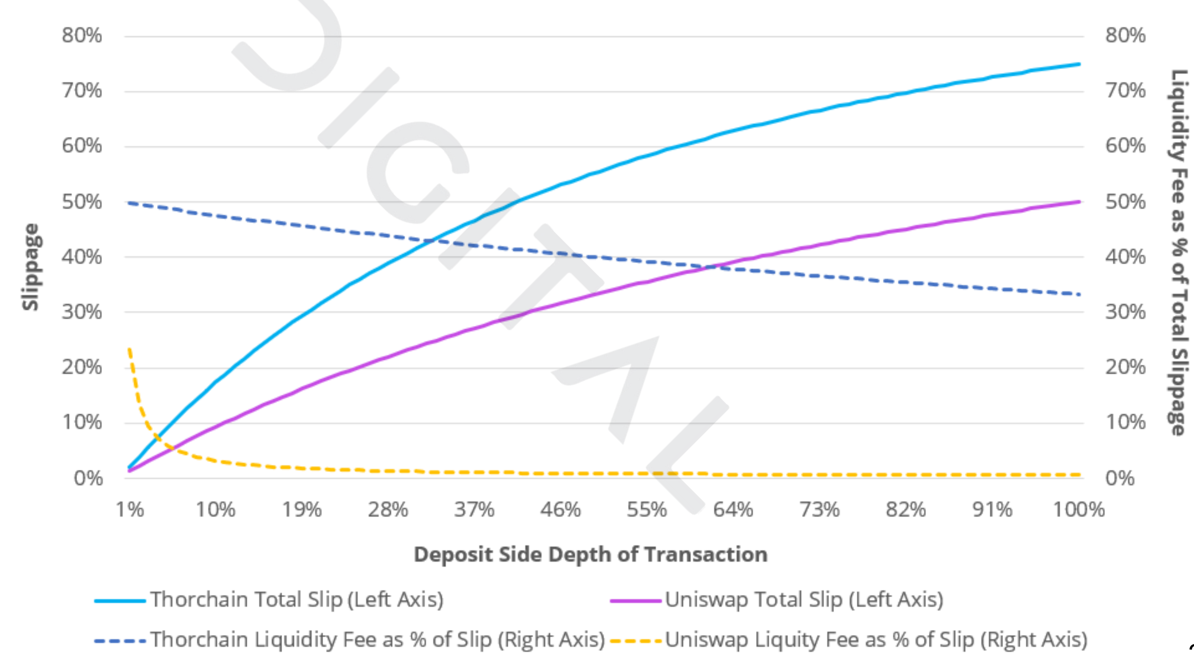

Impermanent & Permanent Loss

The XYK constant product model and THORChain's slip based CLP model vary significantly in some aspects, so doing a pure apples to apples comparison is difficult. When price deviates within the XYK model, the loss is considered impermanent as portfolio value oscillates on curve K with path independence. Prices can continue to shift randomly in any order, but will absolve the portfolio of any losses if they return to parity. With an XYK model, profitability is assessed by comparing liquidity yield with impermanent loss. Since liquidity fees accrue to the pool rather than get distributed, liquidity yield is measured by comparing the ratio of the square root of the constant product and the quantity of liquidity tokens to itself over time, and seeing how much that ratio increases. That growth represents the liquidity yield. The square root of the constant product grows from both new capital entering the pool and the fees that accrue to the pool over time. Liquidity tokens only increase through new capital entering the pool. By dividing them, you cancel out new capital and are left with fees that accrued to the pool.

The CLP model also has a liquidity token that increases in a comparable way, however, the impermanent loss aspect doesn't exist. With a CLP structure, you can model your way to peak losses by following all the ideal trades to maximize the arbitrage scenario, which we've done to create this curve. It's unrealistic to assume this will ever happen, so it serves as a useful floor to understand how significant losses can be. Losses will likely be less since you'll have multiple trades occurring simultaneously, and it's unreasonable to assume each trader will successfully maximize arb, or even attempt to. When a single incorrect trade happens in the arb chain, it comes to the benefit of LPs. What we're able to deduce is that price change, despite not being accompanied by impermanent loss, will likely be less detrimental to CLPs than it will to liquidity providers in the XYK model. Although the losses with an XYK model are impermanent, you'd need prices to revert to Day 1 values in order for that impermanence to actually take effect.

Constant Product Curve

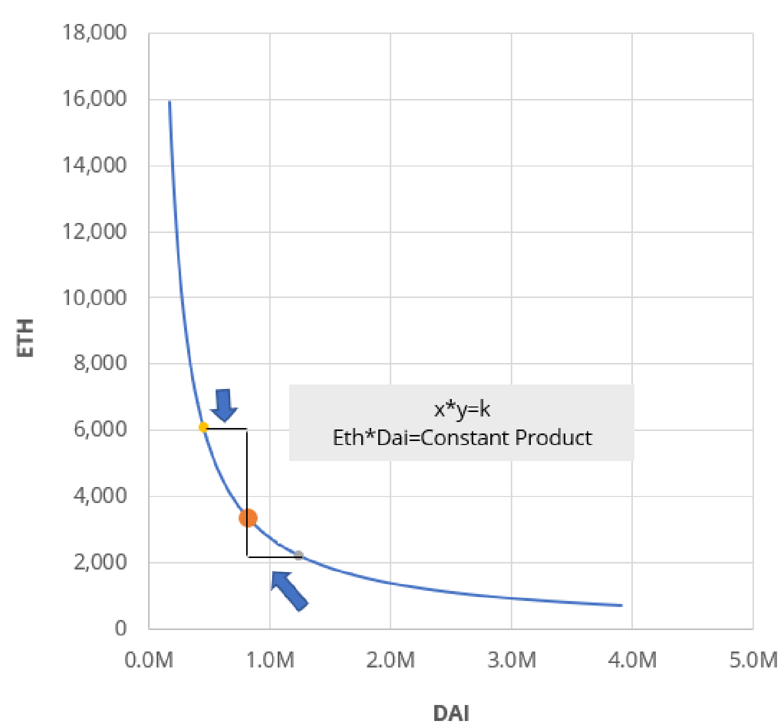

Max Permanent Loss (CLP) vs Impermanent Loss (XYK)

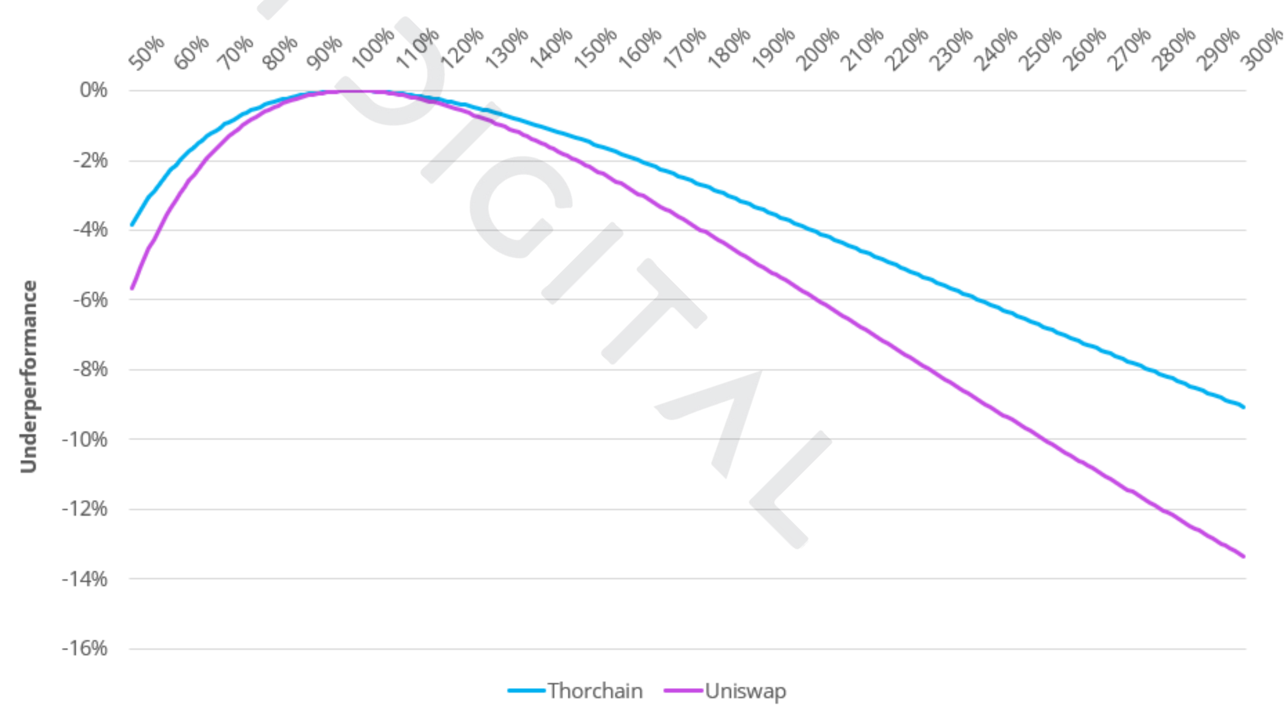

Economic Risks & Long Term Sustainability

DELPHI DIGITAL

In the short term, RUNE's issuance alone will provide the incentives necessary for people to stake and provide capital for liquidity. This will ideally lead to enough depth in these asset pools where trading makes economic sense. The ability to trade anonymously and from the convenience of your own wallet will also be a consistent driver for volume, even when the economics aren't better than ones from a centralized source. As assets are staked for liquidity, the price of RUNE is forced to appreciate, since the 3x ratio between RUNE participating across security and liquidity and non-RUNE staked assets has to hold. The mechanisms in place appear effective enough to allow this ratio to hold.

(1) When the liquidity pools aren't balanced 1:1, there's an arbitrage opportunity in place that'll bring Liquidity Pools back to parity.
(2) When the value bonded by nodes securing the network isn't balanced with the assets in Liquidity Pools, the portion of Liquidity Fees Nodes were previously entitled to will increase significantly, creating strong financial incentives for additional nodes to come online.

At Uniswaps current liquidity total of $25m, this would mean the value of total staked RUNE would have to be $75m. It's reasonable to assume that figure would be higher for that amount of liquidity considering there will always be RUNE that's purely held for speculation, as is the case with any productive token. There will also be RUNE locked up in potential new nodes that are lining up to join the network.

In the long term, as pool sizes begin to grow, there can be concerns around total liquidity fees being enough to support a network of that size. Liquidity fees are used to pay both Liquidity Providers and Node Operators. As pool sizes grow and transaction sizes relative to the pool shrink, that means individuals will be paying less and less in fees. This mechanism is very useful in the short to medium term, but can potentially create long term risks if fees aren't enough to support the network. The portion of total income that's supported by issuance will mitigate this risk, but that reward also declines over time.

A potential mitigant, which was previously mentioned, is that each pool being half RUNE implies that in order to transact across two non-RUNE assets, you'll be paying fees to two pools. This already exists with Uniswap, and will certainly be taken advantage of early on by traders that highly value privacy and access. It will likely become more common place as these pools grow and it's economically attractive to make these transactions. For context, the outputs from the Uniswap vs THORChain analysis was for a pool that averaged nearly $3 million in Total Liquidity (ETH+MKR) throughout the month and facilitated about $22.5 million in volume. A pool of that size can (and already does with Uniswap) support competitive pricing, particularly on some of the smaller trades, although it's certainly pricey for some of the larger trades.

When an individual is deciding to provide liquidity they always have to weigh it against the risks. If this network grows significantly, the technical risk declines as attack vectors become limited. The financial risk of under-performance will always exist, but CLP model helps significantly reduce that, and allows the break even yield for otherwise unproductive assets to decline considerably.

Disclosures

The Research Team may own the tokens represented in this report, and as such this should be seen as a disclosure of any potential conflict of interest. Anyone can contact Delphi Digital for full token disclosures by team member at Team@DelphiDigital.io. This report belongs to Delphi Digital, and represents the opinions of the Research Team.

Delphi Digital is not a FINRA registered broker-dealer or investment adviser and does not provide investment banking services. This report is not investment advice, it is strictly informational. Do not trade or invest in any tokens, companies or entities based solely upon this information. Any investment involves substantial risks, including, but not limited to, pricing volatility, inadequate liquidity, and the potential complete loss of principal. Investors should conduct independent due diligence, with assistance from professional financial, legal and tax experts, on topics discussed in this document and develop a stand-alone judgment of the relevant markets prior to making any investment decision.

Delphi Digital does not receive compensation from the companies, entities, or protocols they write about. The only fees Delphi Digital earns is through paying subscribers. Compensation is not received on any basis contingent upon communicating a positive opinion in this report. The authors were not hired by the covered entity to prepare this report. Delphi Digital did not receive compensation from the entities covered in this report for non-report services, such as presenting at author sponsored investor conferences, distributing press releases or other ancillary services. The entities covered in this report have not previously paid the author in cash or in stock for any research reports or other services. The covered entities in this report are not required to engage with Delphi Digital.

The Research Team has obtained all information herein from sources they believe to be accurate and reliable. However, such information is presented “as is,” without warranty of any kind – whether expressed or implied. All market prices, data and other information are not warranted as to completeness or accuracy, are based upon selected public market data, reflect prevailing conditions, and the Research Team’s views as of this date, all of which are accordingly subject to change without notice. Delphi Digital has no obligation to continue offering reports regarding this topic. Reports are prepared as of the date(s) indicated and may become unreliable because of subsequent market or economic circumstances. The graphs, charts and other visual aids are provided for informational purposes only. None of these graphs, charts or visual aids can and of themselves be used to make investment decisions. No representation is made that these will assist any person in making investment decisions and no graph, chart or other visual aid can capture all factors and variables required in making such decisions.

The information contained in this document may include, or incorporate by reference, forward-looking statements, which would include any statements that are not statements of historical fact. No representations or warranties are made as to the accuracy of such forward-looking statements. Any projections, forecasts and estimates contained in this document are necessarily speculative in nature and are based upon certain assumptions. These forward-looking statements may turn out to be wrong and can be affected by inaccurate assumptions or by known or unknown risks, uncertainties and other factors, most of which are beyond control. It can be expected that some or all of such forward-looking assumptions will not materialize or will vary significantly from actual results.

DELPHI DIGITAL

85 Broad Street
New York, NY, 10004
www.delphidigital.io