GSR

September 2022

Ethereum's Roadmap: A Guide to The Merge and Beyond

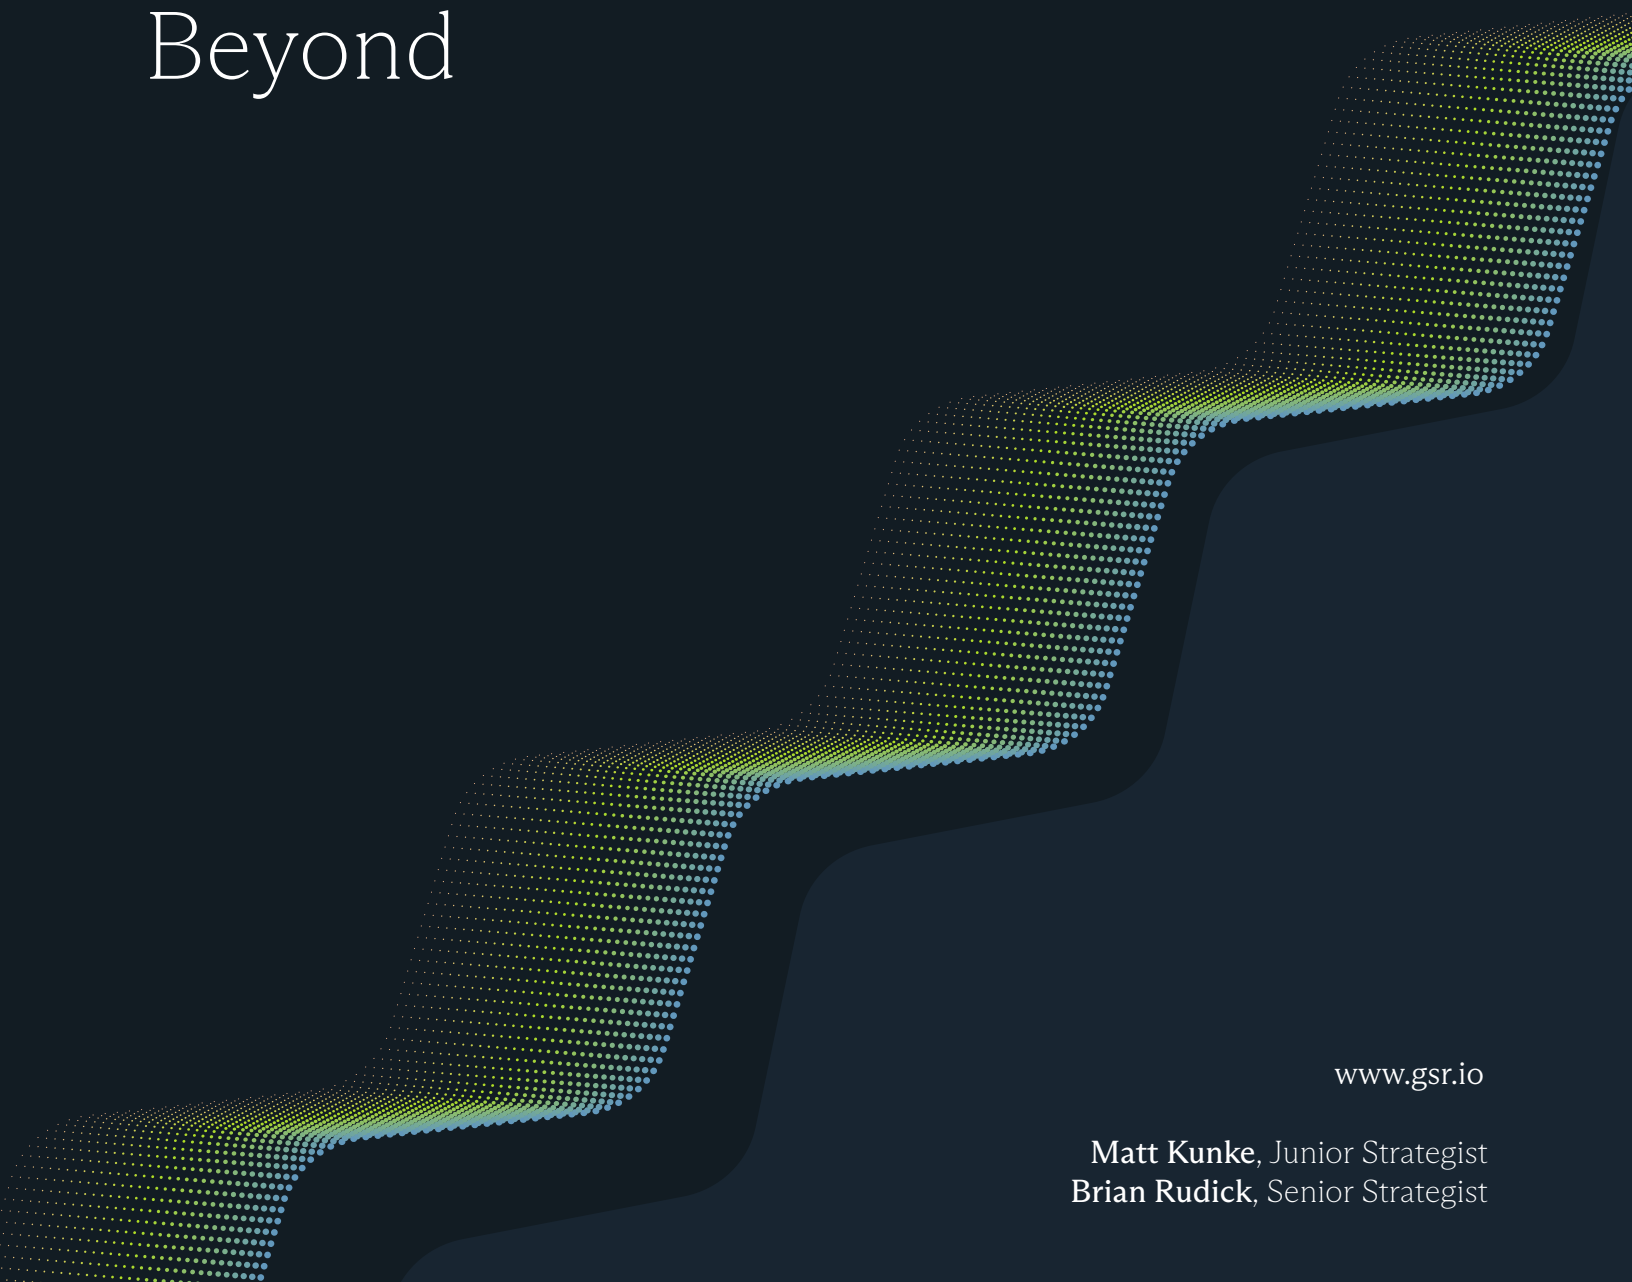

www.gsr.io

Matt Kunke, Junior Strategist
Brian Rudick, Senior Strategist

After years of research and development, Ethereum recently completed its final testnet Merge and confirmed the details of its transition to proof-of-stake, which is expected to occur around September 15th. This report dives into Ethereum's proof-of-stake blockchain and the implications of the coming Merge, as well as previews future upgrades and related topics such as sharding, statelessness, and proposer-builder separation.

A Brief Overview of Ethereum

First conceived in 2013 by Vitalik Buterin and launched in mid-2015, Ethereum revolutionized blockchain technology and cryptocurrencies forever by creating a world settlement layer or decentralized state machine. No longer were blockchains home to a single, siloed application such as Bitcoin or Namecoin, or were developers confined by the limited set of transaction types. The Ethereum network could process general-purpose code, known as smart contracts, adding programmability and arbitrary computation without needing to modify the underlying blockchain. Moreover, common smart contract standards and open-source code enabled applications to interact and build on each other, creating a composable programming environment. And all of this was made possible while still adhering to the central tenets of crypto and blockchain, allowing anyone to permissionlessly interact with code that the network executes in a trustless fashion.

Soon, hundreds of developers began building decentralized applications (dapps) on Ethereum, leveraging smart contracts for a wide range of use cases that now span DeFi, NFTs, DAOs, and more. And despite inspiring the creation of many well-funded, competing smart contract blockchains, Ethereum still commands the largest developer base supporting its industry-leading ecosystem of dapps. This has resulted in over $35b of value locked in the network, and the second highest market cap in crypto, behind that of only bitcoin at just shy of $200b.

Like Bitcoin, Ethereum currently uses a proof-of-work consensus mechanism, allowing the decentralized network of unknown parties to agree on which transactions should go into a block and onto the blockchain. And by requiring block producers to expend something of value - in this case, significant computational and energy resources - Ethereum prevents bad actors from attempting to execute a double spend by reverting a block from the blockchain via an alternative longer chain in a malicious chain reorganization or reorg. And the odds that any individual miner solves the mining puzzle and receives the block reward/miner tips are proportional to the amount of work, or hashrate, contributed, encouraging and rewarding greater participation in the consensus process and further securing the network.

Ethereum, however, has long planned a transition to proof-of-stake, where block producers, known as validators under proof-of-stake rather than miners under proof-of-work, run full nodes, bond or stake the protocol's native token, propose blocks when chosen to do so, and attest to the validity of a particular proposer's block when not selected. Validators are chosen at random to produce a block in proportion to their stake. Importantly, by requiring and choosing block producers based on stake, attempting to revert a finalized block in proof-of-stake is asymmetrically expensive and even more costly than reorging a block under proof-of-work. Validators receive staking rewards for

the work they perform but may be forced to pay a small penalty fee if they perform their duties poorly, or they may even be slashed and removed from the network if they behave maliciously. Such actions provide the network with the ever-important function of consensus and provide the network with its security.

Ethereum's transition to proof-of-stake, known as The Merge, has been a long and arduous one with years of delay, but it now appears to be on the imminent horizon. In fact, Ethereum's difficulty bomb, a feature intended to incentivize the transition to proof-of-stake by exponentially increasing the difficulty of mining ETH, has been pushed back six times since its first implementation in 2015. Despite this, developers have now transitioned several testnets to proof-of-stake, have tested the transition on many other shadow forks, and are now comfortable enough that they announced September 15th as the likely date of mainnet's transition to proof-of-stake. And once complete, proof-of-stake validators assume the network's block production and security responsibilities, replacing the role of proof-of-work miners today.

While perhaps the most prominent and certainly the most topical, The Merge is one of many important technical upgrades designed to enhance Ethereum's scalability, decentralization, and/or security. Last December, Vitalik outlined Ethereum's roadmap, breaking it into five categories known as The Merge, The Surge, The Verge, The Purge, and The Splurge, illustrated in the exhibit below. Importantly, this exhibit reads from left to right, and developers are working on components of each category simultaneously. This series of upgrades was historically referred to as the Ethereum 2.0 roadmap until a rebranding earlier this year, yet the roadmap is still colloquially referred to as Ethereum 2.0 by many in the community. In what follows, we first introduce Ethereum's implementation of proof-of-stake on the Beacon Chain before diving into each category of upgrades outlined in the roadmap.

Exhibit 1: Ethereum's Roadmap

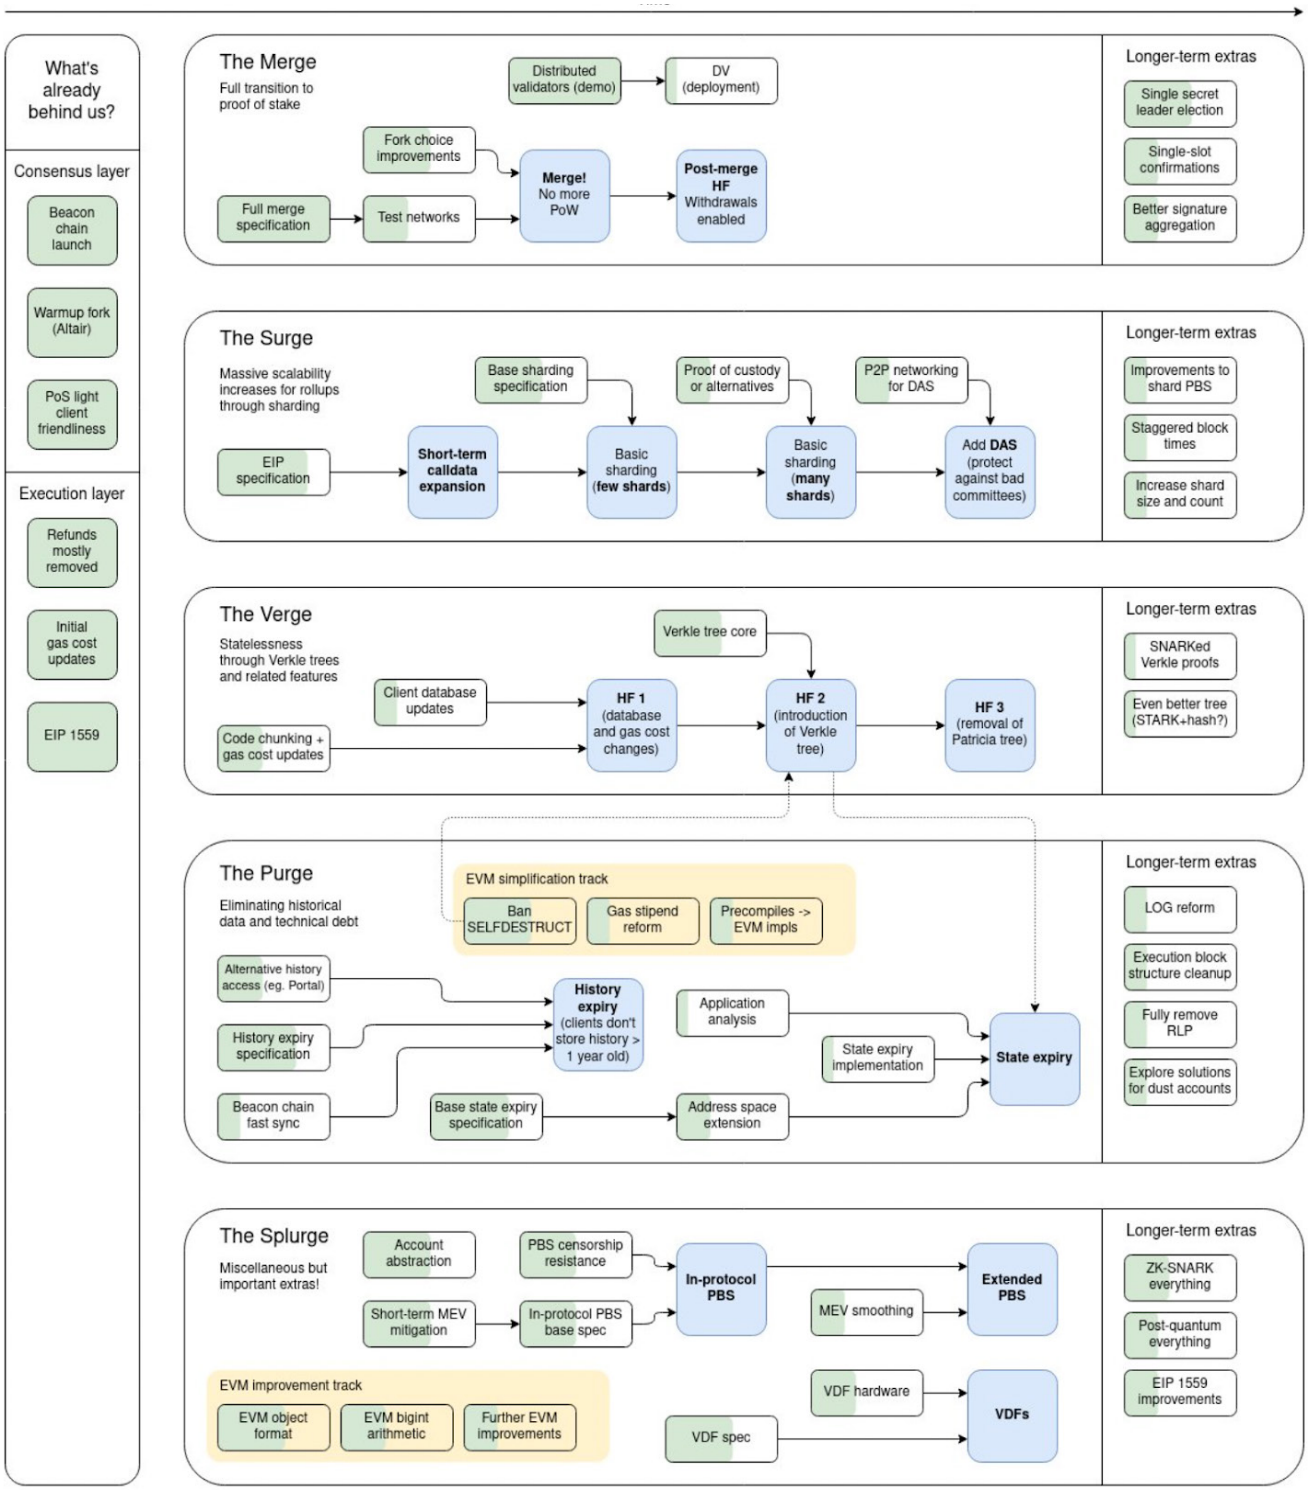

Source: Vitalik Buterin as of December 2021. GSR. Note that the illustration reads by time from left to right, development across the five sections of the roadmap occurring in parallel. However, the roadmap is from year-end and remains fluid. For example, short-term calldata expansion (EIP-4488) was expected to be the first upgrade across each of these phases, but it's still in review and appears unlikely to be implemented before The Merge. Further, it may not be implemented at all if PDS is prioritized.

The Beacon Chain

Launched in December 2020 as the first concrete step in Ethereum's transition to proof-of-stake, the Beacon Chain, also known as the consensus layer, is a distinct proof-of-stake blockchain running in parallel to Ethereum's proof-of-work mainnet. Perhaps the biggest challenge with The Merge is that Ethereum cannot be easily shut down or paused to implement the upgrade, as it exists not on a centralized server but simultaneously on thousands of computers (nodes) that run one of the open source Ethereum software implementations known as a client. This challenge is perhaps best illustrated by an analogy where Ethereum can be viewed as a spaceship (multiple functional layers - consensus, execution, settlement, data availability) with the Beacon Chain being a newly built engine (consensus), and after substantial testing, The Merge will swap out the spaceship's engine mid-flight. Hence, the Beacon Chain was launched as an independent blockchain to eventually swap out Ethereum's existing consensus mechanism while generally retaining the rest of Ethereum's history and other functionality. This approach allows for live testing while permitting everyday Ethereum users and their assets to remain segregated until developers are comfortable executing The Merge.

Nodes operating on Ethereum mainnet today simply need to run an execution layer (EL) client implementation such as Geth or Erigon, while nodes on the Beacon Chain are required to run a different consensus layer (CL) client implementation such as Prysm or Lighthouse. Additionally, since validator account balances on the Beacon Chain are credited by depositing 32 ETH into a smart contract on mainnet that currently functions as a one-way bridge, Beacon Chain nodes must also watch the staking deposit contract on an EL client to be aware of any new validator instances or changes to validator account balances. Many validators outsource this responsibility pre-Merge to third-party node providers such as Infura or Alchemy, however.

Validators on Ethereum have several roles, including attesting to their view of the chain, proposing new blocks, and participating in sync committees that support light client functionality. And while the Beacon Chain has not been processing transactions just yet, as execution is currently taking place on mainnet, validators on the Beacon Chain have been reaching consensus on state and agreeing on the active validator set and their corresponding account balances. While this may seem superficial today, given that mainnet uses proof-of-work consensus currently, the Beacon Chain will be the coordination layer for arriving at distributed consensus for all mainnet transactions after The Merge.

Validators are frequently chosen to vote on their view of the chain head block, where the chain head block is the most recent block in the validator's view of the canonical chain. Each vote is known as an attestation, and the attestation process relies heavily on time in the Beacon Chain, which is divided into periods lasting 6.4 minutes, known as epochs, and each epoch is further partitioned into 32 distinct slots, each lasting 12 seconds. Each slot has one designated block proposer that the protocol randomly selects to propose a block in that particular slot. Hence, new blocks come every 12 seconds unless a selected proposer fails to deliver a block leading to an empty slot, in which case the next block would be expected to arrive 24 seconds after the previous block. Every validator is placed on one beacon committee per epoch, and each beacon committee is randomly assigned to a particular slot and required to attest to their view of the chain head block during their assigned slot. Beacon committees are equally spread amongst the 32 slots in an epoch, resulting in 1/32 of the total validator set attesting to the validity of each slot/block (assuming a

block is proposed in each slot)[1]. By dividing validators into beacon committees, the network cuts down on messaging requirements, allowing for individual attestations to be aggregated in parallel and gossiped at the committee level. In fact, each slot has multiple beacon committees of validators, all attesting to the same information in that particular slot, so the number of aggregated attestations per slot will align with the number of committees per slot in an idealized example. Each beacon committee makes a single attestation per epoch before being disbanded and the process restarting anew in the next epoch. A small set of validators are also chosen at random to join sync committees (which are different from the aforementioned beacon committees), which pay additional rewards to validators and help light clients sync up and determine the head of the chain. Sync committees are particularly lucrative as participating validators receive a reward for each slot, and the selection lasts for 256 epochs, or 8,192 slots before a new committee is selected.

Exhibit 2: Epochs, Slots, and Beacon Committees Illustrated

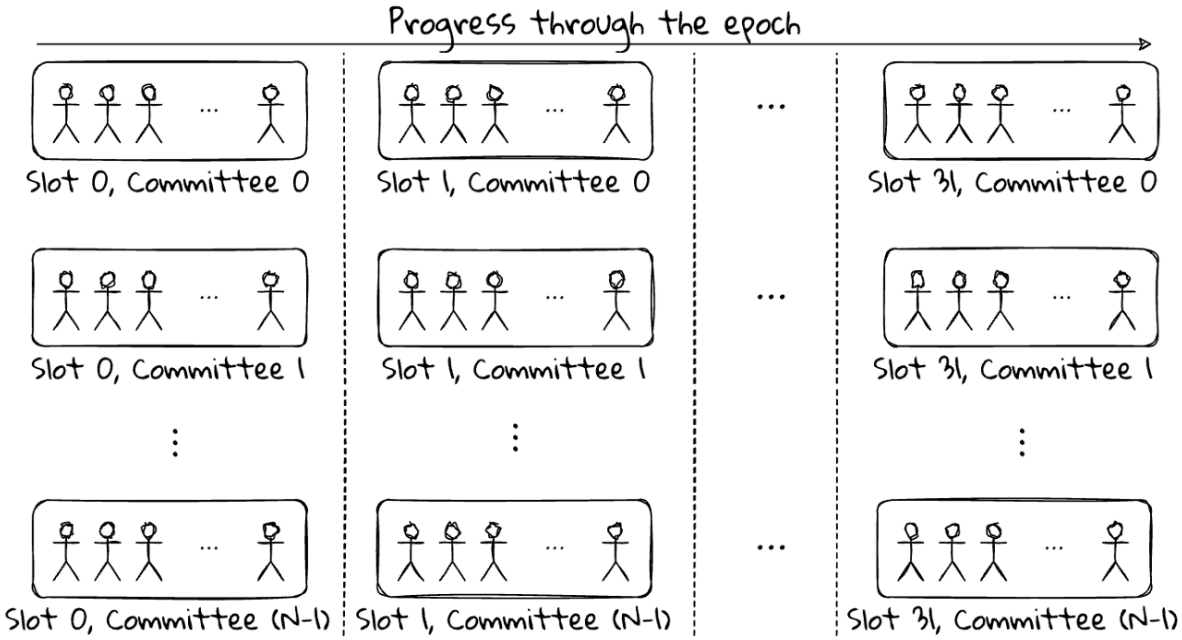

Source: Upgrading Ethereum, GSR.

The Beacon Chain employs a proof-of-stake consensus protocol named Gasper, which the Ethereum team designed internally. Gasper combines the Casper FFG finality mechanism, a practical Byzantine Fault Tolerant (pBFT)-inspired finality gadget for the realization of proof-of-stake, with the LMD GHOST fork-choice rule, which at a high level selects the chain by choosing the branch with the most attestations[2]. By doing so, Gasper combines the low overhead benefits that allow for a high number of participants to support decentralization seen in longest chain systems with the finality benefits of a pBFT-inspired system. In short, Gasper favors liveness over safety, continuing to produce blocks even if finality thresholds aren't reached, which may result in a fork, but it always keeps the chain moving forward and producing blocks (liveness). Alternative approaches favoring safety like Tendermint will not allow for forks (safety), but they cease block production and halt when finality thresholds are not met.

Gasper uses a system of checkpoint attestations of prior blocks, which requires a supermajority of attestation votes and increases the cost of reorganizing the blockchain prior to such checkpoints. Every epoch has one checkpoint, and that checkpoint is a hash identifying the latest block at the start of that epoch[3]. Validators attest to their view of two checkpoints every epoch, and the validator also runs the LMD GHOST fork-choice rule to attest to their view of the chain head block. The two checkpoint blocks are known as a source and a target, where the source is the earlier of the two checkpoint blocks. Generally speaking, the target checkpoint is the validator's view of the block at the start of the current epoch, while the source checkpoint is their view of the most recent 'justified' checkpoint. Simply put, the most recent 'justified' checkpoint is the most recent checkpoint that received more than two-thirds of the stake weight voting for its inclusion in the canonical chain; typically, this will refer to the previous epoch's checkpoint. If more than two-thirds of the total validator stake vote to link two adjacent checkpoint blocks, then there is a supermajority link between these checkpoints, and they both achieve an increased level of security. The target checkpoint becomes 'justified,' and the source checkpoint, which is already justified from a prior supermajority link, becomes 'finalized.' A checkpoint typically receives the necessary votes to become finalized after two epochs, and once a checkpoint is finalized, all previous slots become finalized. Reversing a finalized block would require malicious action by two-thirds of the total validator stake, and resultantly, the protocol guarantees they would be slashed at least one-third of the total network stake[4]. Assuming an ETH price of $2,000 and using the ~13.3m ETH locked in the deposit contract currently, reversing a finalized block would cost the attacker(s) more than $8.8b. This is referred to as economic finality - while a finalized Beacon Chain block can be reversed at a later date, unlike a protocol that achieves absolute finality such as Tendermint, it is impossible to do so without having a prohibitively large amount of stake slashed.

Exhibit 3: Gasper Checkpoint Blocks & Finality Illustrated

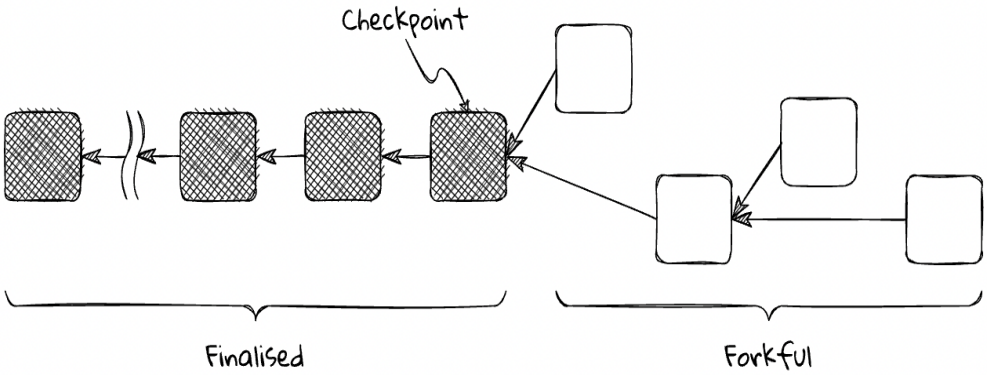

The honest nodes have agreed that the checkpoint and all its ancestor blocks are "final" and will never be reverted. There are therefore no forks before the checkpoint. The chain descending from the checkpoint remains liable to forking.

Source: Upgrading Ethereum, GSR. Note that the checkpoint block illustrated in the graphic represents the source checkpoint. The target checkpoint is unlabeled, and it would live in the forkful section of the illustration as it's not finalized.

Gasper's inclusion of an explicit finality mechanism helps deter reorg attempts by illuminating the cost to reorg a finalized block, which is notably distinct from existing proof-of-work chains that rely on probabilistic finality where an attacker's ability and cost to reorganize the chain is probabilistic in nature. Additionally, proof-of-stake has an asymmetric cost advantage that should disincentivize chain reorgs even more so than proof-of-work. The cost to a miner of attempting a chain reorganization and failing under proof-of-work is the electricity cost of their hashrate and the opportunity cost of coins that could have been mined on the canonical chain. The proof-of-stake reorganization equivalent requires a malicious validator to front as much as two-thirds of the total Ethereum stake, understanding that they will be slashed at least one-third of the total network stake after reorganizing a finalized block. Ethereum researcher Vlad Zamfir notably analogized that it's the equivalent of a miner's entire ASIC farm burning down as soon as it participated in a reorganization.

The Beacon Chain has additional mechanisms in place, such as the inactivity leak, to ensure finality is reached eventually, even if it's temporarily disrupted. Whether the impediment is from validators being offline due to a client issue or a fork caused by a consensus disagreement, the inactivity leak is designed to penalize validators that impede finality by failing to attest to the chain, and it will eventually allow for the chain(s) to finalize as the impeding party accrues quadratically growing penalties until a supermajority is reclaimed. Slashing and penalties are the features that underpin Ethereum's strong economic finality guarantees.

Rewards and penalties are aggregated across slots and paid to validators every epoch. Rewards issued for validating the chain are dynamic and depend on the total amount of ETH staked in the network. Specifically, the total ETH issued to validators in aggregate is proportional to the square root of the number of validators. This mechanism incentivizes validators with larger issuance

as the validator set grows and attracting additional validators becomes less essential. As an example, with the ~410k validators on the network today, about ~602k ETH would be issued annually, representing an average yield from issuance of ~4.6% across validators. However, the average yield from issuance would fall to about 3.3% if the validator count was to double. Note that these numbers simply show the total issuance over the total stake or the average yield paid across all validators, but individual validators will achieve different yields based on their performance, as well as other uncontrollable factors.

Exhibit 4: Annual ETH Issuance by Validator Participation

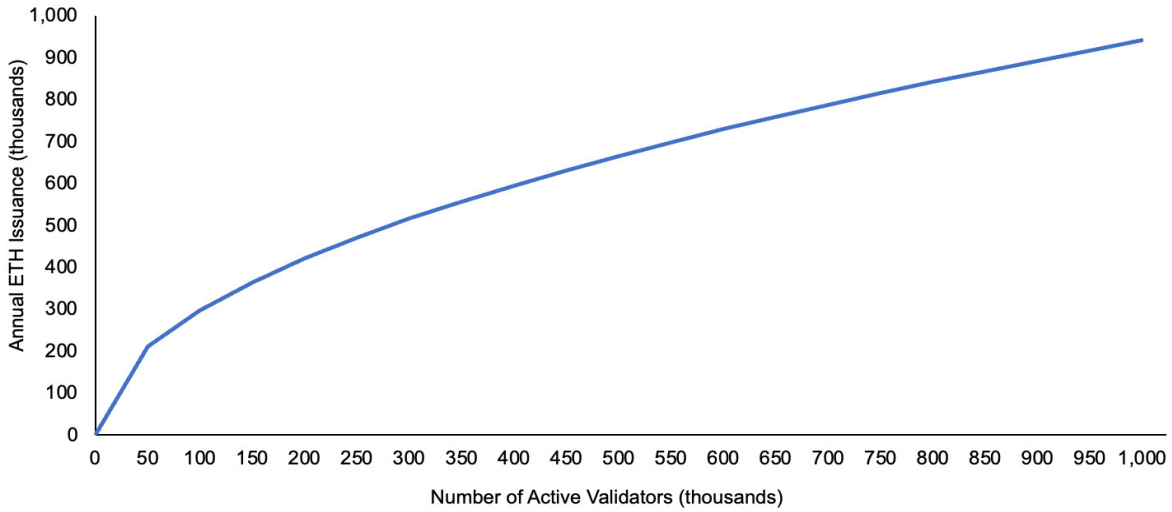

Source: GSR, Upgrading Ethereum. The graphic only denotes ETH issued to proof-of-stake validators, which will account for all gross ETH issuance after The Merge, but this only represents ~10% of total ETH issuance today under proof-of-work. The illustration assumes the Beacon Chain is running optimally, validators are performing their duties perfectly, and all validators have a 32 ETH effective balance. Actual issuance will be lower than illustrated as validators do not behave optimally in practice, but data since the launch of the Beacon Chain indicates that live validator performance is only a few percentage points below optimal.

A substantial portion of validator rewards are derived from attestations, as every validator will make one attestation during each epoch. Rewards from attestations also depend on the validator's timeliness in attesting and their correctness relative to the block proposer's view of the chain. Attesting too slowly or incorrectly will result in rewards turning into penalties. In addition, the rewards realized by individual validators will further vary as incremental rewards accrue to the randomly selected block proposers and sync committee participants. Additionally, the rewards scale with a validator's effective balance and with the total participation rate of other validators in the set. In short, this essentially means that validators with a balance below 32 ETH due to penalties for going offline or being slashed for malicious behavior will have their rewards scaled downward versus validators with a 32 ETH balance. Conversely, a validator's effective balance is capped at 32 ETH, so one's proportional share of rewards does not continue to grow beyond this level as their balance grows, making it suboptimal to hold a balance much higher than 32 ETH. Lastly, rewards between validators could further diverge as validators elect to participate in MEV by leveraging modified software such as Flashbot's MEV-Boost, but this is notably outside of the Beacon Chain protocol for now.

While this section aimed to set the stage for Ethereum's upcoming roadmap, we acknowledge that the full depth and nuance of the Beacon Chain deserves coverage far beyond the scope of this report. For readers interested in a deeper understanding of the Beacon Chain, we found Ben Edgington's Upgrading Ethereum to be the most detailed resource.

The Merge

Arguably the most important upgrade in Ethereum's seven-year history, The Merge formally deprecates Ethereum's proof-of-work consensus mechanism in favor of proof-of-stake. The Merge is actually a sequence of two upgrades, a CL client upgrade known as Bellatrix and an EL client upgrade known as Paris, which encompasses the improvement proposals of EIP-3675 and EIP-4399. Bellatrix will occur on September 6th, and it gives the Beacon Chain logic to be aware that The Merge is coming, while Paris is the actual Merge itself, where the consensus mechanism is switched in real-time. The Merge will result in Ethereum mainnet and the Beacon Chain merging together under a new consensus mechanism while maintaining the network's full transaction history under proof-of-work. After The Merge, Ethereum miners will cease to exist, and validators will leverage Gasper to achieve consensus, decreasing the network's aggregate energy consumption by more than 99.9%, according to estimates.

The Merge will be triggered when the chain reaches a pre-specified terminal total difficulty (TTD) level, which is a measure of the total cumulative mining power used to build the proof-of-work chain since genesis. Once a proof-of-work block is added to the chain that crosses the preset TTD threshold, no additional proof-of-work blocks will be produced from this point on. While prior mainnet upgrades have been triggered at a predetermined block number, The Merge uses TTD to reduce attacks where a bad actor directs hashrate to a forked chain, pulling forward the block number, which would force a chain reorg similar to a 51% attack. Upon hitting TTD, Ethereum EL clients will toggle off mining and cease their gossip-based communication about blocks, with similar responsibilities now being assumed by CL clients. The two distinct blockchains that were historically running in parallel will have merged into the Beacon Chain, and new blocks will be proposed and extend the Beacon Chain as usual, but with transaction data that was historically included in proof-of-work blocks.

Exhibit 5: Ethereum Block Architecture Pre & Post Merge

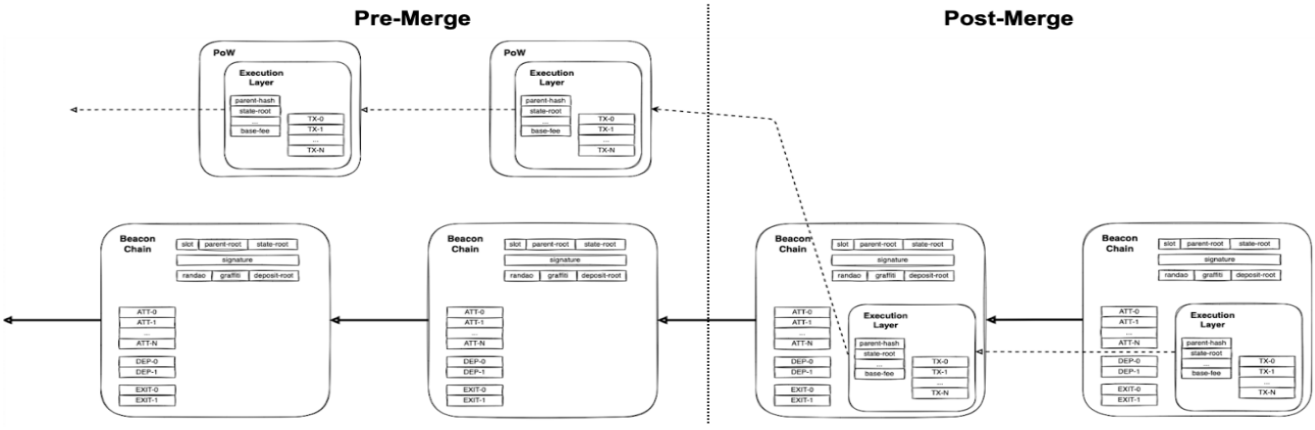

Source: Danny Ryan, Tim Beiko. Annotations by GSR. The Merge requires a CL client upgrade (Bellatrix) and an EL client upgrade (Paris; EIP-3675 & EIP-4399) which will happen ~1.5 weeks apart, yet the graphic illustrates it as a single upgrade for simplicity. We would recommend this post to those interested in a very precise series of events.

One notable challenge associated with The Merge is the sheer number of pairwise combinations between consensus and execution layer clients. Unlike Bitcoin, which has a single reference implementation in Bitcoin Core, post-Merge Ethereum nodes must run an execution client and a consensus client paired together, with the implementations chosen at the discretion of the node operator. Further, Ethereum has multiple distinct client teams independently developing and implementing the EL and CL protocol specifications. Ignoring client implementations with less than one percent of the user base, there are four EL client implementations and four CL client implementations, according to clientdiversity.org. This creates 16 distinct pairs of EL and CL client implementations that all need to interoperate seamlessly. While client diversity certainly makes The Merge more challenging, it enhances Ethereum's decentralization and security by reducing various risks that arise when a network's stake is concentrated in a single client implementation. Additionally, the Beacon Chain explicitly incentivizes client diversity by incorporating a correlation penalty that escalates the punishment when a slashing event occurs that impacts a large portion of the network's stake. The inactivity leak further punishes correlated failures that impede finality. Hence, running a majority client increases one's correlation with other validators, increasing the potential cost of a client bug that results in a large set of validators going offline or being slashed.

Building the Beacon Chain specification and battle-testing the client implementations is no small feat, and Ethereum developers have run through a large number of tests aiming to simulate The Merge in a controlled environment. Around 20 shadow forks, which are simply copies of the state of a network used for testing purposes, have been executed across mainnet and Goerli, allowing developers to trial The Merge through a large suite of live network conditions. Shadow forks work by coordinating a small number of nodes to fork off the canonical chain by pulling their Merge implementation timeline ahead of the live network. These shadow forks provide for a more realistic testing environment than simulating The Merge on newly launched testnets like Kiln that don't have a long history of transactions, and they were an important step before merging existing public testnets. Developers have since merged the Ropsten, Sepolia, and Goerli testnets in June, July, and August, respectively. While many bugs were identified, these tests were widely viewed as successful, and Goerli's Merge was the final test that precipitated developers to specify the TTD level of 58750000000000000000000 to activate The Merge on mainnet. Based on the Ethereum hashrate mining currently, The Merge is likely to occur on September 15th, but the expected date can be monitored in real-time here.

While The Merge is expected to be minimally disruptive to most participants of the Ethereum network, there are a few important changes to be aware of. Importantly and as discussed above, the upgrade will now require full nodes to run an EL client and a CL client. In contrast, transactions and blocks could previously be received, validated, and propagated with a single EL client. Moving forward, both EL and CL clients will have a unique peer-to-peer (p2p) network. The CL client will gossip blocks, attestations, and slashings while the EL client will continue to gossip transactions, handle execution, and maintain state. The two clients will leverage the Engine API to communicate with each other, forming a full post-Merge Ethereum node in tandem. In addition, Ethereum applications are not expected to be materially affected by The Merge, but certain changes like a marginally decreased block time and the removal of proof-of-work-related opcodes like difficulty could impact a subset of smart contracts.

Another byproduct of The Merge is a substantial reduction in the amount of new ETH issued. Currently, about 15,100 new ETH are issued daily, equating to ~4.6% of Ethereum's supply on an annualized basis, with ~13,500 of this going to miners and the remaining ~1,600 going to validators. After The Merge, however, no more ETH will be issued to miners, and total issuance will fall by nearly 90%, with gross annualized issuance representing 0.5% of the supply. Moreover, net issuance may be deflationary, as gas fees burned under EIP-1559 may more than offset the new, lower issuance schedule. In fact, it will only take an average base fee of ~15 gwei to fully offset ETH issuance. For context, the median base fee has averaged ~58 gwei since EIP-1559 was implemented last August, but gas prices have notably fallen and are averaging closer to 13 gwei over the last two months. Further, priority fees, which are equivalent to miner tips and are the unburned portion of gas fees, will begin to accrue to validators, and these fees are expected to increase the staking yield by ~50%. Immediately following The Merge, validators are expected to be able to generate a ~7% nominal staking yield with the potential for a higher real yield depending on the ETH burn dynamics. Lastly, it's important to note that Beacon Chain validators will still not be able to withdraw their stake until the Shanghai upgrade that's anticipated to come about six to twelve months after The Merge. As a result, all new ETH issuance will be illiquid as it will accrue to validator accounts where it cannot be withdrawn or transferred until after the next upgrade. And even then, there are validator exit limits in place to prevent a simultaneous run to the exits after staked ETH becomes liquid.

All told, a successful Merge will result in many changes and positive benefits. Ethereum's current energy consumption is expected to fall by ~99.9% as proof-of-work miners are replaced by more energy-efficient proof-of-stake validators. As touched upon in the Beacon Chain section, block times will decrease from a ~13.3-second average to a constant 12 seconds (assuming no empty slots)[5], and the network will offer stronger economic finality guarantees underpinned by slashing and other validator penalties. From an economic perspective, large expenses paid to miners for security will dissipate, and new ETH issuance will fall by ~90% as block rewards paid to miners cease. And finally, validator staking yields are expected to rise by ~50% as transaction priority fees will begin to accrue to validators, as well as MEV revenue which we will cover in the last section of this report.

The Surge

Another major upgrade is The Surge, which refers to the set of upgrades commonly referred to as sharding that are designed to help Ethereum scale transaction throughput. For traditional databases, sharding is the process of partitioning a database horizontally to spread the load, and in earlier Ethereum roadmaps, it aimed to scale throughput on the base layer by splitting execution into 64 shard chains to support parallel computation, with each shard chain having its own validator set and state. However, as layer two (L2) scaling technologies developed, Vitalik Buterin proposed a rollup-centric scaling roadmap for Ethereum in October 2020, simplifying the long-term Ethereum roadmap by deemphasizing scaling at the base layer and prioritizing data sharding over execution sharding. The updated roadmap aims to achieve network scalability by moving virtually all computation (i.e., execution) over to L2 while making it cheaper for data to be posted back to Ethereum mainnet. Simply put, computation is already very cheap on L2s, and the majority of L2 transaction fees today are driven by the cost of posting the computed data back to mainnet. Hence, improving mainnet's ability to make data cheaply available will be the largest driver in decreasing L2 transaction costs. The updated roadmap retrenched the focus of

Ethereum's mainnet to consensus, settlement, and data availability, allowing L2 platforms to compete in the free market to provide execution.

Currently, rollups post their state roots back to Ethereum using calldata for storage. Calldata is the cheapest form of storage on Ethereum today, but it's still expensive as the data goes through the EVM and is logged permanently in the blockchain's history. While a full primer on rollups is beyond the scope of this piece, rollups do not need permanent data storage but only require that the data is temporarily available for a short period of time. More precisely, they require data availability guarantees ensuring that data was made publicly available and not withheld or censored by a malicious actor. Hence, despite calldata being the cheapest data solution available today, it is not optimized for rollups or scalable enough for their data availability needs.

Ethereum's current plan for sharding is known as Danksharding (DS). However, instituting full Danksharding is complex, leading the community to support an intermediate upgrade offering a subset of the DS features known as Proto-Danksharding (PDS; EIP-4844) to achieve meaningful scaling benefits more quickly. PDS introduces a new Ethereum transaction type called a Blob-carrying transaction which allows for data to be posted in 'blobs.' Blob-carrying transactions are like regular transactions, but they also include an extra data blob attached that allows the protocol to provide data availability guarantees without committing to permanently store that data. This new transaction type will materially increase the amount of data available for rollups to interpret since each blob, which is roughly 125 kB, is larger than an entire Ethereum block on average. Blobs are purely introduced for data availability purposes, and the EVM cannot access blob data, but it can only prove its existence. The full blob content is propagated separately alongside a block as a sidecar. Blob transactions have their own independent, EIP-1559-style gas market, targeting eight blobs per block (~1 MB) up to a maximum of 16, with gas prices adjusting exponentially as the demand for blobs deviates from the target. This segregated fee market should yield efficiencies by separating the cost of data availability from the cost of execution, allowing the individual components to be priced independently based on their respective demand (i.e., an NFT mint on mainnet will not increase the price a rollup pays for data availability). Further, data blobs are expected to be pruned from nodes after a month or so, making them a great data solution for rollups without overburdening node operators with extreme storage requirements.

Despite PDS making progress in the DS roadmap, the name is perhaps a misnomer given each validator is still required to download every data blob to verify that they are indeed available, and actual data sharding will not occur until the introduction of DS. The PDS proposal is simply a step in the direction of the future DS implementation, and expectations are for PDS to be fully compatible with DS while increasing the current throughput of rollups by an order of magnitude. Rollups will be required to adjust to this new transaction type, but the forward compatibility will ensure another adjustment is not required once DS is ready to be implemented. Developers are aiming for PDS to be included in the Shanghai hard fork about six to twelve months after The Merge.

While the implementation details of DS are not set in stone, the general idea is simple to understand: DS distributes the job of checking data availability amongst validators. To do so, DS uses a process known as data availability sampling, where it encodes shard data using erasure coding, extending the dataset in a way that mathematically guarantees the availability of the full data set as long as some fixed threshold of samples is available[6]. DS splits up data into blobs or shards, and every validator will be required to attest to the availability of their assigned shards

of data once per epoch, splitting the load amongst them. As long as the majority of validators honestly attest to their data being available, there will be a sufficient number of samples available, and the original data can be reconstructed. In the longer run, private random sampling is expected to allow an individual to guarantee data availability on their own without any validator trust assumptions, but this is challenging to implement and is not expected to be included initially.

DS further plans to increase the number of target shards to 128, with a maximum of 256 shards per block, materially increasing the target blob storage per block from 1 MB to 16 MBs. While an increased block size isn't an issue for nodes validating the network as they can verify the block efficiently with data availability sampling, it does represent a centralizing force for block builders that will need to compute the blob encoding and distribute the data. This increase in validator requirements would be detrimental to the diversity of the network, so an important upgrade from The Splurge, known as Proposer-Builder Separation (PBS), will need to be completed first. For those interested in a deeper understanding of the long-term DS roadmap, The Hitchhiker's Guide to Ethereum offers thorough coverage.

Exhibit 6: Danksharding Illustrated

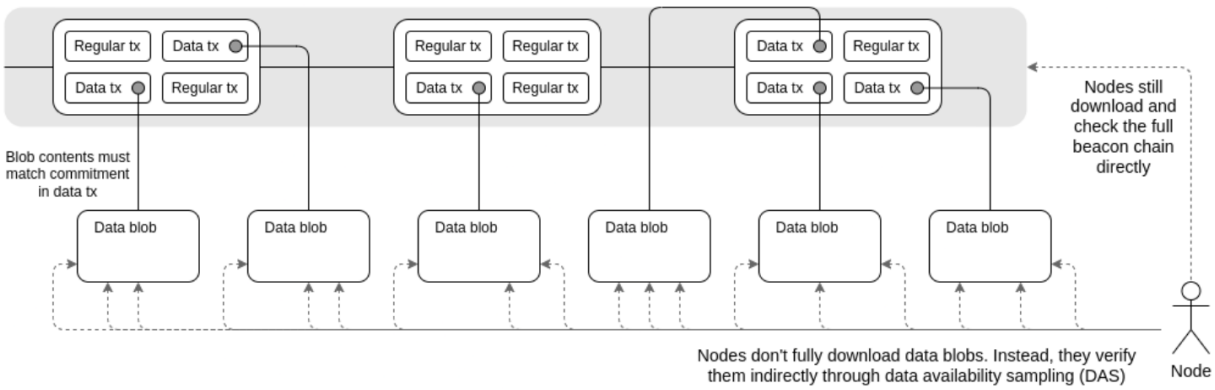

Source: Vitalik Buterin, GSR.

In summary, The Surge focuses on scaling and improving Ethereum's transaction throughput. However, many still misconstrue sharding as scaling Ethereum execution at the base layer, which is no longer the medium-term objective. The sharding roadmap prioritizes making data availability cheaper and leaning into the computational strengths of rollups to achieve scalability on L2. Many have highlighted DS as the upgrade that could invert the scalability trilemma as a highly decentralized validator set will allow for data to be sharded into smaller pieces while statistically preserving data availability guarantees, improving scalability without sacrificing security.

The Verge

The Verge is a series of upgrades that aims to introduce statelessness, a feature that removes the requirement for validating nodes to maintain a copy of Ethereum's state to validate transactions. Ethereum's state is an extensive database encompassing all externally-owned accounts and their balances, smart contract deployments, and associated storage. In addition to its considerable existing size, Ethereum's state is continuously growing as new users join the network and developers deploy new contracts. Presently, all of the data in Ethereum's state is hashed together and compressed into a Merkle-Patricia Tree. And in the current design, Ethereum nodes must store the state to validate blocks and ensure that the network transitions between states correctly. This growing storage requirement increases the hardware specifications to run a full node over time, which could have a centralizing effect on the validator set.

The permanence of state also creates a unique scenario as a user pays a one-time gas fee to send a transaction in exchange for an ongoing cost to the network via permanent node storage requirements. The Verge aims to alleviate the burden of state on the network by replacing the current Merkle-Patricia state tree with a Verkle Tree, a newer data structure first described in 2018. A vital property of both tree structures is that anyone can make a short proof known as a 'witness' that some piece of information is an element of the tree, and anyone can easily verify this proof against the publicly available state root. However, Verkle proofs are much more efficient in proof size compared to Merkle proofs. Unlike a Merkle-Patricia Tree, which requires more hashes as the tree widens with more children, Verkle Trees use vector commitments that allow the tree width to expand without expanding the witness size. For a review of Merkle Trees, see How Bitcoin Works.

The transition to Verkle Trees will allow stateless clients to proliferate as smaller witnesses enable direct block inclusion. Instead of validators maintaining a local copy of Ethereum's state, block builders will provide a Verkle proof giving the portions of the state impacted in a particular block, as well as the proof ensuring the accuracy of these pieces, allowing validators to verify blocks without maintaining Ethereum's state. Stateless clients will enable fresh nodes to immediately validate blocks without ever syncing the state as they would simply request the required block information and proof from a peer. Ethereum aims for 'weak statelessness,' meaning validators can verify blocks without maintaining a copy of Ethereum's state, but block builders will still need the state to construct a block. The assumptions underpinning 'weak statelessness' do not pose material concerns as block builders will be more specialized under PBS and will be able to manage any state growth.

In short, The Verge aims to decrease the hardware requirements of validator nodes by introducing stateless clients that can verify blocks without downloading a local copy of Ethereum's state, which requires an increasingly large amount of solid-state storage to maintain. Enabling nodes to validate the network primarily with RAM will increase validator decentralization.

The Purge

The Purge refers to a series of upgrades aimed at simplifying the protocol by reducing historical data storage and technical debt. Most prominently, it aims to introduce history expiration (EIP-4444) which could potentially come in the months following The Merge. History expiration requires nodes to stop serving historical blocks on the p2p network that are more than one year old, and it would give nodes the option to locally prune this same set of historical blocks, allowing them to be deleted from the node's local copy. Importantly, once a node is fully synced to the head of the chain, validators do not require historical data to verify incremental blocks. Hence, historical data is only used at the protocol level when an explicit request is made via JSON-RPC or when a peer attempts to sync the chain. After EIP-4444, new nodes will leverage a different syncing mechanism, like checkpoint sync, which will sync the chain from the most recently finalized checkpoint block instead of the genesis block.

The deletion of history data is primarily a concern for individual Ethereum-based applications that require historical transaction data to show information about past user behaviors. History storage is viewed as a problem that would be best handled outside of the scope of the Ethereum protocol moving forward, but clients would still offer the ability to import this data from external sources. It's expected that applications will have multiple different solutions to obtain history data outside of Ethereum, including from block explorers like Etherscan, indexing providers such as The Graph, or a more decentralized, Ethereum Foundation-supported protocol like the Portal Network. Removing history data from Ethereum would significantly reduce the hard disk requirements for node operators, and it would allow for client simplification by removing the need for code that processes different versions of historical blocks.

In addition to history expiration, The Purge includes state expiry, which prunes state that has not been touched in some defined amount of time, such as one year into a distinct tree structure, removed from the Ethereum protocol. State expiry is the furthest out of all the upgrades outlined in the roadmap and only becomes feasible after the introduction of Verkle Trees. Expired state assets could always be retrieved by displaying a proof as long as some other party maintains a copy of the chain's history elsewhere. While state expiry is not as essential as it may initially seem after stateless clients are implemented, it will still decrease the strain of dust accounts and other inactive addresses on Ethereum's state.

The Splurge

The Splurge is a catch-all bucket for miscellaneous upgrades that don't fit neatly in any of the previous categorizations. Proposer-Builder Separation (PBS) is the most prominent upgrade in The Splurge as it directly impacts the roadmap for DS and statelessness.

Before expanding on PBS, a brief introduction to block building and MEV is needed, with our in-depth piece on MEV providing more detail for interested readers. In short, MEV is a measure of profit that a miner or validator can extract from block production beyond the block reward and gas fees by including, excluding, and changing the order of transactions in a block. It's commonly measured as the incremental gains achieved by deviating from a basic block building approach that simply orders transactions based on their priority fee. While miners are in a prime position to identify and capitalize on such transactions, as they control which transactions are included and in

what order, the majority of MEV is extracted by independent third parties called searchers that use sophisticated trading strategies to capture MEV. In practice, searchers identify and execute MEV extraction via bots that comb through the pending transactions pool and utilize various strategies such as DEX arbitrage, liquidations, and frontrunning/backrunning sandwich trades. MEV extraction is a fundamentally different skill set than participating in network consensus, and companies such as Flashbots have been created to illuminate, democratize, and redistribute MEV by serving as a neutral, public, open-source infrastructure for permissionless MEV extraction, allowing independent MEV searchers to communicate their bid and granular transaction order preference to mining pools to execute their ordered bundle of transaction. Competition between searchers to extract MEV results in much of the gains accruing to the block proposer in a competitive bidding process.

PBS, as the name implies, separates block builders from block proposers at the protocol level. The validator that is selected to propose the next block in the chain is known as the block proposer, and they outsource block construction (transaction selection and ordering) to a dedicated market of block builders. Under this model, dedicated block builders search for MEV opportunities to build the most profitable block and submit bids to block proposers to propose their block. A proposer's job is as simple as proposing the block of whichever builder offers them the highest fee. This eases the job of validators by selling the computationally difficult optimization problem to a more specialized entity and allowing validators to fulfill their responsibilities with materially lower hardware specifications. Additionally, PBS should redistribute the profit attributable to MEV, as multiple builders compete against each other in an auction, eroding their margins and returning most of the profit to validators. Perhaps ironically, the set-up somewhat resembles the scale economies inherent in proof-of-work. Building a block that maximizes total profit is a difficult problem that's best outsourced to specialized entities who benefit from economies of scale, but verifying the validity of this block remains very easy. This results in more centralized block production, but validation is still trustless and should be even more decentralized since block building responsibilities are delegated elsewhere.

While the specification details of in-protocol PBS are not fully decided at this point, censorship resistance is an explicitly categorized area of focus on the roadmap. The PBS implementation will include a censorship-resistance list (crLists) that the proposer publishes to display their view of censored transactions in the mempool. Proposers cannot force transaction inclusion on builders if they deliver a full block, but if a block is not full, builders will be required to include the proposer's selected transactions. In-protocol PBS will not be available immediately following The Merge, but validators will be able to leverage Flashbot's MEV-Boost for block production as needed in the interim.

Another notable upgrade in The Splurge is account abstraction, with the most prominent proposal being EIP-4337. This proposal lets users employ smart contract wallets as their primary Ethereum account instead of an externally-owned account (EOA), and it does so by leveraging a higher-layer account abstraction approach that avoids any Ethereum protocol changes. Specifically, EIP-4337 creates a separate mempool consisting of a higher-order transaction-like object called a UserOperation. A special set of users known as bundlers would aggregate UserOperations into a transaction that would directly communicate with a particular smart contract, and that transaction would then be included in a block on mainnet. This improves user experience by atomically batching operations into a single transaction that would otherwise require multiple different transactions to

execute on mainnet. Account abstraction would further provide user flexibility to deviate from the ECDSA digital signature algorithm and employ any arbitrary verification logic, such as a quantum-resistant signature scheme. It also simplifies the use of multisigs and social recovery wallets. Lastly, it introduces a form of gas abstraction where gas fees can be paid in ERC-20 tokens, and applications can subsidize the gas fees of their users.

Conclusion

In just 15 days, we will likely witness one of the most significant events in blockchain history, as the first (and apex) smart contract blockchain attempts a nearly impossible feat - to change its consensus mechanism mid-flight. The stakes couldn't be higher with billions at risk, though developers and community members have prepared for this moment for years. In just 15 days, the Beacon Chain will merge with Ethereum mainnet, as proof-of-work is switched off and proof-of-stake takes over. Validators will immediately begin proposing and attesting to blocks, as beacon committees are formed and disbanded at every epoch. Validators, following Gasper, will attest to both checkpoints and chain heads, identifying the canonical chain and introducing the notion of economic finality. And when all is said and done, energy consumption will plummet, finality guarantees will strengthen, ETH issuance will fall, and staking yields will rise, all upon reaching that fateful terminal total difficulty level.

Though a crowning achievement, Ethereum's work will be far from done, with Vitalik himself estimating that the network will be just 55% complete post-Merge. Ethereum, however, is equipped with a thoughtful and well-defined roadmap, and its developers have been hard at work perfecting the various upgrades to bring about their many benefits. Through Danksharding, with its data blobs and data availability sampling, The Surge will make data availability cheaper and distribute the job of checking data availability amongst nodes, providing material scalability benefits on L2. The Verge will achieve statelessness by implementing Verkle Trees, allowing stateless clients to verify blocks without maintaining a local copy of Ethereum's state. The Purge will introduce history expiration and state expiry, archiving history data, pruning untouched state, and generally simplifying the protocol. And The Splurge will add Proposer-Builder Separation, reducing validator hardware requirements and redistributing MEV, as well as account abstraction, increasing wallet choice/functionality and improving user experience. In the end, Ethereum is expected to process ~100,000 transactions per second with vast improvements in its already leading security and decentralization, enabling mass adoption and achieving Ethereum's initial goal of creating a trustless and permissionless world settlement layer for a diverse suite of dapps and beyond.

Footnotes:

1) Notably, validators are attesting to their view of the chain head block for LMD GHOST during their slot, which is generally but not necessarily the block proposed in their slot. In practice, blocks are proposed at the very beginning of each slot, so the block proposed in a given slot will generally be the chain head block that validators are attesting to in that slot, resulting in 1/32 of the validator set attesting to each block. However, if a block proposer does not deliver a block in their assigned slot, the validators in that slot would attest to their view of the chain head block, which would likely be the same chain head that validators in the previous slot attested to. Hence, it's not necessarily always 1/32 of the validator set attesting to each block. In aggregate, validators are delivering one attestation per epoch, but that attestation includes three items: 1) a vote on the source checkpoint, 2) a vote on the target checkpoint, and 3) a vote for the chain head block. Notably, the chain head vote uniquely determines the source and target vote, so strictly speaking, voting on all three is redundant and unnecessary, but it simplifies processing.

2) A fork-choice rule is simply a protocol that determines one's view of the head of the chain based on the available information. Bitcoin and Ethereum today use a rule that selects the longest chain, or more precisely, the chain with the most cumulative chainwork. Gasper, however, follows the chain containing the justified checkpoint that has the greatest block height without ever reverting a finalized block. From here, it essentially counts the accumulated votes from validators for blocks and their descendent blocks (the economically heaviest chain).

3) An epoch's checkpoint is generally going to be the first block in an epoch. However, a checkpoint is a hash identifying the latest block at the start of that epoch, and the block identified by a checkpoint's hash isn't always included in this new epoch because an empty slot can occur at the beginning of an epoch, so "the latest block" would reference the last block in the prior epoch. Checkpoints are known as epoch boundary blocks in other literature, which may help with intuition. Checkpoints are identified by their block root (hash) and an epoch number. A block can theoretically serve as the checkpoint for multiple epochs if the slots throughout remain empty.

4) Since checkpoint finalization requires a two-thirds supermajority, the finalization of two competing checkpoints would require attestations from at least four-thirds of the total stake weight, guaranteeing at least one-third of the total stake attested to two different checkpoints for the same epoch. In a normal scenario where the honest validators' views of the canonical chain are not divergent or split, the attacker would also need to procure two-thirds of Ethereum's total stake to execute this attack in the first place. Imagine an attacker has two-thirds of the total network stake, and honest validators possess the other one-third. Assuming honest validators all attest to the same checkpoint, the attacker could finalize this checkpoint by attesting with half of their stake, or one-third of the total network stake equivalently. The attacker could then unilaterally finalize a competing checkpoint by attesting to it with the entirety of their two-thirds stake, but as guaranteed by the first sentence in this footnote, this would result in them having one-third of the total network stake slashed as this would require them to maliciously reuse one-third of the total network stake that they already used to attest to the first checkpoint. Despite this attack requiring two-thirds of the stake to conduct, only half of the stake was required to attest to competing checkpoints, and this is the only provably malicious act that the protocol can detect, so the one-third is all that can be slashed. It could easily be argued that this attack should warrant the malicious actor being slashed the full two-thirds, but this is not detectable in-protocol and would require social slashing through a user-activated soft fork. It may be possible to revert a finalized block with less than two-thirds stake in the event that honest validators are partitioned on their view of the canonical chain, but regardless, any attacker is always guaranteed to be slashed one-third of the total network stake under all circumstances which seems to be a sufficiently large distinctive.

5) Based on the frequency of empty slots in the Beacon Chain historically, block times are expected to decrease below the ~13.3 second-average on mainnet currently, but the average block time will be greater than 12 seconds if there are any empty slots. Notably though, block times will be a constant 12 seconds if a block is proposed in each slot. Unlike proof-of-work, where block times are probabilistic in nature based on how quickly a miner solves the mining puzzle, all slots are exactly 12 seconds apart, and block times will only deviate from 12 seconds if a proposer fails to propose a block in their assigned slot. This deviation would naturally only come in multiples of 12 seconds.

6) It also requires KZG Polynomial Commitments to prove that the original data was encoded correctly. This is beyond the scope of this report and covered thoroughly in the sourced Delphi piece.

Authors:

Matt Kunke, Junior Strategist | Twitter, Telegram, LinkedIn
Brian Rudick, Senior Strategist | Telegram, LinkedIn

The authors would like to thank Justin Drake of the Ethereum Foundation for providing helpful comments, as well as Ben Edgington of ConsenSys for answering many questions on the Beacon Chain.

Sources:
Ethereum.org, All Core Devs Update, ETH Roadmap FAQ - Tim Beiko, Validator FAQ, Endgame - Vitalik Buterin & Bankless, Endgame - Vitalik Buterin, The Ethereum Merge - Tim Beiko & Epicenter, The Hitchhiker's Guide to Ethereum - Delphi Digital, Upgrading Ethereum - Ben Edgington, Ethereum 2.0 Knowledge Base, Combining Ghost & Casper, Bellatrix & Paris - Etherworld, Client Diversity - Dankrad Feist, EIP-3675 - Upgrade to Proof-of-Stake, EIP-4399 - Supplant Difficulty Opcode, Danksharding - Dankrad Feist, Dive Into Danksharding - Ethereum Foundation & Bankless, Proto-Danksharding - Protolambda, EIP-4844 - Proto-Danksharding, EIP4844.com, Data Availability Sampling - Paradigm, A Rollup-centric Ethereum Roadmap - Vitalik Buterin, Addressing Common Rollup Misconceptions - Polynya, Rollups + Data Shards - Polynya, A Theory of Ethereum State Size Management - Vitalik Buterin, Verkle Tree Integration - Vitalik Buterin, Verkle Trees - Vitalik Buterin, EIP-4444 - History Expiration, crList Proposal, EIP-4337 - Account Abstraction, EIP-1559 - Fee Market Change, Watchtheburn.com

About GSR

GSR has nine years of deep crypto market expertise as a market maker, ecosystem partner and active, multi-stage investor. GSR sources and provides spot and non-linear liquidity in digital assets for token issuers, institutional investors, and leading cryptocurrency exchanges. GSR employs over 300 people around the globe, and its trading technology is connected to 60 trading venues.

GSR's main service areas are: otc trading; market making; proprietary and algorithmic trading; client execution; structured products; risk management solutions; and portfolio investments.

For more information or if we can help with anything, please see gsr.io or contact us at gsr@gsr.io.

Required Disclosures