Pakistan Blockchain Institute

PAKISTAN BLOCKCHAIN INSTITUTE

Class-2

MODULE-3

ETHEREUM 2.0 EXPLAINER

Raja Rizwan Saleem
Lead Blockchain Trainer

Edversity.

SHARDING

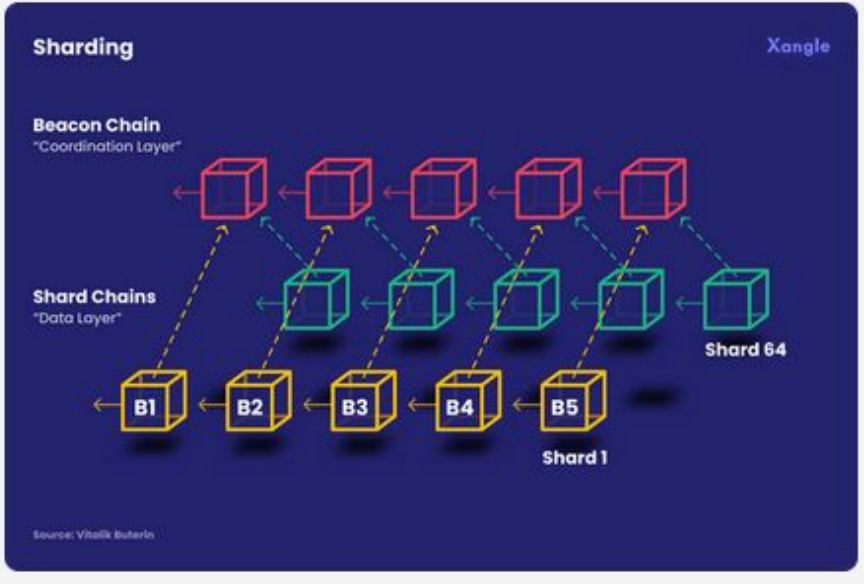
Sharding
Xangle
Beacon Chain
"Coordination Layer"
Shard Chains
"Data Layer"
Shard 64
B1
B2
B3
B4
B5
Shard 1
Source: Vitalik Buterin

Blockchain Trilemma

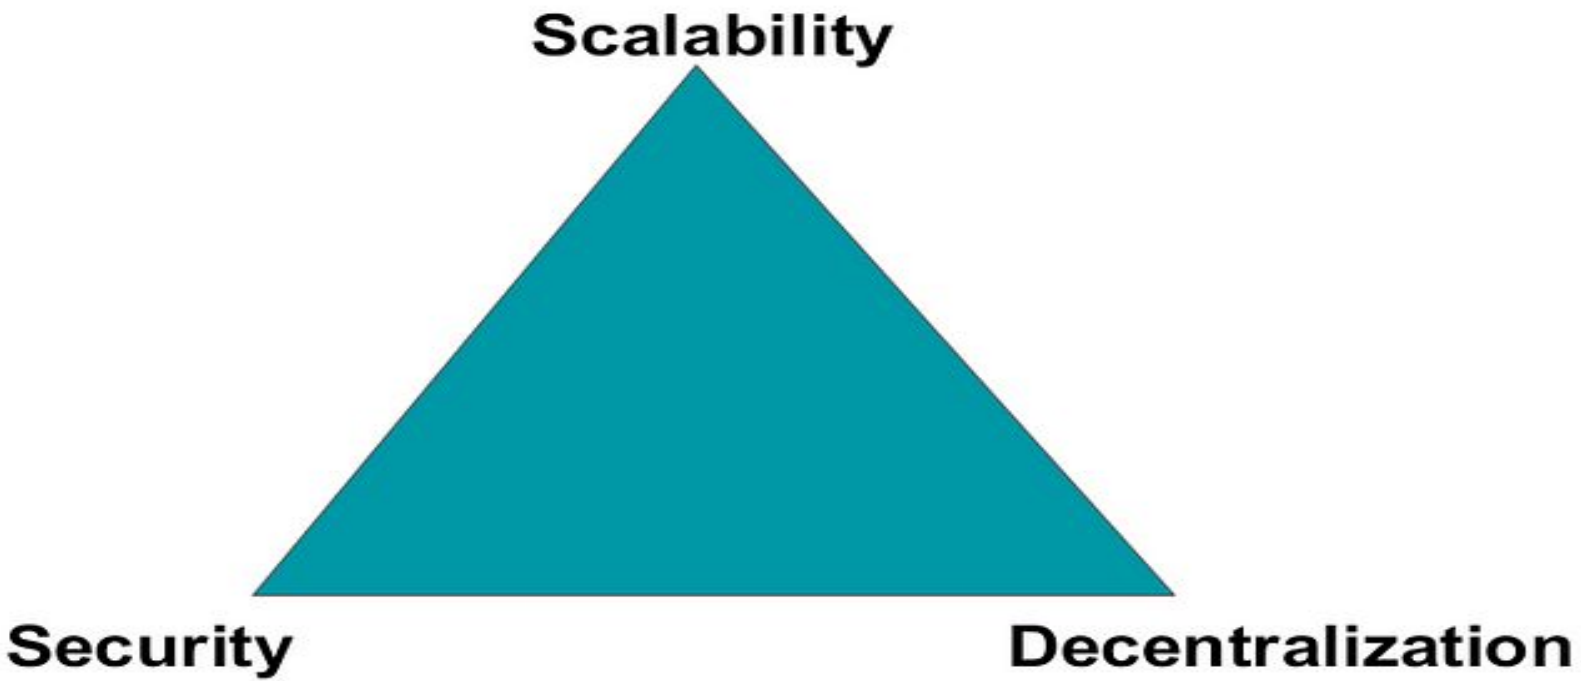

Sharding

- Sharding is actually very similar to partitioning, except that it's used on systems where data doesn't have one central location but rather can be spread out over many machines depending on its complexity.
- The word "sharding" is derived from "dividing". In a blockchain, sharding refers to the process of dividing the network into smaller blockchains. The nodes within each individual blockchain are identical and hence, operate in a similar fashion.
- Combined with other techniques such as pruning, sharding results in the blockchain handling more transactions at higher speeds while also reducing storage requirements.
- It also makes it possible for nodes to only process transactions that pertain to their area of interest (e.g., if you want updates about Bitcoin Cash instead of Bitcoin). This greatly increases security and avoids overloading any single node with information that is irrelevant to it.

Sharding

- Sharding is a scaling technique used in blockchain networks to increase their capacity to process transactions and handle a larger number of users and applications.
- The concept of sharding is akin to breaking a database or network into smaller, more manageable parts, or "shards," each capable of independent processing.
- This approach aims to alleviate the limitations of blockchain's single-chain architecture, where every transaction needs to be processed by all nodes in the network.
- In the context of blockchain, particularly Ethereum 2.0, sharding involves dividing the network into smaller, interconnected chains, or shards, which work together to process transactions.
- Each shard operates independently, processing its own transactions and executing smart contracts.

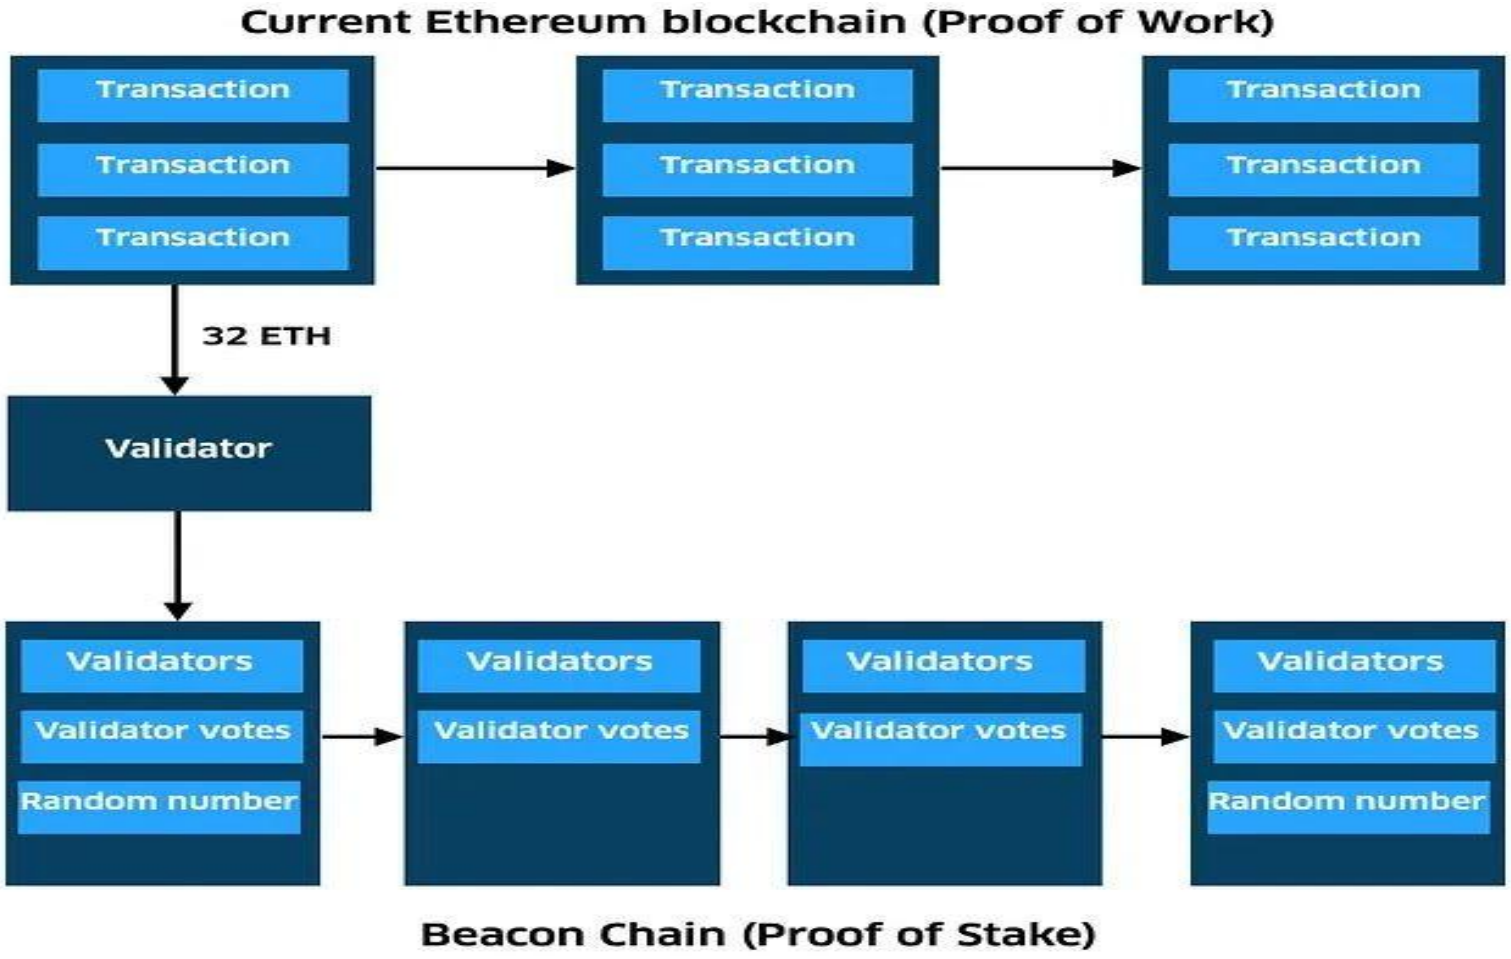
Current Ethereum blockchain (Proof of Work)
Transaction
Transaction
Transaction
Transaction
Transaction
Transaction
Transaction
Transaction
Transaction
32 ETH
Validator
Validators
Validator votes
Random number
Validators
Validator votes
Validators
Validator votes
Validators
Validator votes
Random number
Beacon Chain (Proof of Stake)

Sharding

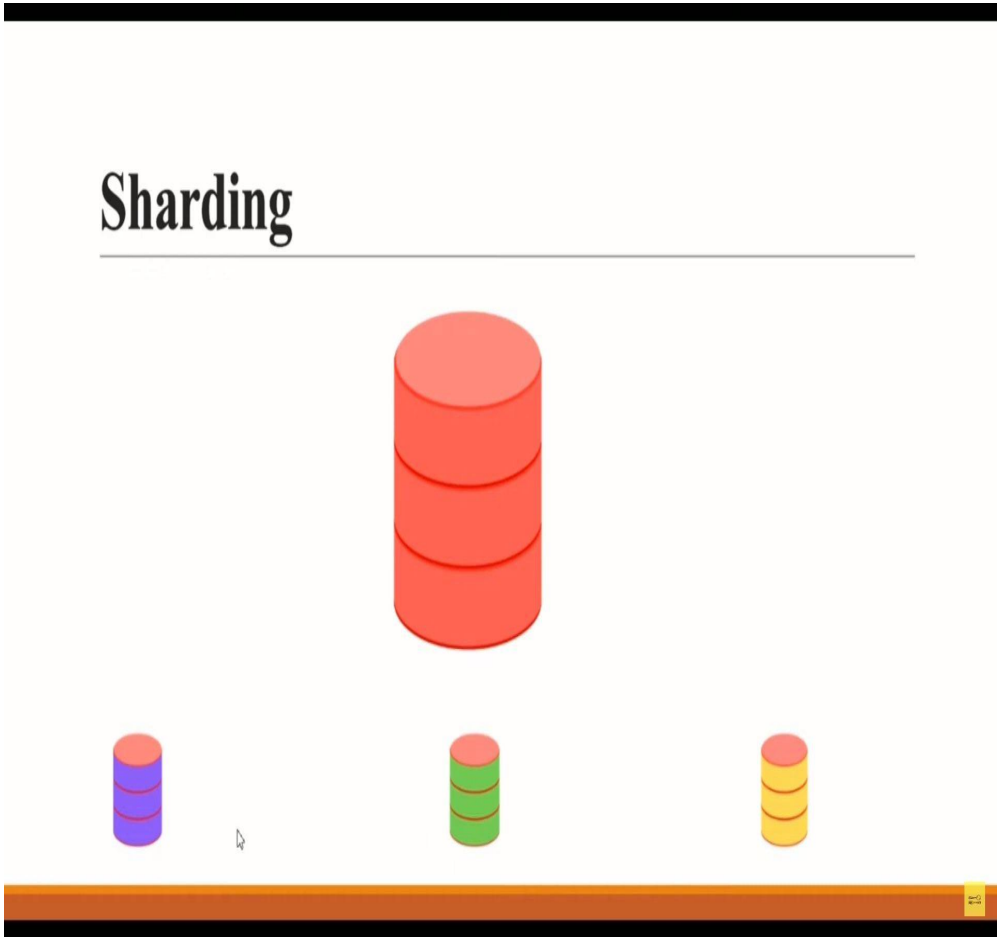

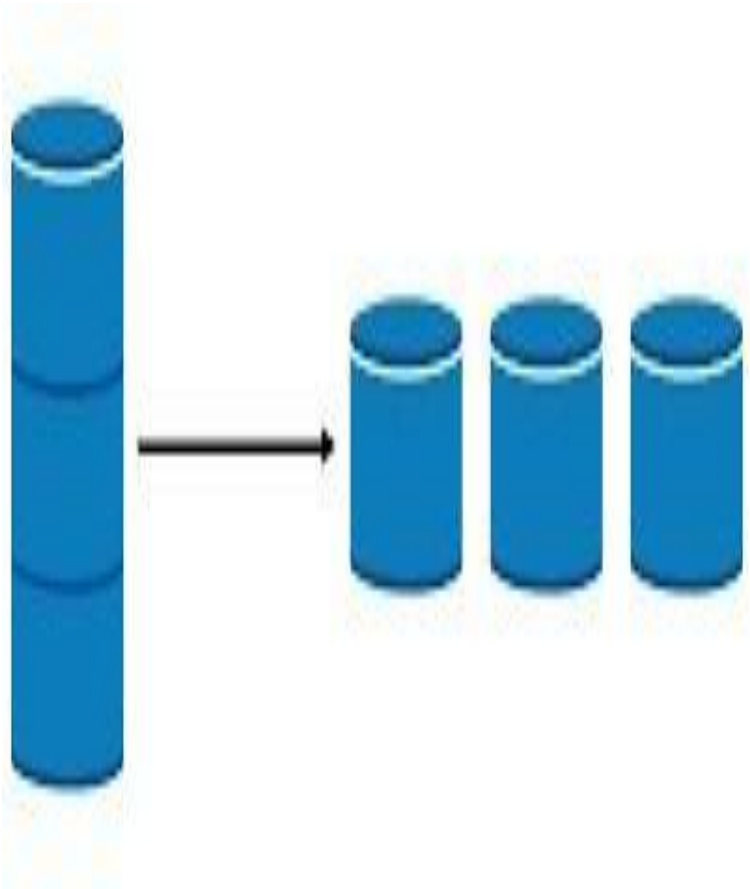

Sharding

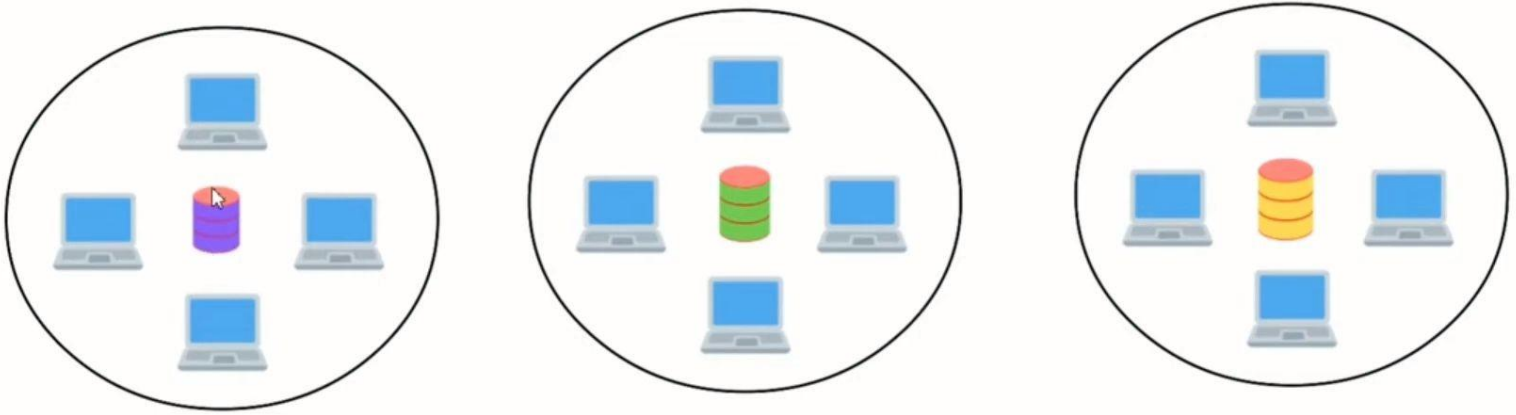

Network A

Network B

Network C

What is Shard?

A "shard" means a "small part of the whole." In database management, a shard is a subset of a large database hosted on a separate server. While each shard contains chunks of data, they all form one logical dataset.

Using our previous example, we could have on shard, "Shard 1," for city residents with surnames starting with 'A', "Shard 2," for those with surnames starting with 'B', and so on.

If you combine these logical shards, you'd get a single dataset of records for all city residents.

Sharding

splitting a blockchain into multiple pieces,
or shards, and storing them in different places

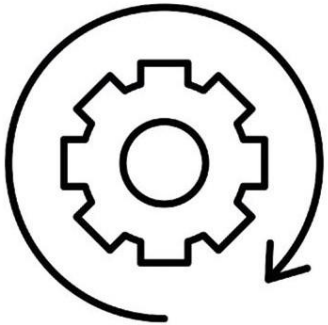

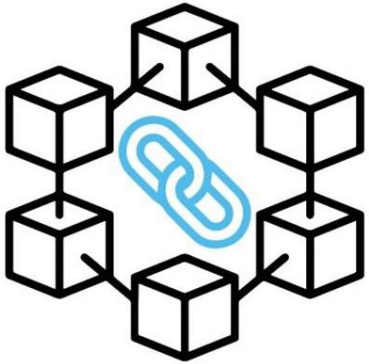

Using sharding, its possible for nodes to function
without having to maintain all of that data at once

parallel execution model

transactions will be processed simultaneously
and in parallel on each shard

network nodes will process only certain specific operations and not the all as they were doing in the linear model

horizontal partitioning

rows of the same database being distributed across multiple nodes

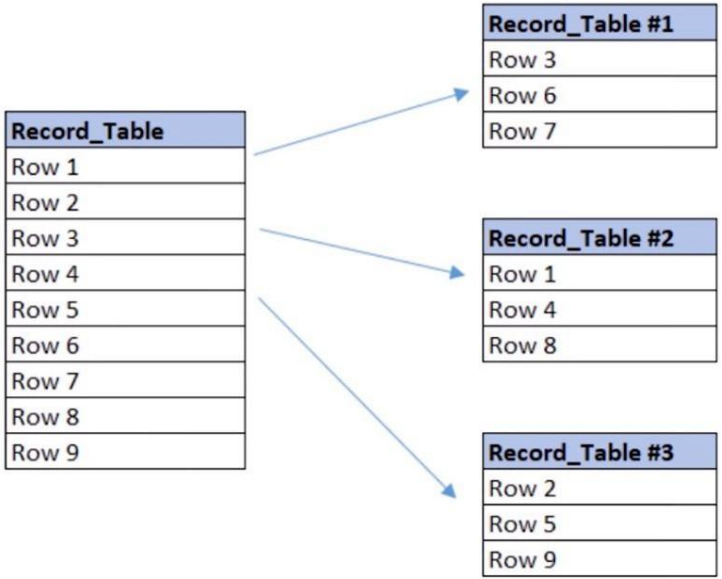

partitions

each partition has the same schema and columns, but all the respective rows are entirely different

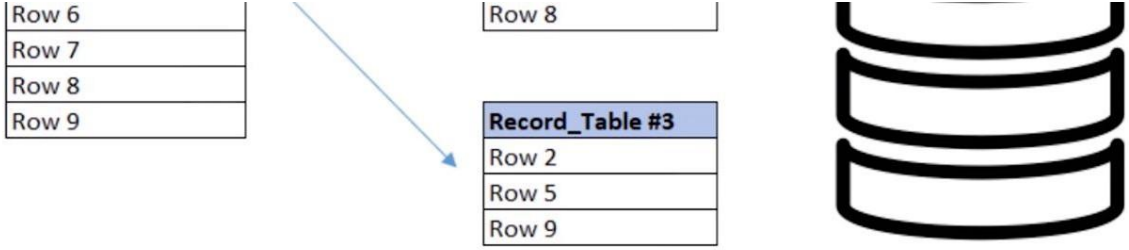

vertical partitioning

different information in a separate databases

A	B	C	D	E
1	2	3	4	5

A	D
1	4

B	C	E
2	3	5

How Sharding is Done

Proof of stake (or PoS)

nodes validate transactions based on
the amount of tokens they have staked

stakers dealing with different shards
of the same blockchain, and accordingly
processing a network transaction

Proof of Work (or PoW)

because network nodes face difficulty validating transactions with only the information from a single shard, and not the whole network

Full Node

archive a copy of the blockchains entire history on itself

SHARDING

full nodes no longer have to store or process the entirety of the networks activities

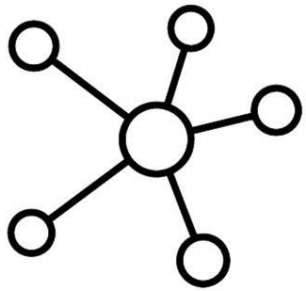

Edversity.

SUBSCRIBE

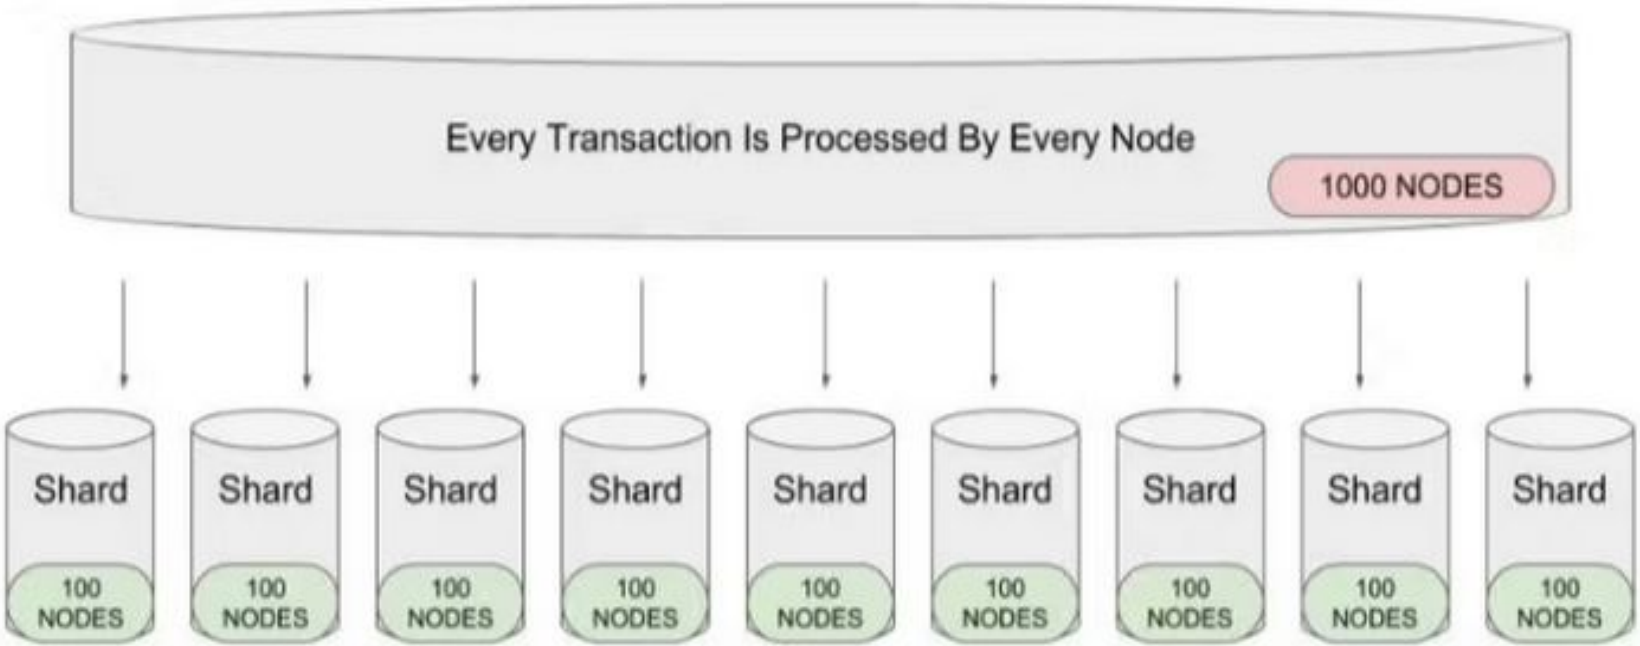

1000 nodes can be divided into 10 shards (100 nodes each) to achieve 10x performance.

Ethereum Sharding Terminology

State

State refers to the information about a system at any point in time. In Ethereum, state is a description of the network at a particular time—contract code, accounts, address balances, etc. Every new transaction alters Ethereum's state.

Merkle Tree

A Merkle tree or root is a cryptographic mechanism that stores large amounts of information via hashes. Merkle trees/roots are essential for Ethereum's security, as they allow nodes to quickly verify if a piece of data is part of the larger structure.

Ethereum Sharding Terminology

Collation

A collation is a group of transactions conducted on a shard chain, similar to a block in proof-of-work (PoW). Collations are submitted to the main chain and linked together to form the blockchain.

Collation Header

The collation header is similar to a block header in proof-of-work consensus. A collation header contains metadata about the information inside the collation such as:

- The single shard that the collation belongs to
- The root hash of the parent collation
- The Merkle root of all transactions in a collation
- The pre-state root and post-state root
- Signatures of notaries

Ethereum Sharding Terminology

Ethereum Sharding Terminology

Notaries

Notaries are validators randomly assigned to a shard chain to vote on proposed collations (blocks). These votes are called "attestations" and prove collation validity. Every collation needs at least ⅔ of collators to sign off on it before being added to the consensus chain.

Proposers

A proposer is a collator (or validator) selected to create a collation and propose it for validation. The proposer has the same duties as a miner in PoW blockchains.

Committees

A committee is a collection of validators or notaries that attest the validity of shard blocks. These committees are randomly shuffled at intervals, so validators cannot predict which committee(s) they'll be in.

Sharding A-Big Picture

1. **Shard Creation:**
   In a sharded blockchain network, a central chain, often referred to as the "Beacon Chain," manages and coordinates the operation of the shards. The Beacon Chain ensures that shards are synchronized, validators are assigned to shards, and consensus is maintained across the network.

2. **Shards' Independence:**
   Each shard operates like a mini-blockchain, containing its own subset of validators and transactions. Shards can process transactions and execute smart contracts without needing consensus from the entire network, which significantly improves scalability.

Sharding A-Big Picture

3. **Parallel Processing:** Shards work in parallel, enabling multiple transactions and smart contracts to be processed simultaneously across different shards. This parallel processing capability boosts the overall throughput of the blockchain network.
4. **Cross-Shard Communication:** Although shards operate independently, there's still a need for communication and coordination among them. Cross-shard communication mechanisms allow transactions that involve multiple shards to be processed smoothly. This is essential for maintaining consistency and integrity.

Sharding Benefits

- **Scalability:** Sharding greatly increases the network's transaction processing capacity, making it more suitable for mainstream adoption and high-demand applications.
- **Reduced Latency:** With transactions distributed across multiple shards, the overall network can process transactions more quickly, reducing latency.
- **Lower Fees:** Increased throughput and reduced congestion can lead to lower transaction fees, enhancing the user experience.
- **Energy Efficiency:** By processing fewer transactions per shard compared to a single-chain architecture, sharding can reduce energy consumption.

How might Ethereum PoS Sharding have worked in practice?

Imagine Ethereum has 10,000 validators and 100 shard chains.

Through a pseudorandom protocol, eligible validators, who have deposited ETH in the Validator Manager Contract (VMC), and are assigned to shards 1-100.

In Shard 1, a validator (proposer) is selected to group new transactions into a collation.

Other validators (notaries) download the collation and verify the validity of transactions.

If two-thirds of notaries attest to the collation, it is submitted to the main chain via the VMC.

It's important to note that the entire collation isn't added to the Beacon Chain—it would be difficult and time-wasting to verify collations from every shard.

Instead, the validator nodes on the main chain simply check the attestations (signatures) for each collation to determine its validity.

How might Ethereum PoS Sharding have worked in practice?

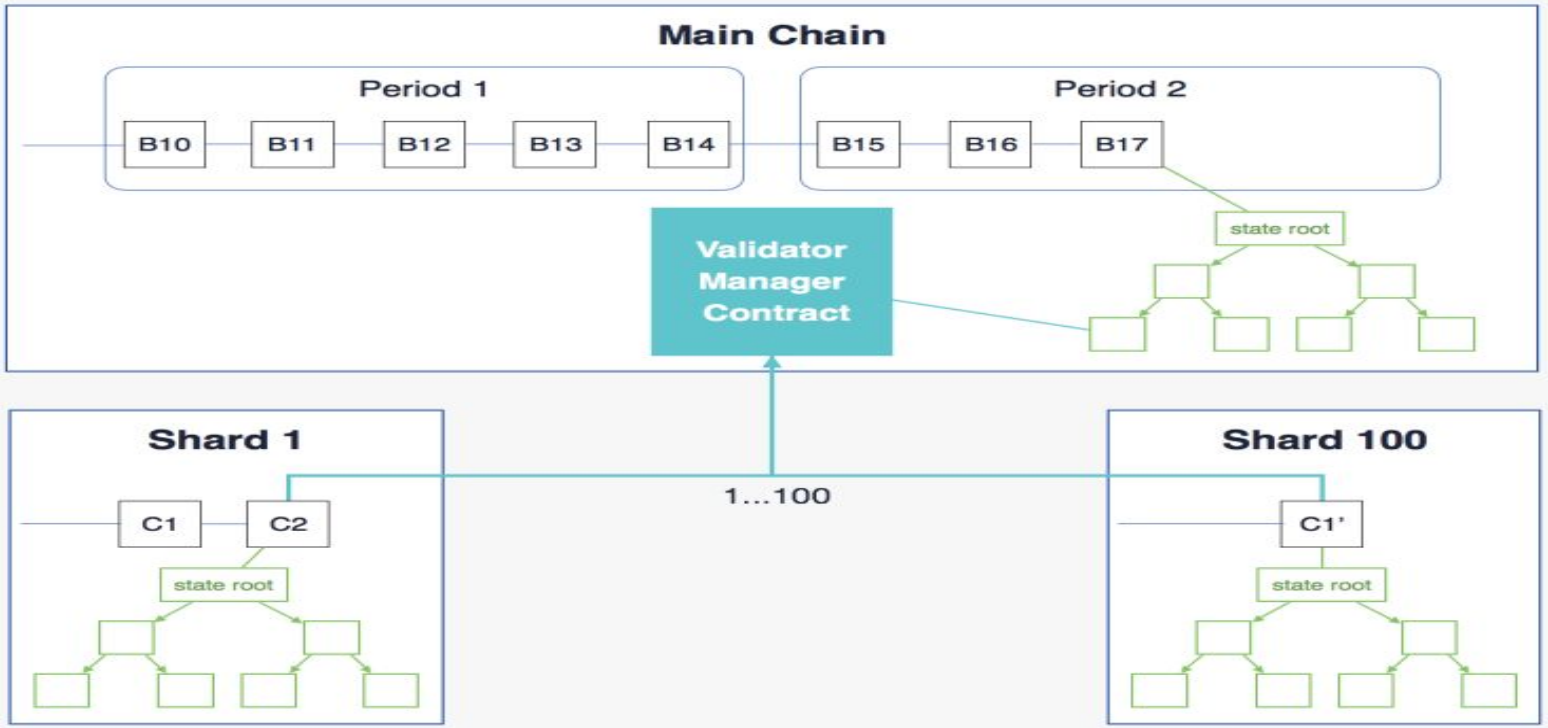

Major benfits

- Transactions per second increase.
- Powerful and expensive computers will not be needed.
- More validators will join.
- Energy consumption will reduce.

Sharding Challenges

- Ensuring secure cross-shard communication
- Managing data availability across shards.
- Maintaining overall network security.

Sharding

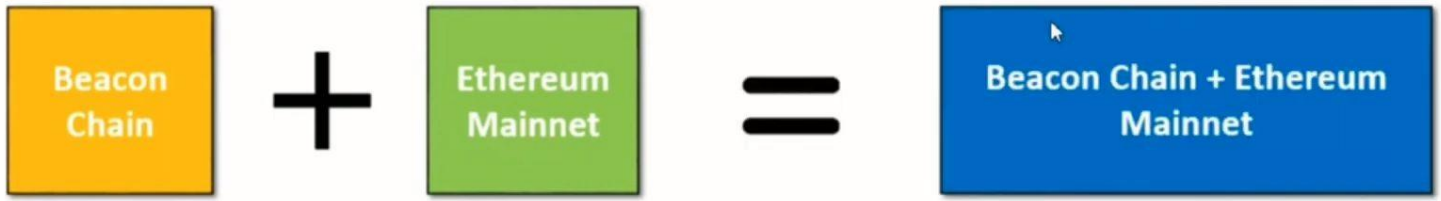
Beacon Chain
+
Ethereum Mainnet
=
Beacon Chain + Ethereum Mainnet

Sharding

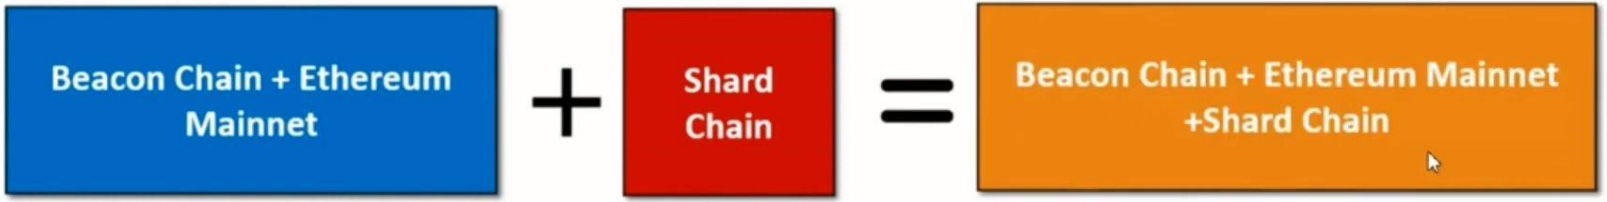
Beacon Chain + Ethereum Mainnet
+
Shard Chain
=
Beacon Chain + Ethereum Mainnet +Shard Chain

Cross-shard communication

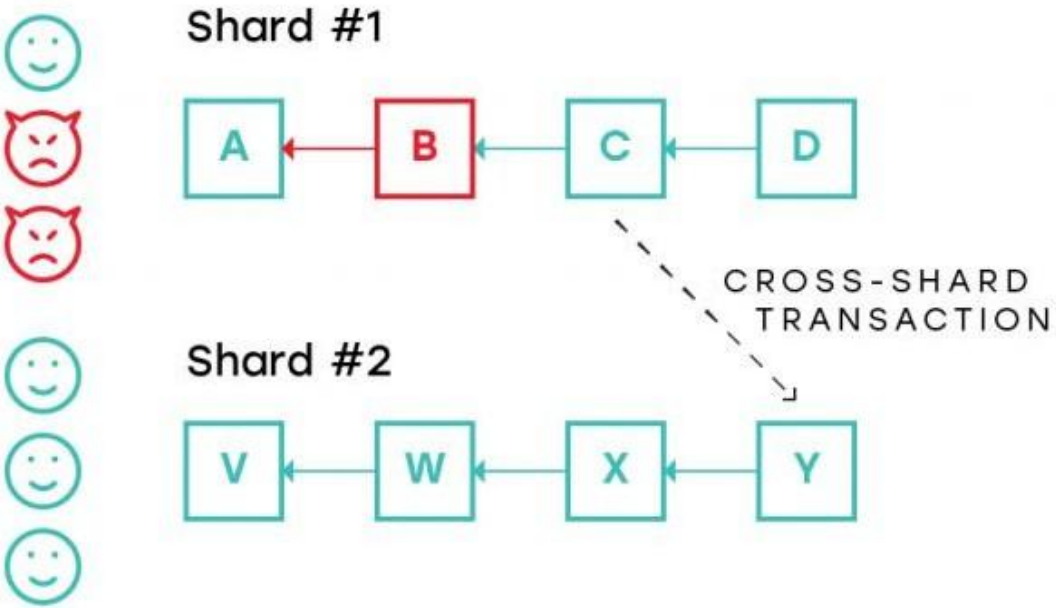

What is a Shard Chain

In the context of blockchain networks, a shard chain would contain a portion of the data and handle a portion of the transaction processing responsibilities.

Shard chains are like a collection of mini-blockchains that operate independently, and to preserve security, each shard chain submits a record of transactions to the main chain (Beacon Chain) at regular intervals through the Validator Manager Contract (VMC).

Because each shard chain will have a unique transaction history and a set of nodes to validate new transactions, multiple shard chains can run simultaneously to bolster network latency and throughput through parallel processing.

What is a Shard Chain

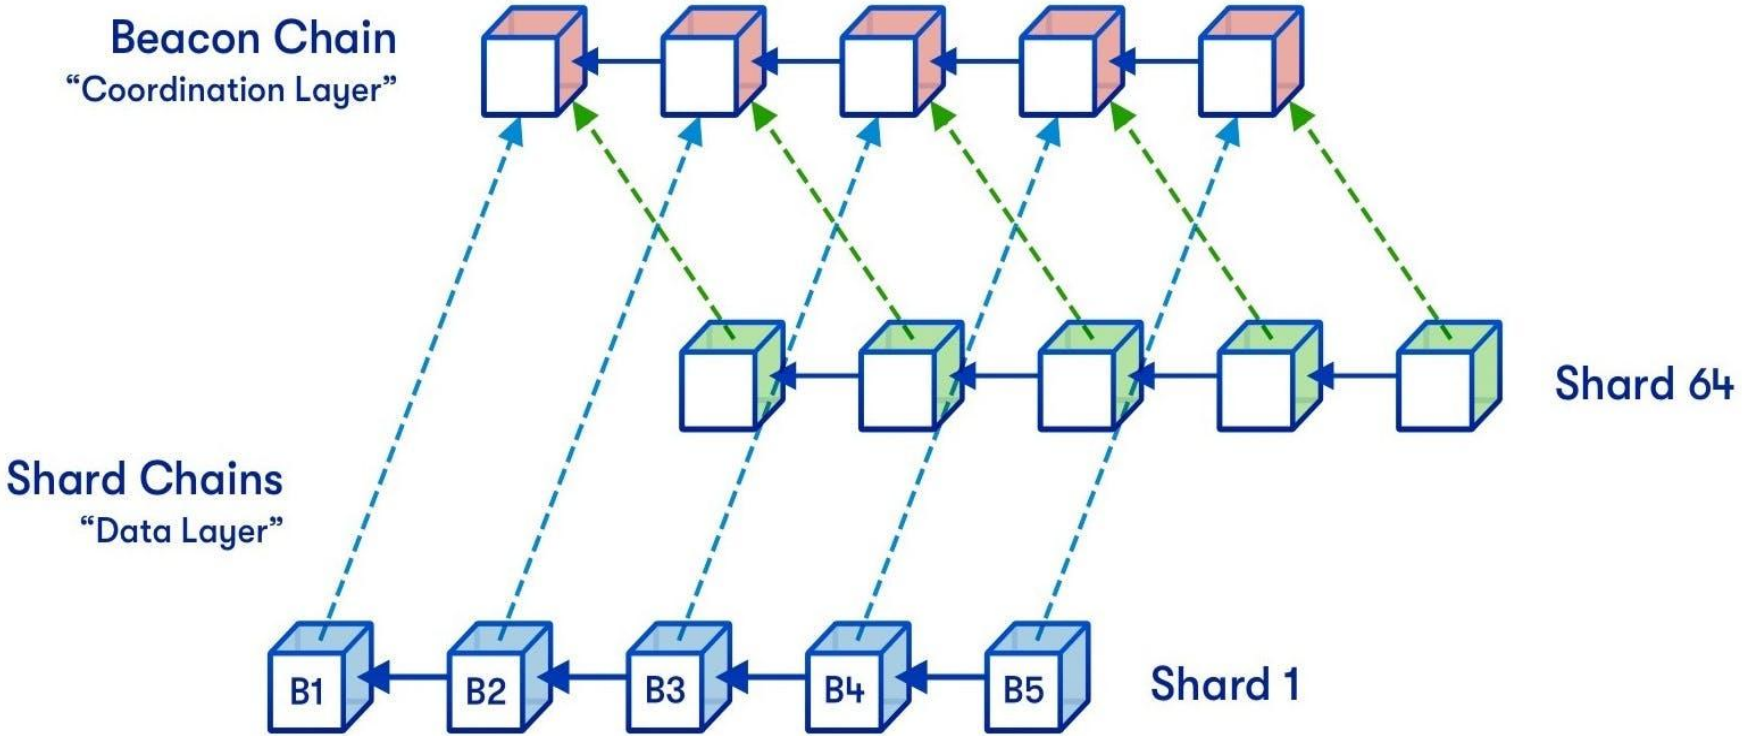

TRANSACTION EXECUTION POS

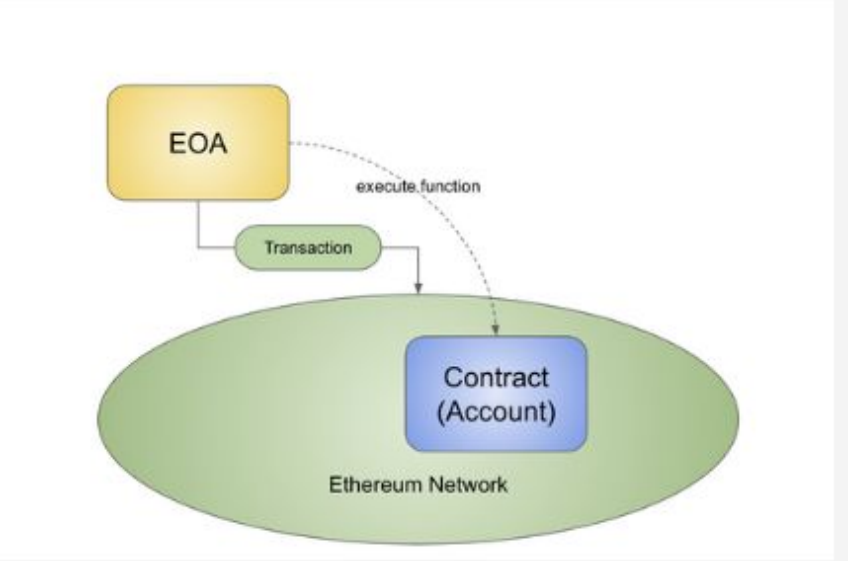

HOW A TRANSACTION GETS EXECUTED IN ETHEREUM POS

1. A user creates and signs a transaction with their private key. This is usually handled by a wallet or a library such as ether.js, web3js, web3py etc but under the hood the user is making a request to a node using the Ethereum JSON-RPC API. The user defines the amount of gas that they are prepared to pay as a tip to a validator to encourage them to include the transaction in a block. The tips get paid to the validator while the base fee gets burned.
2. The transaction is submitted to an Ethereum execution client which verifies its validity. This means ensuring that the sender has enough ETH to fulfill the transaction and they have signed it with the correct key.
3. If the transaction is valid, the execution client adds it to its local mempool (list of pending transactions) and also broadcasts it to other nodes over the execution layer gossip network. When other nodes hear about the transaction they add it to their local mempool too. Advanced users might refrain from broadcasting their transaction and instead forward it to specialized block builders such as Flashbots Auction. This allows them to organize the transactions in upcoming blocks for maximum profit (MEV).

HOW A TRANSACTION GETS EXECUTED IN ETHEREUM POS

4. One of the nodes on the network is the block proposer for the current slot, having previously been selected pseudo-randomly using RANDAO. This node is responsible for building and broadcasting the next block to be added to the Ethereum blockchain and updating the global state. The node is made up of three parts: an execution client, a consensus client and a validator client. The execution client bundles transactions from the local mempool into an "execution payload" and executes them locally to generate a state change. This information is passed to the consensus client where the execution payload is wrapped as part of a "beacon block" that also contains information about rewards, penalties, slashings, attestations etc. that enable the network to agree on the sequence of blocks at the head of the chain

HOW A TRANSACTION GETS EXECUTED IN ETHEREUM POS

5. Other nodes receive the new beacon block on the consensus layer gossip network. They pass it to their execution client where the transactions are re-executed locally to ensure the proposed state change is valid. The validator client then attests that the block is valid and is the logical next block in their view of the chain (meaning it builds on the chain with the greatest weight of attestations as defined in the fork choice rules). The block is added to the local database in each node that attests to it.

6. The transaction can be considered "finalized" if it has become part of a chain with a "supermajority link" between two checkpoints. Checkpoints occur at the start of each epoch and they exist to account for the fact that only a subset of active validators attest in each slot, but all active validators attest across each epoch. Therefore, it is only between epochs that a 'supermajority link' can be demonstrated (this is where 66% of the total staked ETH on the network agrees on two checkpoints).

End of Module-3 (Class-2)

Thank You!