DFINITY

Internet Computer Consensus

PODC, July 2022

Jan Camenisch, Manu Drijvers, Timo Hanke,
Yvonne-Anne Pignolet, Victor Shoup, Dominic Williams

DFINITY Foundation, Switzerland

Outline

- What is the Internet Computer?
- Internet Computer Consensus
- Measurements

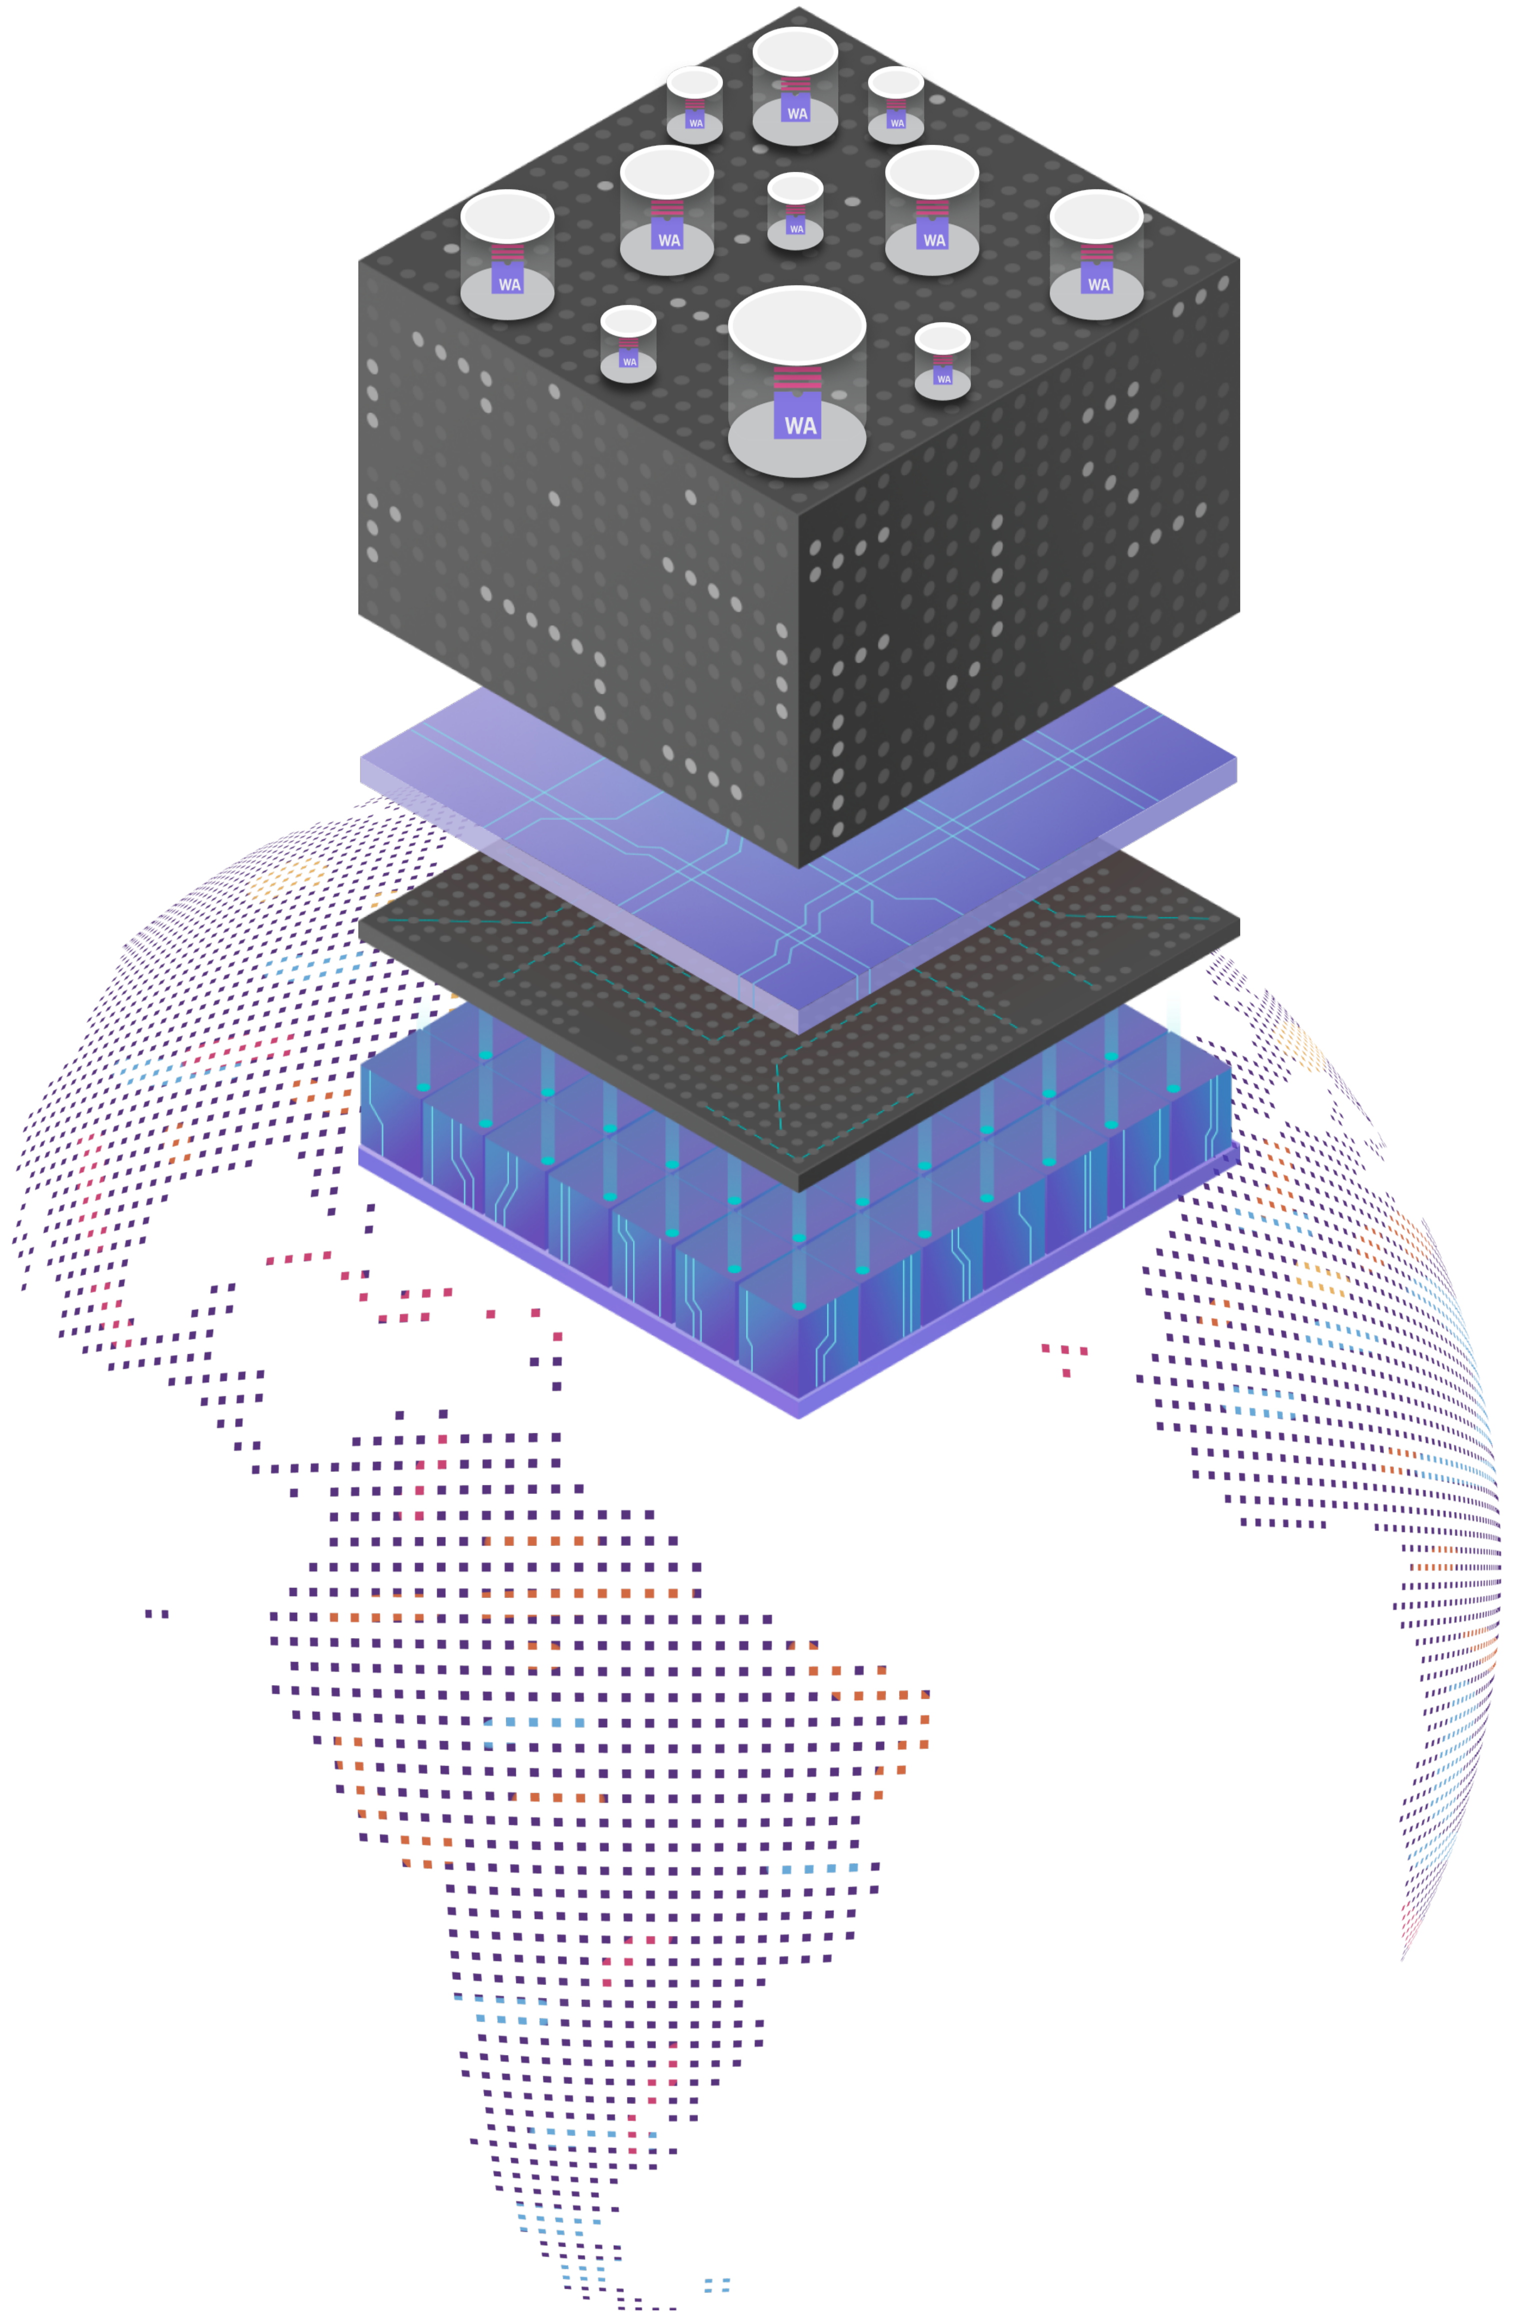

What is the mission of the Internet Computer?

What is the mission of the Internet Computer?

Platform to run any computation, using blockchain technology for decentralisation and security

Nodes in Independent Data Centers

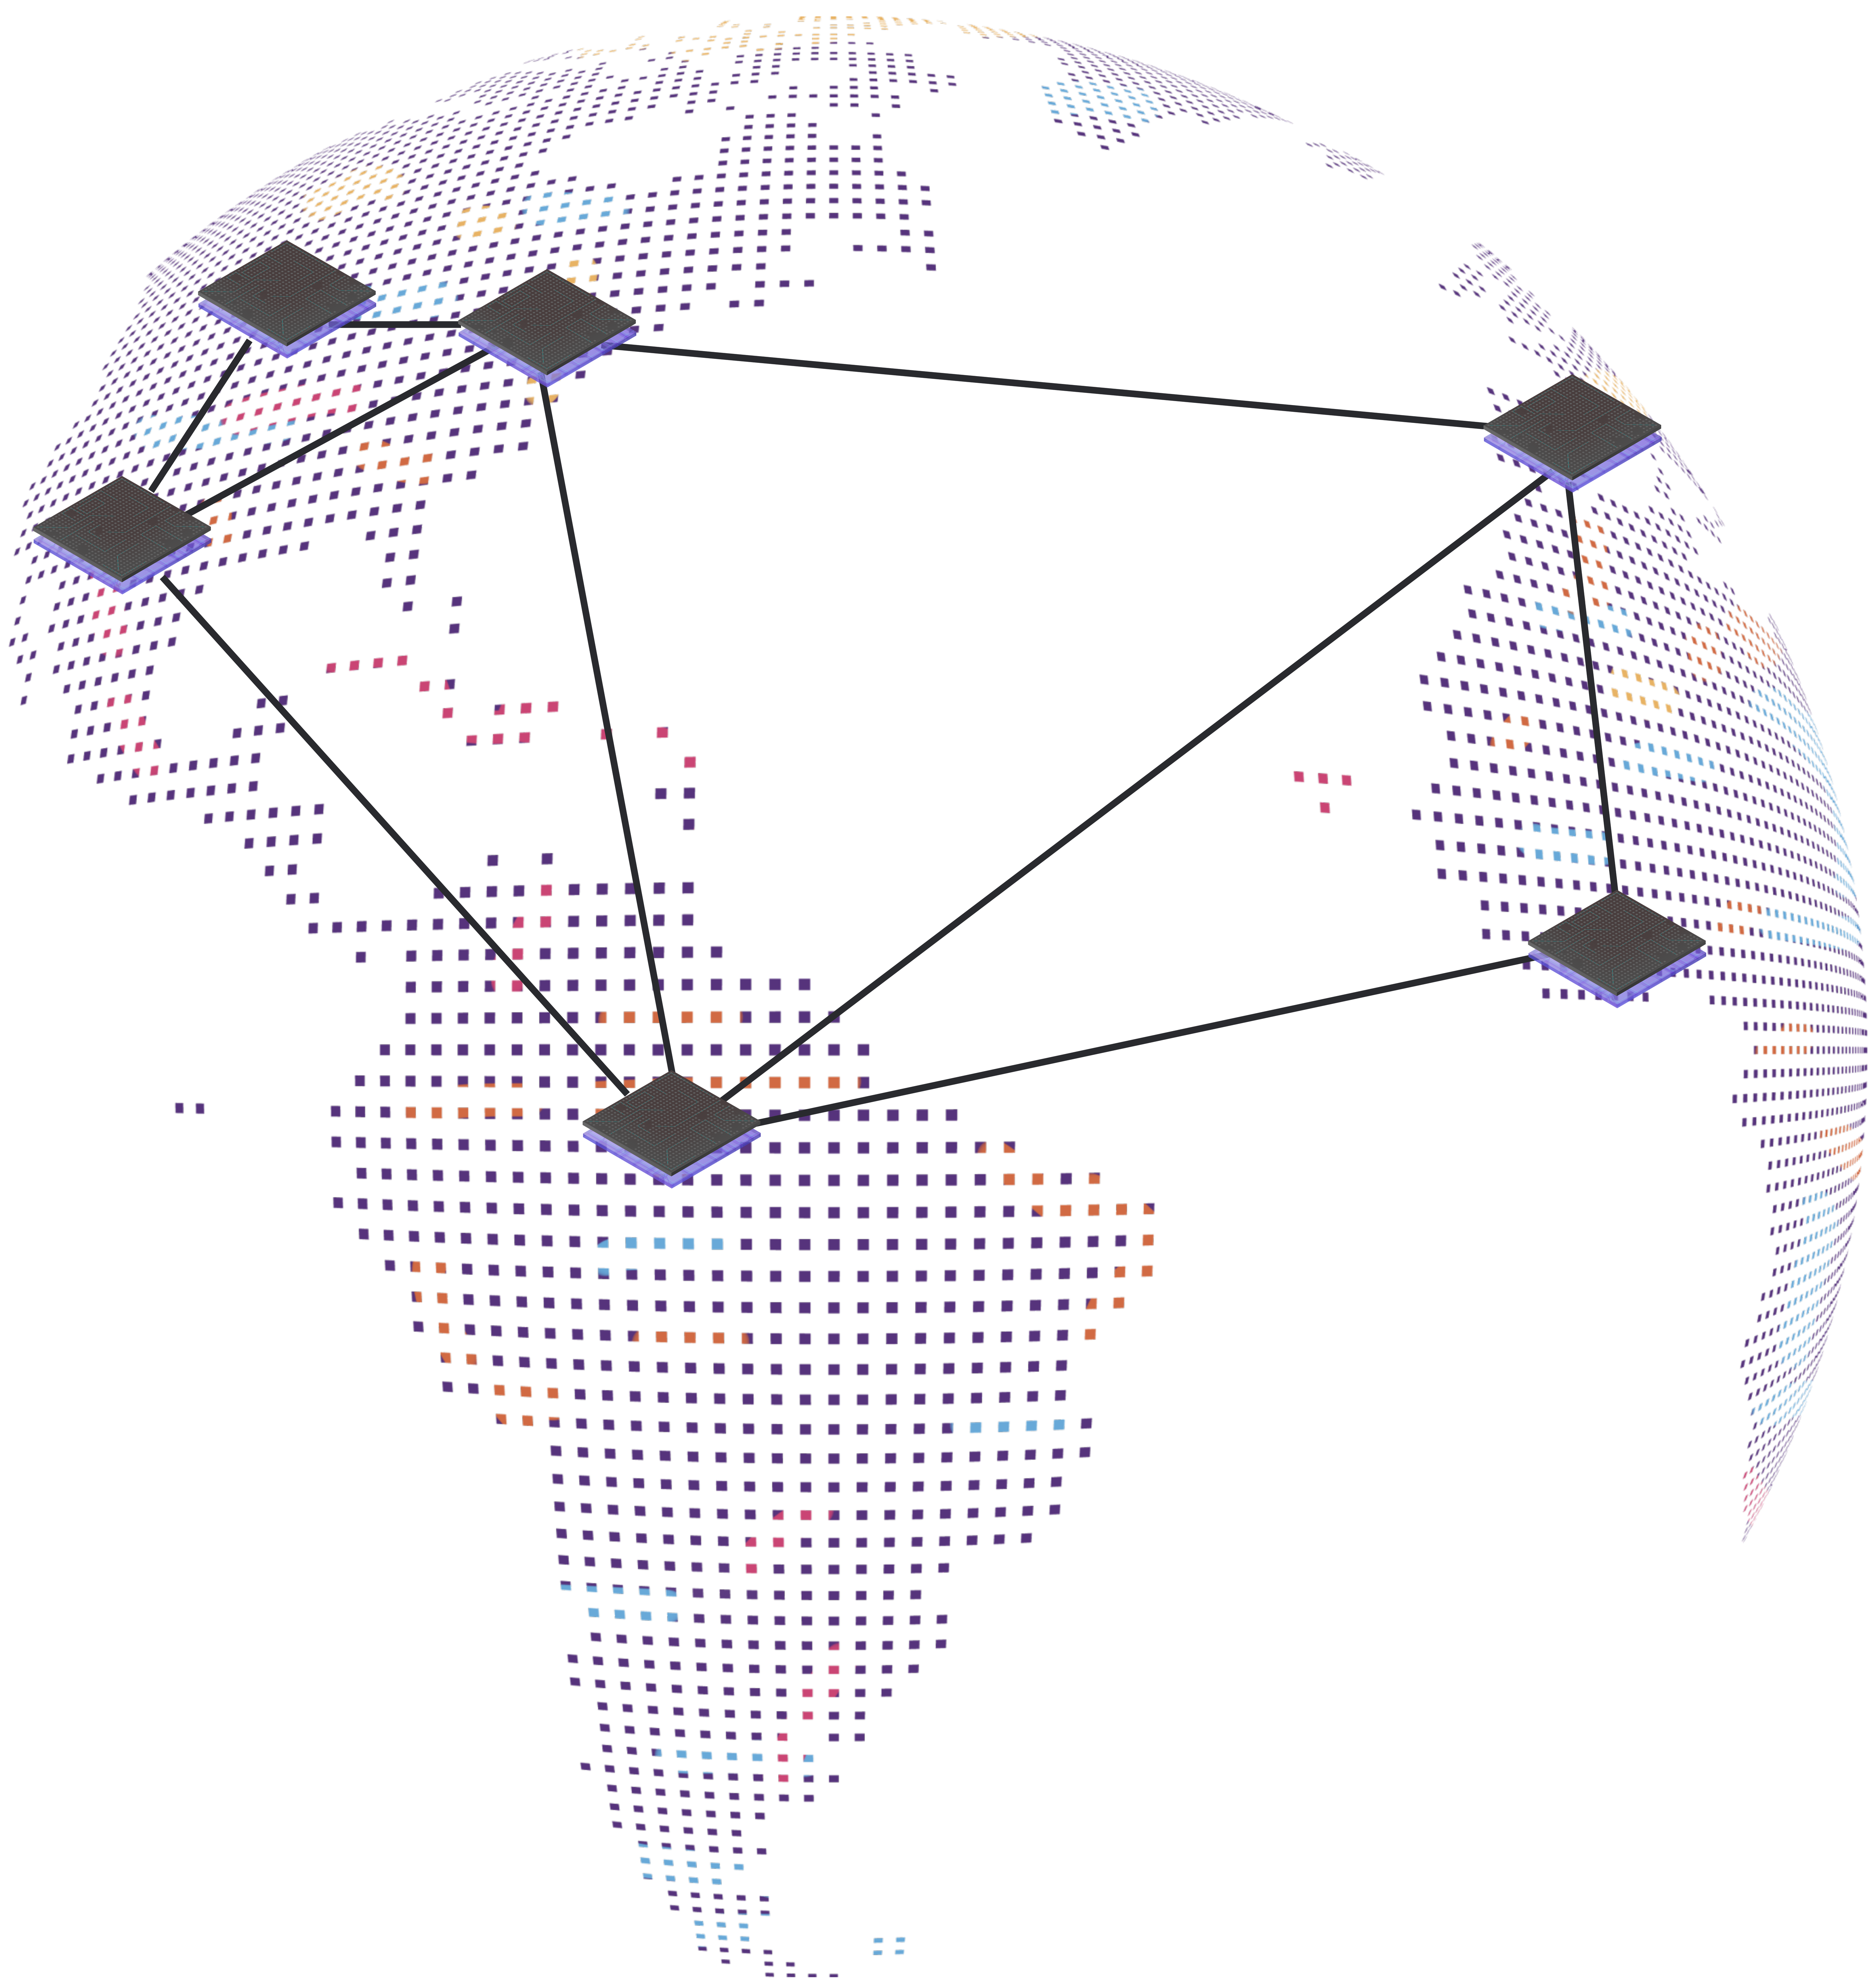

https://dashboard.internetcomputer.org/

Canister Smart Contracts: Combination of Data and Code

Developers and users interact directly with Canisters on the IC

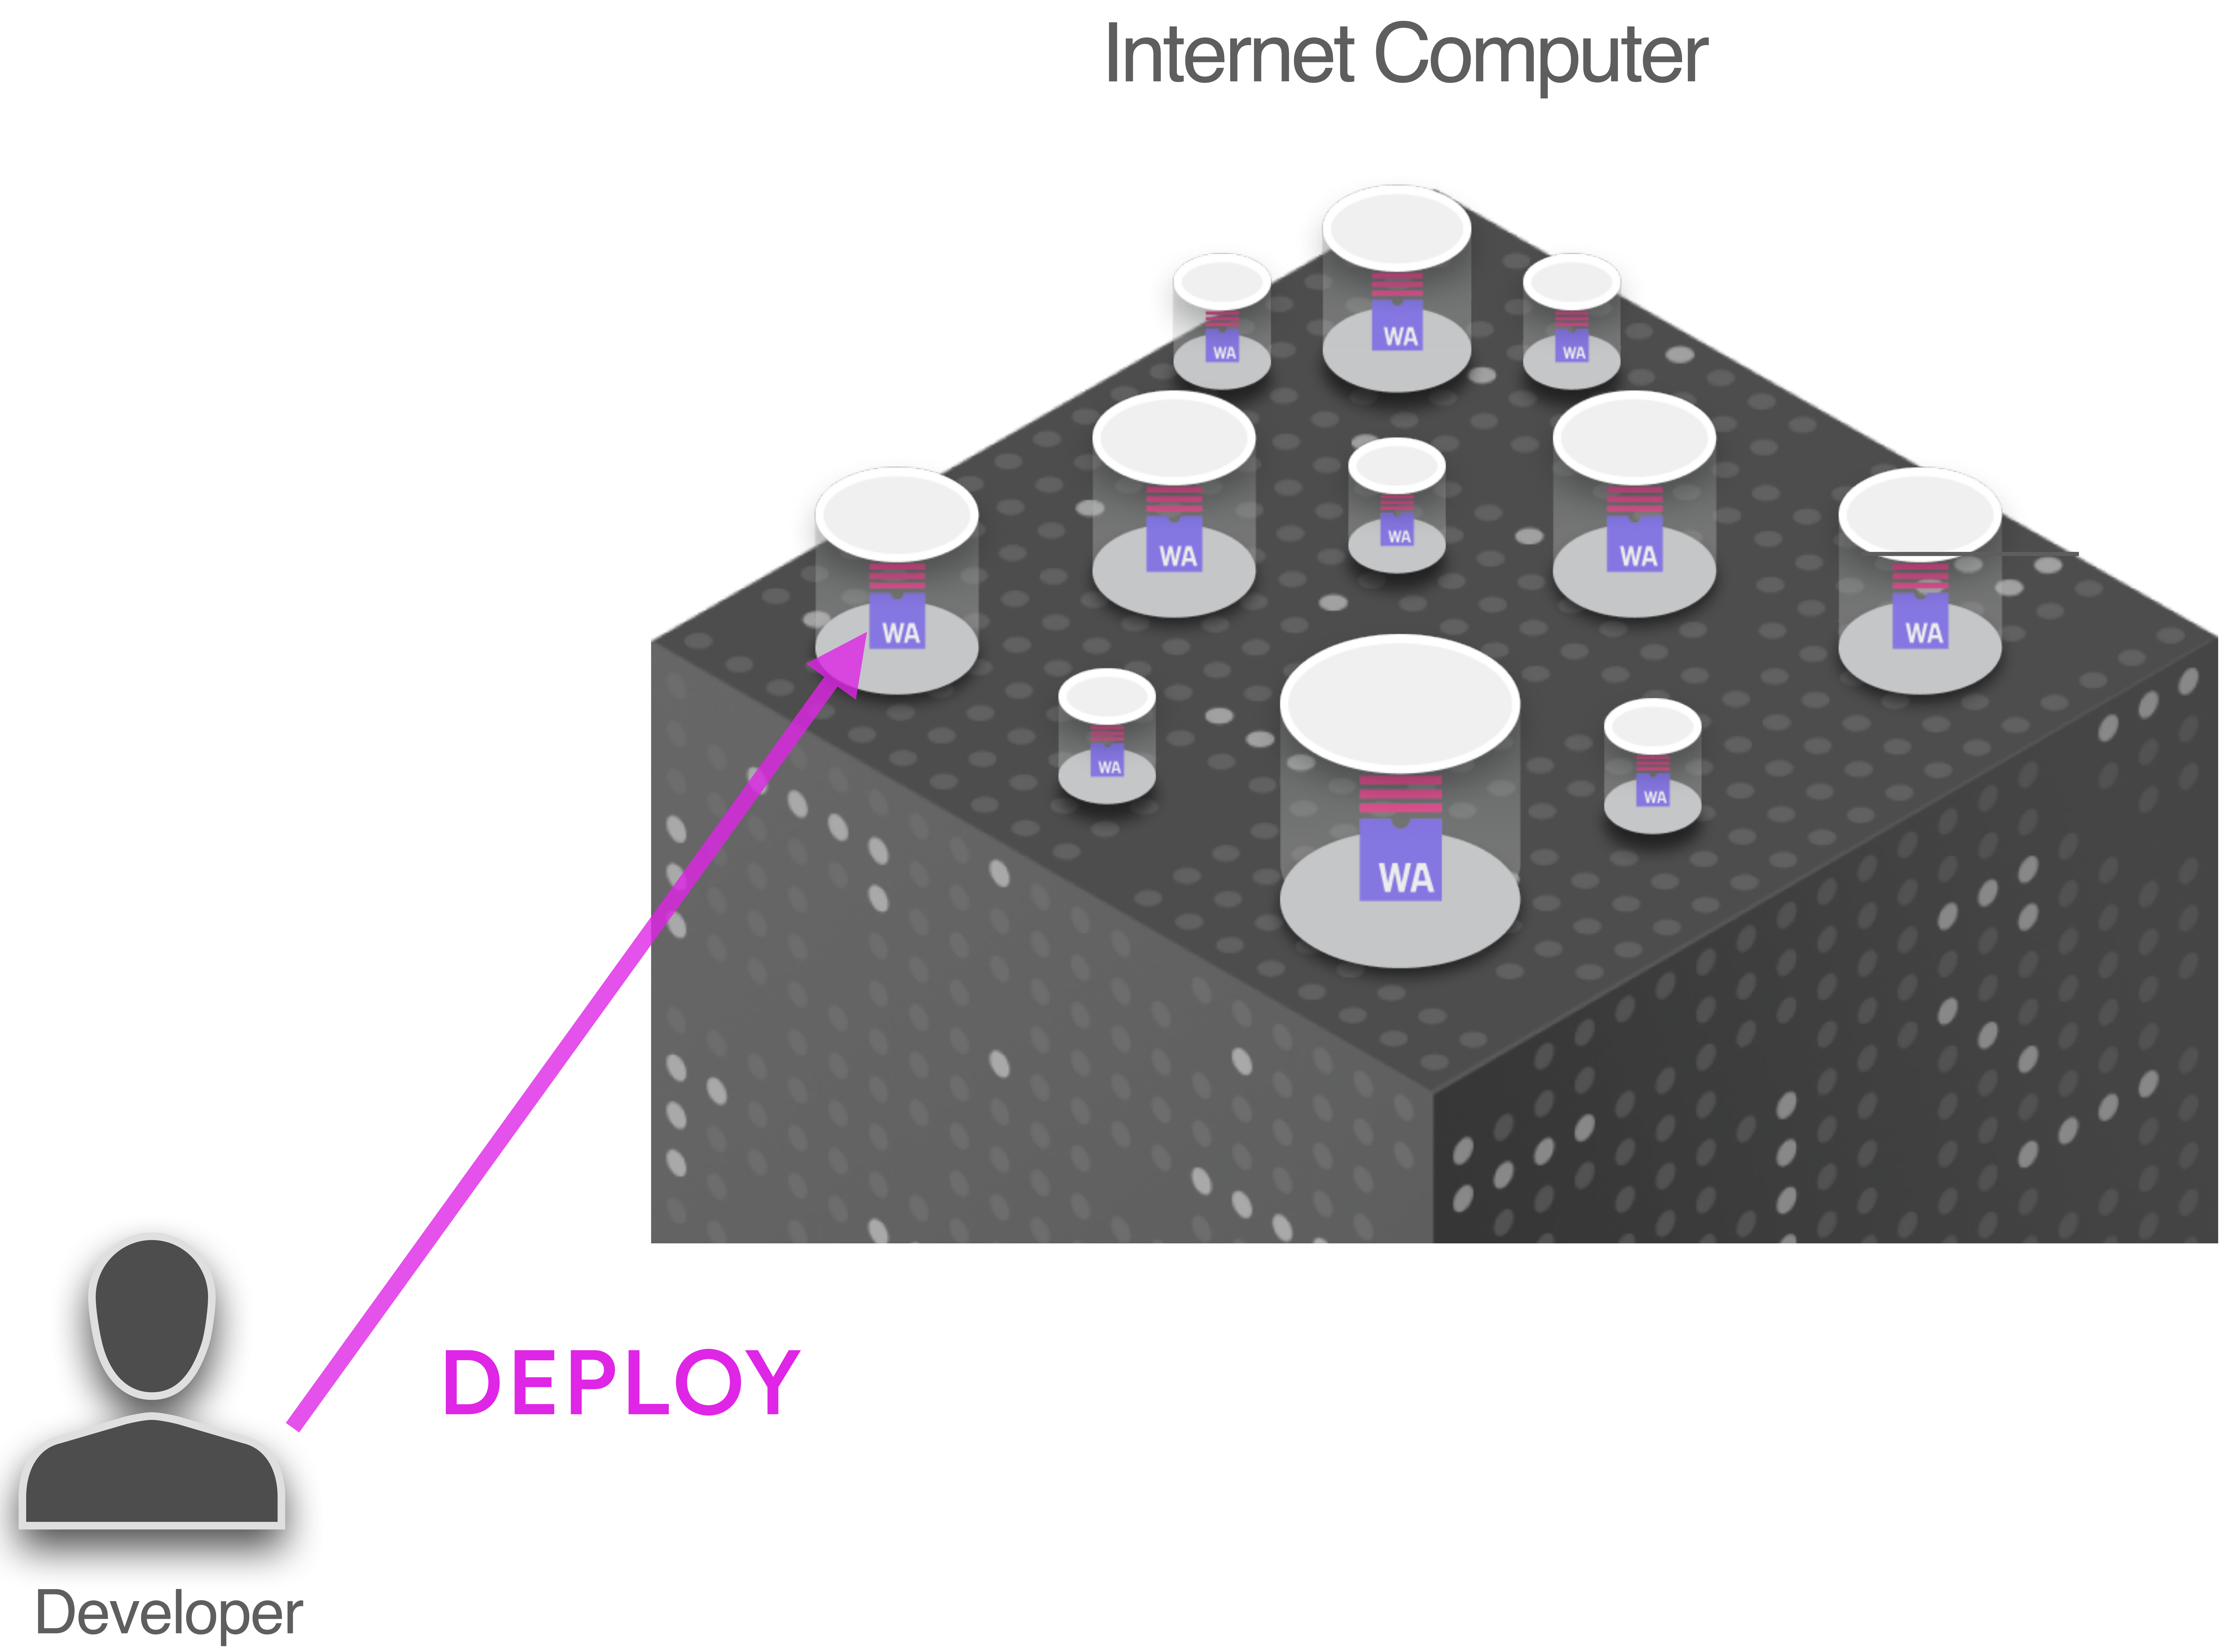

Developers and users interact directly with Canisters on the IC

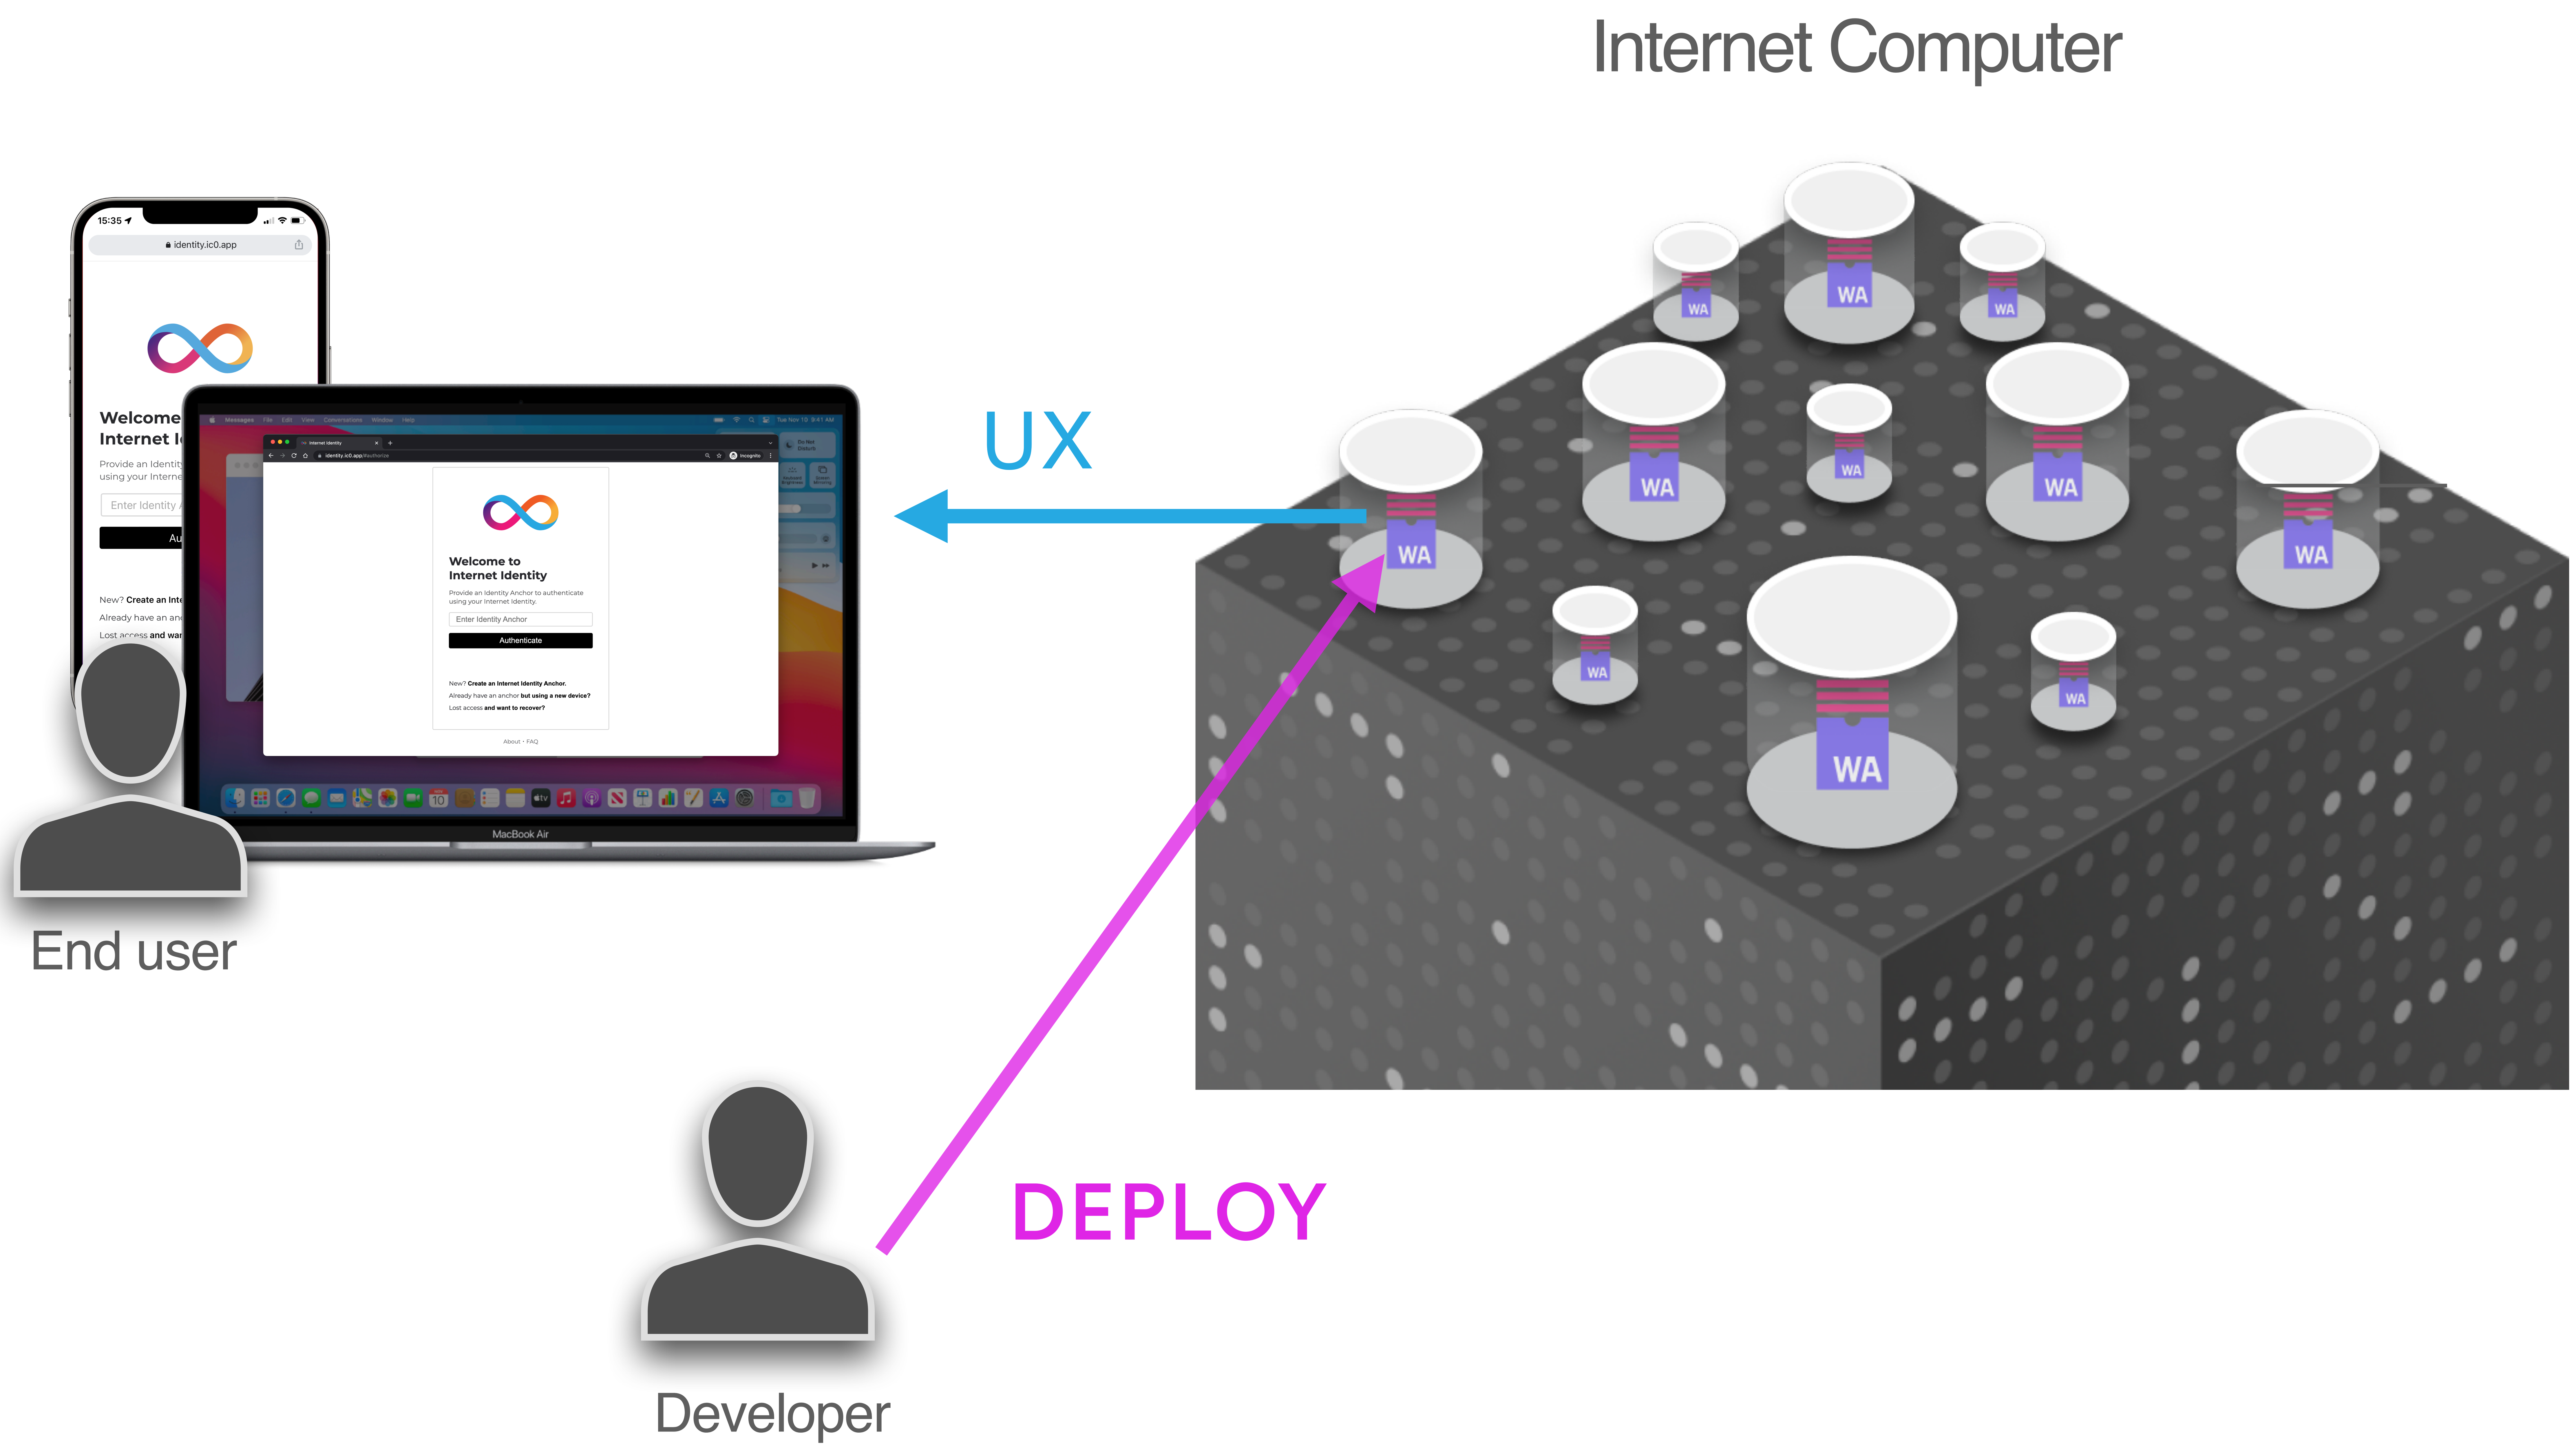

More than 60 000 Canisters deployed

FLEEK
Fleek brings decentralized web-hosting to the Internet Computer. With thousands of webpages deployed, Fleek enables anyone to deploy their content on Web3.0

fleek.co

#Infrastructure #Tools

DISTRIKT
Distrikt is a completely decentralized, community-owned professional network. Users of the platform will vote on upgrades, and no user data will ever be mined or sold. Create your account, secured by Internet Identity today.

19 000 users

#Social #Dapp

ORIGYN
The Origyn Foundation is blending luxury goods, with NFTs by providing digital verifications for physical objects. Only possible on the Internet Computer.

www.origyn.ch

#Enterprise #NFT

OPENCHAT
Decentralized messaging has been a pipe-dream for decades. With the advent of the Internet Computer, real-time messaging is now possible on a blockchain.

50 000 users

7e6iv-biaaa-aaaaf-aaada-cai

#Social #Dapp

INTERNET IDENTITY
Internet Identity guarantees that your data isn't visible, tracked, or mined. The blockchain authentication system enables users to sign in to dapps on the Internet Computer and sites across the web anonymously and securely.

1 000 000+

identity.ic0.app

#Authentication #Dapp #Infrastructure

IC ROCKS
IC.Rocks is a complete "block explorer" for the Internet Computer – built by the community. Tracking everything from transactions, to network upgrades, to cycles, IC.Rocks enables anyone to explore the inner-workings of the Internet Computer.

ic.rocks

#Infrastructure #Explorer

NNS DAPP
The NNS front-end dapp allows anyone to interact with the Internet Computer's Network Nervous System with a user-friendly UI. Served completely end-to-end through blockchain, this dapp allows you to manage ICP, stake neurons, participate in voting, and earn rewards.

#Dapp #Infrastructure #Wallet #NNS

DANK
Dank is the first Decentralized Bank built on the Internet Computer, developed by Fleek. Through a collection of Open Internet Services for users and developers, Dank makes cycles management seamless.

dank.ooo

#Infrastructure #DeFi

TONIQ LABS
Toniq Labs is the creator of Entrepot NFT marketplace, Stoic Wallet, Exponent, and Rise of the Magni, Cronic NFTs and more. Try out their projects that range from NFTs to wrapped cycles to games built on, and for, the Internet Computer blockchain.

igpeu-waaaa-aaaad-qaava-cai

#Infrastructure #Dapp

CANLISTA
The Internet Computer community canister registry. Find, publish and extend applications and services built on the Internet Computer. Log in with Internet Identity.

k7gat-daaaa-aaaae-qaahq-cai

#Infrastructure #Dapp

AGRYO
Agryo is the global risk intelligence provider that enables financial institutions to assess and manage financial risks in the crop field level for underwriting agriculture insurance, loans, and trade finance globally; as well as meet sustainability goals.

www.agryo.com

#Enterprise #DeFi

SUDOGRAPH
Sudograph is a GraphQL database for the Internet Computer. Its goal is to become the simplest way to develop applications for the IC by providing flexibility and out-of-the-box data management.

i67uk-hiaaa-aaaae-qaaka-cai

#Infrastructure #Dapp #Tools

PLUG
Plug Wallet, built and open sourced by Fleek, is a browser extension that allows you to access your ICP, Cycles and other tokens – as well as log into Internet Computer dapps with one click. Download it here.

100 000 users

plugwallet.ooo

#Wallet #Infrastructure #NFT

REVERSI
Reversi is one of the first canister smart contracts deployed to the Internet Computer and is a completely decentralized multiplayer game. Play against a friend (or foe) in real-time, from any browser, anywhere in the world.

ivg37-qiaaa-aaaab-aaaga-cai

DFINITY EXPLORER
DFINITY Explorer, a project started in 2018, is an open-source, community-built dashboard and explorer for the Internet Computer, providing live information and statistics about the network, governance, and the ICP utility token, including account and transaction information.

www.dfinityexplorer.org

NNS CALCULATOR
The Network Nervous System Calculator is a calculator that allows anyone to edit variables and estimate voting rewards based on number of proposals voted on, length of stake, accumulated maturity, and more.

networknervoussystem.com

The Wall · Uniswap Front End on the IC · DFinance · Rise of the Magni · Welcome Into the Metaverse · HexGL · ICPunks · Entrepot · Departure Labs · Axon

https://internetcomputer.org/showcase/

Scalability: Nodes and Subnets

Nodes are partitioned into **subnets**

Canister smart contracts are assigned to different subnets

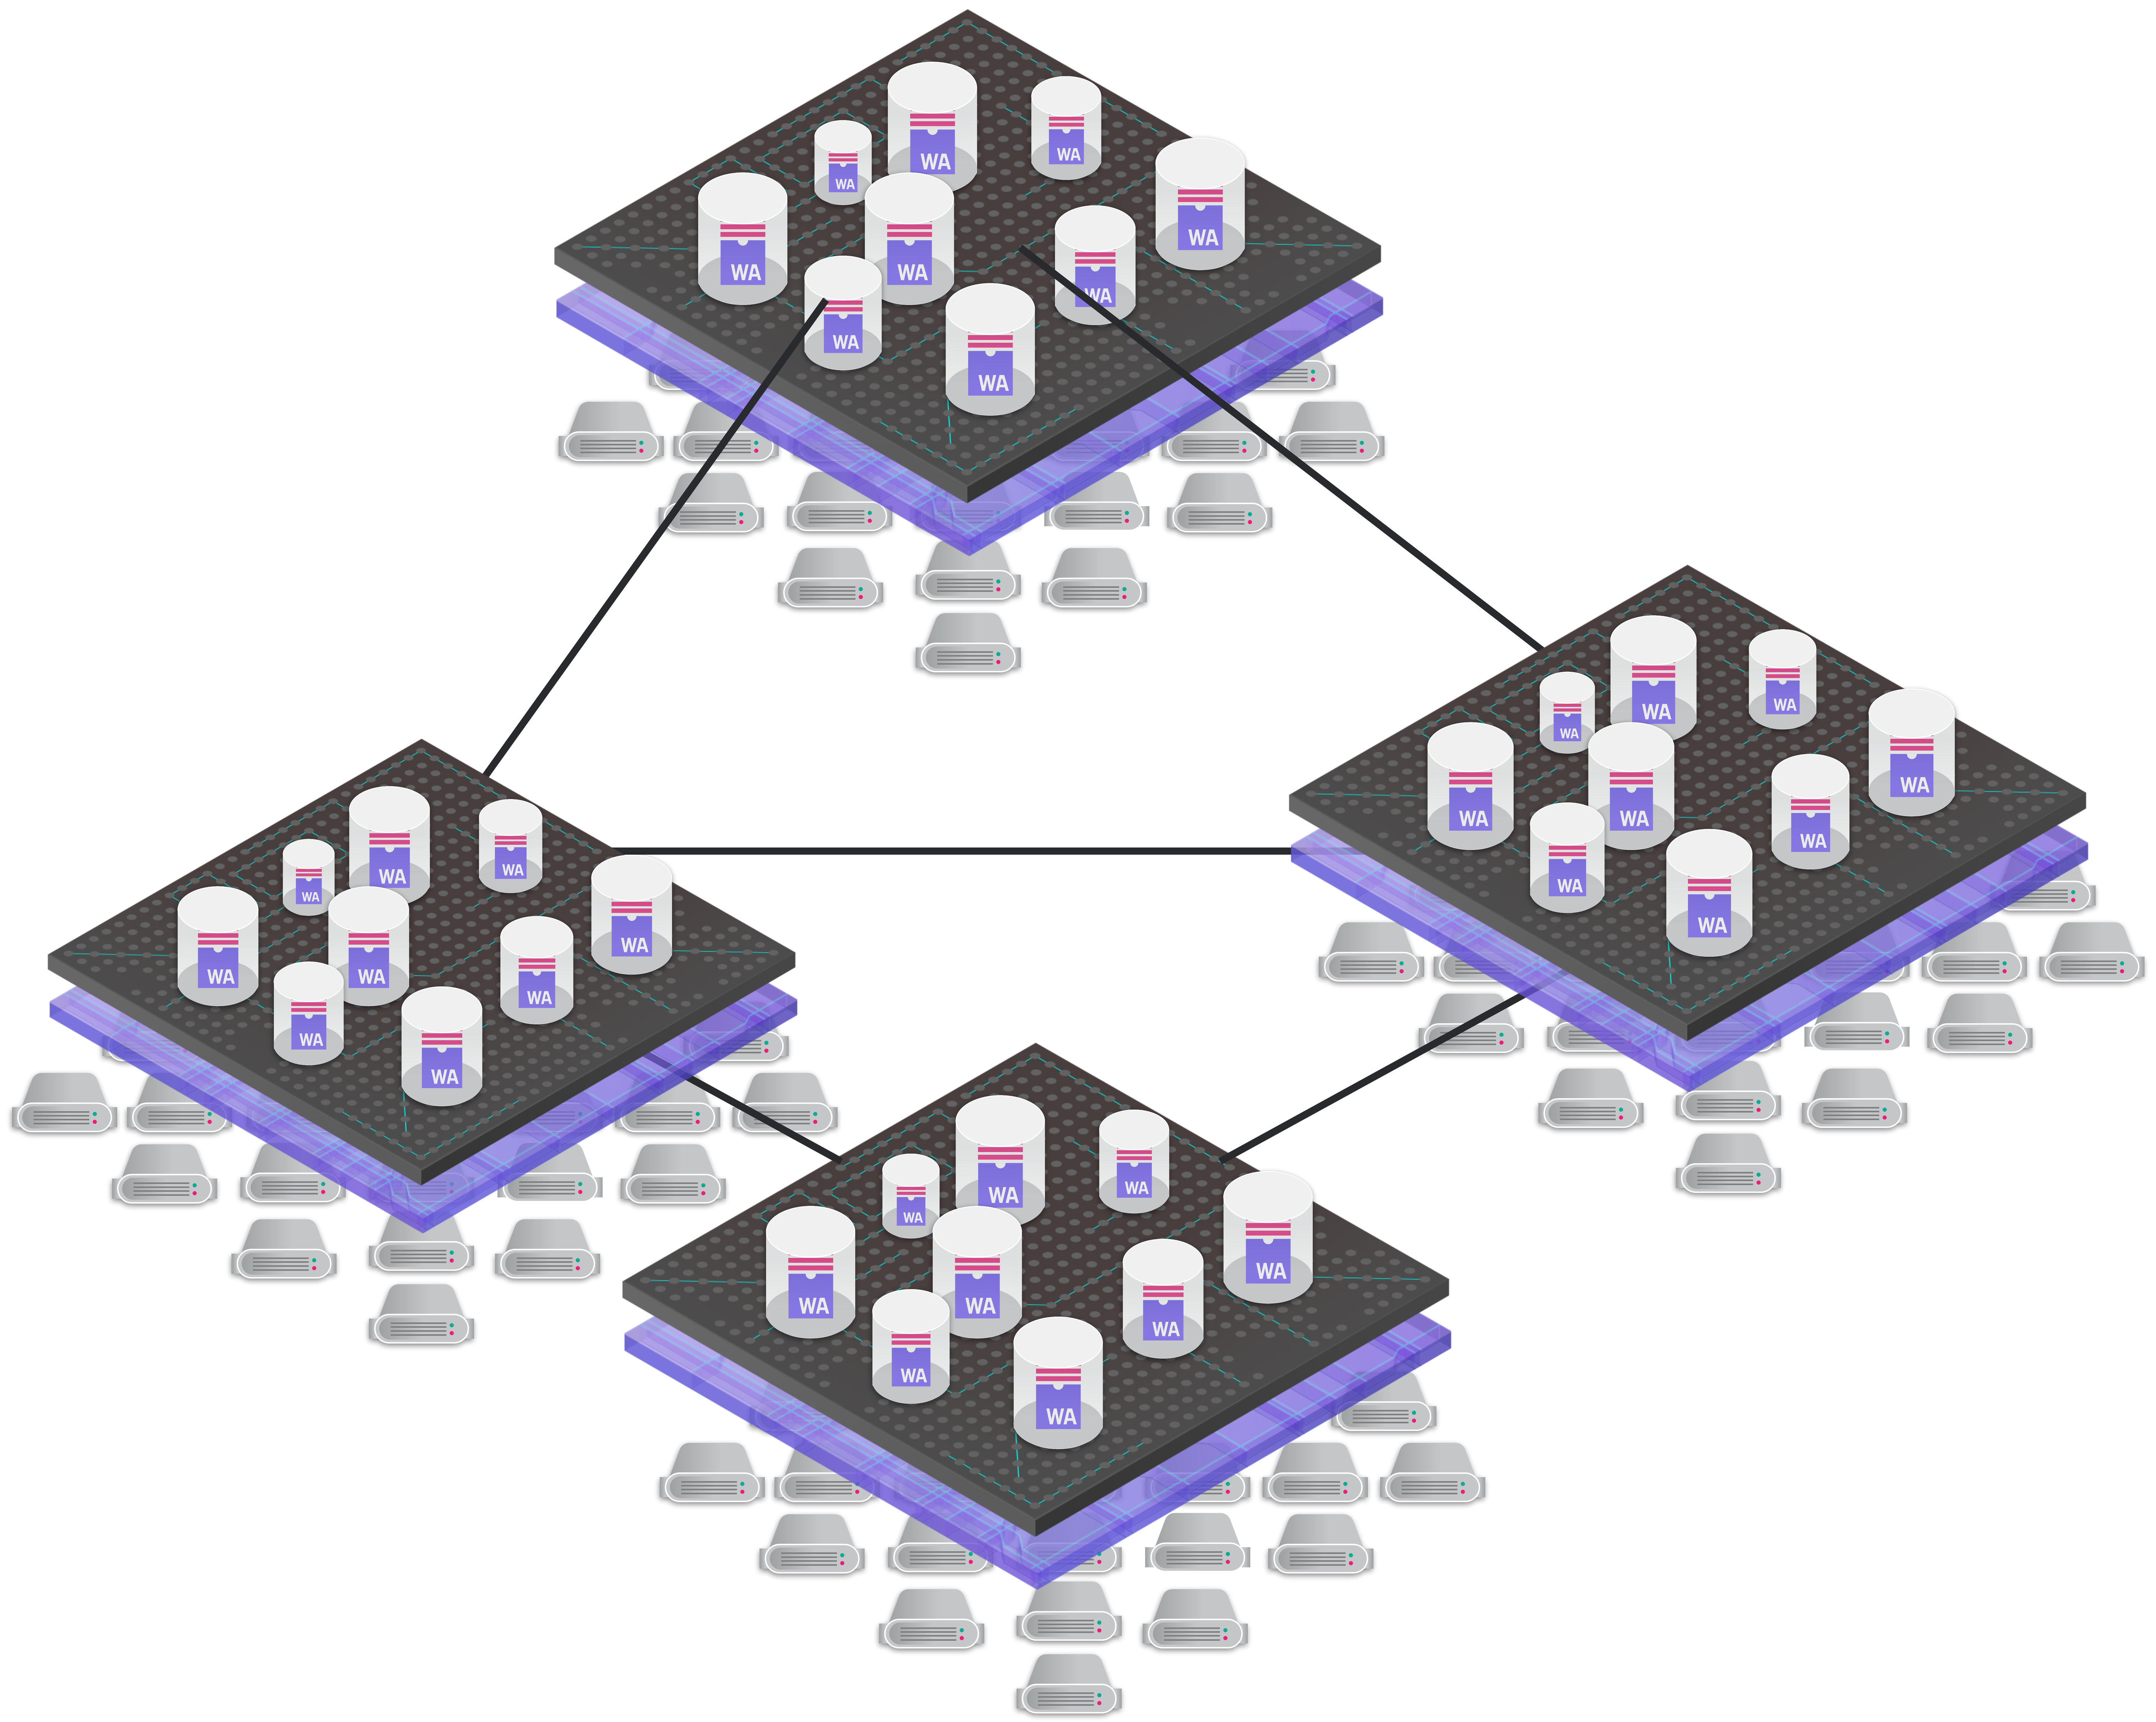

Scalability: Nodes and Subnets

Nodes are partitioned into **subnets**

Canister smart contracts are assigned to different subnets

One subnet is special: it host the **Network Nervous System (NNS)** canisters which govern the IC

ICP token holders vote on
- Creation of new subnets
- Upgrades to new protocol version
- Replacement of nodes
- …

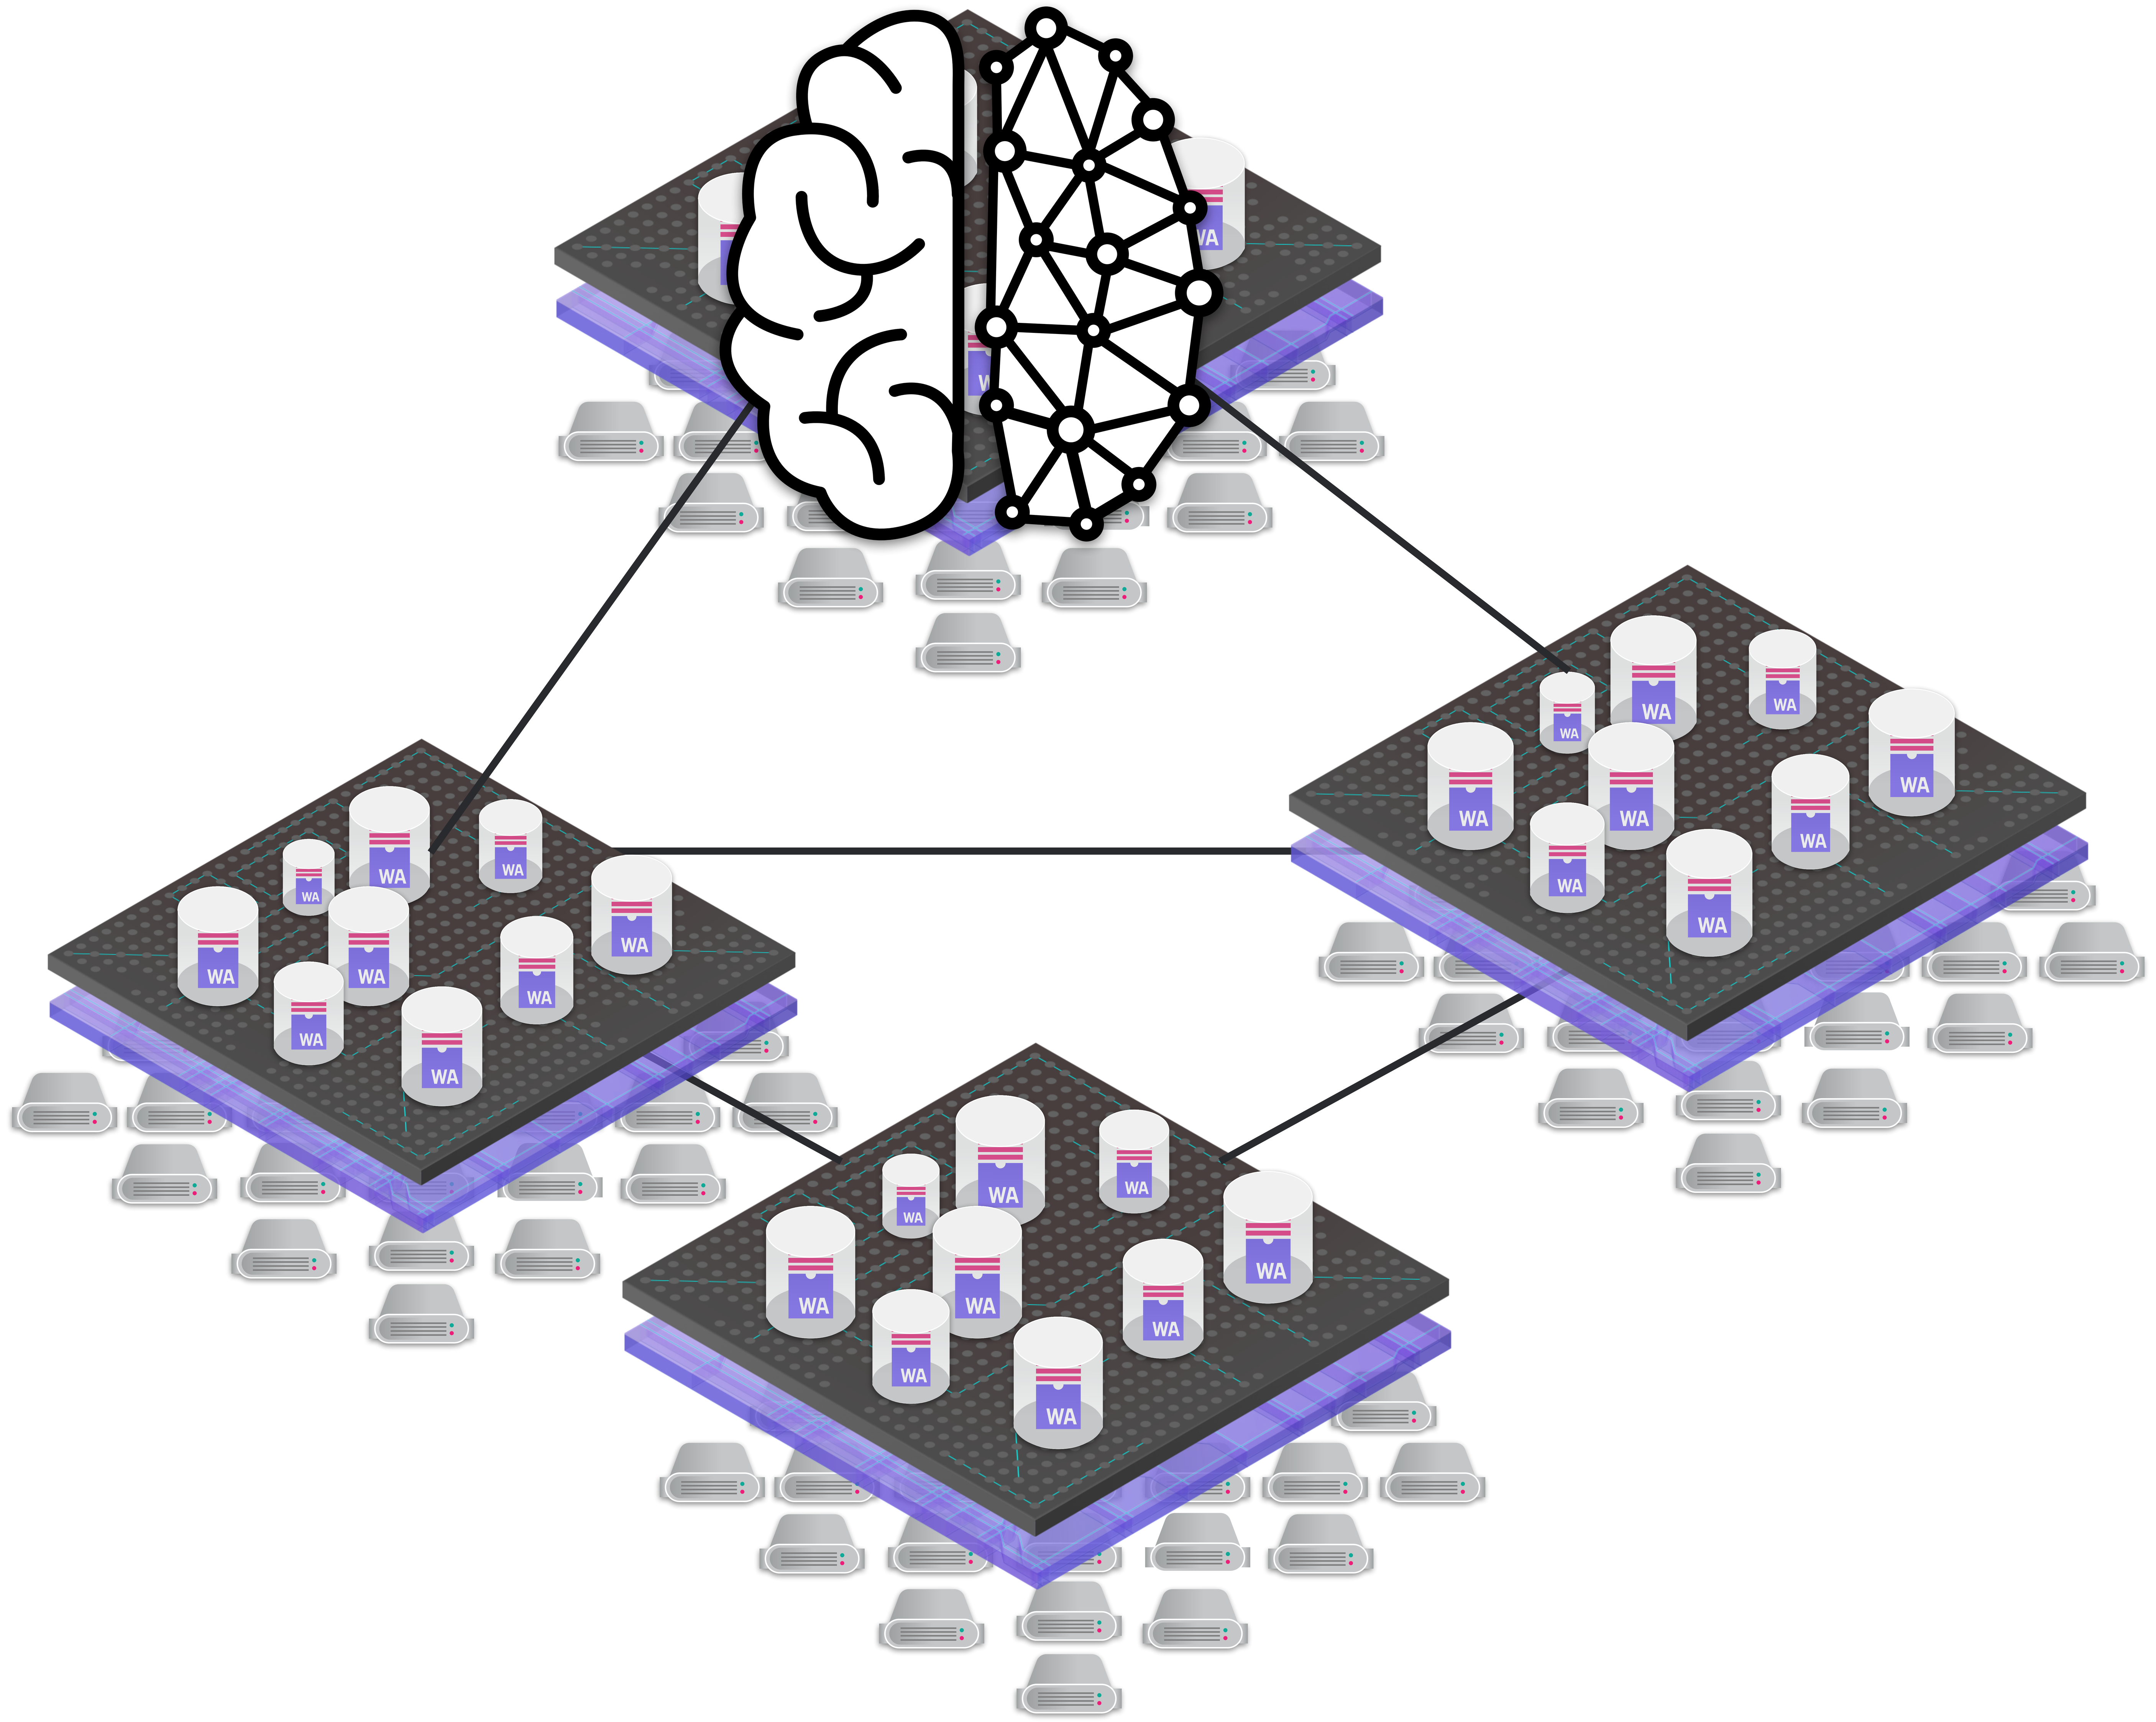

Each Subnet is a Replicated State Machine

State:
- canisters and their queues

Inputs:
- new canisters to be installed,
- messages from users and other canisters

Outputs:
- responses to users and other canisters

Transition function:
- message routing and scheduling
- canister code

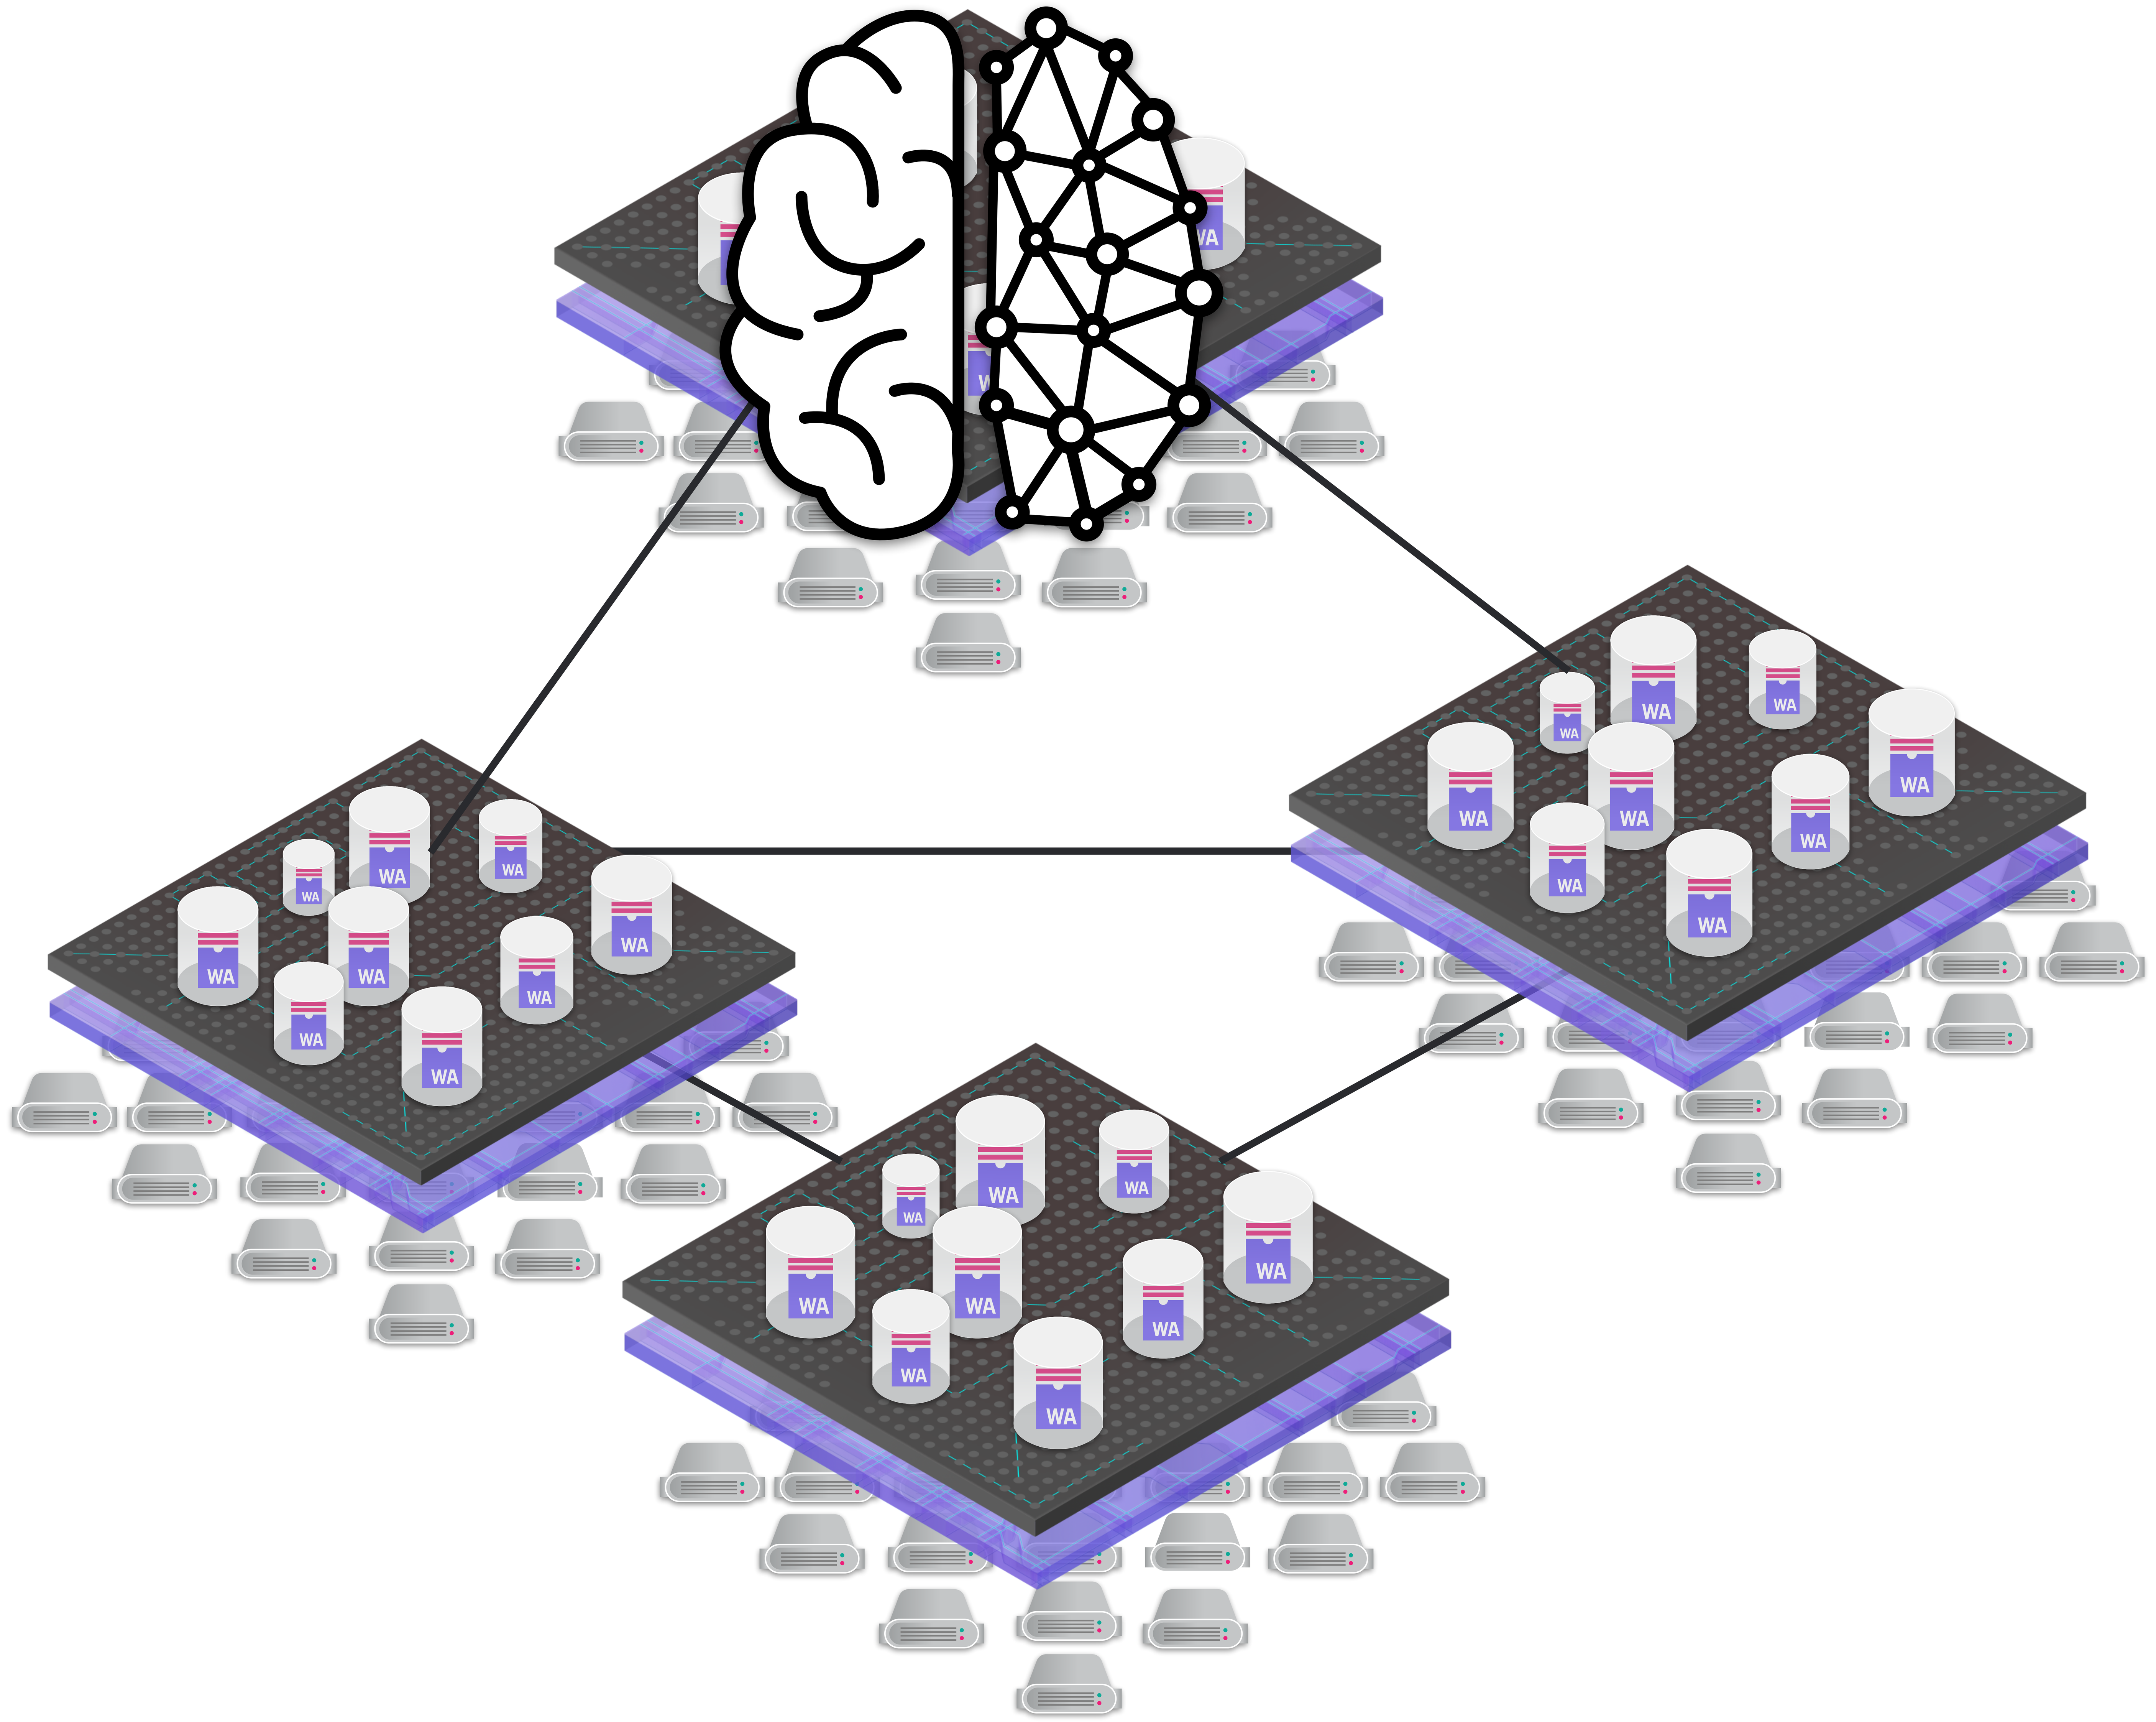

Consensus on the Internet Computer

Requirements

High Througput
- Several thousands of messages per second

Low Latency
- ~1 second (+ user network latency) for state changes

Robustness
- Tolerate bad communication links between nodes as well as Byzantine node behaviour
    - Safety under asynchrony
    - Liveness under short intervals of synchrony
- Graceful degradation
    - “slow path” is a simple variation on “fast path”

Simplicity
- Facilitate fast implementation and debugging

Consensus Properties

Messages are placed in **blocks**. We reach agreement using a blockchain.

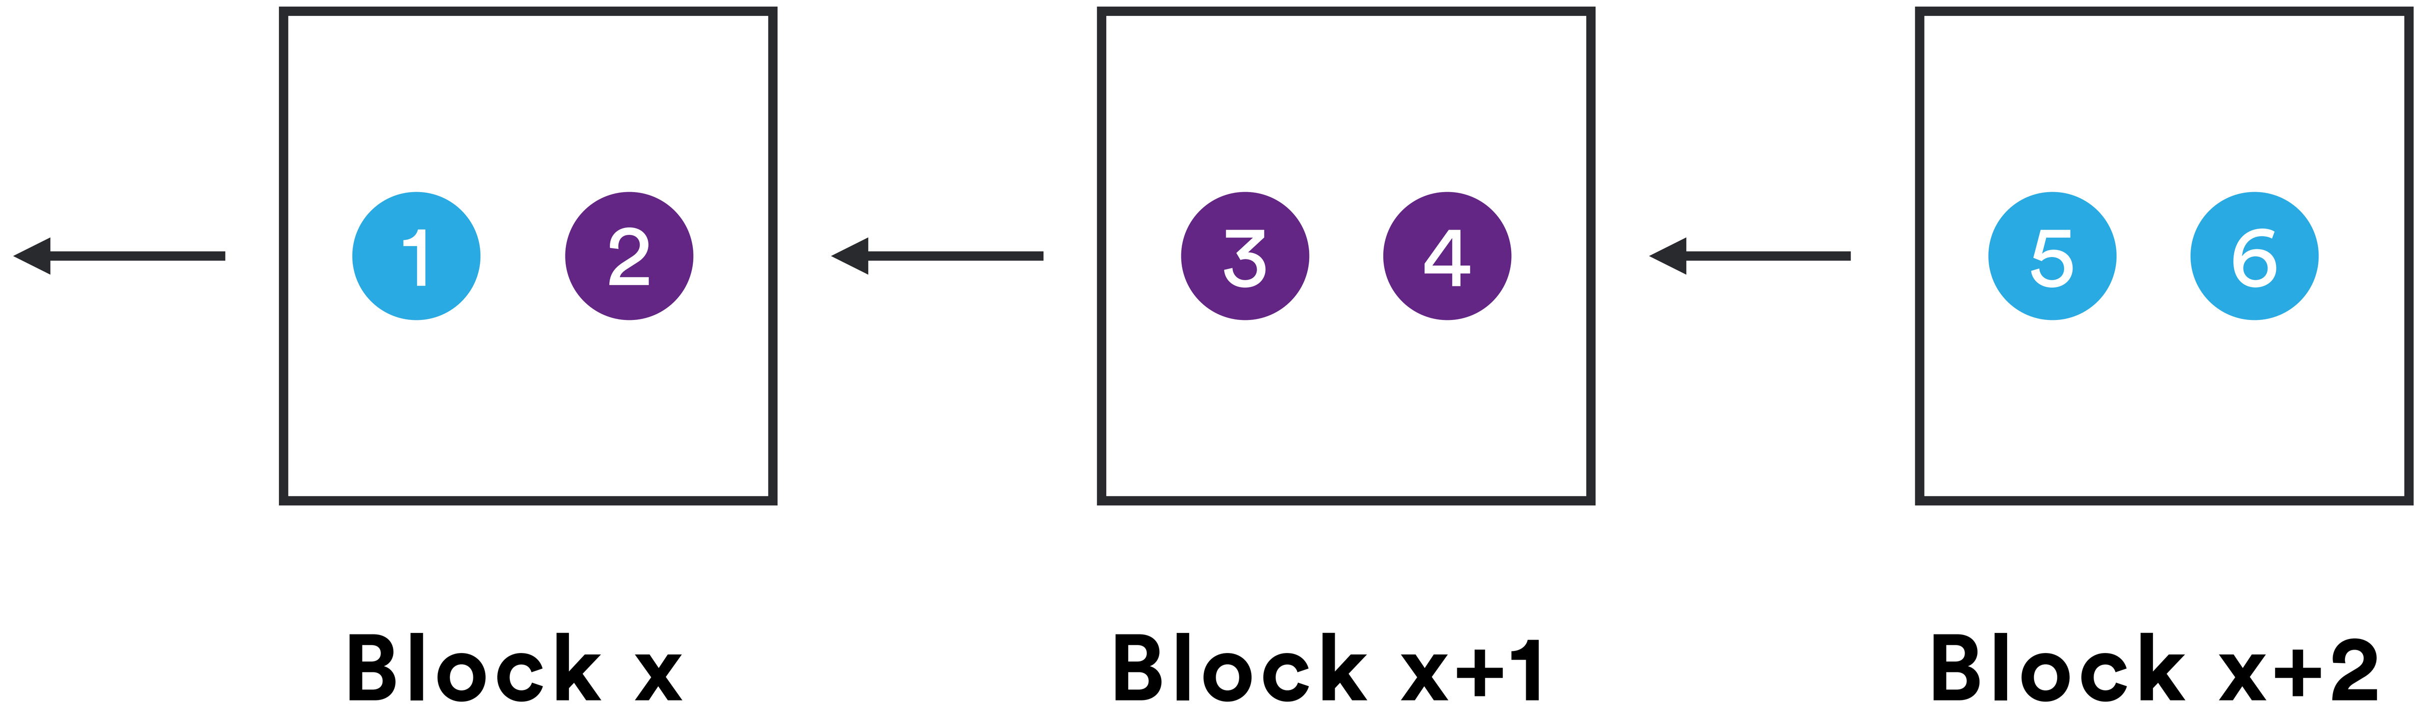

We use ***n* = 4, *f* = 1** in examples

The following properties must hold even if up to $f < n/3$ nodes misbehave

- **Safety**: For any *i*, if two honest nodes think that the *i*-th block is agreed upon, they must have the same block
- **Liveness**: For any *i*, at some point every honest node will consider the *i*-th block is agreed upon

Block Making

● Message (user → canister) ● Message (canister → canister)

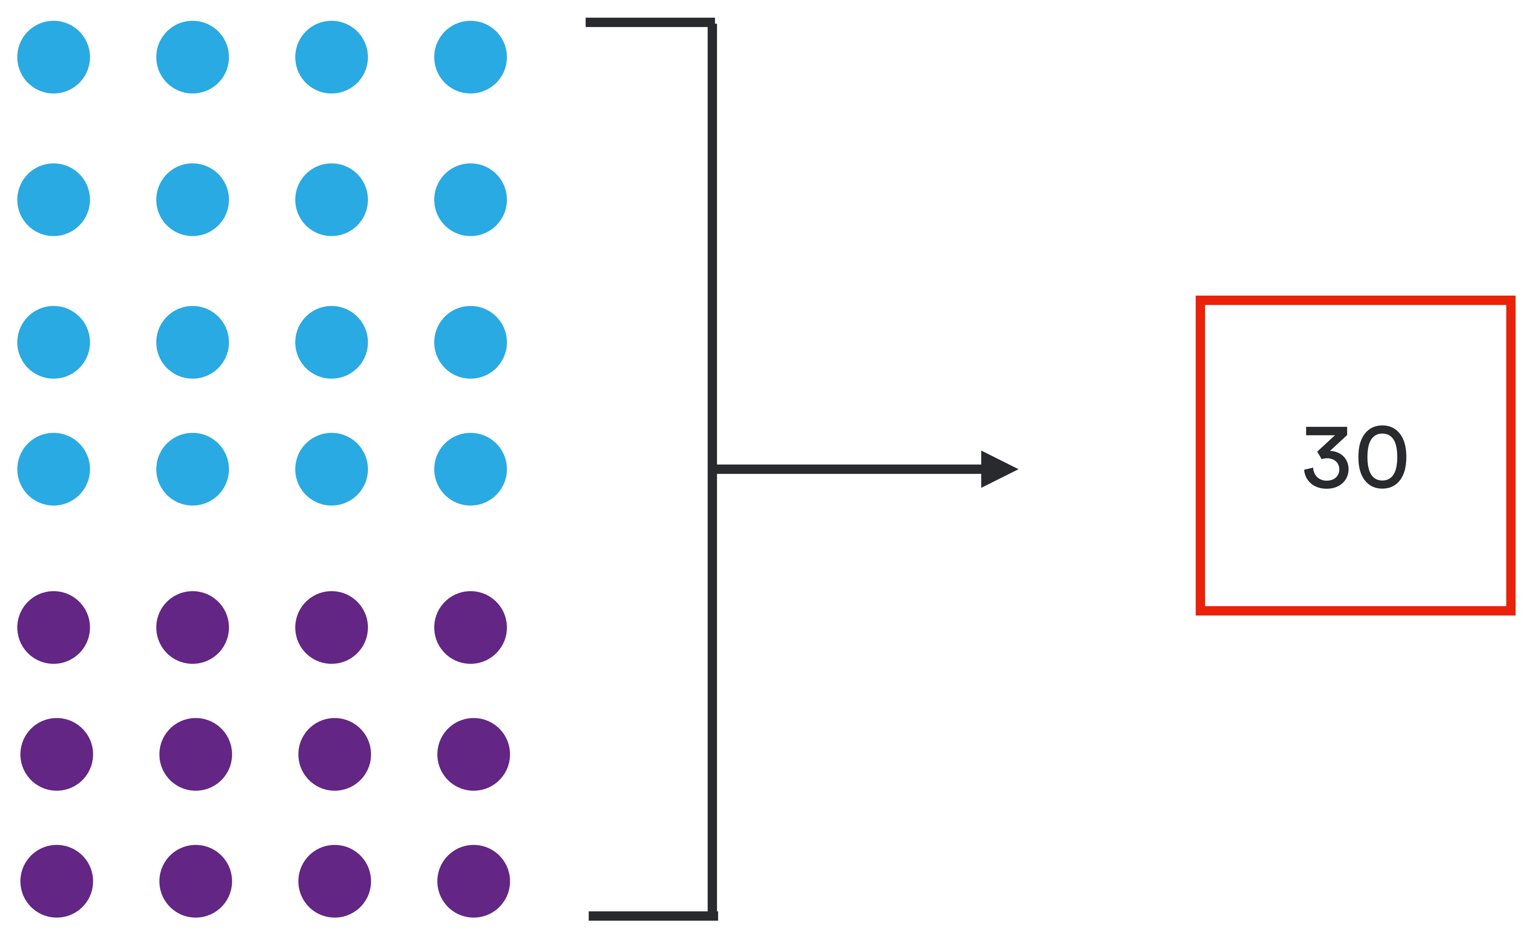

Node selects available messages and combines them into a block together with reference to predecessor and meta-data and broadcasts it

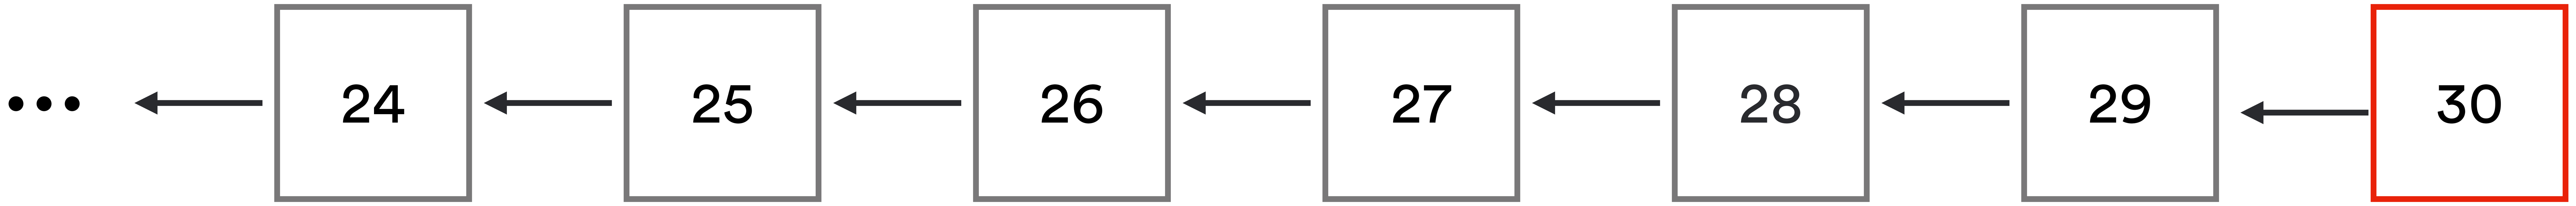

Block Making

● Message (user → canister) ● Message (canister → canister)

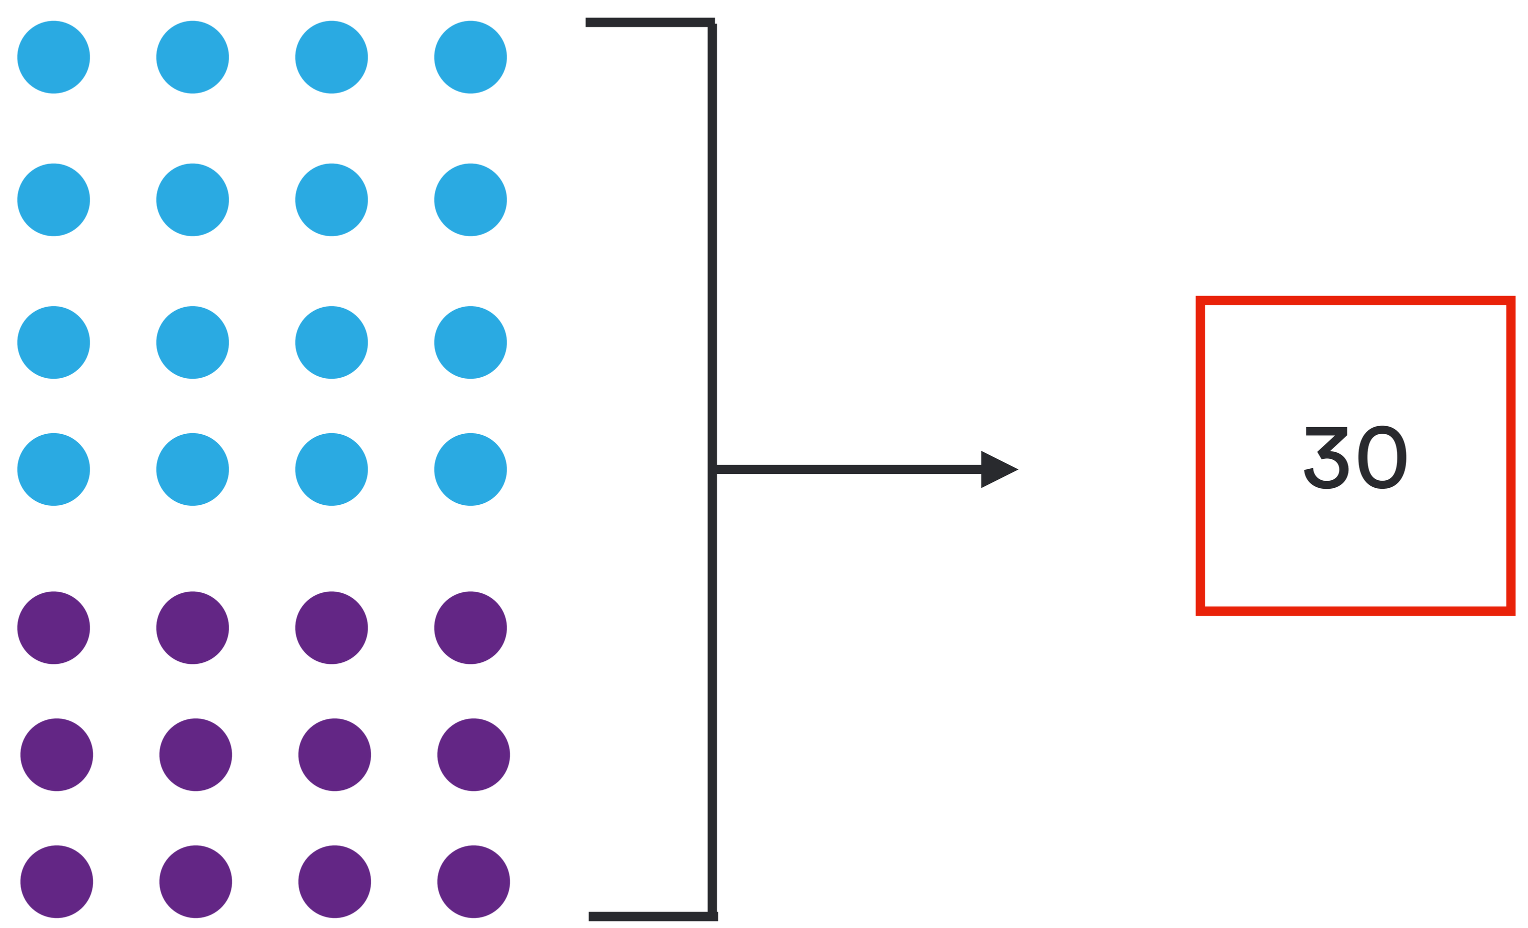

Node selects available messages and combines them into a block together with reference to predecessor and meta-data and broadcasts it

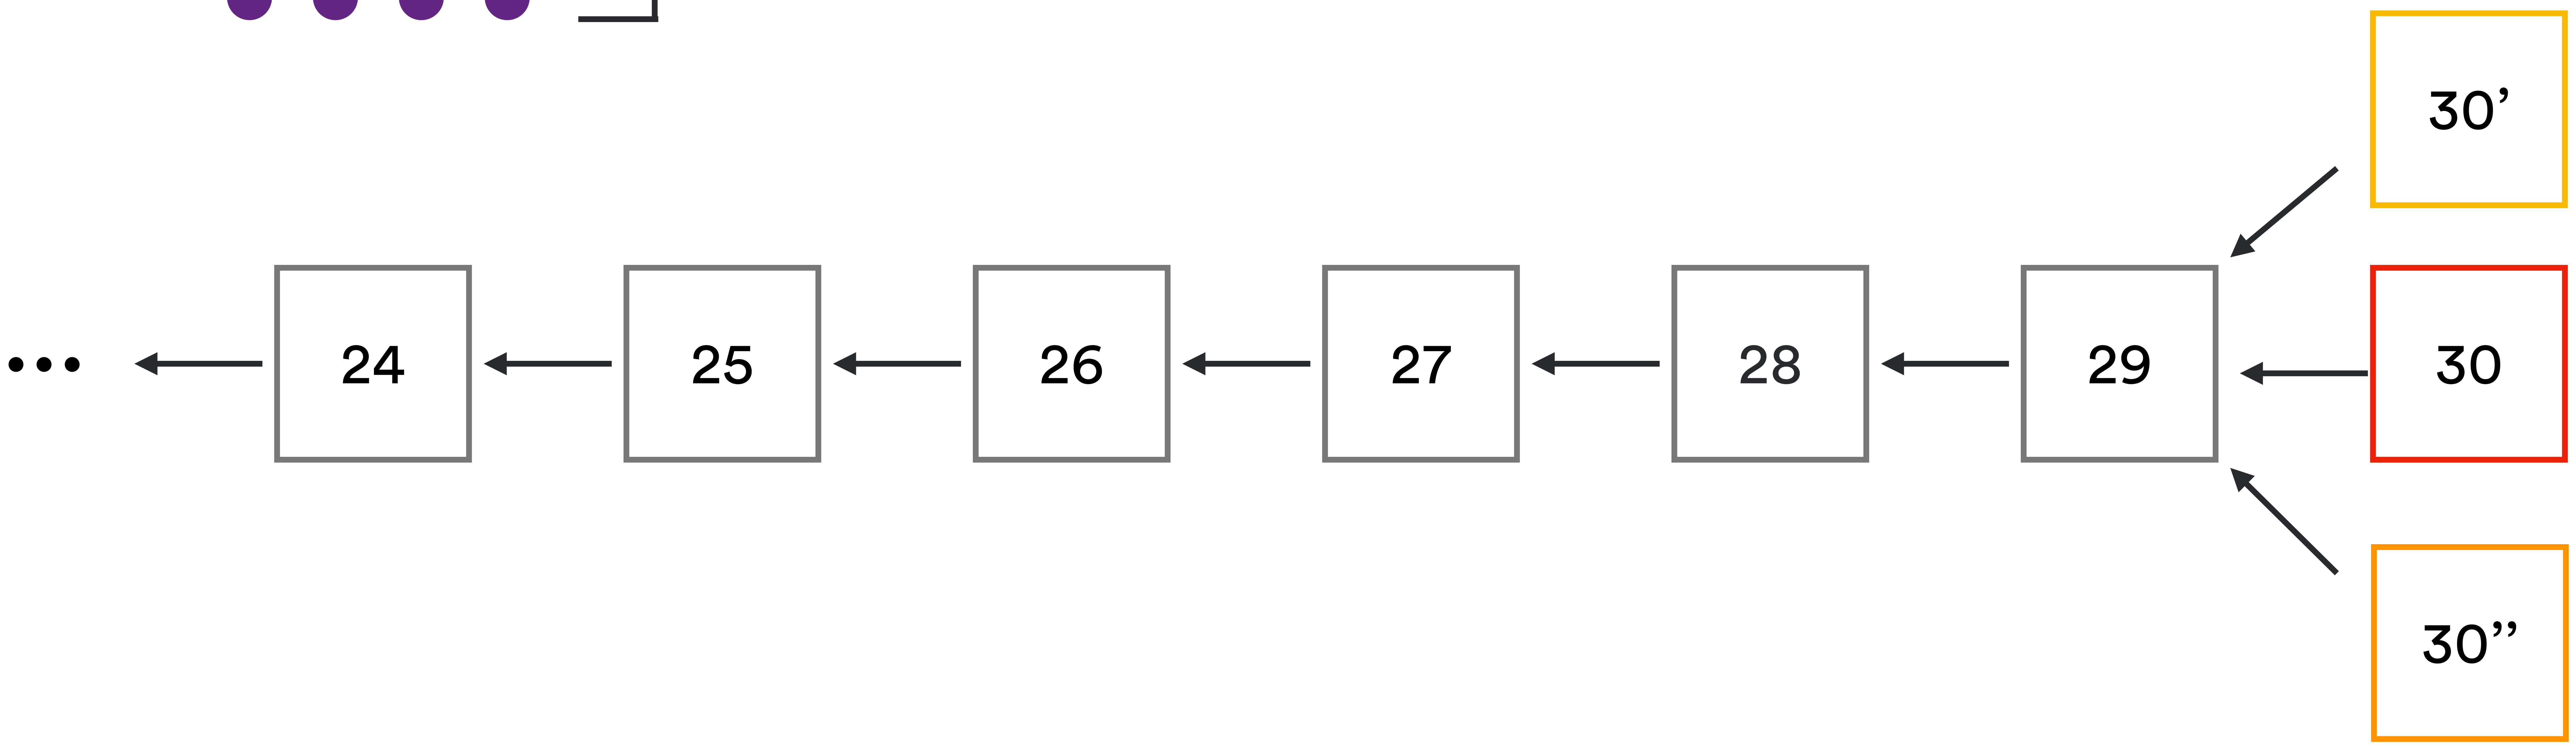

Notarization

The notarization process ensures that a *valid* block is known for every round

Step 1

Node 1 receives a block proposal for height 30, building on some notarized height 29 block

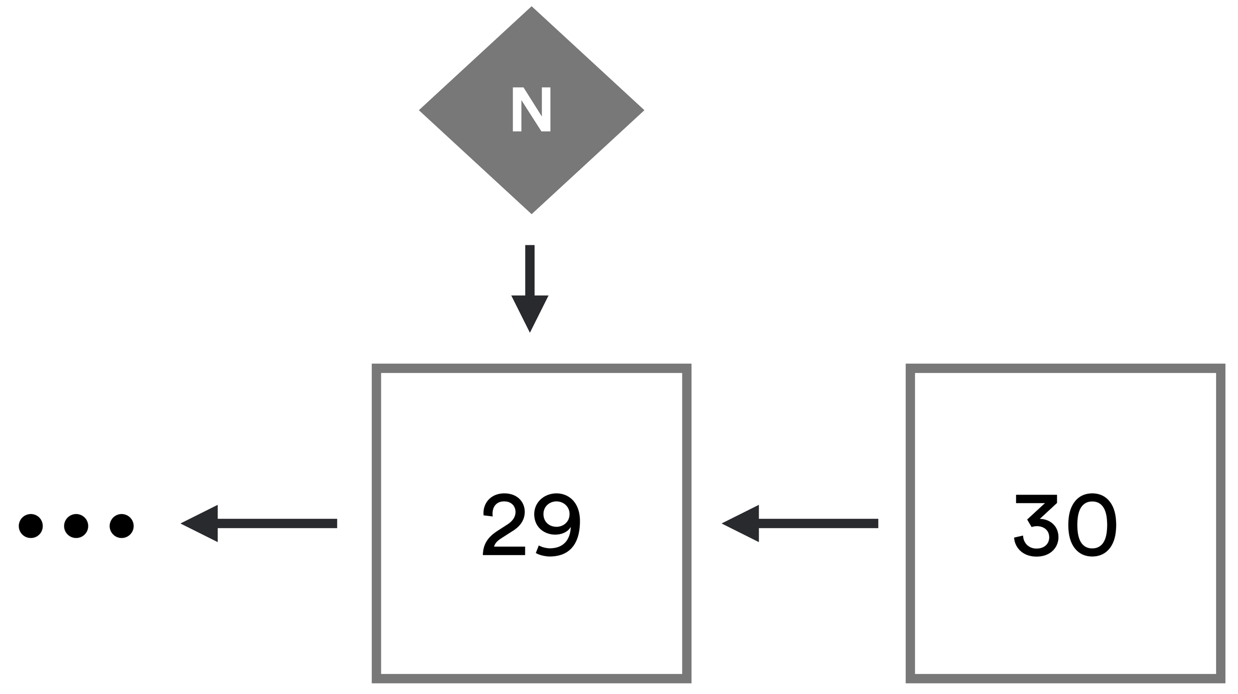

Step 2

Node 1 sees that the block is valid, signs it, and broadcasts it together with its *notarization* share

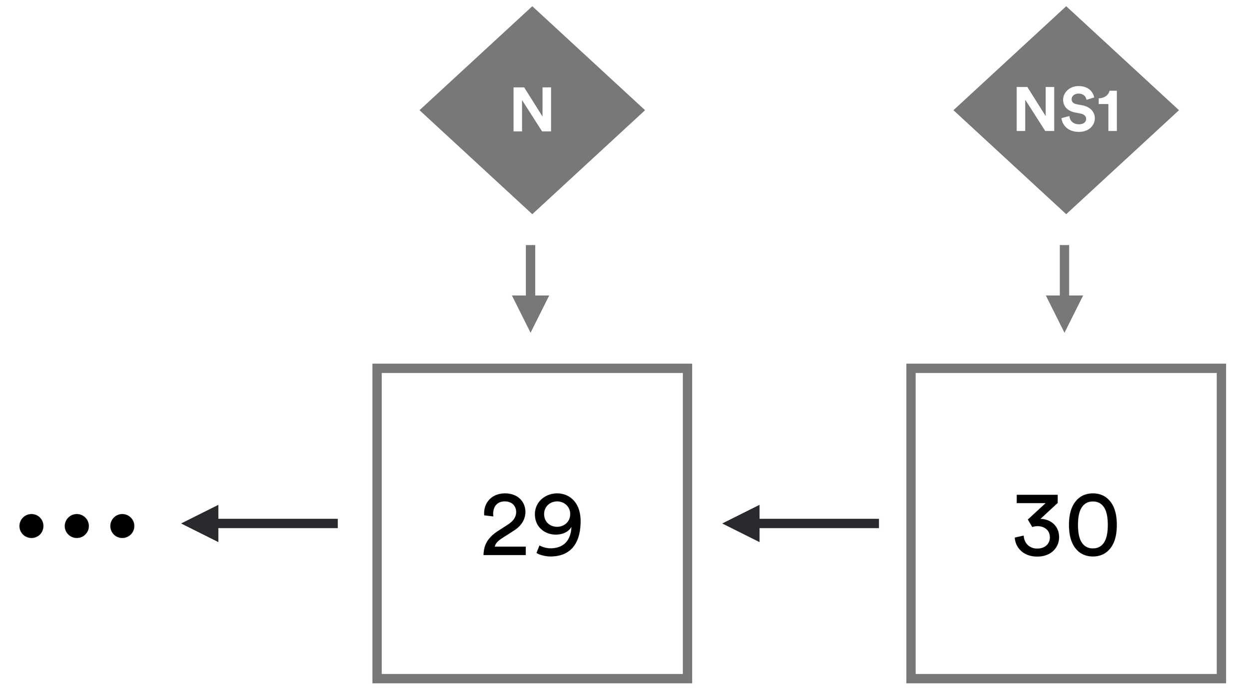

Step 3

Node 1 sees that nodes 3 and 4 also published their notarization shares on the block

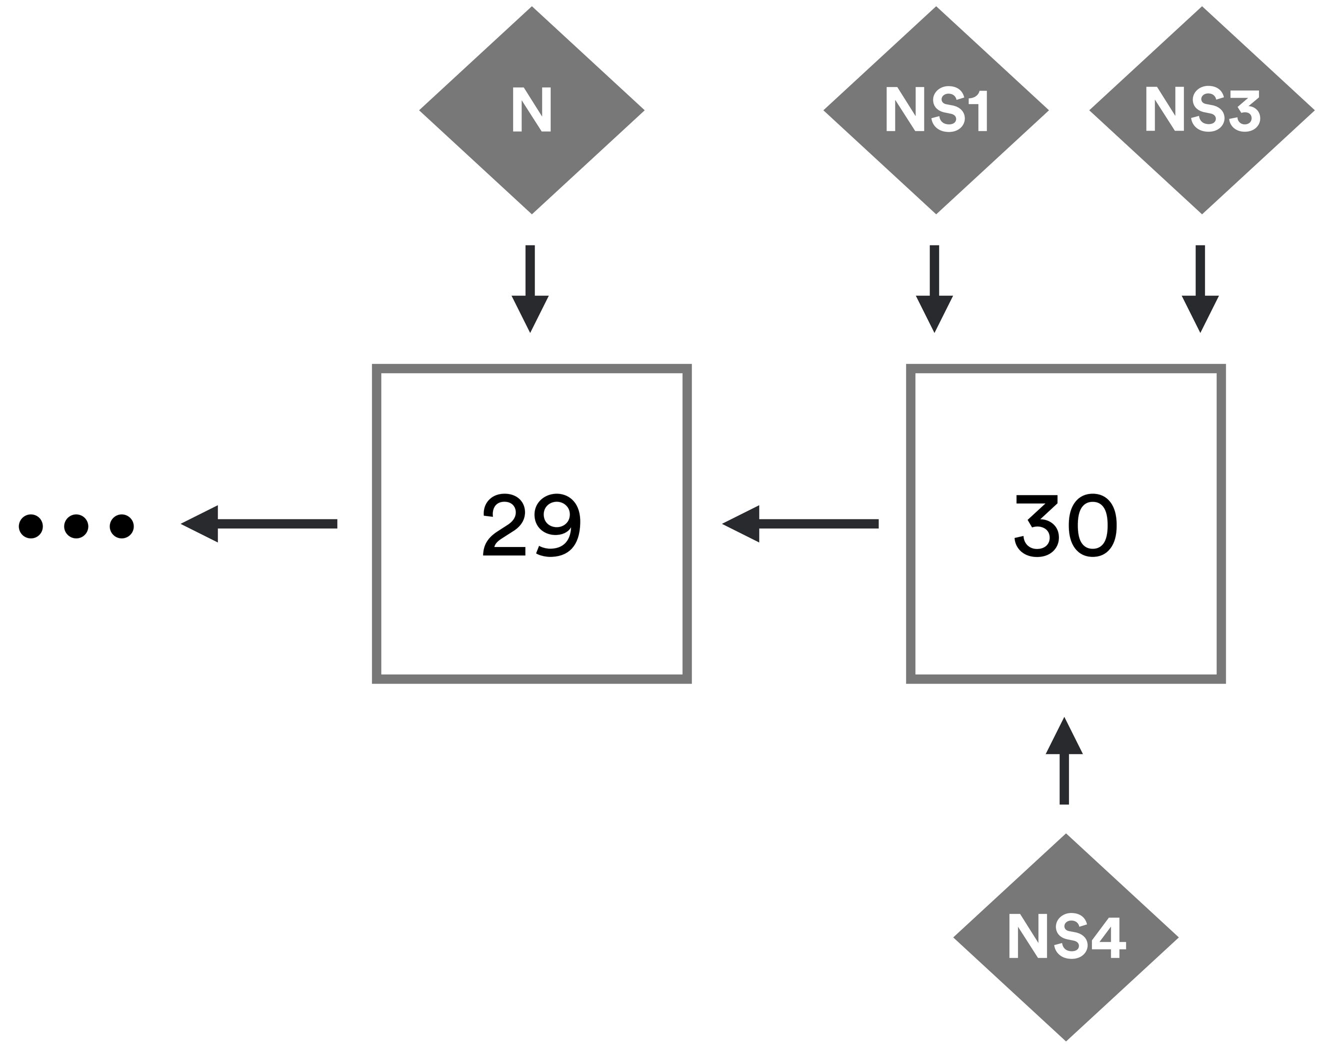

Step 4

3 notarization shares are sufficient approval: the shares are aggregated into a single full notarization. Block 30 is now notarized, and nodes wait for height 31 blocks

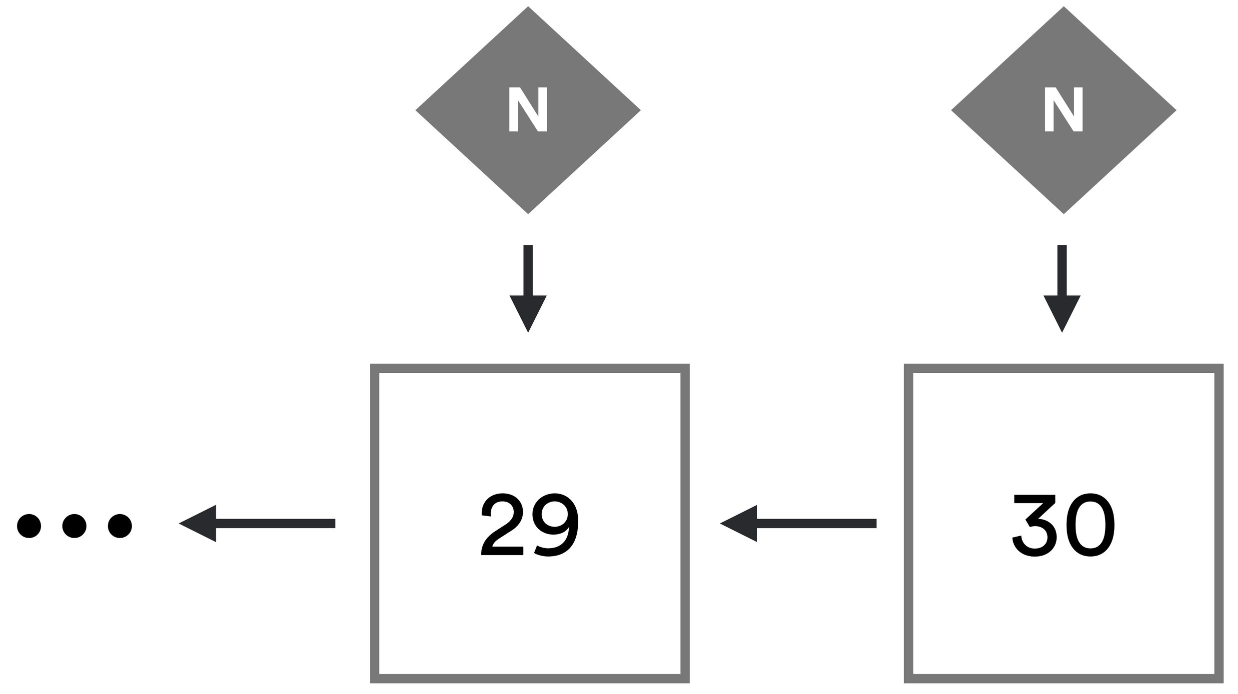

Notarization

Nodes may notary-sign multiple blocks to ensure that at least one block becomes fully notarized

Step 1

Node 1 receives a block proposal for height 30, building on some notarized height 29 block

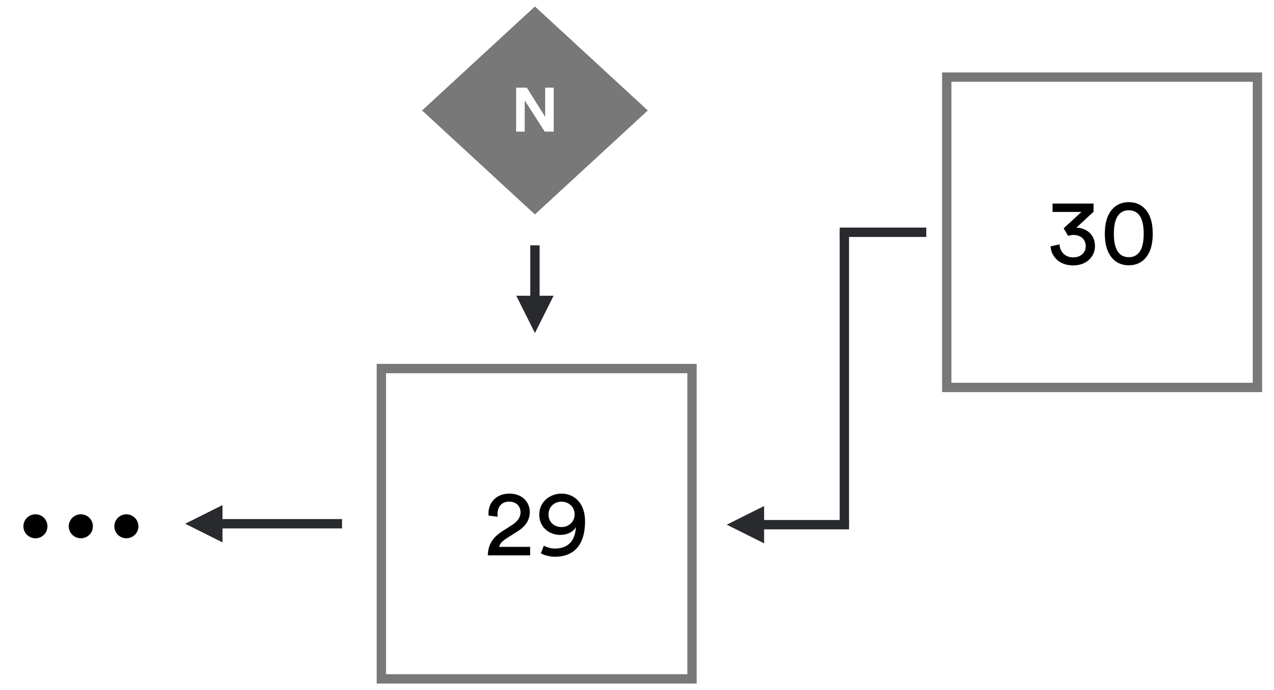

Step 2

Node 1 sees that the block is valid, signs it, and broadcasts it together with its *notarization* share

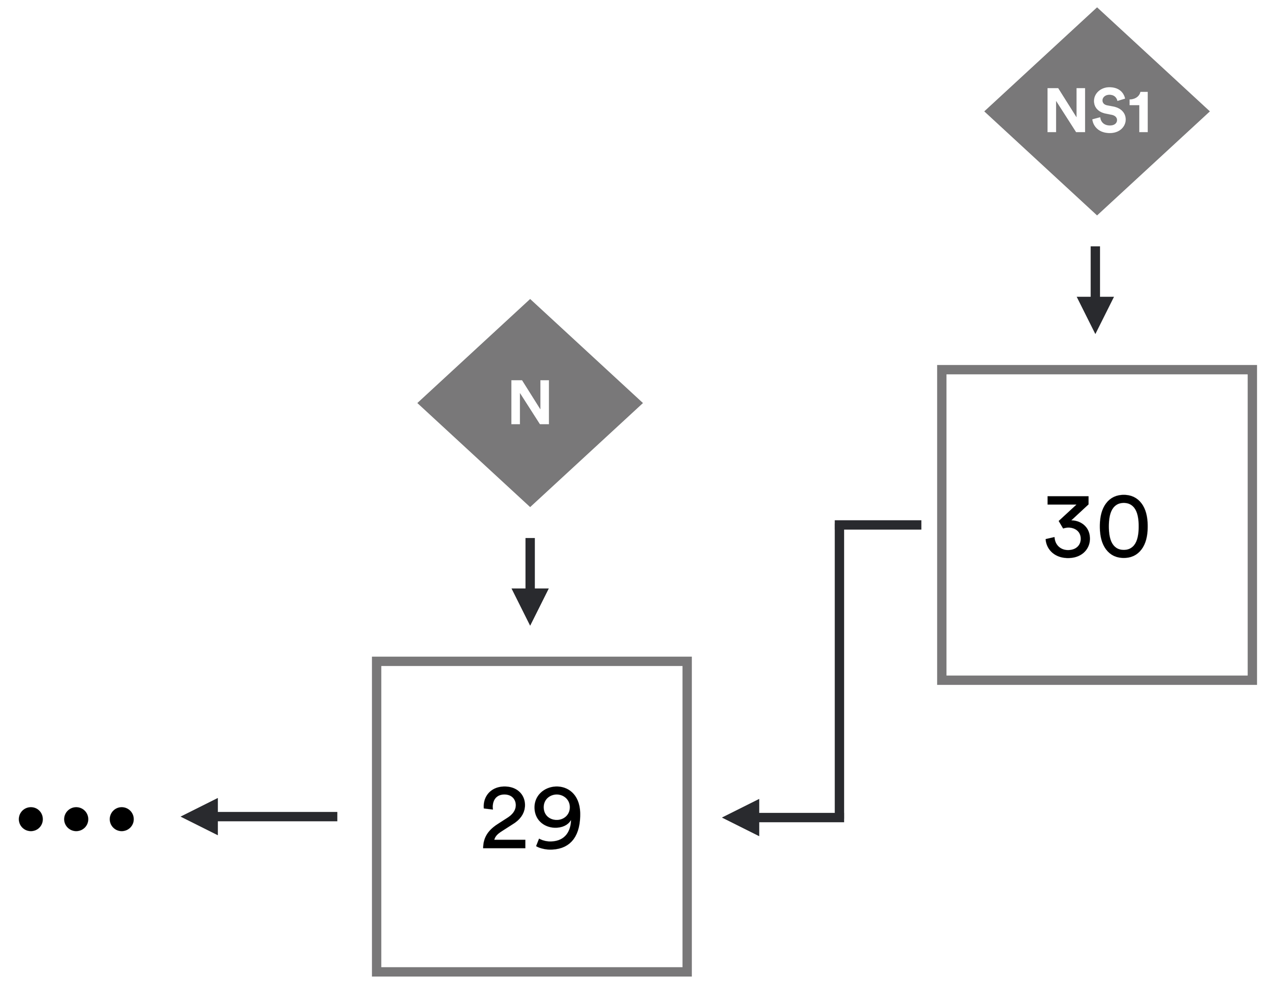

Step 3

Nodes 1 sees another height 30 block, which is also valid, and it broadcasts it together with another notarization share

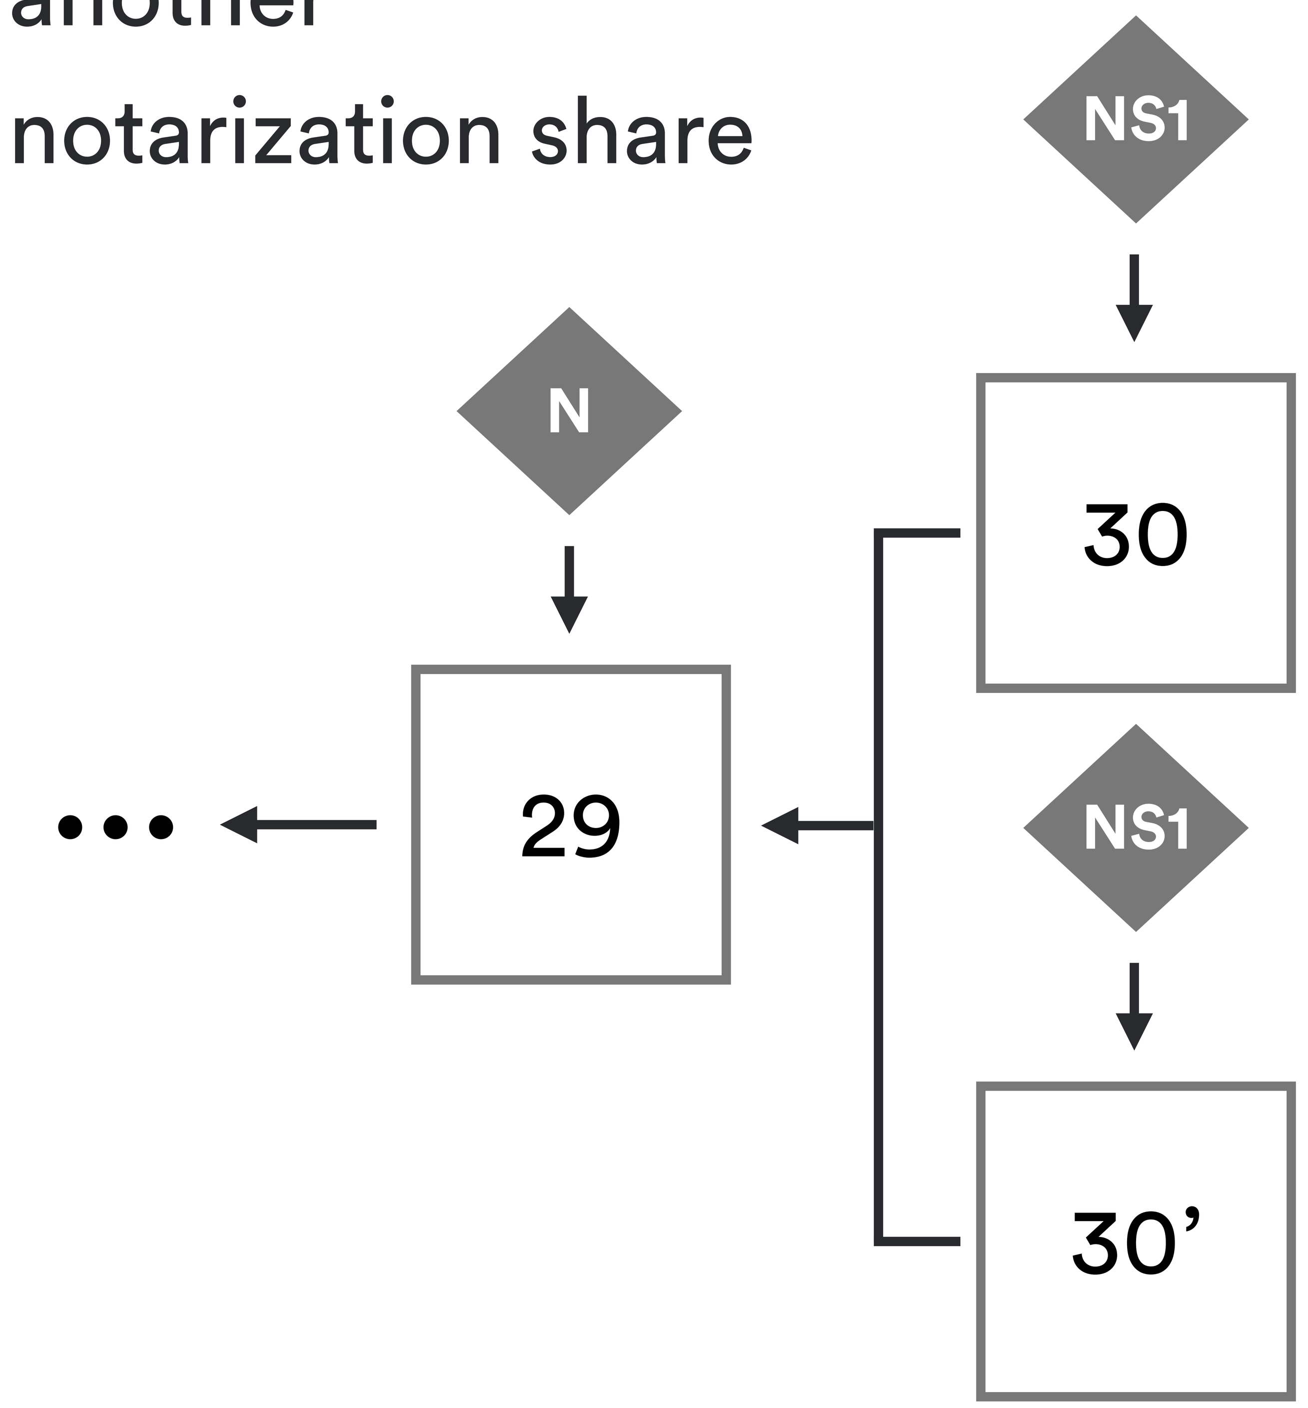

Step 4

Both height 30 blocks get enough support to become notarized

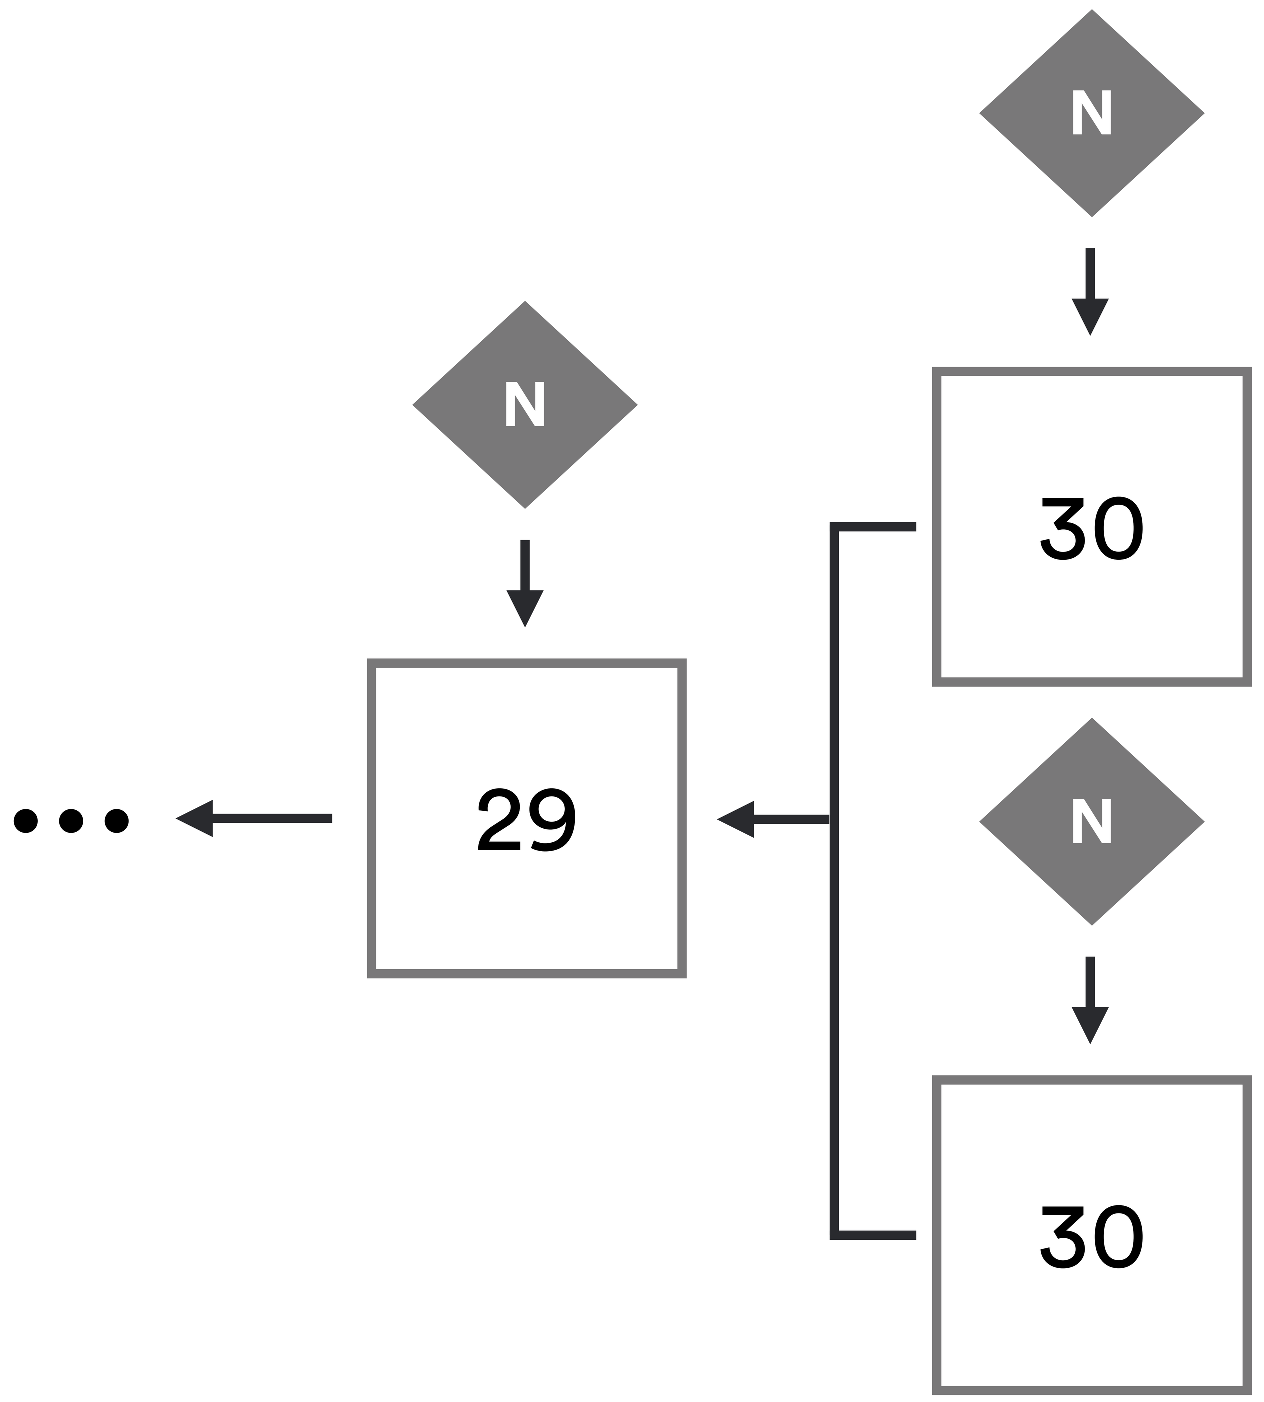

Notarization

Multiple notarized blocks may exist at the same height, at least one per height

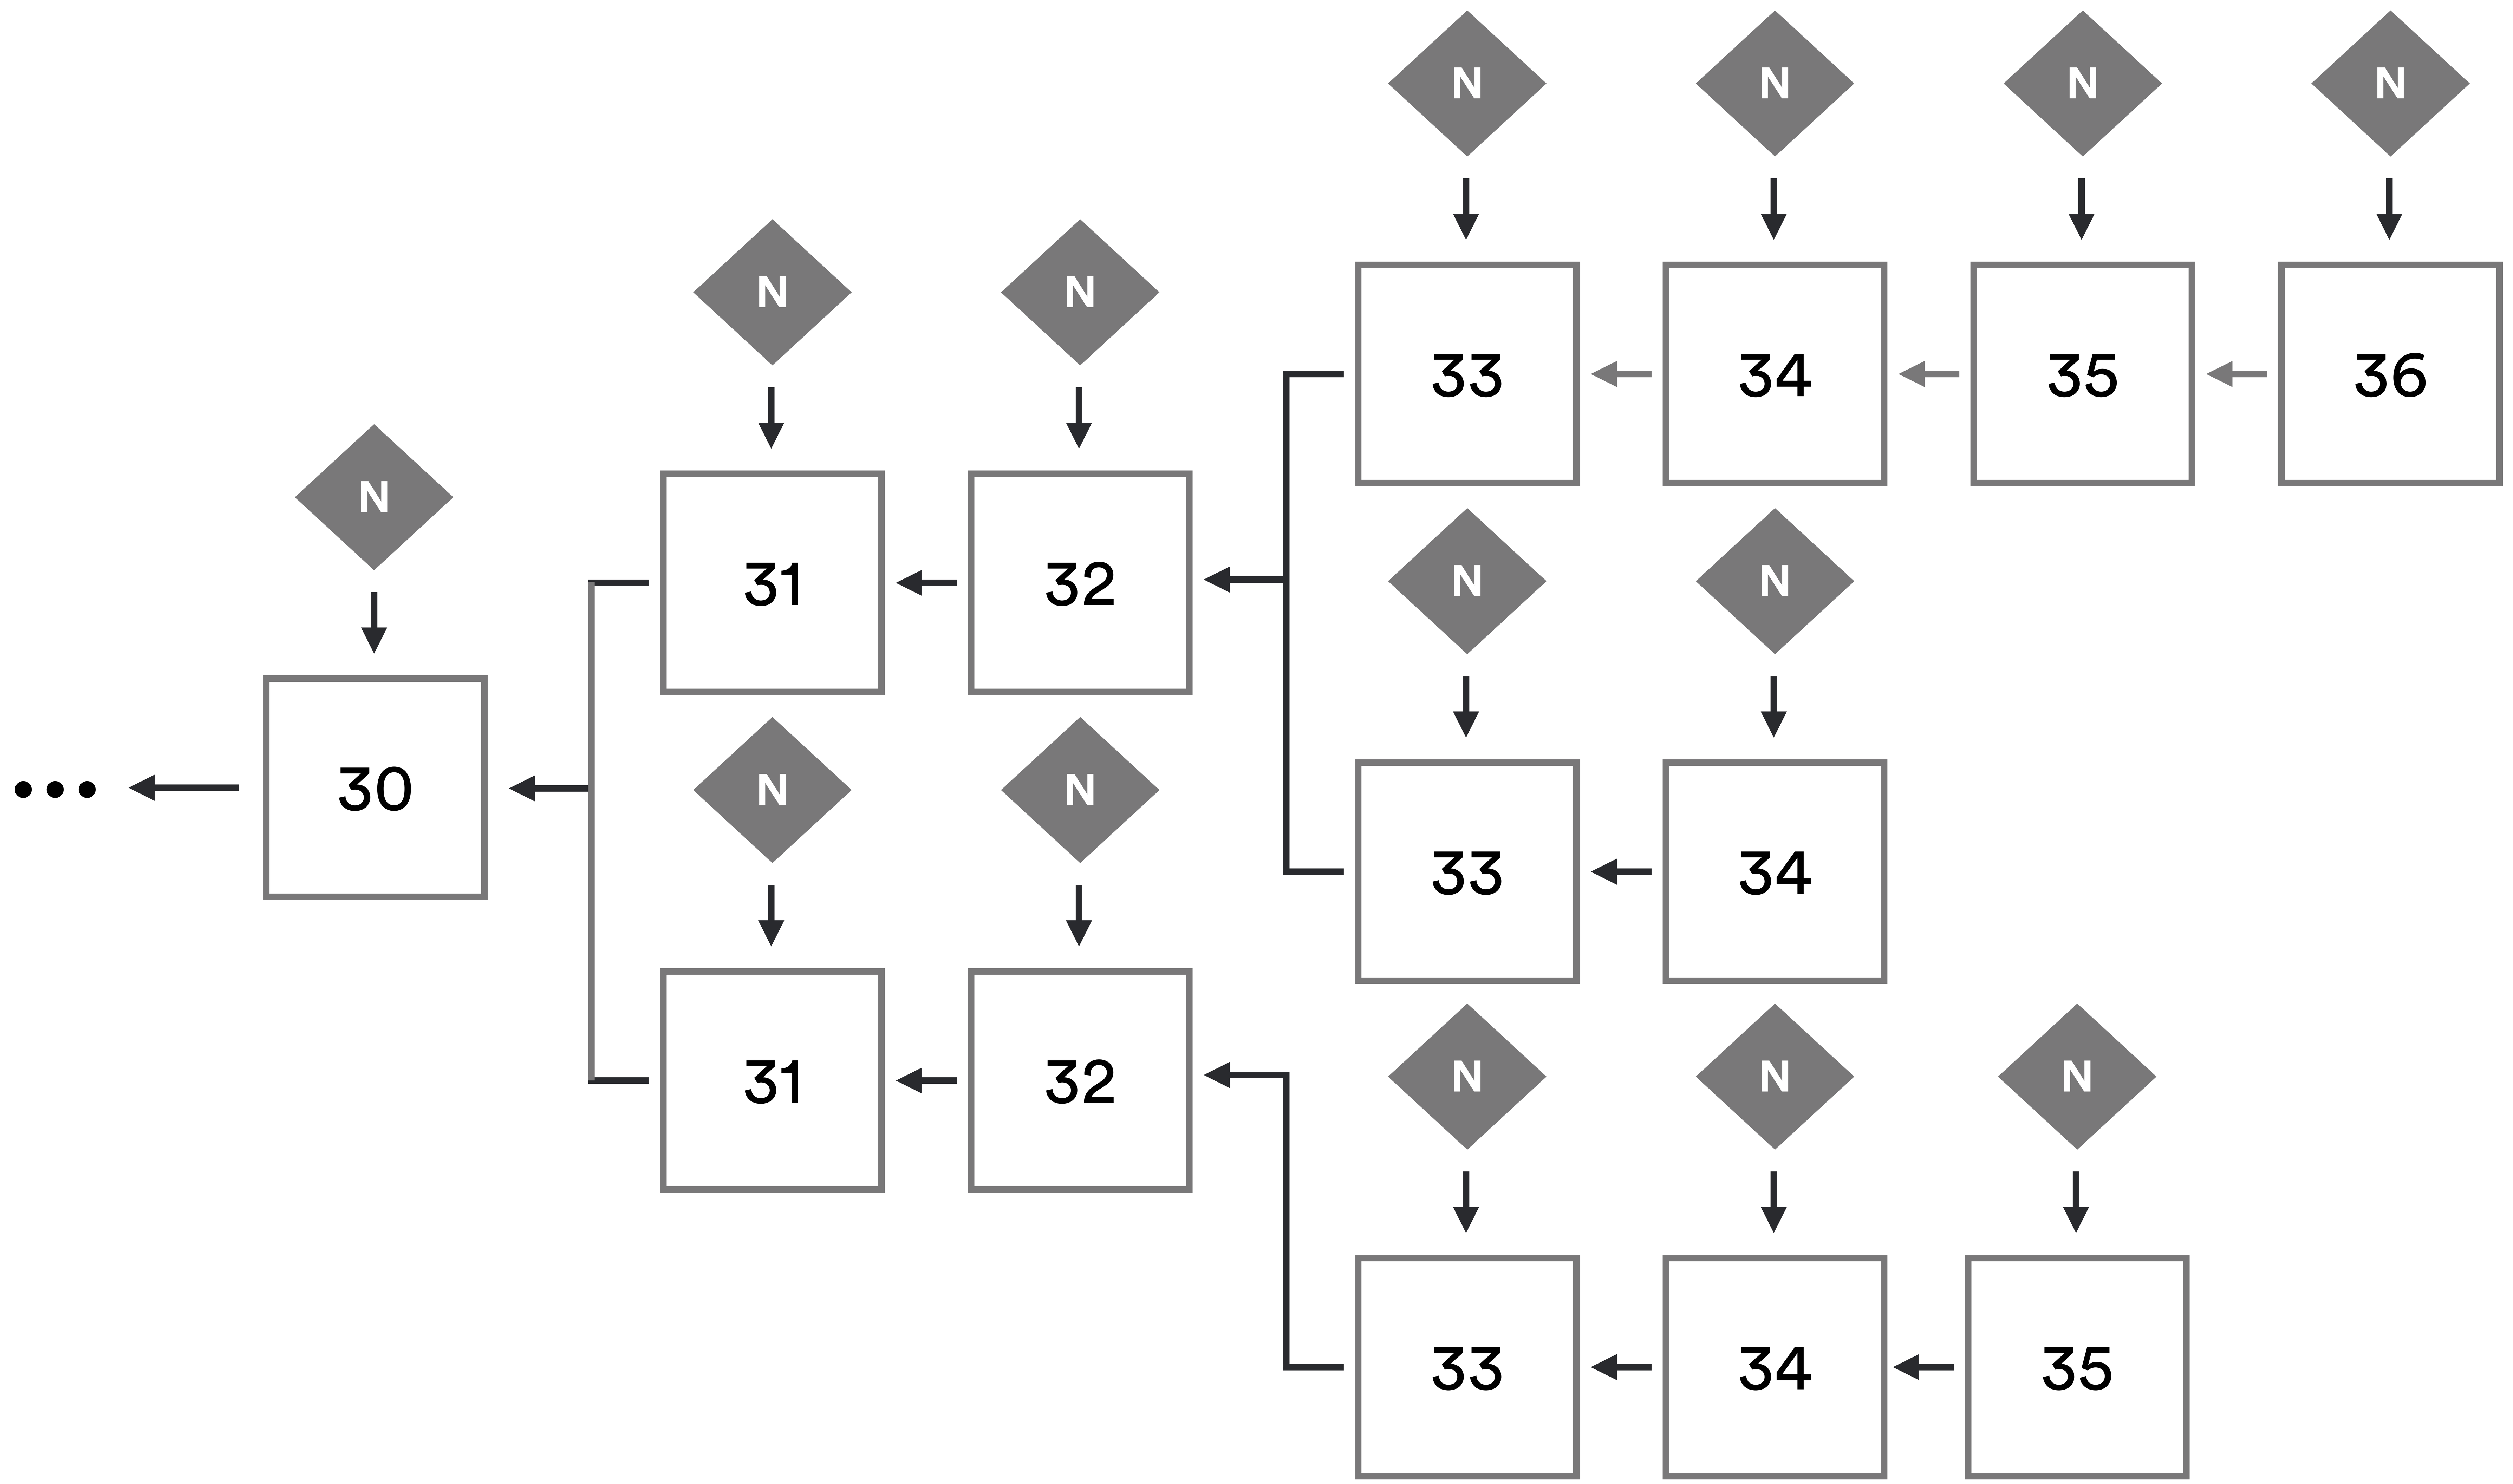

Random Beacon

At every height, there is an unpredictable random value shared by the nodes

BLS-Threshold Signatures

- Pseudo random (not predictable, no last actor bias)
- Non-interactive distributed key generation
- Non-interactive independent signature share creation
- Unique: for every message *m* there exist one signature, regardless of the threshold group

Random Beacon

At every height, there is an unpredictable random value shared by the nodes

Step 1

Node 1 has Random Beacon 29 and wants to help constructing Random Beacon 30

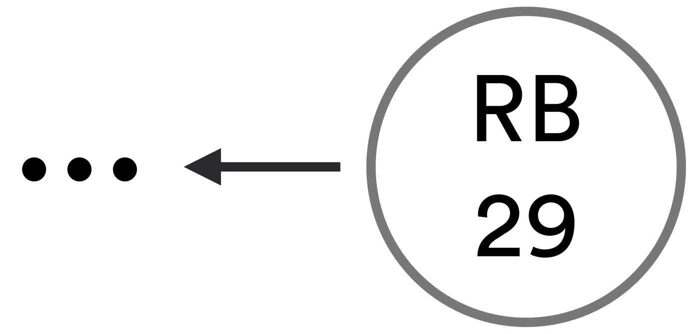

Step 2

Node 1 signs RB29 using a threshold signature scheme, yielding a share of random beacon 30

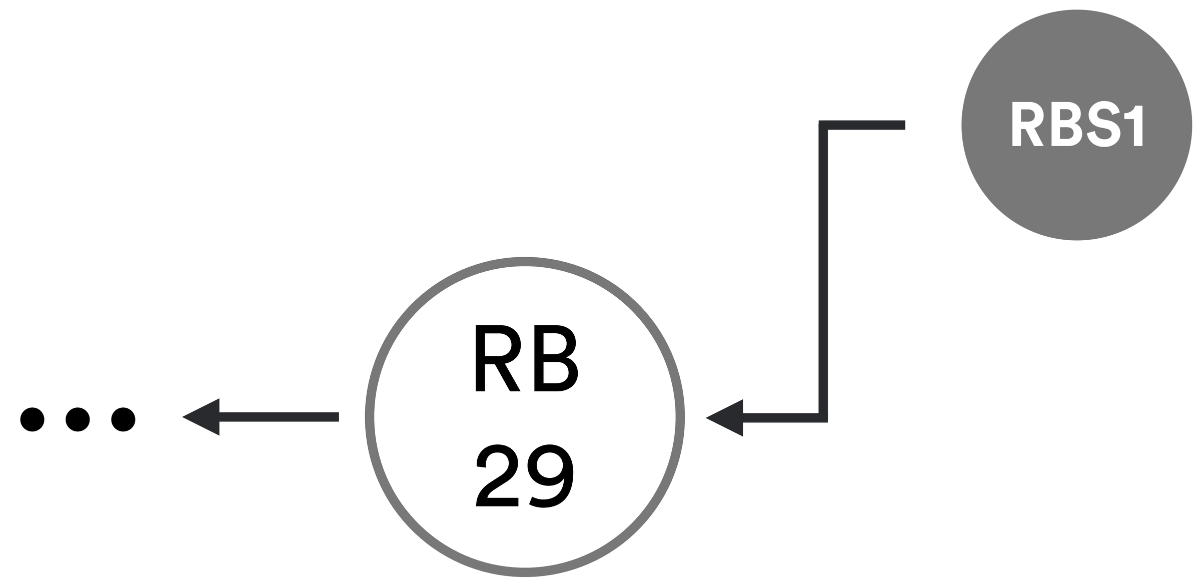

Step 3

Nodes 1 sees that node 2 also published a share of Random Beacon 30

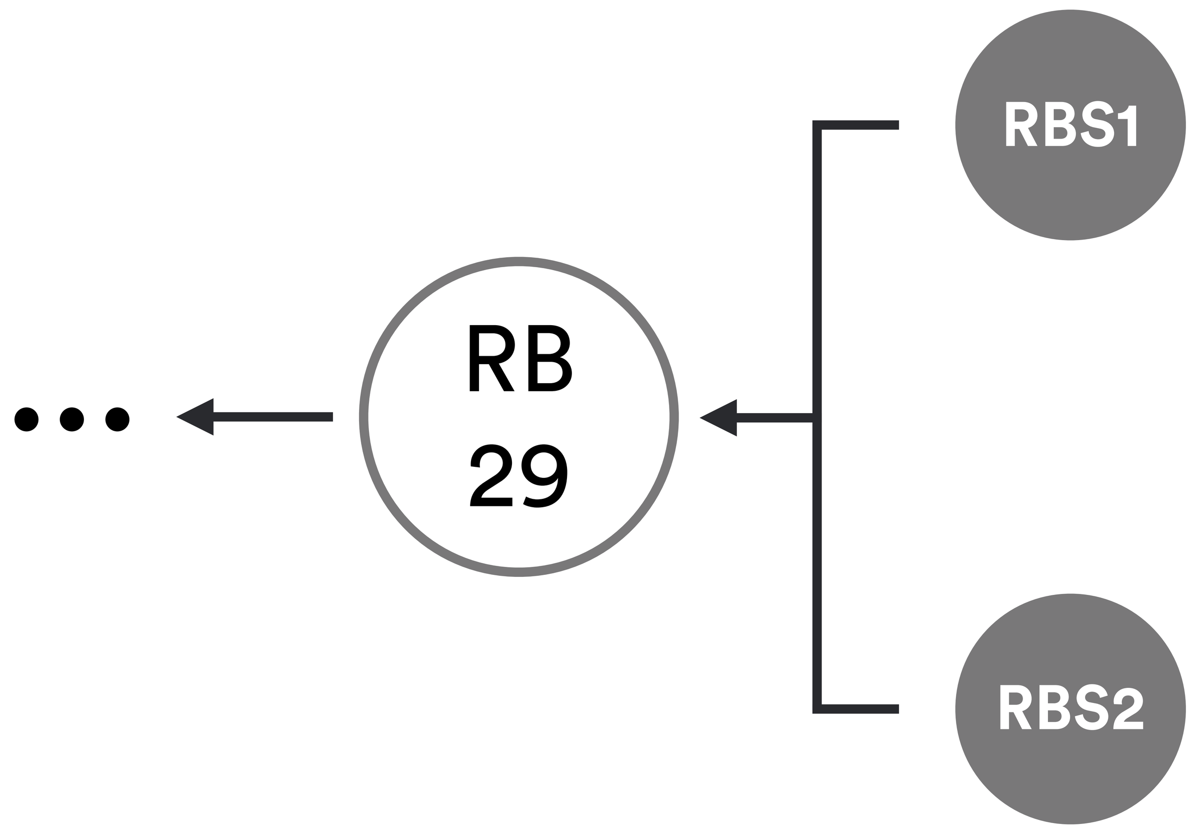

Step 4

2 random beacon shares are sufficient to reconstruct a full threshold signature, which is Random Beacon 30

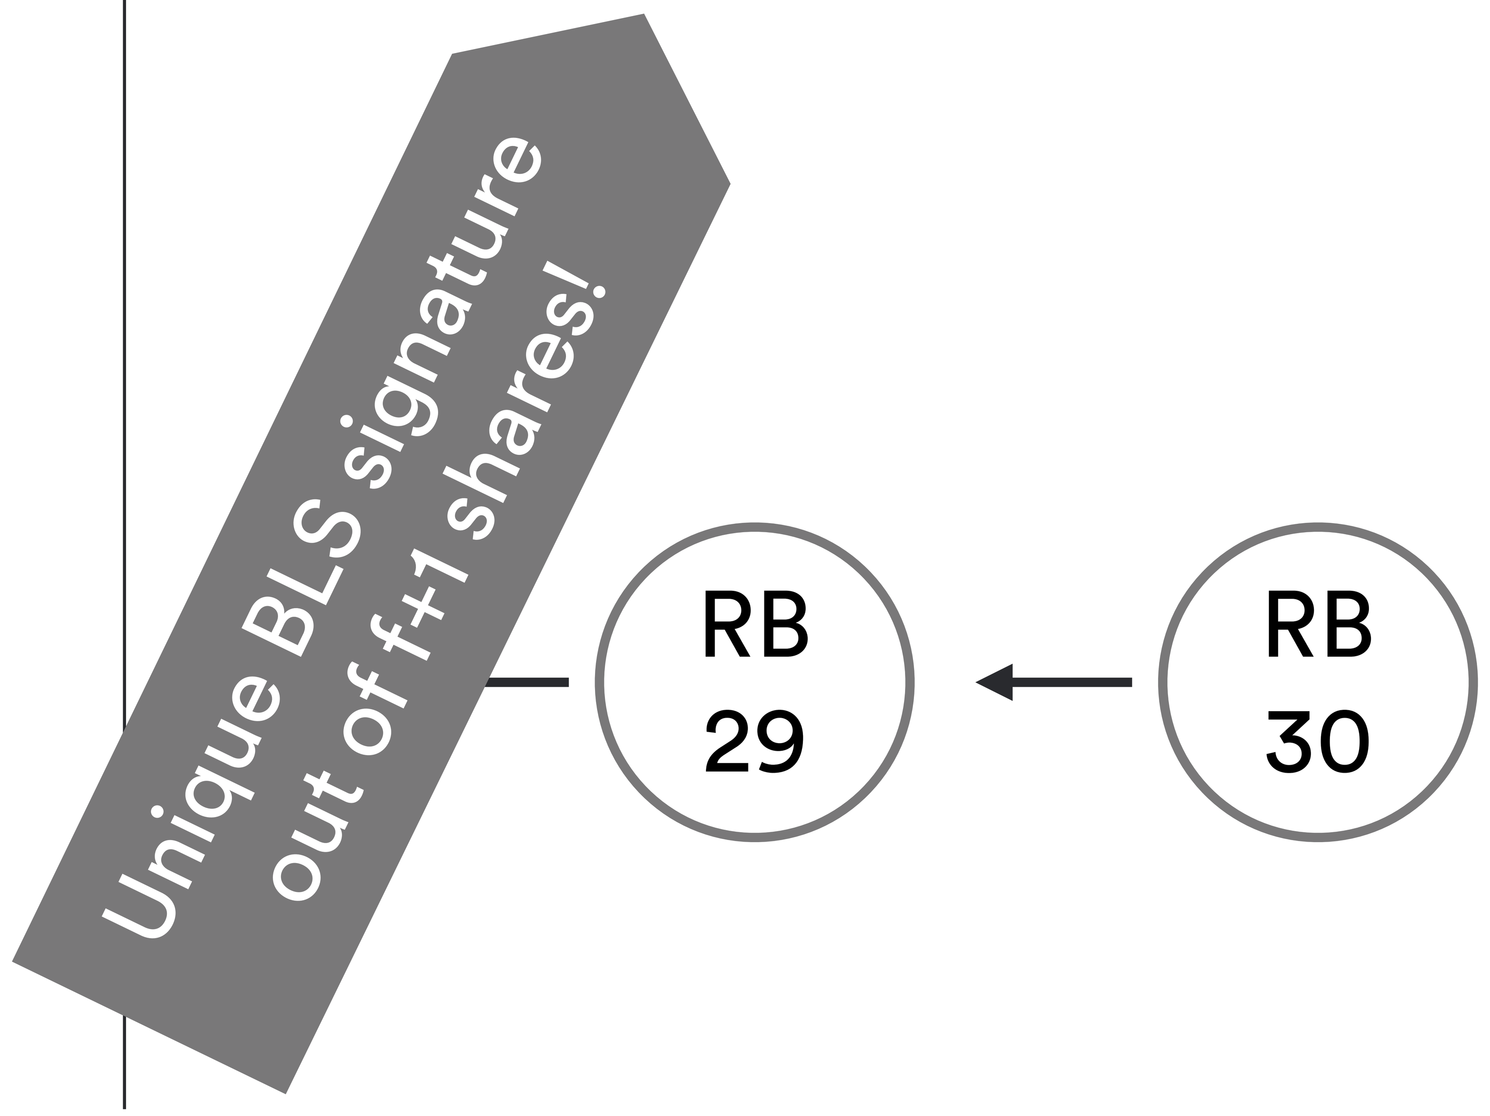

Block Maker Ranking

The Random Beacon is used to rank block makers

High Priority

	Round 24	Round 25	Round 26	Round 27	Round 28	Round 29
Rank 0	Node 1	Node 4	Node 2	Node 3	Node 4	Node
Rank 1	Node 4	Node 3	Node 3	Node 1	Node 2	Node 1
Rank 2	Node 3	Node 1	Node 4	Node 4	Node 1	Node 3
Rank 3	Node 2	Node 2	Node 1	Node 2	Node 3	Node 4

Low Priority

Notarization with Block Maker Ranking

Rounds are divided into time slots defining when block maker proposals are considered

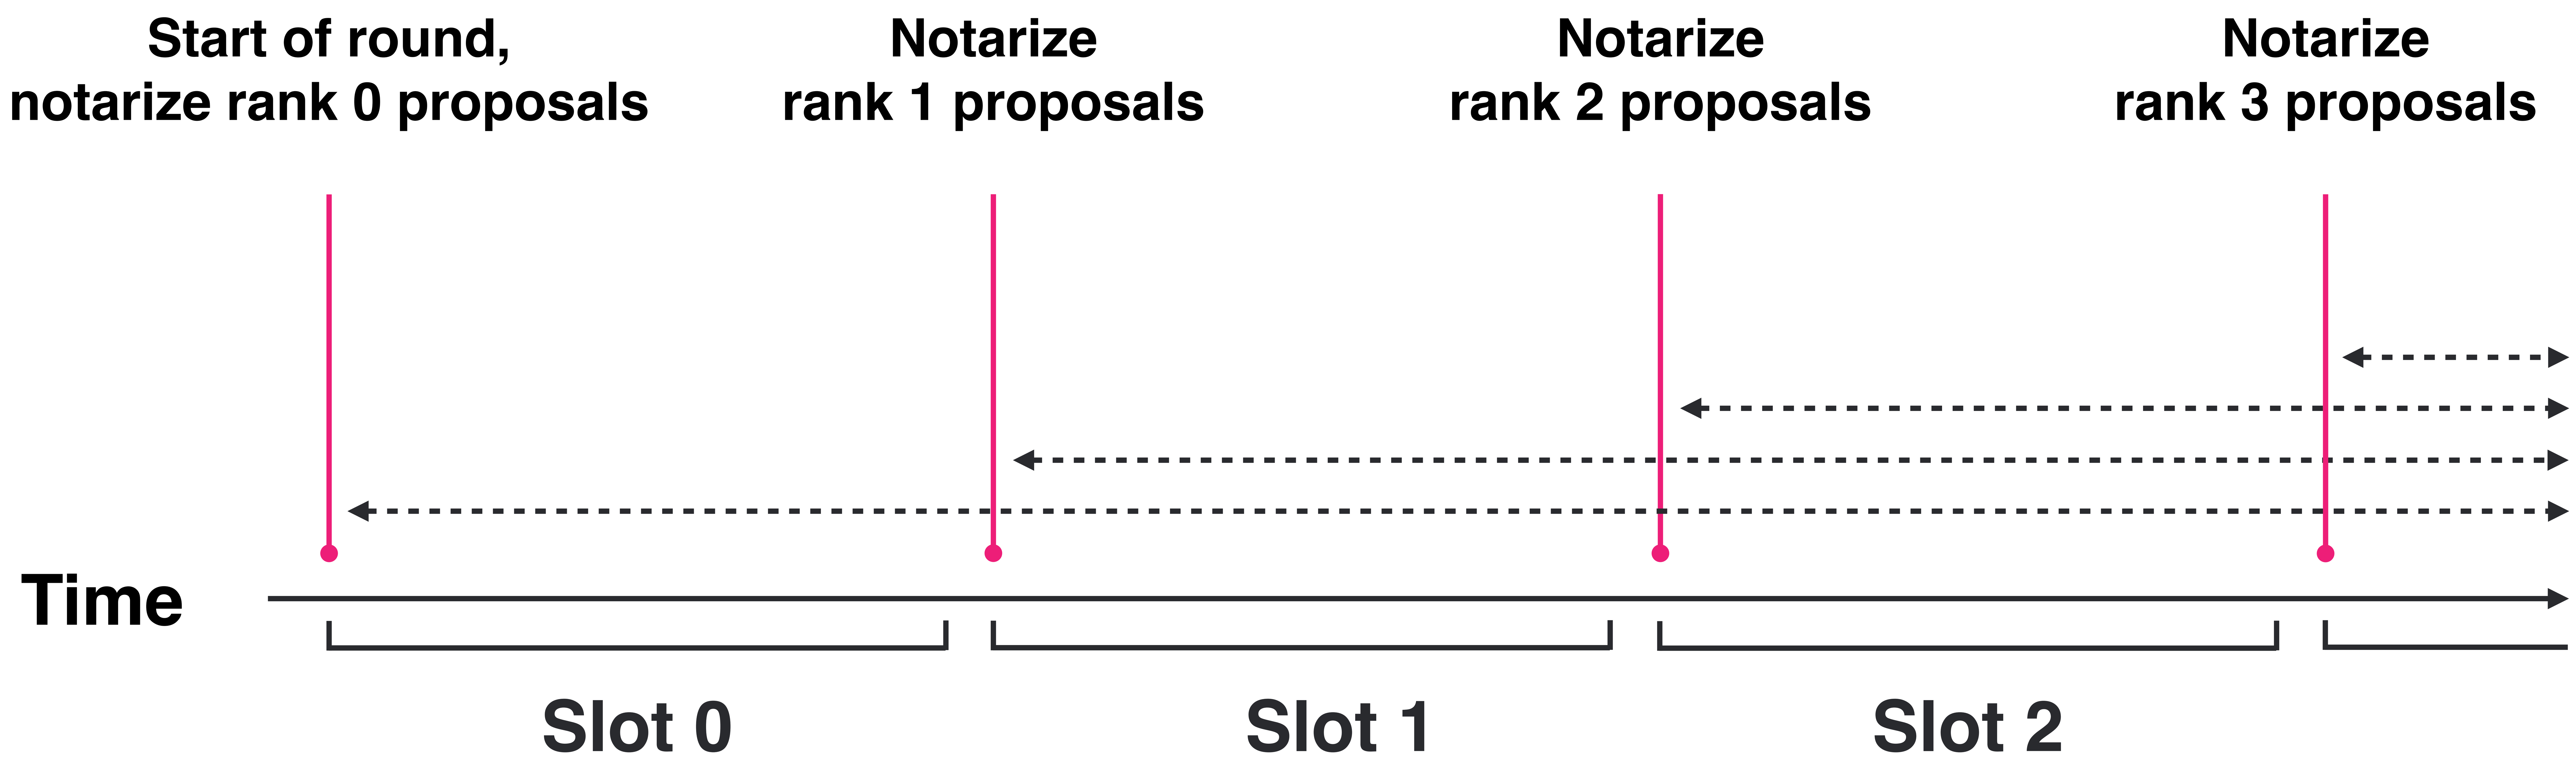

Notarization with Block Maker Ranking

The block ranks can reduce the number of notarized blocks

Step 1

Node 1 receives a rank-1 block proposal for height 30, building on some notarized height 29 block

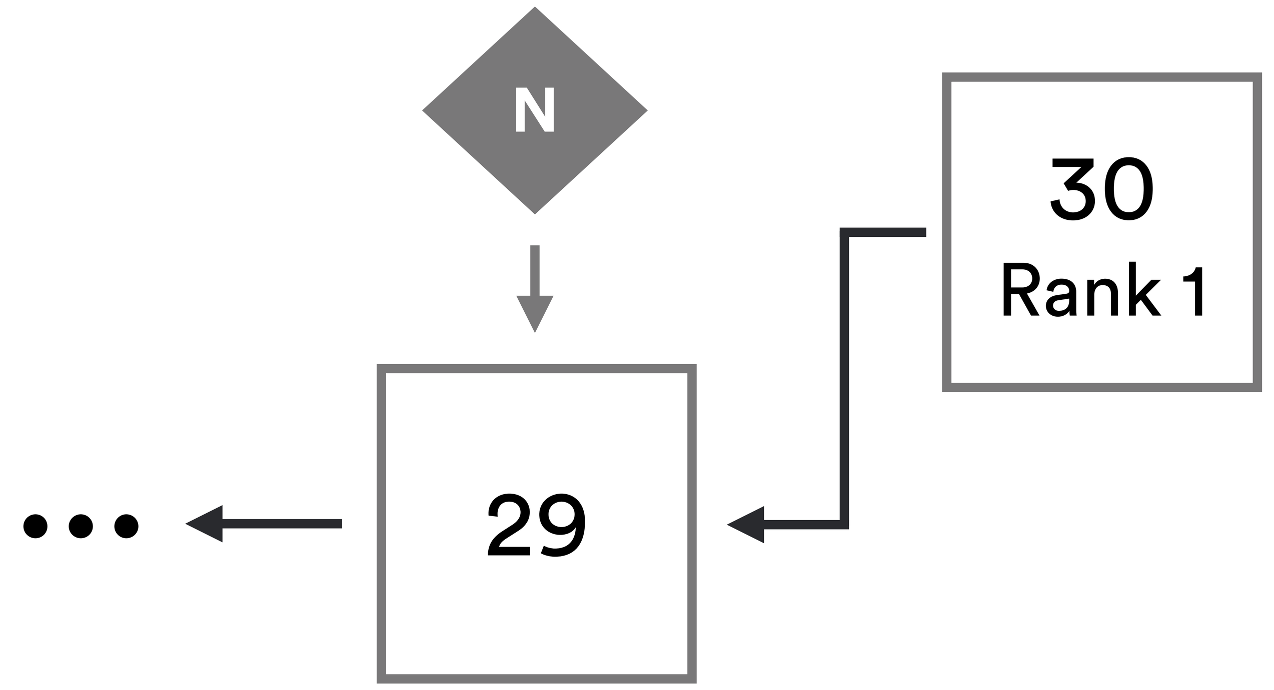

Step 2

Node 1 is still in time slot 0, not willing to notary-sign a rank-1 block yet

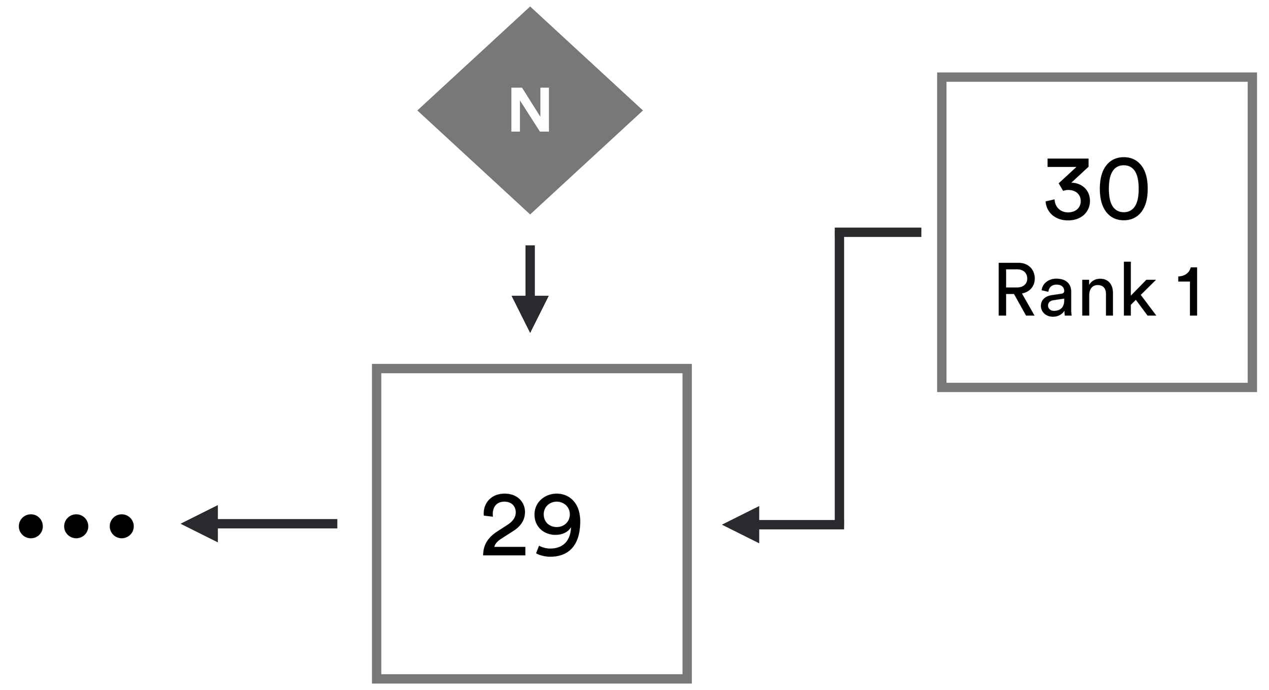

Step 3

Nodes 1 sees a valid rank-0 height 30 block, and it broadcasts it together with a notarization share

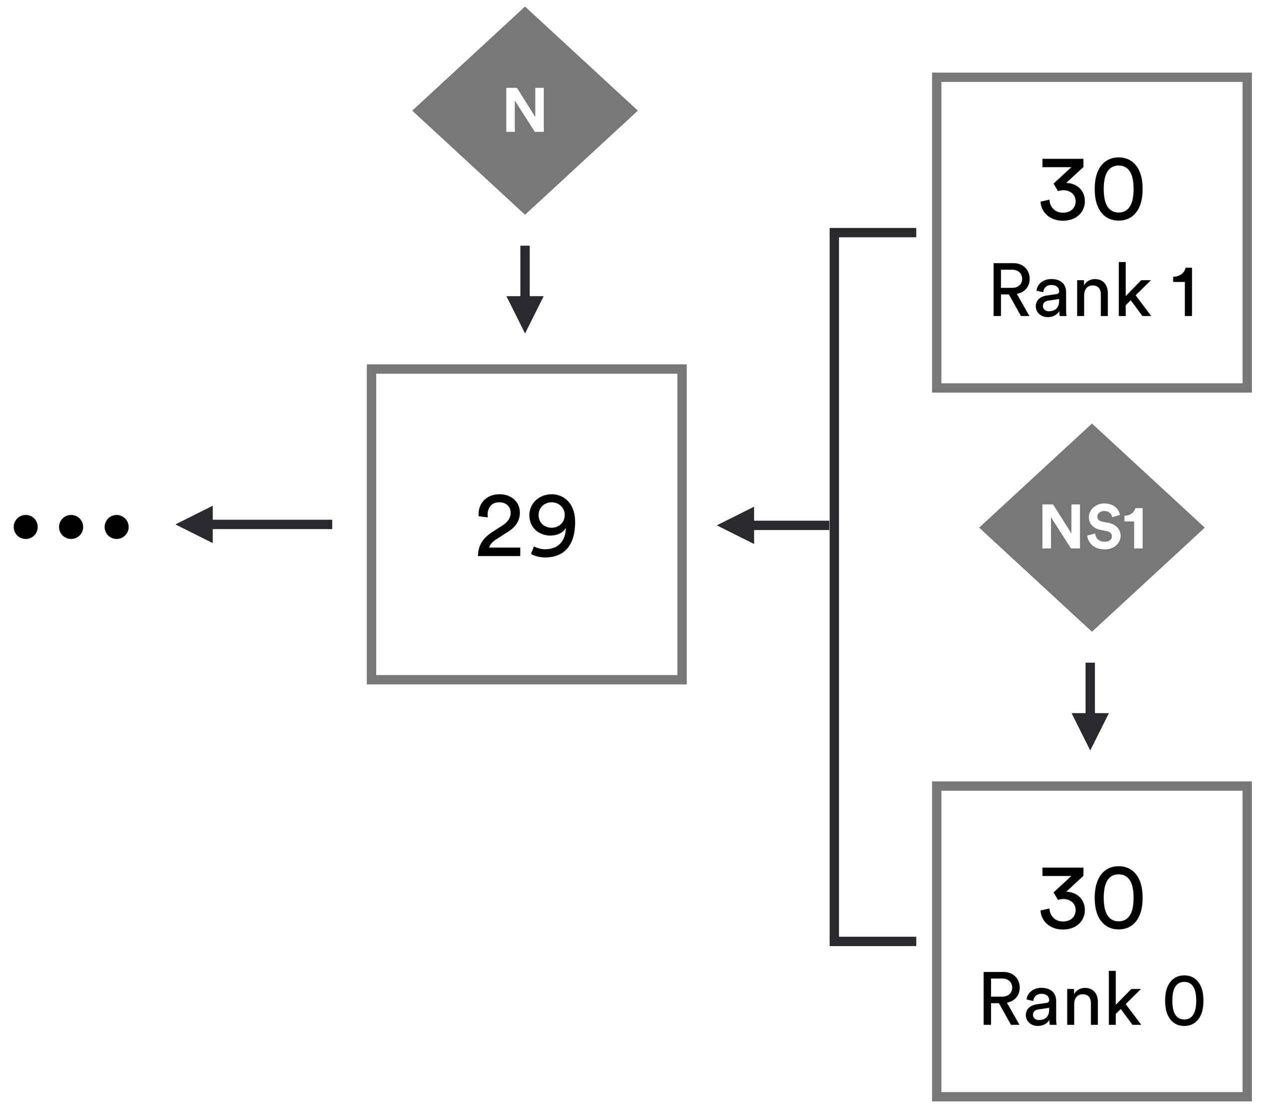

Step 4

Eventually, only the rank 0 block becomes notarized

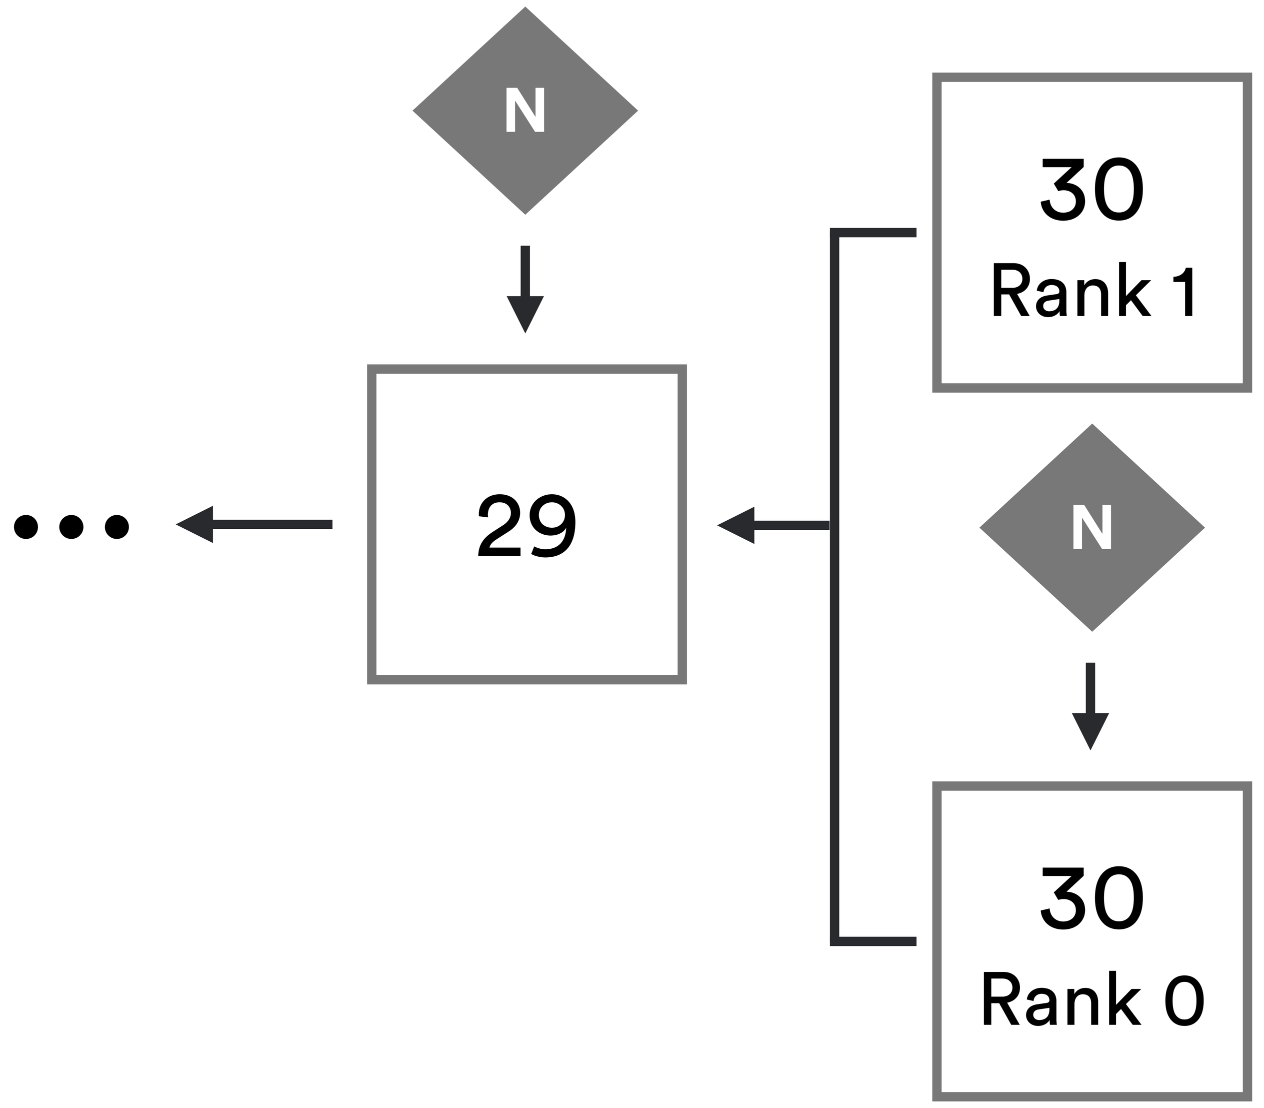

Notarization with Block Maker Ranking

One notarized block *b* at a height *h* = Agreement up to *h*

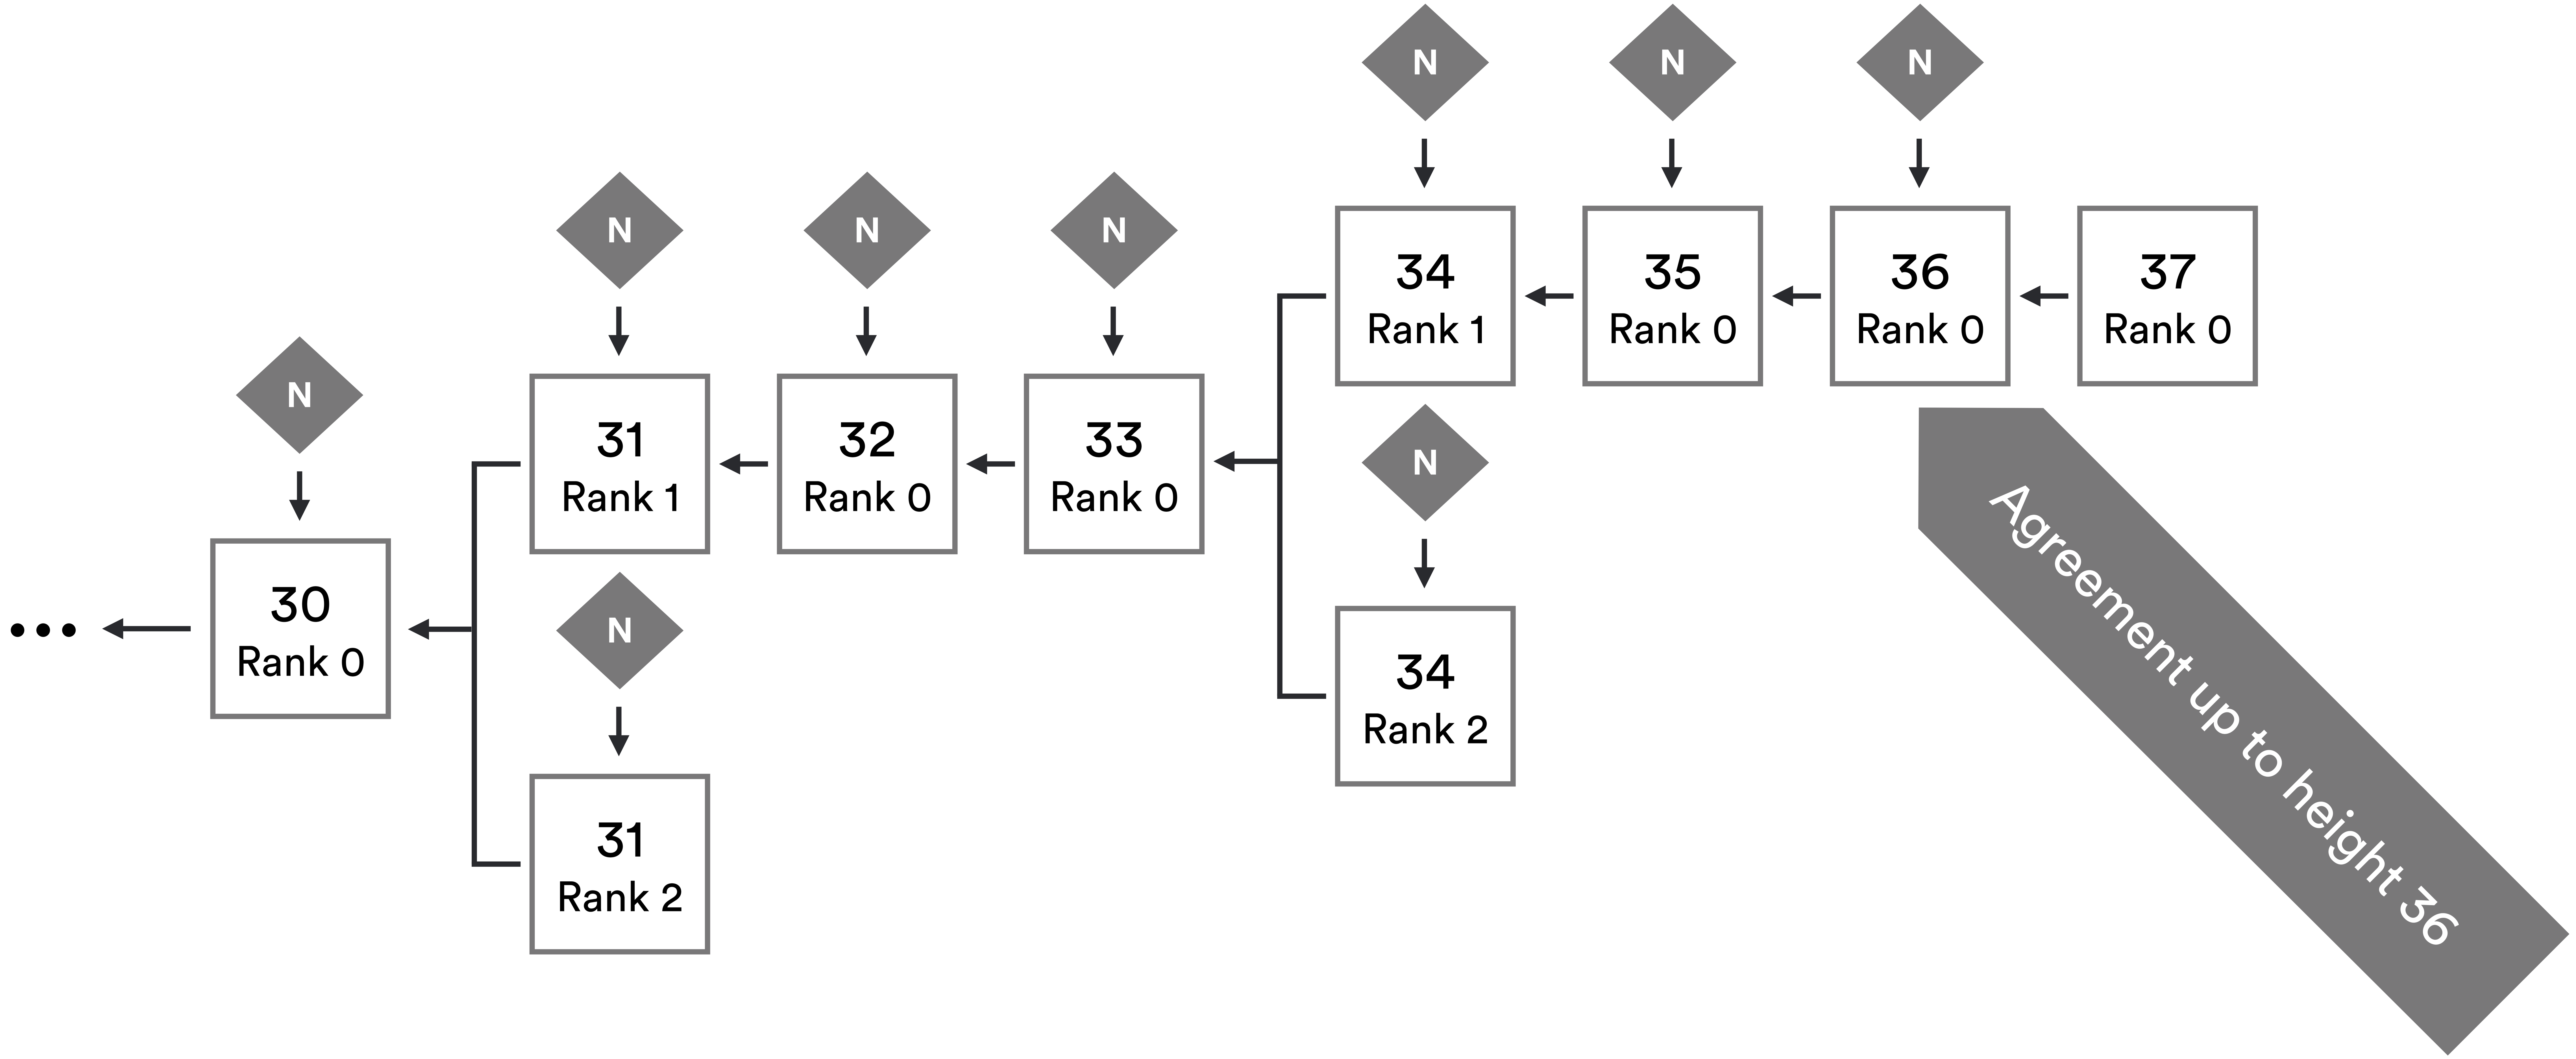

How can we detect this…?

Notarization with Block Maker Ranking

Synchronous communication → Forks can be removed

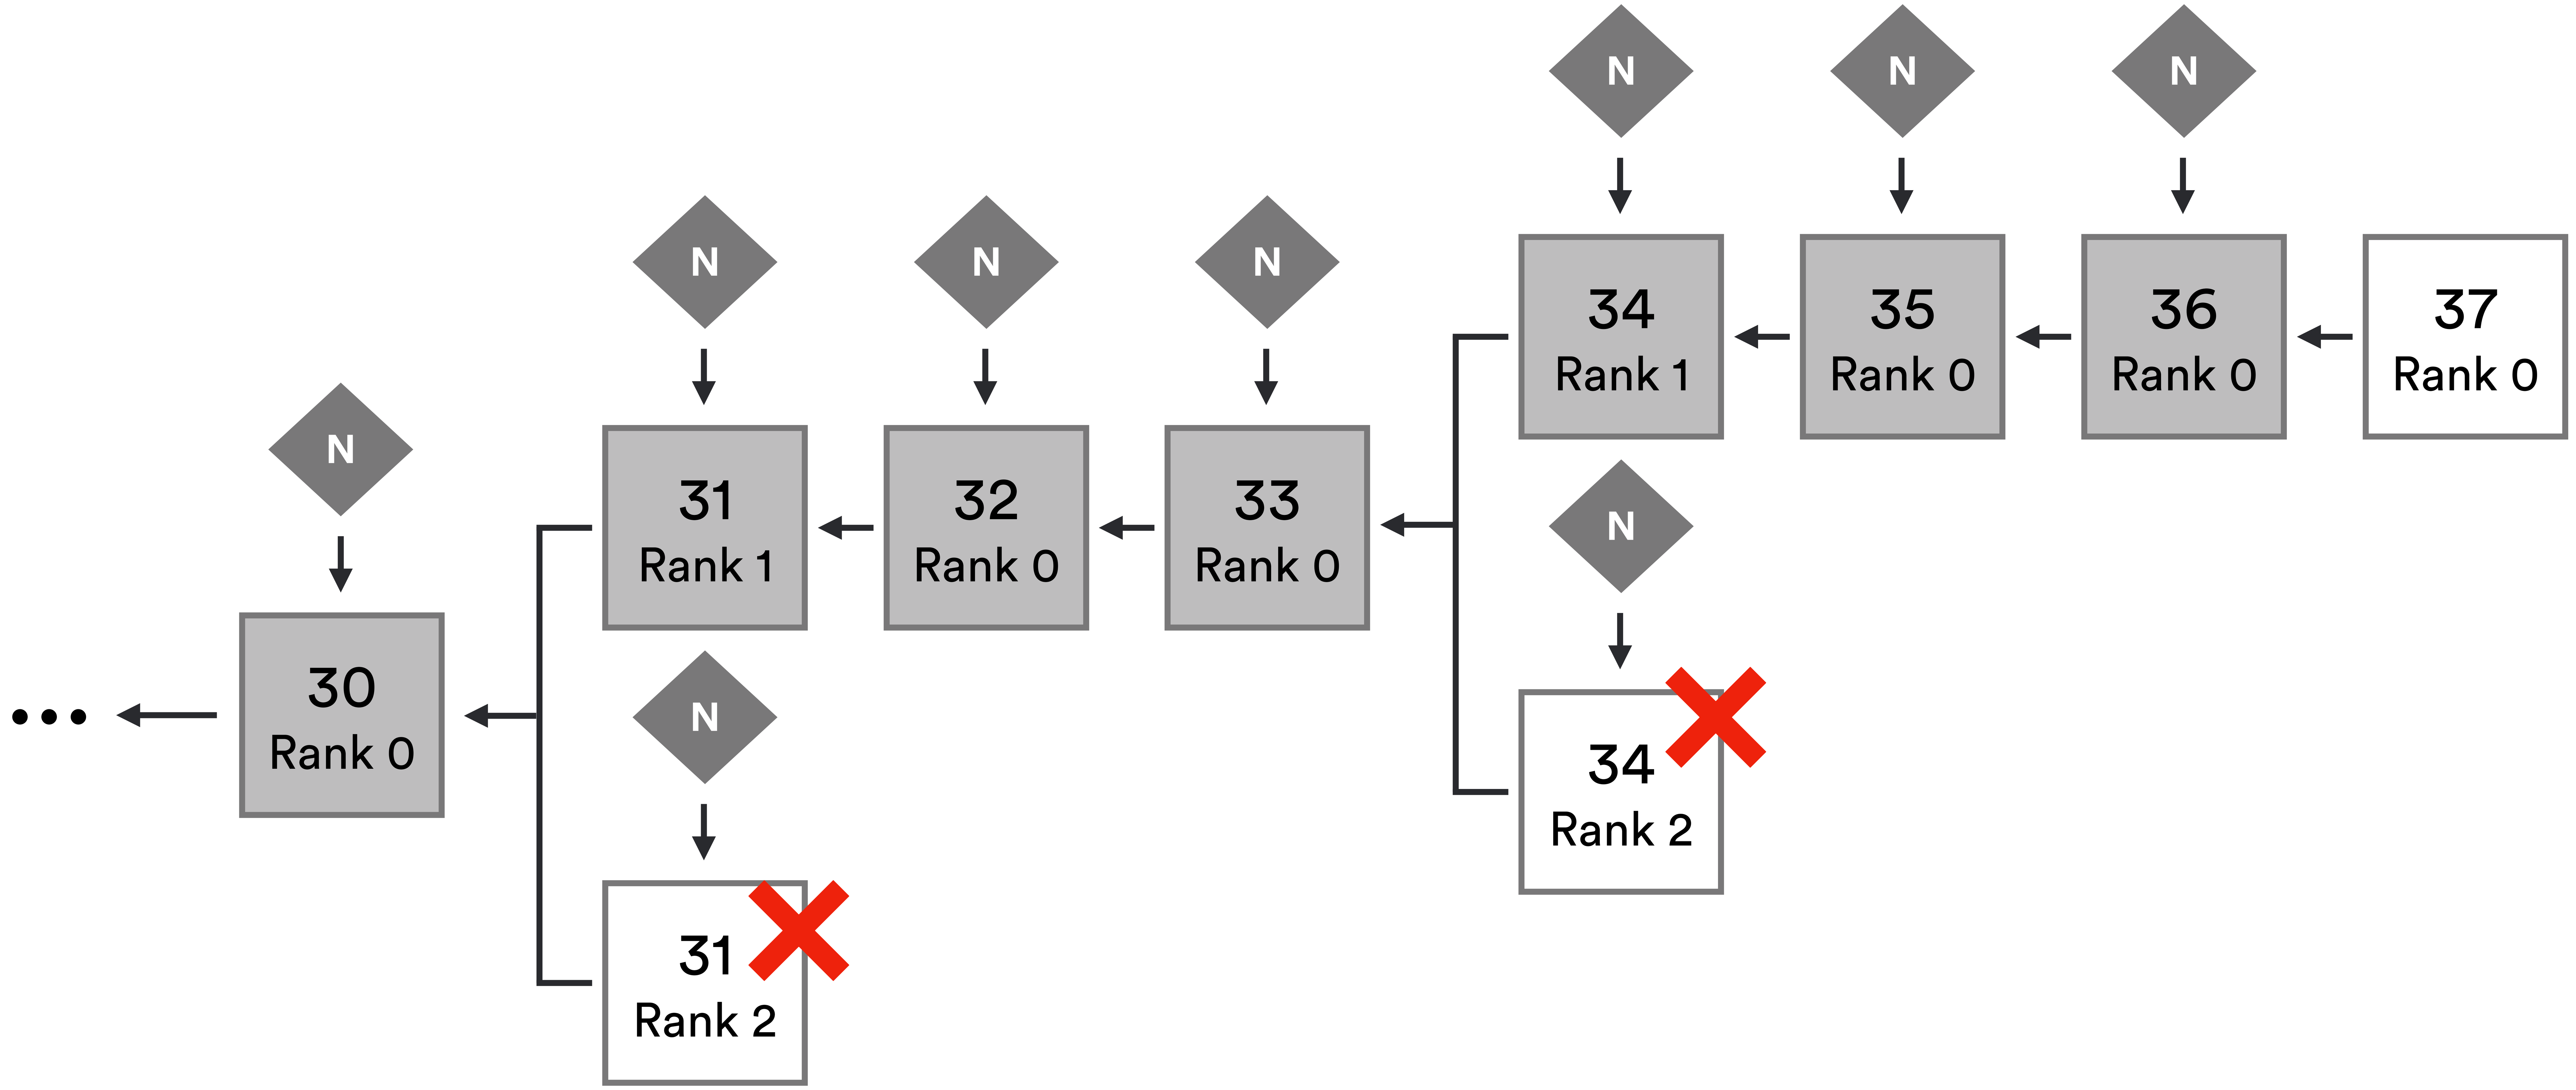

Notarization with Block Maker Ranking

Asynchronous communication → Forks cannot be removed!

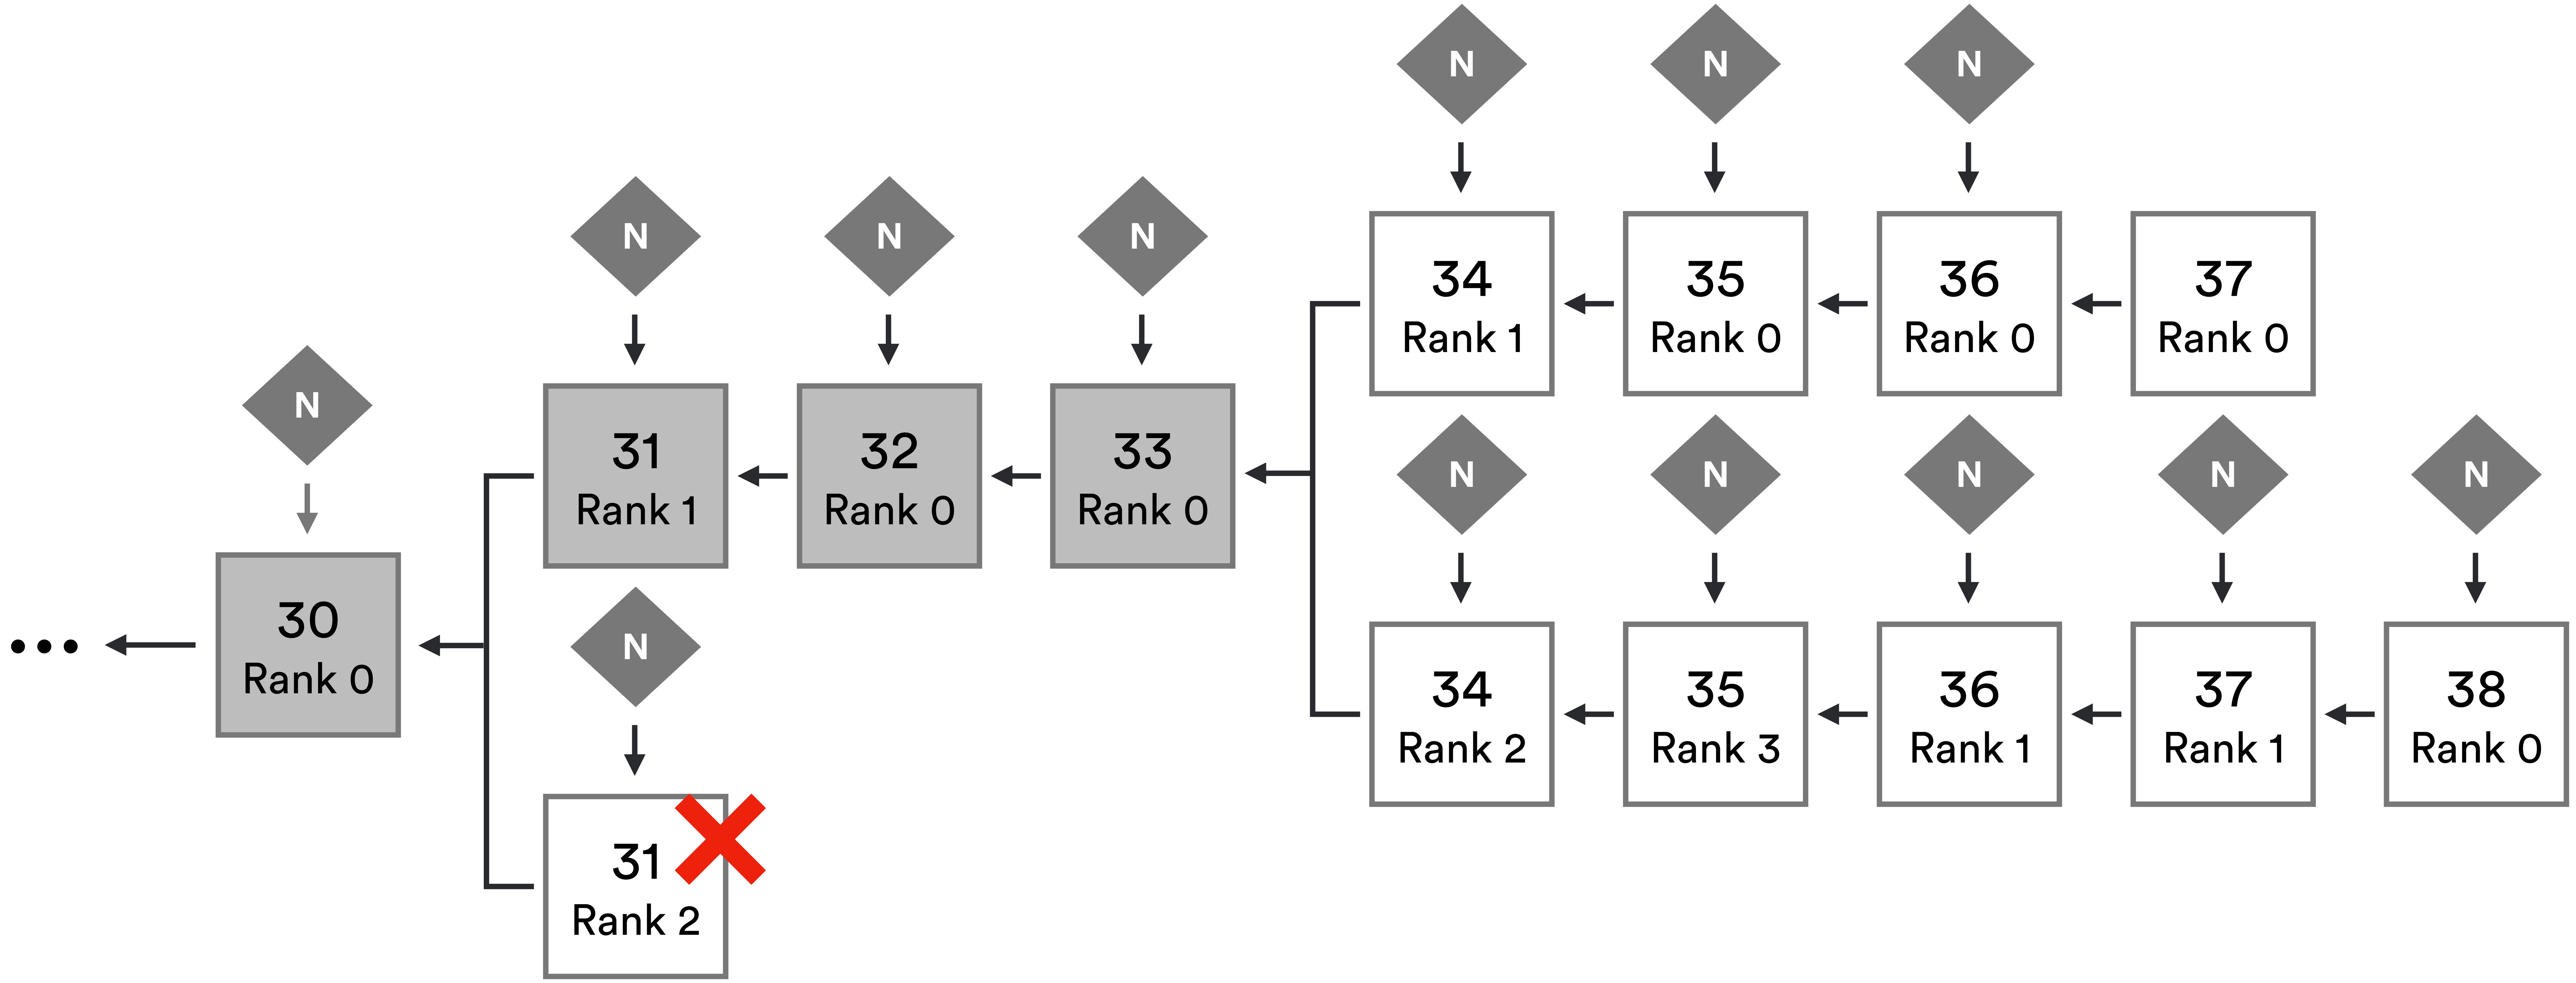

Finalization

Nodes create finalization shares if they did not notary-sign any other block at that height

Step 1

Node 1 notary-signs block *b* at height 30

Step 2

Node 1 observes that block *b* is fully notarized and will no longer notary-sign blocks at height ≤ 30

Step 3

Since node 1 did not notary-sign any other block than block b, it creates a finalization-share on *b*

Step 4

Nodes 2 and 4 also cast finalization shares on block *b*

Step 5

3 finalization-shares are sufficient approval: the shares are aggregated into a single full finalization

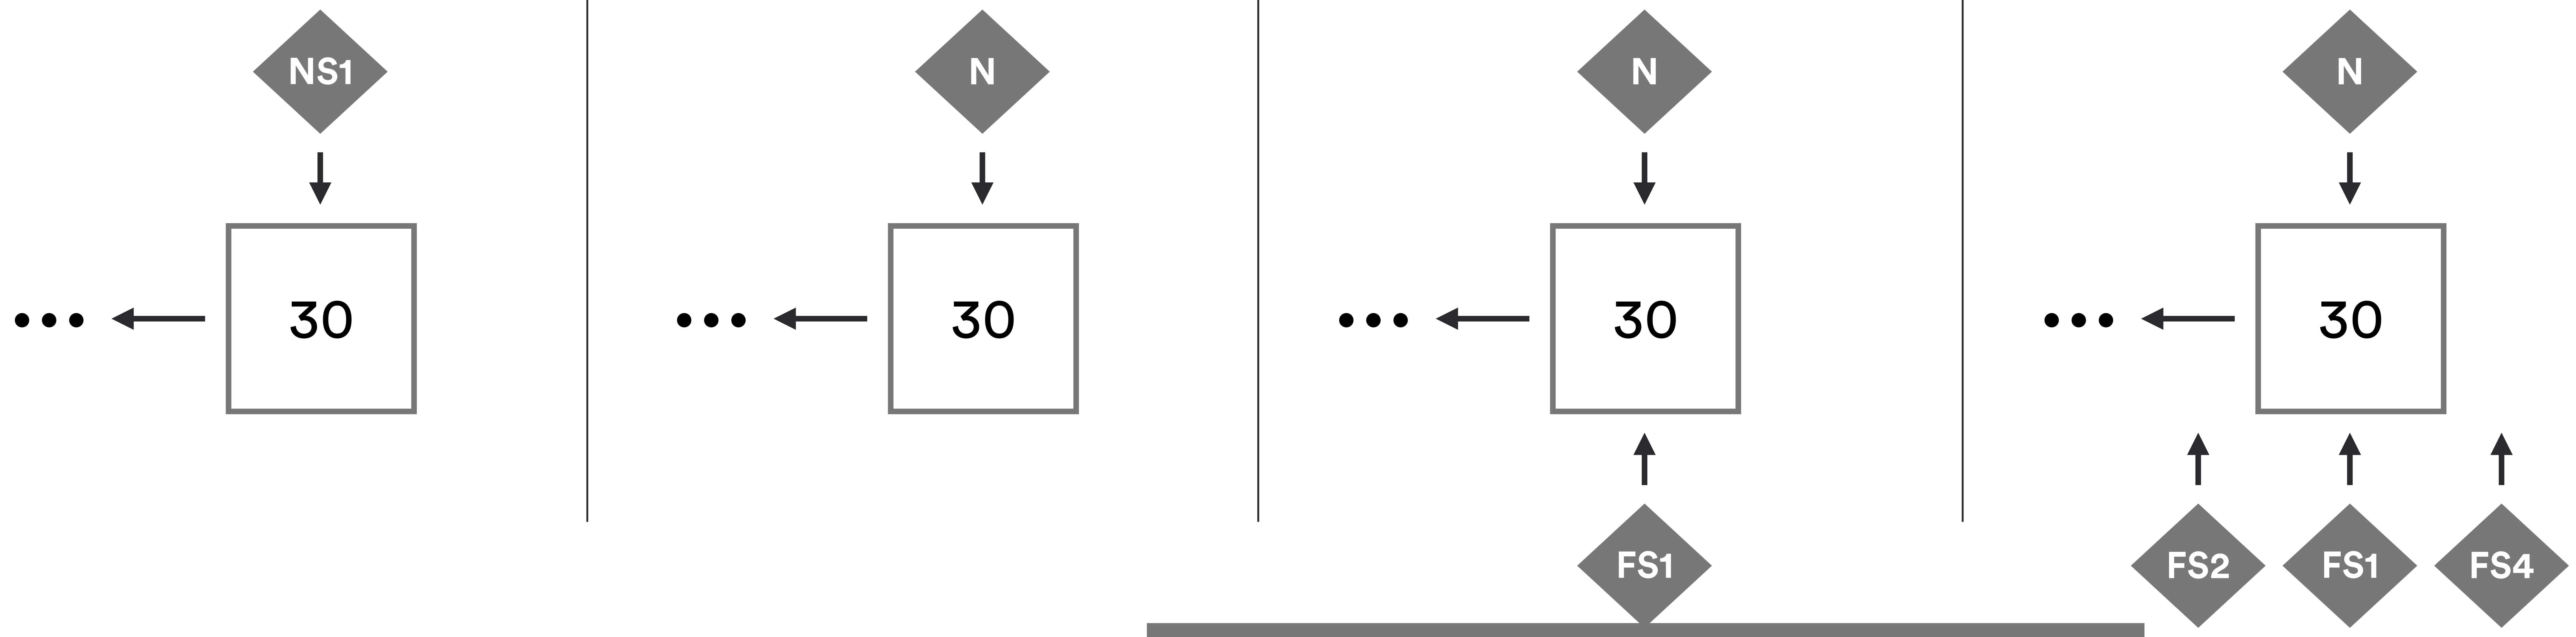

Node 1 did not notary-sign any height 30 block other than b

Finalization

Finalization on block *b* at height *h* = Proof that no other block is notarized at height *h*

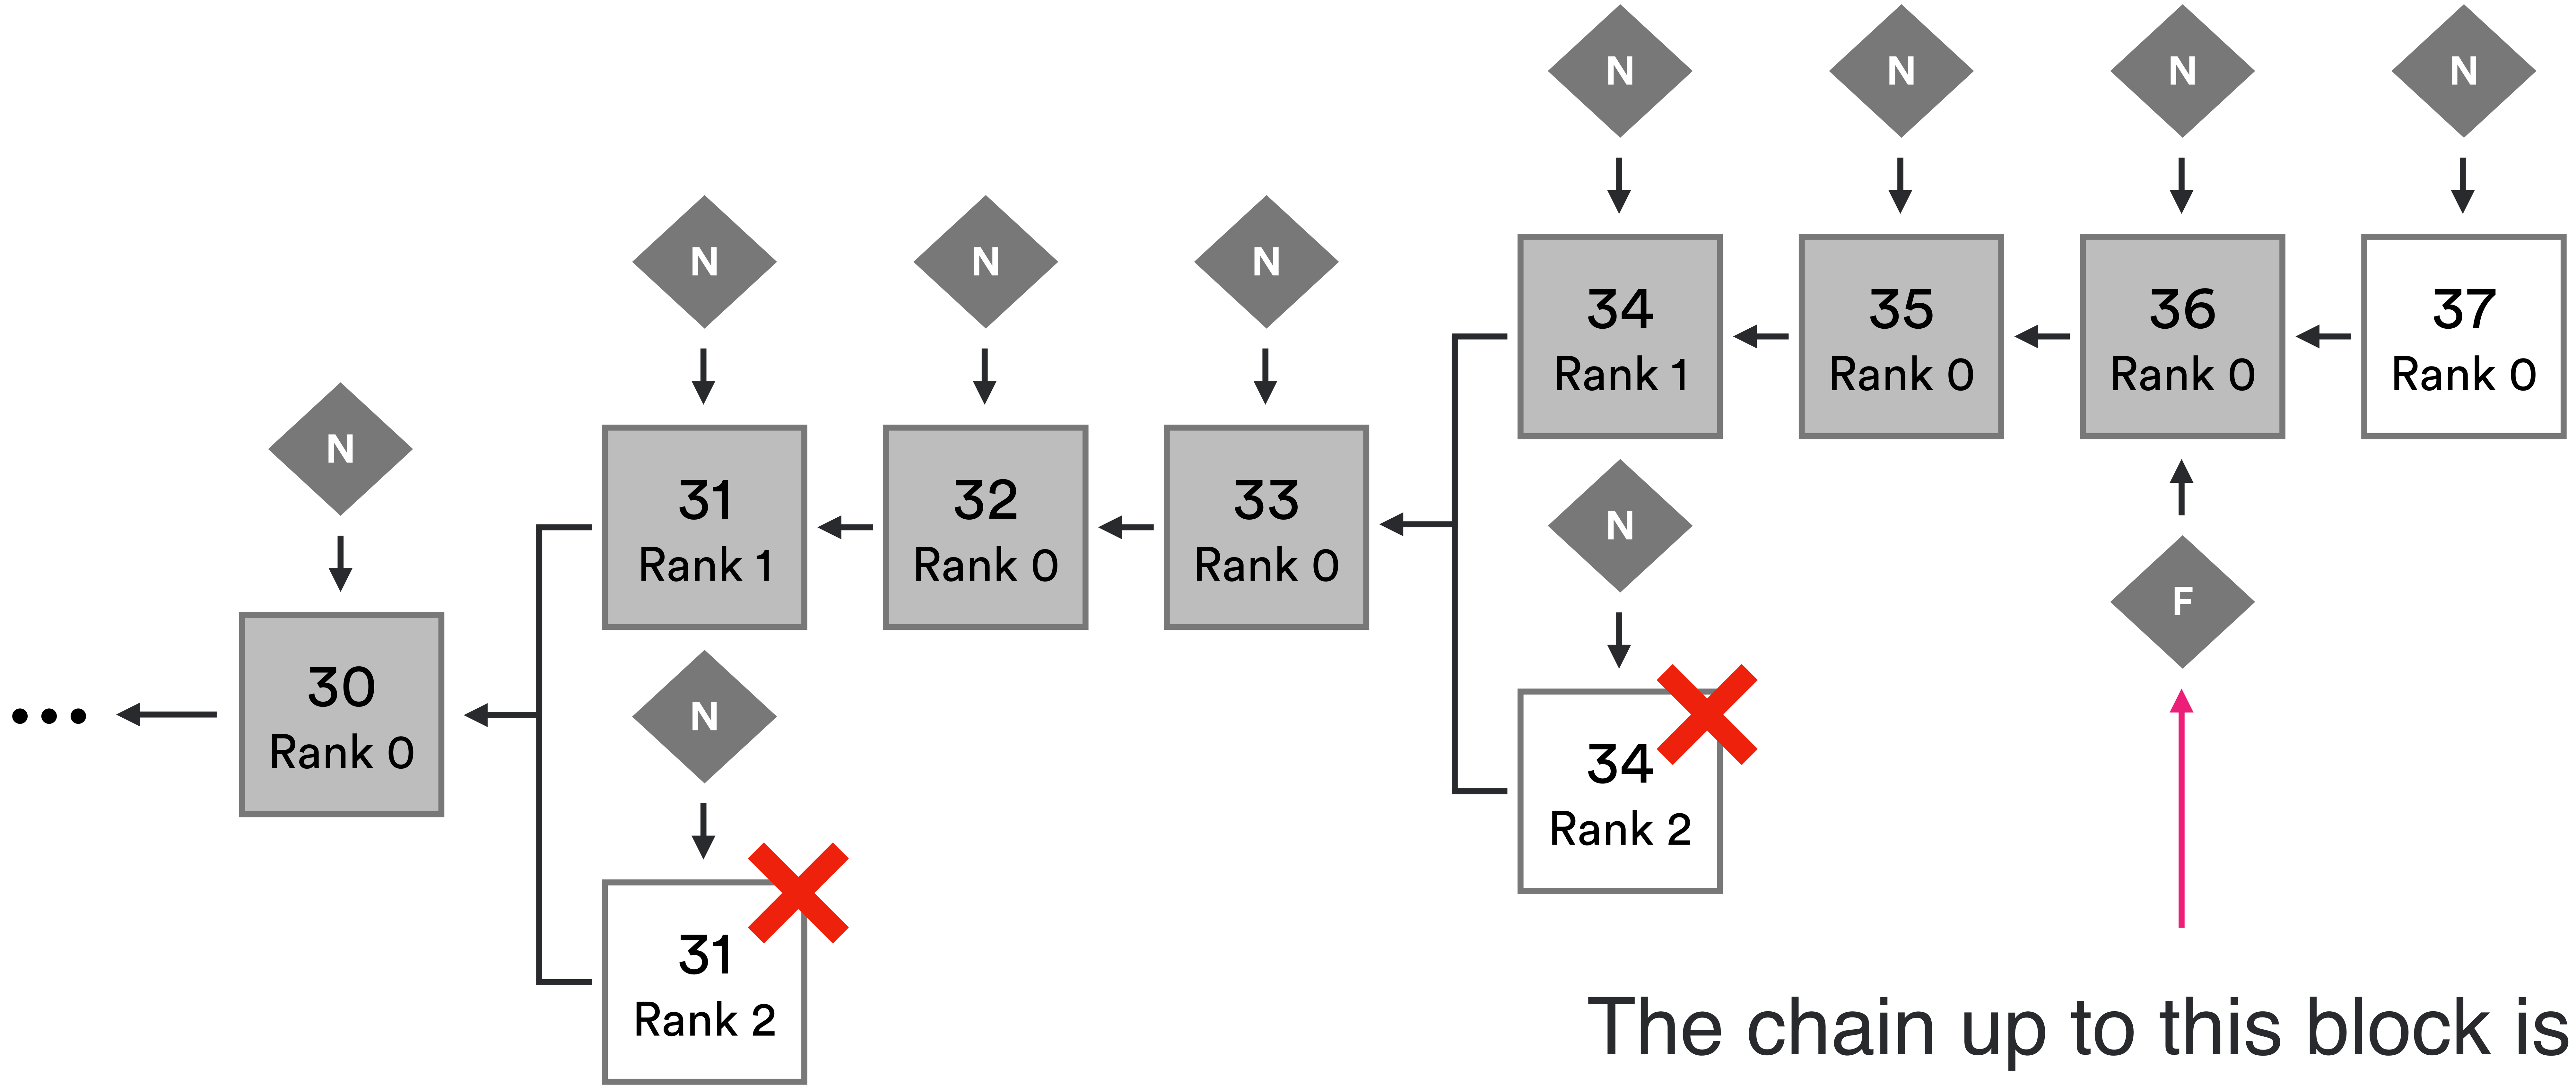

The chain up to this block is final

Algorithm Summary

1. Block tree building with notarization threshold signatures
   => at least one block per round

2. Random beacon from BLS threshold signature chain to rank block makers
   => reduce message and bit complexity

3. Recursive tree pruning with finalization threshold signatures
   => exactly one block per round

Safety

If block *b* at height *h* is finalized, then there is no finalized block $b' \neq b$ at height *h*.

Proof sketch:

1. A full finalization on *b* requires ***n-f*** nodes to finality-sign (by construction)
2. At least ***n-2f*** of the **n-f** nodes that finality-signed *b* must be honest (by assumption that ≤ ***f*** nodes are corrupt)
3. An honest node that finality-signed *b* did not notary-sign any other block at height *h* (by construction)
4. At least ***n-2f*** nodes did not notary-sign any height *h* block other than *b* (by 2. & 3.)
5. A full notarization requires ***n-f*** notarization-shares (by construction)
6. The ***n-(n-2f)*** **<** ***n-f*** remaining nodes that may have notary-signed a block *b'* are not sufficient to reach the notarization threshold of ***n-f*** (by 4. & 5.)

Liveness

The communication network is δ -synchronous at time T if all messages sent by honest nodes by time T are delivered by $T+\delta$

Assume that:

(i) $k > 1$, the first honest node P to enter round k does so at time T

(ii) Node Q with rank 0 in round k is honest;

(iii) the communication network is δ -synchronous at times T and $T+\delta$;

(iv) slot 0 lasts at least 2δ.

Then when all round-k messages from honest nodes have been delivered to all honest nodes, each honest node will have Q's round-k proposed block as a finalized block.

Measurements

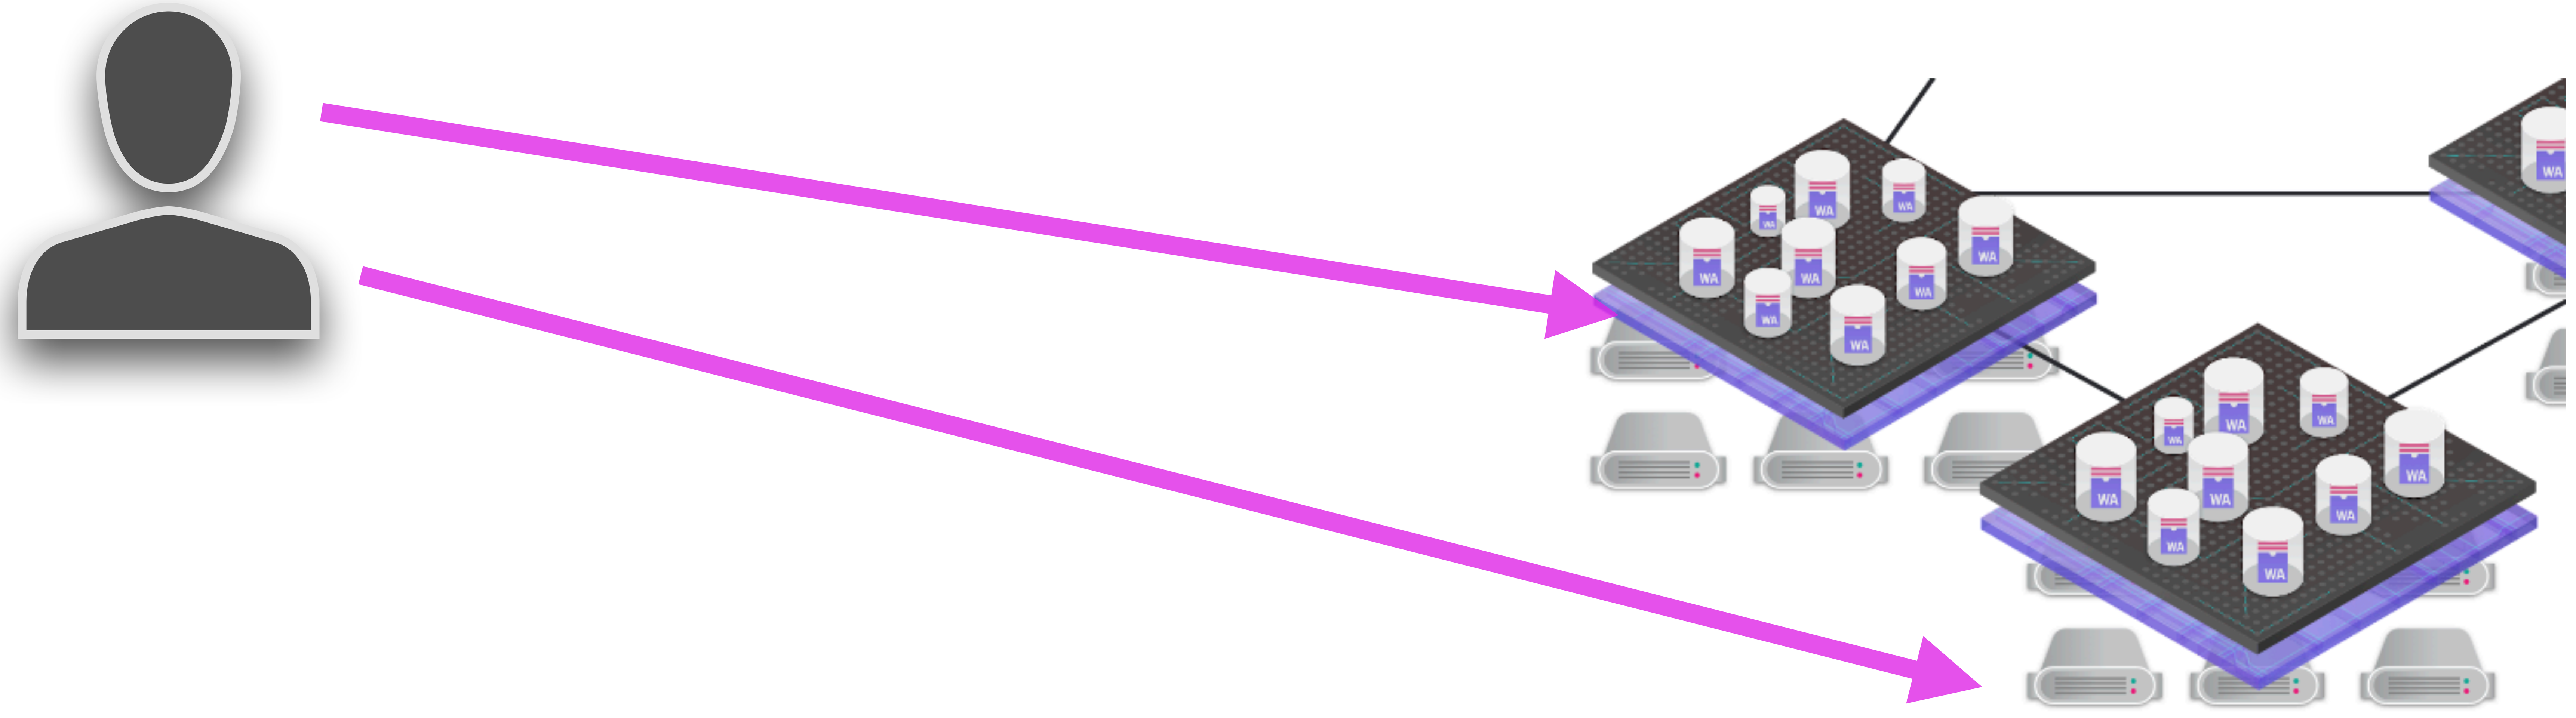

	without load	with load	with load and node failures
13 node subnet	1.09 blocks/s 1.64 Mb/s	1.10 blocks/s 4.72 Mb/s	0.45 blocks/s 4.39 Mb/s
40 node subnet	0.41 blocks/s 4.63 Mb/s	0.41 blocks/s 7.32 Mb/s	0.16 blocks/s 5.06 Mb/s

Average block rate and sent traffic.

Wanna know more?

- Full version with proofs link
    - Includes protocol variants + analysis for message complexity, latency, …

- Internet Computer Wiki link
- Technical Library: here (videos of talks) and here (blogposts)

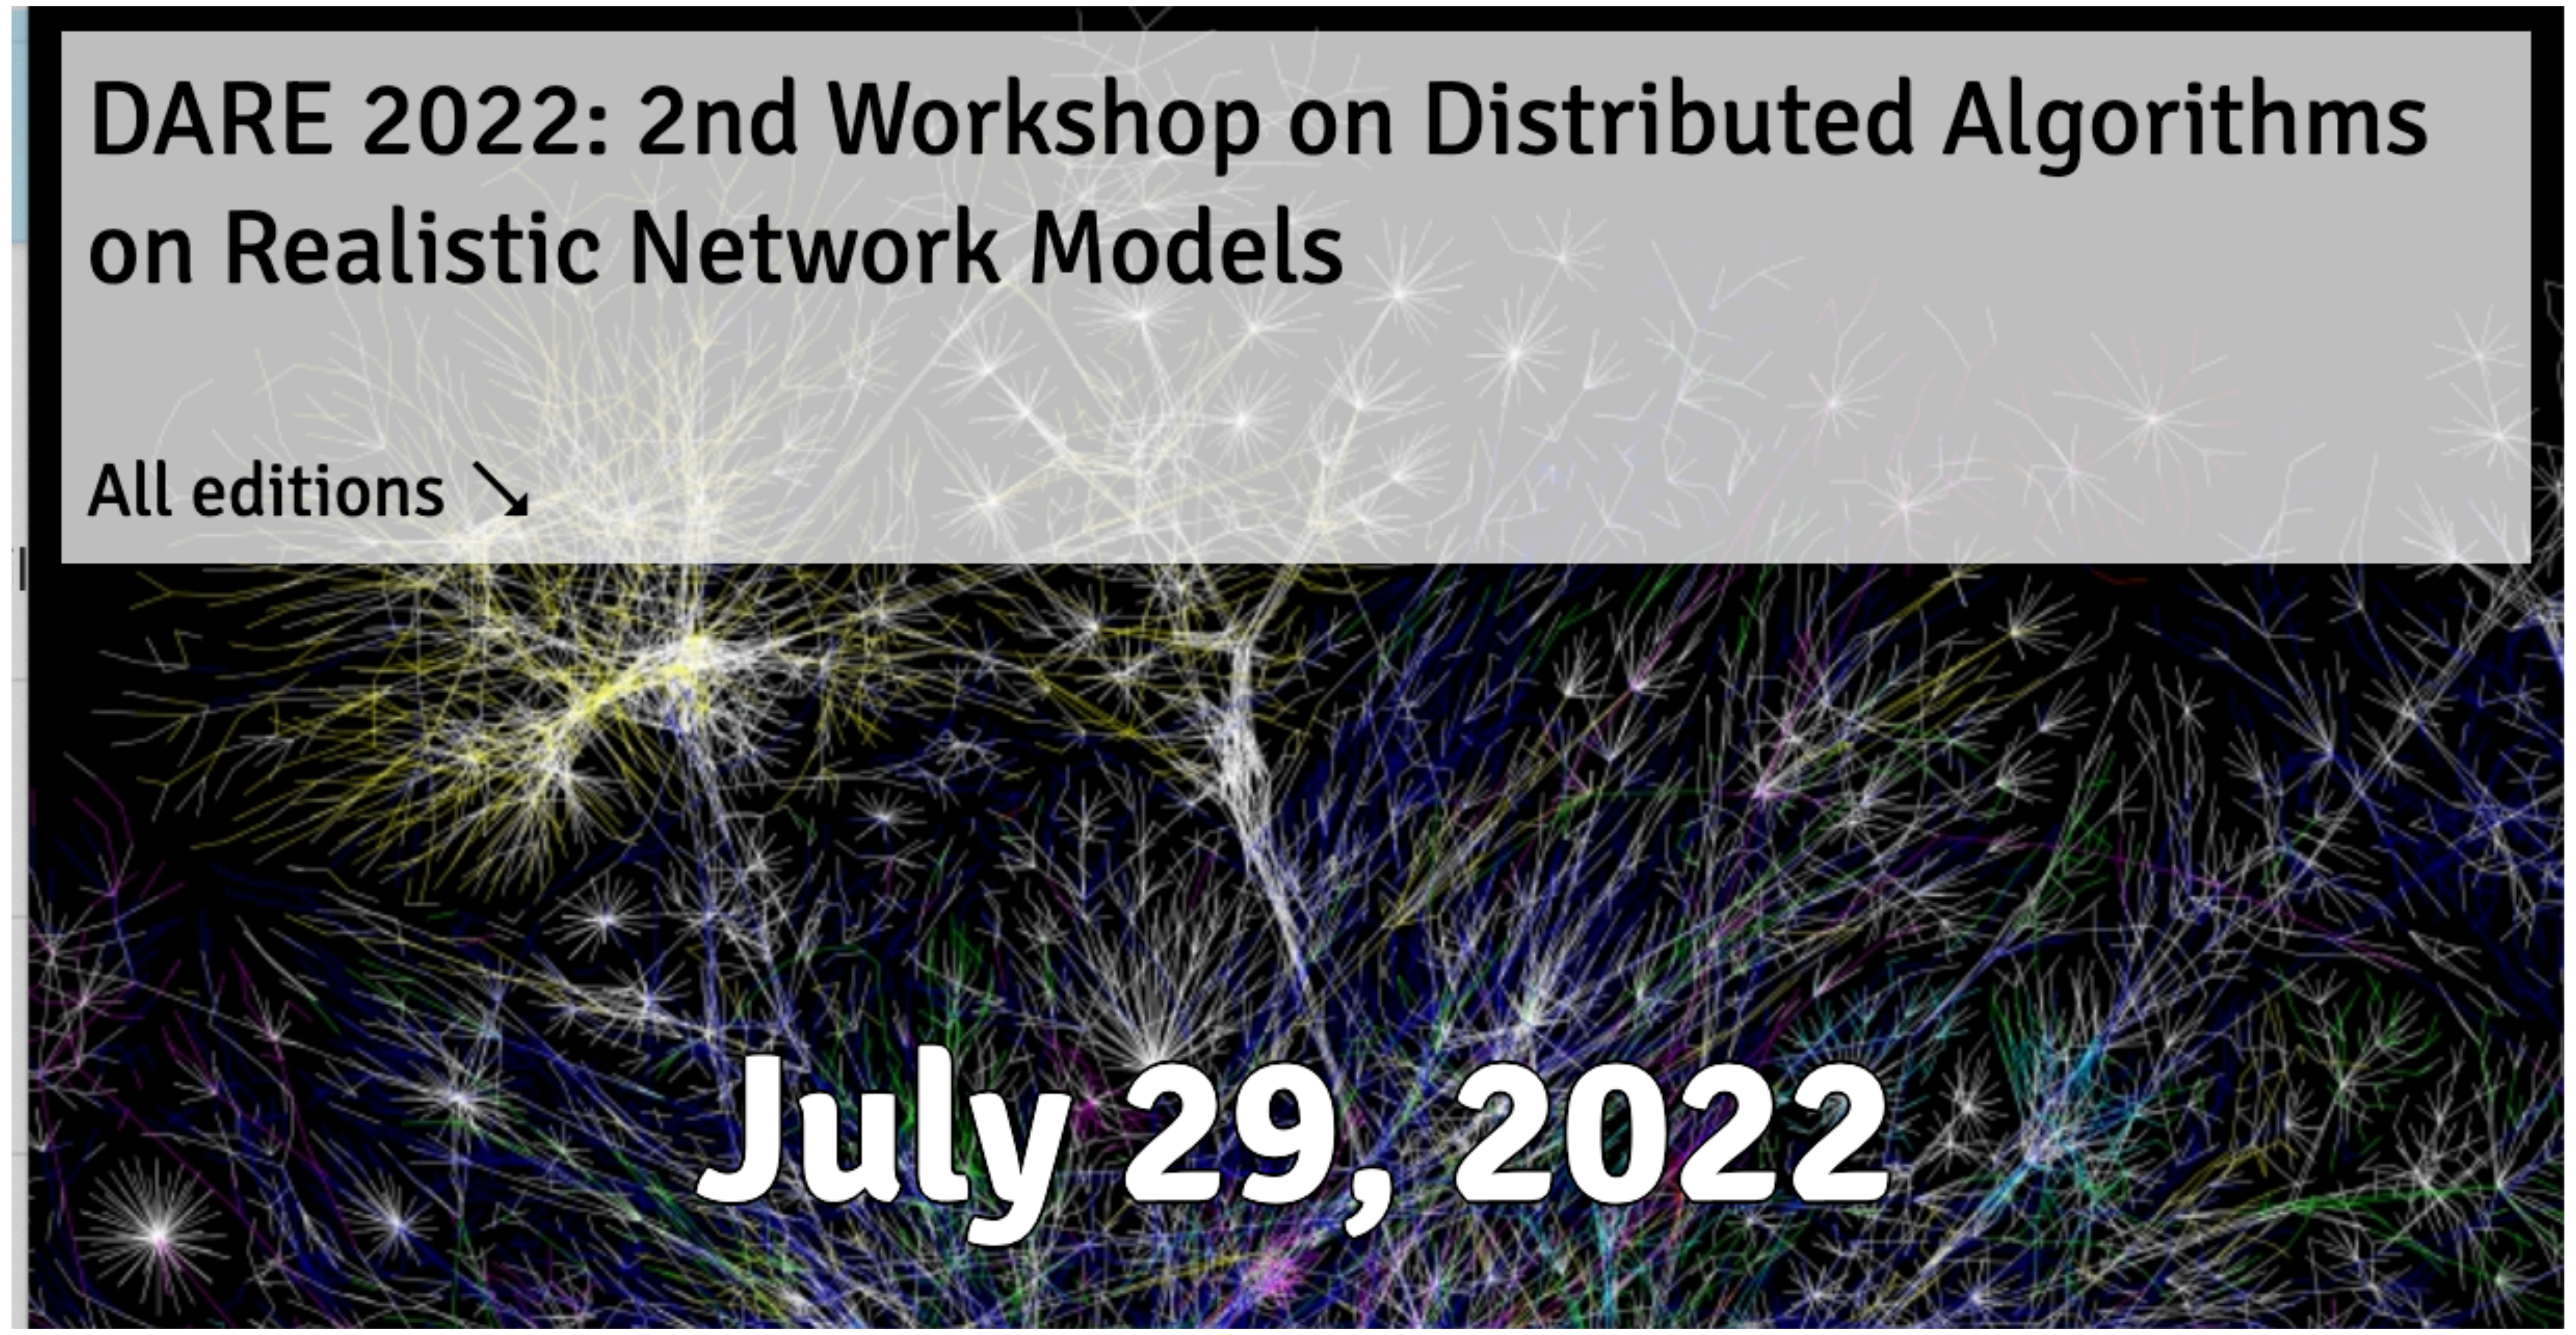

Friday 10:30 - 11:15

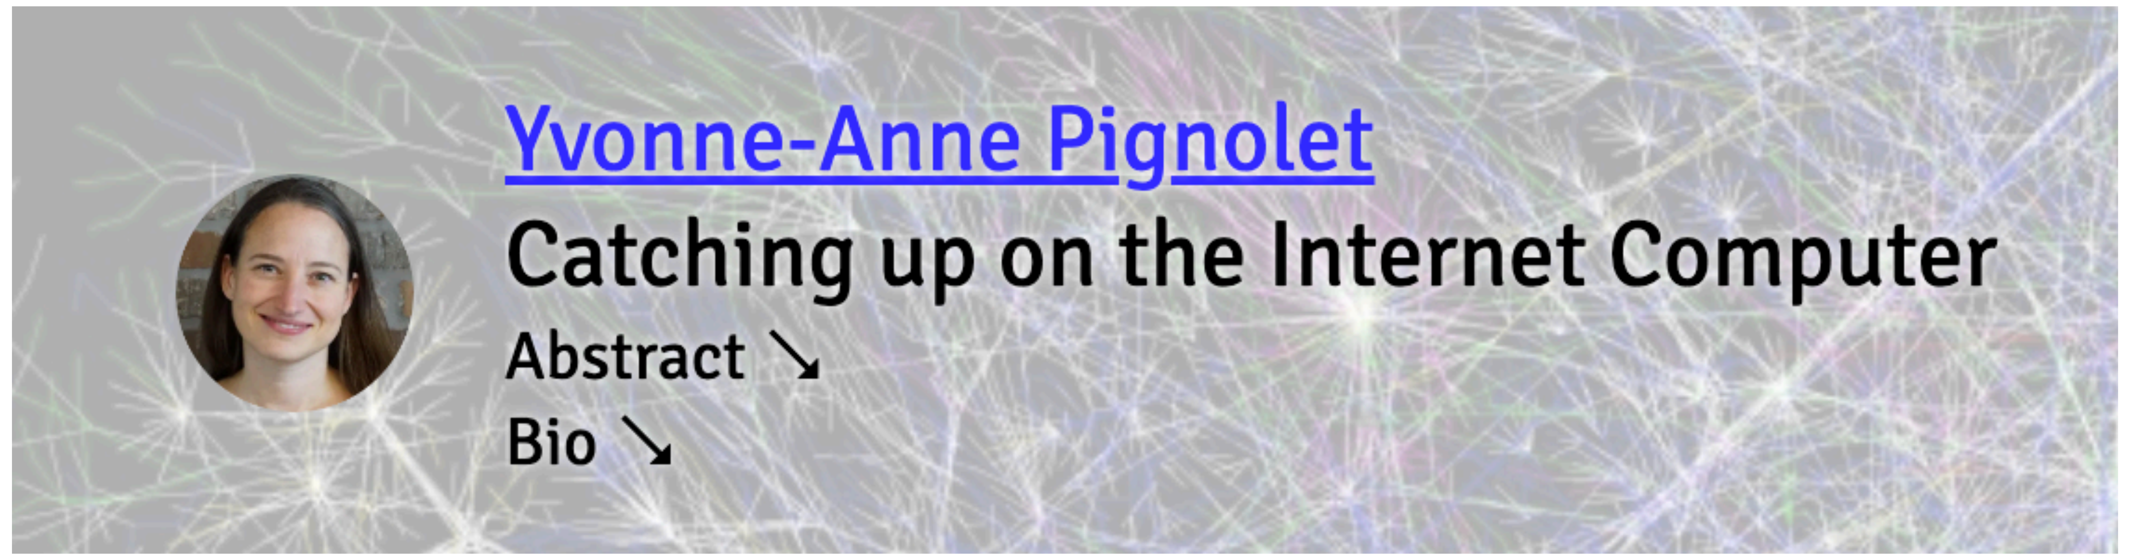

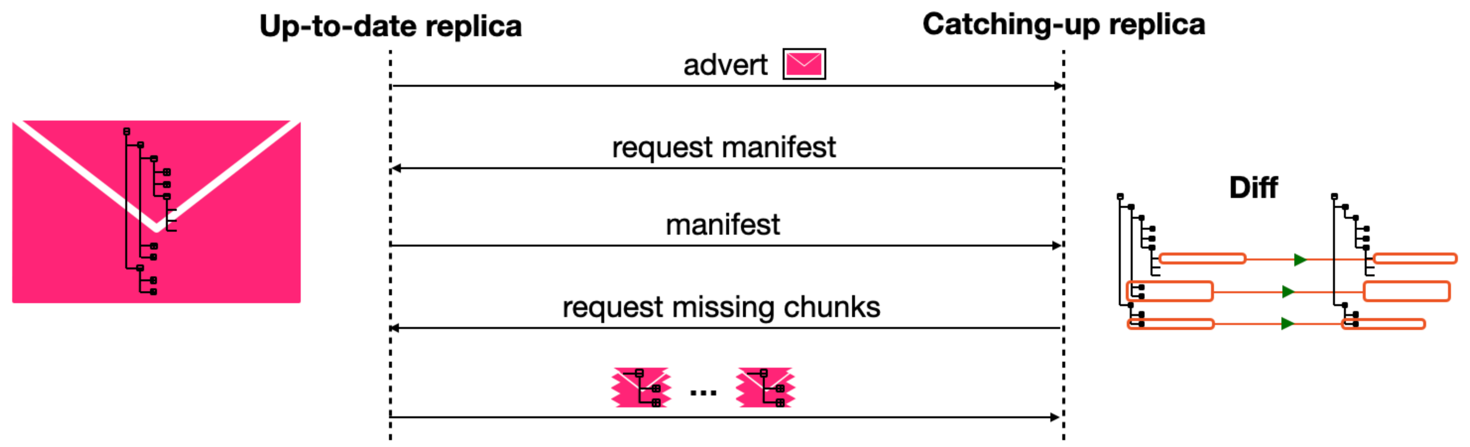

DFINITY

yvonneanne@dfinity.org

We are hiring!
www.dfinity.org/careers