

ANODE / H002

BETTER SYSTEMS FOR COLLECTIVE ACTION

THE DAO DELEGATION HANDBOOK

How Delegated Governance Power Is Allocated, Exercised and Controlled

ONE RIGHT

MANY INTERMEDIARIES

DELEGATE

DREP

VALIDATOR

COUNCIL

STEWARD

CUSTODIAN

AGENT

A handbook for delegated discretion

Research and writing: Enrique Alles Fernandez

V2 / September 2026

Research current through: 25 August 2026 unless otherwise noted.

Published by: Anode Research Studio

WHAT THIS HANDBOOK STUDIES

Delegation is usually described as a way to move voting power from a tokenholder to a delegate. That description is technically correct but institutionally incomplete. What moves is not only a vote: it is some amount of discretion - the ability to interpret information, choose among alternatives and act without asking the principal for instructions each time.

Once framed that way, several questions become unavoidable. What problem requires an intermediary? What authority is actually being transferred? Who receives it, how much judgment can they exercise, what limits apply, how is their behavior observed, and what happens when the arrangement stops working?

The handbook follows those questions across token delegation, treasury allocation, DReps, expert bodies, stewards, institutional custody workflows and early agentic systems.

HOW TO READ IT

The chapters move from diagnosis to control. First: why delegate at all, and what object of authority is being delegated? Next: how power is allocated, what mandate guides the agent and what incentives shape behavior. Then: how accountability, revocation, concentration and lifecycle should be evaluated.

The later chapters widen the lens to specialization, professional governance intermediaries, institutional holders and AI-assisted delegation before ending with a practical architecture canvas.

Cases are used as evidence of mechanisms, trade-offs and redesigns - not as proof that one governance model should be copied universally.

Central question

How can a principal reduce the cost of collective decision-making without losing control of the authority it owns?

The unit of analysis is the authority relationship

Throughout the handbook, the label attached to an actor is secondary. A "delegate" may hold broad political discretion; a steward may hold narrow operational discretion; a council may have a veto; a custodian may only provide an execution path. The useful comparison is therefore source -> scope -> discretion -> execution -> revocation.

THE CENTRAL CLAIM

DELEGATION IS NOT THE TRANSFER OF A VOTE.

IT IS THE TRANSFER OF

discretion.

The governance problem begins after authority leaves the principal.

Contents

Executive Summary	6
1. Why delegate at all?	9
2. What exactly is being delegated?	13
3. Who gets the power?	16
4. What does a delegate owe?	21
5. Who pays for governance?	24
6. Can power actually be taken back?	29
7. Does delegation decentralize power?	33
8. The delegate lifecycle	37
9. When delegates stop being the answer	41
10. The governance intermediary	45
11. Institutional delegation	48
12. Delegation after AI	52
13. When not to delegate	55
14. The Delegation Architecture Canvas	57
Worked example - applying the canvas	59
Appendix A - Comparative system map	63
Appendix B - Failure and redesign registry	66
Appendix C - Method and source notes	67


How to use it

Protocol designers can start with Chapters 1-3 and 13-14; delegates with Chapters 4-8; institutional holders with Chapters 10-11; researchers with Appendices A-C. The worked example on pages 59-61 shows how the full framework fits together.

HOW TO READ THIS HANDBOOK

Read the system, not the title

A protocol can call someone a *delegate* while giving them almost no discretion; another can avoid the word entirely while granting a steward the ability to change live risk parameters within hard-coded bounds. The label is secondary.



This five-layer reading is used throughout the handbook. Two systems can both support delegation and still produce opposite political contracts.

FACT

Documented in a contract, proposal, official documentation, program report or directly observable governance record.

OBSERVATION

A pattern or calculation that follows from available data but does not establish causality.

INTERPRETATION

A research inference supported by multiple observations. It is presented as analysis, not fact.

SPECULATION

A forward-looking possibility. Used sparingly and labeled explicitly.

Baseline technical references: OpenZeppelin Votes/Governor ^[1] and Snapshot delegation ^[2].

Six terms used throughout

PRINCIPAL

The holder, DAO, institution or body that owns or creates the authority.

AGENT

The person, organization, council, contract or software system that receives authority.

DISCRETION

Freedom to interpret information and choose an action without fresh instructions each time.

MANDATE

The policy, constituency, purpose or rule that guides the agent's discretion.

REVOCATION

The path by which authority can be withdrawn, narrowed, replaced or allowed to expire.

CONTESTABILITY

Whether a capable alternative can realistically replace the incumbent in practice.

EXECUTIVE SUMMARY

Delegation became an institutional layer

The 2026 delegation landscape is no longer well described by the simple relation **tokenholder → delegate → vote**.

The principal may be a retail holder, fund, DAO treasury, foundation, staker or protocol body. The agent may be a professional delegate, DRep, validator, expert council, steward, custody workflow or software agent. The delegated object may be a general vote, one referendum track, a veto, or a narrow operational function.

Delegation solves coordination and information costs by moving them.

It does not make those costs disappear. The new question is whether the intermediary layer is legible, contestable and controllable.

2024

Public delegate platforms and early compensation pilots.

2025

Treasury delegation, professionalization, scorecards and performance programs.

2026

Delegator incentives, optimistic approvals, domain specialization, bounded stewards, institution-ready custody workflows and agentic experimentation.

NEXT

Less emphasis on maximum participation; more emphasis on where discretion should sit and how it can be withdrawn.

Twelve findings that frame the problem

The rest of the handbook develops these observations into a design framework. They are starting points for analysis, not universal laws.

01	Delegation covers voting, representation and execution - distinct governance primitives.	02	Participation is not always the objective; quorum security, anti-capture and specialization increasingly matter.
03	ENS moved incentives toward the principal: 90% of its 2026 pilot pool is for delegators, 10% for active delegates. [10][11]	04	Optimism Season 8 showed that governance can remain operational with fewer explicit interactions when optimistic approvals preserve oversight capacity. [3]
05	Uniswap recalled ~12.5M UNI of bootstrap treasury delegation after the original quorum rationale weakened. [8][9]	06	ZKsync increased active delegation by ~205M ZK in its staking pilot while also surfacing a concentration problem in delegate discovery. [13][14]
07	Aave separates unpaid delegates from paid contributors/service providers while using bounded stewards for operational authority. [4]	08	Public rationales improve observability; they do not by themselves create accountability or constraint.
09	Technical redelegation can coexist with behavioral inertia, custody friction and low principal attention.	10	Professional delegation creates economies of scale and a new cross-DAO conflict surface.
11	Institutional custody increasingly permits voting/delegation without moving the underlying assets. [26][28][29]	12	AI is making the artifacts many programs reward cheaper to produce; judgment is becoming harder to infer from visible output.

Taken together, these findings shift the question from "How do we get more people to vote?" to a harder one: where should discretion sit, what should constrain it, and how can authority move again when the arrangement no longer serves its purpose?

SIX DISTINCTIONS

DELEGATION

≠

PARTICIPATION

DELEGATION

≠

REPRESENTATION

REVOCABILITY

≠

CONTESTABILITY

TRANSPARENCY

≠

ACCOUNTABILITY

PROFESSIONALIZATION

≠

INDEPENDENCE

MORE DELEGATION

≠

**MORE
DECENTRALIZATION**

Why the distinctions matter

LEFT SIDE OF THE EQUATION

Delegation != participation: moving power does not guarantee that it will be used well or at all.

Delegation != representation: a technical proxy can exist without a clear political mandate.

Revocability != contestability: removing an agent is weak protection when there is no credible replacement.

RIGHT SIDE OF THE EQUATION

Transparency != accountability: visibility matters only when somebody can evaluate behavior and change the state of the relationship.

Professionalization != independence: paid expertise can improve continuity while also creating conflicts and dependency.

More delegation != more decentralization: power can become easier to exercise while concentrating in fewer intermediaries.

A good delegation system has to survive all six distinctions.

THE PRIOR QUESTION

1. Why delegate at all?

Before choosing delegates, define the governance problem that delegation is supposed to solve.

Delegation is attractive because governance is costly. Reading proposals, understanding technical trade-offs, maintaining signing infrastructure, coordinating internally and tracking execution all consume attention. But the same word - delegation - is used to solve very different problems.

ATTENTION

Principals cannot rationally monitor every decision. Delegation compresses information and decision costs.

EXPERTISE

Some decisions require technical, economic, legal or domain-specific knowledge that is unevenly distributed.

QUORUM / ACTIVATION

Idle voting power may make decisions fragile or capture cheaper. ENS explicitly frames active delegation as a security objective. ^[10]

OPERATIONS

Recurring parameter changes may be too frequent for full governance. Bounded stewards can move execution closer to operators while preserving limits. ^{[4][6]}

The design mistake: treating "low voting participation" as the diagnosis rather than a symptom. A system with few votes may be broken - or it may have deliberately moved routine decisions to optimistic or specialized processes.

THE TRADE

Delegation reduces some costs by creating new ones.

WHAT MOVES AWAY FROM THE PRINCIPAL

Continuous proposal reading, routine coordination and some domain analysis can move to an agent. The principal buys attention compression and, in some designs, specialist judgment.

WHAT MOVES BACK TO THE PRINCIPAL

The principal still has to choose the agent, define the mandate, monitor conflicts and decide when to replace it. Delegation relocates governance work; it does not eliminate responsibility.

Delegation is a cost transfer

The principal saves attention by creating an agency relationship. The question is whether the new relationship costs less than the problem it replaces.

NET VALUE OF DELEGATION

attention saved + expertise gained + latency reduced - agency cost - monitoring cost - concentration cost

COSTS IT CAN REDUCE

Proposal reading and synthesis.
Internal coordination and signing overhead.
Technical or domain expertise gaps.
Decision latency when routine issues recur.

COSTS IT CREATES

Selecting an agent and defining the mandate.
Monitoring conflicts, inactivity and judgment.
Replacing the agent when trust changes.
Concentration and dependency on intermediaries.

Low turnout by itself is therefore a weak diagnosis. If voting is low because principals are overloaded, a representative may help. If voting is low because most actions are routine, the better answer may be optimistic approval, automation or narrower expert authority. Delegation is valuable only when it solves a specific cost or risk more efficiently than the alternatives.

Pre-delegation test

Name the burden being reduced, the new risk being introduced, the actor who will monitor it and the concrete path to reverse the decision. If those answers are vague, the architecture is not ready.

DIAGNOSIS BEFORE DELEGATION

The same symptom can require opposite designs

Observed problem	Possible cause	Delegation response	Alternative response
Low turnout	Information overload	General delegate	Triage / fewer votes
Quorum misses	Inactive voting power	Delegator incentives / treasury bootstrap	Quorum redesign
Technical review bottleneck	Generalists lack expertise	Domain delegate / expert council	Optimistic approval + veto
Slow parameter updates	Governance latency	Bounded steward authority	Hard-coded automation
Institution cannot sign	Custody / authorization workflow	External delegate	Custodian-enabled internal vote
Forum quiet	Little material work	None	Do not manufacture activity

Design principle: do not create a representative layer simply because another mature DAO has one. Delegate only after identifying the specific cost or risk the intermediary is meant to reduce.

How to use this diagnostic

Start with the observed symptom, then ask what actually causes it. Low turnout caused by information overload is a different problem from low turnout caused by a quiet governance calendar. The table is meant to prevent a familiar failure mode: selecting a delegate architecture before diagnosing the cost, risk or bottleneck it is supposed to solve.

CURRENT / 2026

Case: Optimism - less voting can still mean more capacity

CURRENT EVIDENCE

Optimism Season 8

-45%

OVERALL INTERACTIONS

Season 8 vs Season 7

7 / 23

OPTIMISTIC APPROVALS

30.4% of proposals

6

LARGE DELEGATES

needed to equal quorum at mid-season

Season 8 recorded materially fewer votes, rationales and forum interactions than Season 7. Yet the report argues that governance capacity remained stable because more decisions used optimistic approvals: routine actions could progress unless sufficient oversight objected. ^[3]

This is a useful counterexample to activity-maximization. A delegate who is required to vote on every routine action may look more active while adding less marginal security than a watcher who intervenes only when a proposal crosses a risk threshold.

Implication: measure participation capacity, escalation quality and veto readiness - not only the raw count of votes or forum posts.

Primary source: Optimism Token House participation and incentives, Season 8. ^[3]

What this case changes

Optimism is useful because it separates governance capacity from visible activity. If routine actions can advance by default while credible actors retain a veto or escalation path, then fewer interactions need not mean weaker control. The design question becomes whether the right actors are watching the right events and can intervene fast enough when something material changes.

THE OBJECT OF AUTHORITY

2. What exactly is being delegated?

A vote, a representative mandate and an execution permission are not the same thing.

A useful starting point is to stop classifying systems by the title of the agent and instead classify them by the authority object.

VOTING POWER

Cast a principal's weight on proposals.

REPRESENTATION

Exercise broad judgment on behalf of a principal or constituency.

DOMAIN AUTHORITY

Exercise judgment only within a specified topic or track.

VETO / OVERRIDE

Stop or challenge an otherwise optimistic action.

OPERATIONAL AUTHORITY

Change live parameters inside predefined boundaries.

EXECUTION AUTHORITY

Translate an approved decision into a transaction or operational act.


Each authority object creates a different risk. General voting delegation creates agency and representation risk. Operational delegation creates scope-creep and key-control risk. Veto delegation creates monitoring and availability risk. The same accountability toolkit cannot be applied blindly to all three.

A simple authority test

Ask what the agent can do without returning to the principal. If the answer is only "cast this pre-decided vote," discretion is narrow. If the agent may interpret proposals, balance competing interests and decide how to vote, the relationship is representative. If it can change parameters or execute transactions, the control problem becomes operational as well as political.

THE FIVE-LAYER CONTRACT

A common grammar for delegated power



ERC20Votes

Holder → address → general voting power → onchain Governor → redelegation. ^[1]

POLKADOT

Holder → expert → one referendum track → OpenGov → undelegate / redelegate. ^[15]

AAVE STEWARD

DAO → bounded operator → specific parameter range → steward transaction → DAO revocation / hard limits. ^{[4][6]}

RESERVE OPTIMISTIC

Vote locker / allocator → delegate watcher → veto-oriented oversight → optimistic governor → withdrawal / program change. ^{[18][19]}

If a governance design cannot answer these five layers, its delegation contract is politically underspecified even if the smart contract is technically clear.

One example in plain English

An AAVE holder delegating voting power gives another address room to exercise judgment across governance proposals. An Aave steward is different: the DAO grants a narrower operational permission, such as changing specified parameters within hard limits. Both are delegation, but the first relies heavily on mandate and political exit; the second can be constrained much more directly in code.

CURRENT / 2026

Case: Aave - representation is not execution

CURRENT FRAMEWORK

Aave Governance Framework v2

Representative layer

Aave defines delegates as community members who receive voting power from others or self-delegate and vote on behalf of those who entrusted them with that power. The July 2026 Governance Framework v2 states plainly: **delegates are not compensated.** [4]

The earlier Orbit program had paid recognized delegates through H1 2026. [5] The change matters because it separates the political status of "delegate" from paid contributor or service-provider work.

Operational layer

Aave also uses stewards and service providers with narrow execution authority. Risk stewards can adjust specified parameters inside quantitative limits and cooldowns; governance can reassign or revoke roles. [4][6]

This is delegated discretion made partly machine-legible: scope and limits can be encoded rather than left entirely to social expectation.

Lesson: professional governance does not require one all-purpose paid representative. Representation, expertise and execution can be decomposed into separate roles.

ROLE DECOMPOSITION

The point is decomposition, not terminology.

Aave matters because it shows that representation and operational execution can be governed by different contracts. Broad voting discretion is controlled through delegation and political exit; steward authority can be constrained ex ante by function, magnitude, cooldown and revocation. When roles have different discretion, they should have different accountability systems.

From authority type to power allocation

Once representation and execution are separated, the next question is not yet how the agent behaves. It is how that agent obtained authority in the first place. Organic delegation, treasury allocation, incentives and default representation can produce the same onchain voting weight while carrying different legitimacy and dependency risks.

ALLOCATION IS GOVERNANCE

3. Who gets the power?

Delegation is not only about how an agent behaves. It is about how the agent acquired authority in the first place.

Onchain voting weight can look identical after allocation while carrying very different legitimacy. Ten million tokens delegated organically by thousands of holders are not institutionally equivalent to ten million allocated by a foundation, treasury or incentive interface.

ORGANIC

Individual principals choose the agent with their own voting power.

TREASURY / FOUNDATION

A central or collectively controlled allocator bootstraps a delegate with voting power it controls.

INCENTIVIZED

The principal receives an economic reason to delegate to an eligible agent.

DEFAULT / EMBEDDED

Representation follows another relationship, such as validator staking, unless the principal overrides it.

Key distinction: voting power is a quantity; legitimacy has provenance.

Same voting power, different story

Imagine two delegates with 10 million votes each. One received them from thousands of independent holders; the other received almost all of them from one treasury program. Their onchain weight is identical, but their political dependence is not. Provenance tells the principal whether authority reflects distributed choice, allocator selection, an incentive or a default relationship.

Selection is constitutional design

Allocation decides whose judgment gets amplified before the first vote is cast. Identical voting power can have very different political provenance.

ORGANIC

Independent principals choose the agent with their own power. Strongest signal of direct principal choice; vulnerable to low-information selection and inertia.

TREASURY / FOUNDATION

A DAO or institutional allocator bootstraps an agent. Useful for entry and quorum support; creates allocator dependence and legitimacy questions.

INCENTIVIZED

The principal is paid or otherwise rewarded to delegate to an eligible agent. Can activate idle power; may turn political capital into a yield-sensitive allocation.

DEFAULT / EMBEDDED

Representation follows another relationship, such as staking, unless the principal overrides it. Low attention burden; higher risk of stale or passive consent.

The selection layer includes more than a smart contract. Eligibility thresholds, ranking rules, featured lists and treasury allocation processes determine who gets seen and who can accumulate enough power to become credible. ZKsync's 2026 pilot is a clear example: the featured-delegate interface had to be adjusted after concentration emerged. [13][14]

Disclosure recommendation

POWER PROVENANCE

Four allocation mechanisms, four failure modes

Allocation	Strength	Primary risk	Useful metric
Organic delegation	Direct principal choice	Inertia / low-information selection	Redelegation rate; unique principals
Treasury / foundation	Breaks entrant chicken-and-egg	Allocator dependence	% VP from institutional allocator
Delegator incentive	Activates idle voting power	Yield-driven concentration	New active VP; concentration delta
Default representation	Low attention burden	Passive consent / stale delegate	Override rate; inactive represented VP

Discovery design sits upstream of all four. A featured list, ranking algorithm or eligibility threshold can materially influence where power lands. ZKsync had to adjust its featured delegate interface after concentration emerged during its 2026 staking pilot. ^{[13][14]}

For this reason, delegate discovery should be treated as part of constitutional design rather than as a neutral frontend concern.

What a principal should ask before selecting

Who made this agent visible? What threshold did it have to cross? How much of its power is organic? What would happen if a treasury program or incentive disappeared? And can a new entrant realistically compete? These questions turn discovery from a frontend concern into part of the governance architecture itself.

EXECUTED / JUNE 2026

Case: Uniswap - bootstrap, mature, unwind

EXECUTED 1 JUN 2026

Return 12.5M UNI to the timelock

2022-23

~12.5M UNI delegated from treasury to Foundation + active delegates while quorum and proposal-threshold access were fragile.

2026

DUNI-era turnout and organic delegation increased; proposal argued the original bootstrap rationale had weakened.

1 JUN 2026

45.2M votes FOR; proposal executed and ~12.5M UNI recalled to the governance timelock.

The important story is not that treasury delegation was "bad". It is that it was **purpose-bound**. The mechanism solved an early participation constraint, then became less defensible once the environment changed. The proposal also identified a structural misalignment: the Franchiser gave voting power without requiring matching economic exposure. ^{[7][8][9]}

Design pattern: bootstrap delegation should carry an explicit sunset or review condition tied to the problem it was created to solve.

SUNSET DESIGN

A bootstrap mechanism should define its own end.

A useful sunset does more than name a date. It identifies the condition being solved, the review point, who has authority to withdraw the power and what transition occurs afterward. Without those elements, a temporary intervention can become an incumbent institution simply because removing it requires more coordination than leaving it in place.

PILOT / SEASON 1 PAUSED

Case: ZKsync - discovery became a power allocator

PILOT / PAUSED

ZKnomics Staking Pilot Season 1

Season 1 linked staking rewards to delegation to an active delegate. The interface used an active/featured delegate filter to reduce discovery friction. The pilot attracted 355M ZK at peak and increased active delegation by ~205M ZK from announcement to season end. ^[13]

Top 3 pilot recipients

~59.1%



Top 6 pilot recipients

~76.5%



Administrators removed several very large delegates from the featured list mid-season and recommended a hard cap for future featured status. They also noted that growth in active delegation did not automatically produce proportional growth in proposal participation. ^{[13][14]}

Lesson: a delegation interface does not merely reveal preferences. It can shape them.

INTERFACE POWER

Discovery is part of the allocation mechanism.

Featured lists, ranking rules and eligibility filters reduce search costs for principals, but they also influence which agents accumulate additional power. That makes interface design constitutionally relevant. A neutral-looking discovery layer can amplify incumbents, channel incentive-driven delegation or make alternative agents effectively invisible.

The allocation lesson

The Uniswap and ZKsync cases point in opposite directions but answer the same question. Bootstrap power can be useful and incentives can activate dormant voting power; both can also become allocation mechanisms with their own concentration effects. That is why the next chapter turns from who receives power to what the recipient is actually expected to do with it.

THE MANDATE PROBLEM

4. What does a delegate owe?

The word "representative" hides a hard question: representative of whom, and for what?

DAO delegation rarely provides proposal-by-proposal instructions from each principal. The typical arrangement is closer to a revocable grant of discretion: the principal selects an agent and the agent applies its own judgment until authority is changed or withdrawn.

PRINCIPAL MANDATE

Vote in the interests or expressed preferences of the delegators who supplied the power.

ECOSYSTEM STEWARDSHIP

Exercise judgment for the health, security or long-term interests of the protocol/ecosystem.

CONSTITUENCY

Represent a defined stakeholder group, region, user class or professional perspective.

INDEPENDENT JUDGMENT

Apply a published policy or analytical framework without taking instructions from individual delegators.

A system should not pretend these mandates are interchangeable. Conflict starts exactly when the interests of a large delegator, the broader protocol and the delegate's own policy diverge.

CONFLICT TEST

Mandate collisions are normal, not exceptional.

The hard case is not when the delegate, delegator and protocol all want the same thing. It is when they diverge. A serious mandate should explain which interests guide the agent, what happens when a major delegator disagrees, when conflicts require abstention or disclosure and when the right response is to return authority rather than improvise a new mandate.

When mandates collide

Conflict is not an edge case. It is the moment when a delegate's claimed mandate becomes observable.

01 / LARGE PRINCIPAL VS ECOSYSTEM

A major delegator prefers a narrow economic outcome while the delegate believes the proposal creates material protocol risk. Does the delegate follow the principal, apply its published policy, abstain or return the power?

02 / CROSS-DAO CLIENT CONFLICT

A professional intermediary represents power in related protocols while also serving clients or service-provider roles. The conflict may be real even when every individual vote is public.

03 / POLICY VS NEW EVIDENCE

A published voting policy points one way, but a novel proposal reveals risks the policy never anticipated. Blind consistency can be as irresponsible as unexplained deviation.

A serious mandate therefore needs a conflict procedure, not only a values statement. The agent should know in advance when it will disclose, consult, abstain, deviate from policy, refuse an instruction or return authority.

Working principle

Independence is a procedure, not an adjective: pre-commit where possible, disclose material conflicts, explain deviations and preserve the principal's ability to exit.

REPRESENTATION

A delegate statement is not a mandate unless it constrains behavior

Mandate model	What the agent optimizes	Failure test
Instructional	Explicit principal preference	Can instructions be expressed and verified?
Representative	Interests of delegators / constituency	Who counts as the constituency?
Stewardship	Protocol / ecosystem health	Who defines "health" when interests conflict?
Policy-based	Published voting principles	Is policy specific enough to predict action?
Expert	Domain-specific technical judgment	Is the scope actually limited to the domain?

Minimum disclosure: a serious delegate should make legible (1) represented stakeholders, (2) decision policy, (3) conflicts and client relationships, (4) when it will abstain, and (5) what could cause it to return or refuse authority.

Disclosure does not eliminate conflicts. It gives principals information with which to exercise exit.


A mandate is a decision rule for hard cases

A delegate statement has value only if it helps a principal predict behavior when interests diverge. The mandate should therefore be specific enough to explain whose interests guide the agent, what principles outrank short-term preferences, when conflicts trigger abstention or disclosure, and when the agent should return authority instead of stretching the mandate.

THE ECONOMICS OF REPRESENTATION

5. Who pays for governance?

"Delegate compensation" is too narrow a category. The real design question is who bears the cost of governance work - and what behavior the payment is buying.



Programs can pay for voting coverage, public rationales, availability, research, communications, proposal authorship, ecosystem work or simply the willingness to maintain operational readiness. Those are different services.

The measurement trap: the easiest outputs to verify are often the least informative about judgment. A DAO can count votes and rationales; it cannot cleanly observe the counterfactual value of a risk identified, a bad proposal improved, or a conflict handled well.

PAYMENT DESIGN

Start with the job, then choose the payment.

A voting representative, an optimistic watcher, a domain specialist and an operational steward incur different costs and create different value. Paying all of them through one activity score encourages the measurable surface of the work to become the job itself. Compensation should follow the actual service being purchased, not the historical label attached to the participant.

WHAT ARE YOU BUYING?

Compensation design changes the role

Model	What it rewards	Typical distortion
Fixed retainer	Availability / continuity	Pay-work mismatch in quiet periods
Per vote	Coverage	Mechanical voting / threshold gaming
Rationale bonus	Public legibility	Performative or AI-generated prose
VP threshold	Market validation	Incumbency / chicken-and-egg
Performance score	Multi-dimensional observable activity	KPI optimization
Delegator incentive	Demand for active representation	Yield-driven concentration
Scoped SP contract	Deliverable / expertise	Contractor capture / role overlap

Aave's 2026 framework is particularly instructive: delegates are defined as unpaid, while paid work lives in contributor and service-provider roles. ^[4] ENS takes another path, paying mainly the delegators who move voting power to active representatives. ^[10]

Metrics change behavior

Compensation is not neutral measurement. Once a DAO pays for a metric, participants rationally optimize toward it. Paying per vote can turn abstention into a penalty even when abstention is appropriate; paying for rationales can reward volume rather than judgment; paying for VP can entrench incumbents. The safer sequence is job -> evidence -> metric -> payment, not metric -> role.

Design compensation from the job backwards

There is no universal "delegate salary" because different delegation architectures purchase different forms of work.

ROLE	PRIMARY JOB	COST SHAPE	POSSIBLE PAY UNIT	BAD PROXY
General delegate	Broad judgment	Lumpy proposals	Availability + coverage	Raw vote count
Optimistic watcher	Exception detection	Readiness + incidents	Readiness / response SLA	Number of votes
Domain specialist	Scoped expertise	Review-driven	Scoped review / term	General forum activity
Steward / operator	Bounded execution	Continuous operations	SP contract / SLA	Rationale volume
Delegator incentive	Activate power	Balance + duration	Time-weighted delegation	Delegate output

The point is not that each row has one correct payment formula. The point is that the metric should correspond to the service actually being purchased. Arbitrum RAD links rewards more closely to proposal cadence; Aave now separates the delegate status from paid service-provider work; ENS is experimenting with paying mainly the principals who activate delegation. [4][10][11][12]

Design rule

Do not use one scorecard to pay representation, research, ecosystem advocacy, technical review and execution. Role ambiguity becomes incentive ambiguity.

ACTIVE / JULY 2026

Case: Arbitrum - pay follows proposal cadence

ACTIVE

Rewarding Active Delegates (RAD)

35

ENROLLED DELEGATES

July 2026

21

DELEGATES PAID

met monthly rules

\$23k

TOTAL REWARDS

July envelope

July included five in-scope governance actions. The monthly participation threshold was 75% (4/5 votes). 21 delegates qualified for payment. Of 117 vote instances, 70 had an on-time public rationale; missing a timely rationale reduced the reward for that proposal by 10%. ^[12]

The program also reduced per-proposal budgets for Q3 2026 after ARB's price decline. This matters because RAD is not a fixed salary: compensation is increasingly tied to the actual proposal mix and market context.

What RAD buys: measurable voting coverage plus a public reasoning trail - not proof that the underlying judgment was correct.

ACTIVE PILOT / JUL-SEP 2026

Case: ENS - pay the principal to choose active representation

ACTIVE PILOT

Delegation Incentives Program

90%

DELEGATOR POOL

10%

DELEGATE
POOL

ENS funded a three-month pilot with up to 90,000 ENS plus 5 ETH for sponsored delegation gas. The monthly pool scales with growth in voting power delegated to active delegates. Active delegates are defined through recent voting participation; rewards are viewpoint-neutral. ^{[10][11]}

The design is notable because it subsidizes **demand for representation**. The principal is rewarded for moving idle voting power to an active agent. Time-weighted balances and caps are intended to reduce short-lived or concentrated reward capture.

Open question: does paid delegation create durable governance attention, or does it create yield-sensitive political capital that leaves when the subsidy ends? The pilot had not completed at the research cutoff.

Two incentive logics, one underlying question

Arbitrum pays the representative for observable work; ENS directs most of its pilot rewards to principals who place power with active representatives. The mechanisms are different, but the design question is the same: what behavior is the DAO trying to buy, and what distortion might the payment create? The next chapter asks what happens after that paid or unpaid authority has been exercised.

CONTROL AFTER DELEGATION

6. Can power actually be taken back?

A system is not accountable because actions are visible. Accountability begins when observation can lead to a consequence.

Onchain governance often has strong technical exit. ERC20Votes holders can redelegate; Snapshot supports undelegation and space-specific overrides; track-based systems can withdraw authority selectively. ^{[1][2][15]}

But practical exit is more demanding. The principal has to notice poor performance, identify a better agent, clear custody or authorization workflows and actually submit the change. Stale authority can therefore survive despite perfect formal revocability.

Control test: if an agent repeatedly fails, what concrete event reduces its power? If the answer is only "people can see the dashboard," the system has observability, not accountability.

PRACTICAL EXIT

Technical revocation has four human frictions.

- Attention: someone must notice that the current agent should be changed.
- Alternative discovery: the principal needs a credible replacement, not merely a delete button.
- Authorization: custody, multisig or internal approvals may be required before redelegation can occur.
- Coordination: institutions and large holders may need policy review before the power actually moves.

FROM INFORMATION TO CONSEQUENCE

The Delegation Accountability Stack

1. MANDATE	What authority is granted, for how long, under what policy?
2. DISCLOSURE	Identity, affiliations, conflicts, represented stakeholders.
3. OBSERVATION	Votes, rationales, execution logs, incidents, response time.
4. EVALUATION	Who judges performance and by which criteria?
5. CONSEQUENCE	Reward reduction, non-renewal, power reduction, removal.
6. EXIT	Redelegation, revocation, replacement, expiry.

A rationale sits at layer 3. A performance review adds layer 4. A real ability to remove voting power or execution authority creates layers 5-6.

Accountability theater occurs when the visible layers are elaborate but the consequence layers are weak.

Why rationales are not accountability

A rationale can explain what an agent did and sometimes why. It does not, by itself, answer who evaluates the decision, what standard applies or what happens after repeated failure. The stack matters because each layer adds a different control: disclosure makes behavior legible; evaluation makes it reviewable; consequence and exit make the relationship enforceable.

Accountability should end in a state change

A dashboard is useful evidence. It becomes accountability only when repeated failure can change the agent's authority, compensation or mandate.

1	OBSERVE Missed material vote, conflict, stale rationale or failed execution.
2	EVALUATE A defined principal or body applies published criteria.
3	RESPOND Warning, remediation period, reward reduction or tighter scope.
4	CHANGE STATE Non-renewal, reduced voting power, revoked role or expired authority.
5	REPLACE A new agent can realistically acquire the mandate.

The most common failure is to overbuild the visible layers and underbuild the consequence layers. Public voting histories, rationales and scorecards can create strong observability while leaving the underlying authority untouched. The design should therefore work backward from the state change the principal is genuinely willing and able to use.

Three different concepts

Observability = we can see what happened. Accountability = someone can judge it. Constraint = that judgment can materially change the agent's power.

THE MARKET FOR REPRESENTATION

Revocable is not the same as contestable

Revocability asks whether authority can be removed. Contestability asks whether a capable alternative can realistically replace the incumbent. A healthy delegation market needs both.

HIGH REVOCABILITY

**Easy to remove
hard to replace**

Exit exists, but discovery or incumbency blocks competition.

**Easy to remove
easy to replace**

Healthy contestable representation.

LOW REVOCABILITY

**Hard to remove
hard to replace**

Entrenched intermediary power.

**Hard to remove
easy to enter**

Entry exists, but effective exit remains weak.

LOW CONTESTABILITY

HIGH CONTESTABILITY

Why technical exit is not enough

Minimum VP thresholds, featured rankings, reputation lock-in, treasury allocator preferences, identity continuity and the cost of monitoring alternatives can all make replacement difficult. A principal may be able to undelegate instantly and still have no realistic alternative to choose. Contestability therefore asks whether authority can move toward a better agent in practice, not merely whether the smart contract permits removal.

Practical test

If the principal can remove the incumbent but cannot discover, authorize and empower a credible replacement, the system is revocable but not contestable.

DISTRIBUTION AFTER THE INTERMEDIARY

7. Does delegation decentralize power?

Delegation can disperse attention while concentrating decision power.

Counting delegate profiles is almost meaningless by itself. A registry can contain hundreds of names while most effective voting power sits with a small group. Conversely, a system with fewer agents may preserve stronger override, scope limits or domain separation.

ACTIVATION

How much votable power is actually able and likely to participate?

CONCENTRATION

How much of active delegated VP sits with the largest agents?

QUORUM DEPENDENCE

How many top agents are arithmetically sufficient to equal the quorum threshold?

PROVENANCE

How much authority is organic vs treasury/foundation/incentive allocated?

CONTESTABILITY

Can new agents win meaningful power?

PERSISTENCE

How sticky is authority once assigned?

Decentralization has more than one layer

A system can decentralize participation while concentrating decision power. It can also concentrate voting power while preserving strong override, narrow scopes or short mandates. For that reason, "number of delegates" is a weak standalone measure. A useful assessment separates activation, concentration, provenance, contestability and the practical ability to replace an agent.

DELEGATION HEALTH

Measure a dashboard, not a magic score

Dimension	Candidate measure	Caveat
Activation	active delegated VP / votable supply	"active" must be defined consistently
Concentration	CR1 / CR3 / CR5 / CR10; HHI; effective N	requires full synchronized distribution
Quorum resilience	minimum k agents whose VP $\geq$ quorum	not evidence of coordination
Contestability	new entrant VP; redelegation rate	entity resolution is difficult
Provenance	organic vs allocator-origin VP	wallet labeling can be uncertain
Accountability	actual removals / sanctions / expiries	reporting alone is insufficient
Lifecycle	delegate survival / churn	quiet periods distort activity metrics

Do not publish cross-DAO Gini or HHI figures unless the denominator, snapshot timestamp, identity rules and treatment of inactive/self-delegated power are standardized.

Why a dashboard is more useful than one score

Activation, concentration, contestability, provenance and accountability can move in different directions. Compressing them into a single health score hides those trade-offs and invites false precision. A dashboard forces the designer to state the denominator, the time window and the control problem each metric is supposed to illuminate.

MEASURED PILOT DATA

Case: ZKsvnc - activation and concentration can rise together

355M ZK

MAX STAKED

87% of 400M cap

+205M ZK

NET ACTIVE DELEGATION


Feb 2 → May 11

173

DELEGATES RECEIVING VP

through the pilot

The top six recipients received roughly 271.6M ZK through the staking pilot. Relative to the 355M ZK peak staked, that is ~76.5%; the top three account for ~59.1% on the same comparison. ^[13]



This is **not** a claim about all ZK governance concentration. It is a clean measurement of concentration inside one incentive mechanism. The distinction matters.

Interpretation: the pilot succeeded at activation and simultaneously created a distribution problem. Mechanism success and distribution success are separate variables.

ORGANIC VS ALLOCATED

Power provenance should be visible on every delegate dashboard

ORGANIC VP

Selected by independent tokenholders / stakers without programmatic allocation.

ALLOCATED VP

Treasury, Foundation, program, chain or incentive-origin authority.

The two forms can be equally valid for execution but answer different legitimacy questions. Treasury delegation is particularly useful for breaking the entrant chicken-and-egg: a capable newcomer may need VP to become visible, yet visibility is required to win VP.

The cost is allocator dependence. When the allocator withdraws, the delegate learns whether it built an independent constituency or only occupied a program slot.

Uniswap's 2026 recall makes this distinction concrete. The treasury bootstrapped active delegates during a low-participation era, then withdrew the delegated UNI once the community judged the intervention no longer necessary. [8][9]

Recommended disclosure: show every delegate's VP by provenance, not just the aggregate total.

DEPENDENCY

Provenance is also a dependency metric.

A delegate with 80% of its power supplied by one treasury allocator faces a different continuity risk from a delegate with the same total power distributed across many independent principals. The dashboard should make that dependence visible because the probability, speed and legitimacy of a future power loss are not the same.

POWER MOVES - OR IT STAGNATES

8. The delegate lifecycle

Delegation should be studied as a market over time, not as a leaderboard snapshot.

The central lifecycle question is not "How do I become a delegate?" It is whether authority can move toward better agents and away from weaker ones without a crisis.

ENTRY

How does a new agent become discoverable and eligible?

INITIAL POWER

Does the agent earn organic VP, receive bootstrap VP, or buy attention through incentives?

OPERATION

What observable work and decisions occur during the mandate?

REVIEW

Who notices inactivity, conflicts, poor judgment or organizational change?

GROWTH / DECLINE

Can high-quality agents gain authority and weak agents lose it?

EXIT


What happens to authority when an individual, team or brand disappears?

What healthy movement looks like

The lifecycle is healthy when power can move before a crisis. New agents can become visible; strong agents can gain trust; weak or inactive agents can lose authority; and organizational changes are surfaced quickly enough for principals to react. A static leaderboard can look orderly while hiding a market in which almost nothing is actually contestable.

DELEGATE LIFECYCLE

A healthy delegation market has a complete exit loop



ENTRY RISK

Thresholds improve signal but can entrench incumbents. Treasury allocation can open a door while creating allocator dependence.

MID-LIFE RISK

A delegate may remain visible and compensated after its real voting power, team capacity or attention has changed.

EXIT RISK

Authority can remain attached to an inactive identity if principals do not notice retirement, brand changes or key-person transitions.

Lifecycle design is where technical revocability becomes institutional contestability.

The loop is the point

Discovery and eligibility matter only if they connect to review and exit. Without that loop, delegation becomes a one-way allocation event: power arrives, then inertia keeps it in place. The practical test is whether a principal can move from "this agent is no longer right for me" to a replacement without unusual coordination, bespoke governance or a security incident.

Organizations need succession rules

A wallet, brand or delegate profile can outlive the team that originally earned the delegation.

PERSONNEL CHANGE

Who must be disclosed when the analysts, signers or decision-makers materially change?

MERGER / CLIENT TRANSFER

Can governance power follow a commercial relationship automatically? Usually it should not be assumed.

KEY / WALLET CHANGE

Can signing authority move without silently changing the agent the principal selected?

WIND-DOWN

Who announces the shutdown, freezes action, returns roles and helps principals redelegate?

BRAND OR ENTITY CHANGE

Does a rebrand preserve the mandate, or should principals be asked to reauthorize?

Lifecycle controls matter because delegation is usually attached to an address while trust is attached to people, processes and reputation. Those layers can diverge. A serious organizational delegate should publish the identity of the operating team, key-management model, continuity plan and conditions under which authority should be returned rather than transferred.

Exit principle

The principal should never discover a material change in who exercises its governance power only after the next vote.

HISTORICAL EVIDENCE

Two lifecycle warnings: lost delegation and ghost representation

LIDO / 2025

Delegation can disappear fast

During Lido's pilot, three of seven public delegates saw delegation fall close to zero. Two remained active; one stopped participating and became ineligible for incentives. [24]

The case separates **political capital** from **governance labor**: compensation may keep work going after VP falls, but it cannot substitute indefinitely for delegated trust.

Lido's Q1 2026 report later recorded seven eligible delegates with 100% voting participation and a \$75K pool. [23]

SAFE / JAN 2025

Registries can overstate representation

Safe's analytics thread reported 40 active delegates, 66 inactive delegates and 115 "ghost" delegates in January 2025; only 27 delegates voted on the month's single proposal. [25]

A large delegate registry is therefore not evidence of a competitive representative market. The relevant quantities are active power, actual participation and the ease with which principals can migrate away from stale agents.

Lifecycle closes the accountability loop

The Lido and Safe cases show why a delegation market cannot be evaluated from registry size or one participation snapshot. Political trust, operational activity and organizational identity can all change at different speeds. The next step is to ask whether some work should be delegated more narrowly so that failure in one role does not compromise the entire representative relationship.

SPECIALIZATION

9. When delegates stop being the answer

As governance matures, the scalable solution is often to delegate *less broadly*, not to create more general-purpose delegates.

Three design patterns are becoming increasingly important:

DOMAIN DELEGATION

The principal assigns different experts to different decision classes. Polkadot allows delegation by referendum track. [15]

OPTIMISTIC OVERSIGHT

Routine actions pass unless a watcher or veto body objects. Optimism and Reserve illustrate different versions of this pattern. [3][18][19]

BOUNDED EXECUTION

A steward can execute specified parameter changes within coded limits, while governance retains residual control. Aave is the clearest current case. [4][6]

Mature governance often decomposes the delegate role: general political judgment, technical expertise, monitoring and execution do not need to live in the same agent.

RESIDUAL CONTROL

Mature systems often move work without moving ultimate authority.

Specialization scales when the principal can narrow the task and retain a residual check: direct override, veto, charter removal, hard-coded limits or a sunset. The goal is not to create more governance bodies. It is to put recurring or expert work closer to the actor best able to perform it while preserving a credible route back to the principal.

When specialization beats representation


Broad delegates are useful when decisions require recurring political judgment across many topics. They are less useful when the task is narrow, technical and repeatable. In those cases, a domain body, optimistic watcher or bounded steward can reduce the amount of discretion being outsourced while preserving a clear residual check for tokenholders.

CURRENT

Case: Polkadot - one holder, multiple specialists

CURRENT ARCHITECTURE

OpenGov multirole delegation



OpenGov lets a holder specify a different delegate for each referendum class. The holder can delegate one track, multiple tracks or none; undelegation is track-specific. ^[15]

This design rejects the assumption that a principal has one political identity represented by one generalist. Expertise can be modular.

Design implication: if governance issues have meaningfully different knowledge requirements, domain-specific delegation can preserve principal choice while reducing the scope of each agent's discretion.

In plain terms

Polkadot allows a holder to say, in effect, "I trust one person on treasury questions and somebody else on technical questions." That matters because expertise is not universal. Domain delegation reduces the need to hand one generalist an all-purpose political mandate and gives the principal a more precise way to express where outside judgment is actually useful.

ACTIVE / EXPANDED AUG 2026

Case: Reserve - from voting delegate to optimistic watcher

ACTIVE / EXPANDED

Reserve Delegate Program

A three-month pilot concluded in July 2026 after delegates voted on 31 proposals across two Yield DTFs, with 100% delegate voting participation in the reported trial. The review concluded that no proposal in the covered DTFs could have reached quorum without delegate participation. ^[18]

The more interesting development came after the pilot. The program was renewed for six months and expanded to AI DTFs using **optimistic governance**. In those systems, proposals pass unless vetoed; not voting is acceptance. Delegate value therefore shifts from constant voting toward monitoring for unaligned or harmful proposals. ^{[18][19]}

Why this matters: the optimal delegate KPI changes when the architecture changes. "100% voting participation" cannot be a universal standard if the system intentionally minimizes votes.

What changes under optimistic governance

When a proposal passes unless somebody objects, the delegate job changes. Constant voting is no longer the core output; monitoring, exception detection and timely veto become more valuable. A program that kept rewarding raw vote count after such an architectural change would be measuring the old job rather than the current one.

ROLE DECOMPOSITION

Generalists. specialists and operators solve different problems

Role	Primary job	Discretion	Best control
General delegate	Broad proposal judgment	High, across topics	Mandate + redelegation + disclosure
Domain delegate	Judgment in one decision class	High, narrow scope	Track scope + redelegation
Expert council	Technical / security filter	Medium-high, institutional	Election / charter / veto or override
Optimistic watcher	Detect exceptions / veto harmful action	Event-driven	Clear trigger + response SLA
Steward	Operate parameters inside bounds	Low-medium, coded scope	Hard limits + cooldown + revoker
Executor	Implement an already-made decision	Low	Exact transaction scope + audit trail

Do not pay all six roles with one delegate scorecard. The object of accountability is different.

Role decomposition changes what should be measured

A general delegate, expert council, optimistic watcher, steward and executor do not create the same risk. Their accountability should follow their discretion: broad judgment needs mandate and political exit; narrow operational authority can often be controlled through bounds, cooldowns, logs and a revoker. This decomposition also changes how professional governance firms should be evaluated across DAOs.

PROFESSIONALIZATION

10. The governance intermediary

Professional delegation is creating a reusable governance-intermediation layer across protocols.

The valuable asset is not simply voting power. Professional operators accumulate research process, forum context, signing infrastructure, reputational capital, conflict procedures, institutional relationships and the ability to monitor multiple governance venues.

That creates economies of scale - and new risks that single-DAO frameworks do not capture.

BENEFIT

Specialization, continuity, faster information processing, operational security and a repeatable governance process.

RISK

Cross-DAO influence, correlated judgment, client conflicts, shared research assumptions and key-person / organizational continuity.

The right unit of analysis is the intermediary, not the brand name. Firms can merge, shut down, rebrand or transfer teams while delegated addresses and relationships persist.

FIRM-LEVEL RISK

Professionalization changes the unit of governance risk.

When one organization operates across several DAOs, the relevant control problem extends beyond a single voting wallet. Shared research assumptions, client relationships, signers, policies and personnel can create correlated behavior across systems. Protocol-by-protocol disclosures remain useful, but they do not fully describe the intermediary-level conflict surface.

What a governance intermediary actually sells

The reusable asset is a process: monitoring, research, policy interpretation, signing operations, institutional memory and the ability to act consistently across venues. That scale can lower the cost of governance for principals. It also means the firm - its clients, staff, methods and conflicts - becomes part of the risk surface rather than a neutral wrapper around individual votes.

INTERMEDIARY RISK

Cross-DAO power cannot be measured by adding token balances

Ten million ARB + one million UNI + two million OP is not a meaningful aggregate influence metric. Tokens, quorums and decision scopes are incomparable.

A serious cross-DAO concentration study would instead normalize each intermediary's influence by variables such as:

ACTIVE VP SHARE

Share of active delegated voting power within each system.

QUORUM INFLUENCE

Whether the intermediary materially changes the minimum coalition needed to satisfy quorum.

PIVOTALITY

How often its vote could change an outcome under observed turnout.

ROLE OVERLAP

Whether the same entity is delegate, service provider, grantee, council member or client representative in related systems.

Public disclosures today are mostly protocol-by-protocol. Cross-system client and conflict disclosure remains immature.

Research gap: there is no reliable public census of normalized cross-DAO governance influence by professional intermediary.

Normalization before comparison

Professional intermediaries cannot be compared by adding raw token balances across protocols. A better approach asks how much of each system's active power they control, whether their participation changes quorum coalitions, how often their vote is pivotal and which other paid or governance roles overlap with that influence.

Cross-DAO conflict is a network problem

Professional governance firms accumulate roles across systems. The relevant conflict surface is the network of relationships, not one delegate profile.



RISK SURFACES

- Correlated judgment across related protocols.
- Commercial or client conflicts.
- Information advantages from service-provider work.
- Key-person and organizational continuity.

MINIMUM DISCLOSURE

- Material governance roles.
- Economic exposure and paying relationships where disclosure is permitted.
- Conflict / recusal procedure.
- Who signs, who decides and what happens when the team changes.

Measurement warning

Cross-DAO influence cannot be measured by adding token balances. Comparison should normalize active VP share, quorum influence, pivotality and role overlap within each system.

FROM CUSTODY PROBLEM TO POLICY PROBLEM

11. Institutional delegation

For institutional holders, the technical ability to exercise governance rights is increasingly available. The harder question is who is authorized to decide.

Coinbase Prime documents governance delegation for assets including AAVE, GTC, UNI, COMP, ARB and OP while assets remain in custody. Anchorage provides governance voting from qualified custody and internal approval workflows, including Aave delegation. [26][27][28][29]

This weakens an old explanation for non-participation: "the tokens are in custody, so we cannot vote." In many cases the remaining constraint is organizational policy.

Institutional governance is becoming an authorization and stewardship workflow: what rights exist, which events matter, who may decide, who may execute, and what evidence is retained.

THE PRINCIPAL STILL OWNS THE POLICY

Outsourcing execution is not outsourcing responsibility.


An institutional holder still needs to decide which assets have material governance rights, which events require escalation, who can instruct or override an external representative, what conflicts are prohibited and what evidence must be retained. Custody infrastructure can make governance technically possible; it does not answer those policy questions.

The hidden work is internal

For an institution, delegation usually begins before any onchain transaction. Someone must decide which governance rights are material, which proposals require escalation, who is allowed to instruct an external representative and what evidence is needed for audit or review. Custody can enable execution; it cannot substitute for this internal policy layer.

OPERATING MODEL

The institutional delegation workflow



INTERNAL

Research and vote inside the institution. Highest direct control; highest internal attention cost.

EXTERNAL

Delegate discretion to a specialist under a mandate. Lower attention cost; introduces agent and conflict risk.

ABSTAIN

Do not exercise the right by policy. Can be rational when governance is immaterial or legal/operational cost is disproportionate.

Policy comes before the execution model

An institution can vote internally, delegate externally or abstain, but each posture still needs an explicit policy. The workflow is therefore circular rather than linear: observed outcomes and incidents should feed back into the policy that determines future materiality thresholds, approval rights and delegation choices.

INSTITUTIONAL MARKET

What the evidence supports - and what it does not

SUPPORTED

- Qualified/institutional custody can support governance voting and delegation for multiple major assets. [26][28][29]
- Professional multi-protocol governance operators exist.
- Governance can be separated from asset custody through authorization and key workflows.
- Institutions therefore face a real policy choice: vote, delegate, monitor or abstain.

NOT YET PROVEN

- A large, transparent market of crypto proxy-adviser mandates equivalent to ISS / Glass Lewis.
- A public standard for fiduciary-style digital-asset voting policies across asset managers.
- A reliable dataset showing how many funds outsource voting and under what contractual mandate.

Do not import the traditional proxy-voting analogy too literally. The useful comparison is the operating problem - many rights, limited attention, policy, execution and audit - not legal equivalence.

What the evidence supports as an operating model

The available evidence is enough to support a workflow - policy, authorization, decision authority, execution and audit - even though it is not enough to claim that crypto already has a mature proxy-advisory market. The practical task is therefore to define the mandate clearly and keep the principal in control of material exceptions.

A minimum institutional delegation mandate

Delegating governance rights does not outsource ownership of the policy. The principal still needs to define the scope, escalation rules and revocation path.

1. SCOPE Assets, chains, governance venues and proposal classes covered.	2. POLICY Decision principles; prohibited actions; when abstention is required.
3. AUTHORITY LEVEL Research only, recommendation, vote, delegation or bounded execution.	4. MATERIALITY / ESCALATION Which events require principal approval before action.
5. CONFLICTS Required disclosures, recusal rules and return-of-authority triggers.	6. EXECUTION Custody workflow, signers, approvals, transaction controls and evidence.
7. REPORTING Votes, rationales, exceptions, incidents and periodic review.	8. TERMINATION Who can revoke, how quickly, what happens to keys/roles and records.

The mandate can be short. It cannot be vague. "Vote in the best interests of the ecosystem" is not enough for an institution that must explain who had authority, which policy applied, what controls existed and how the action was reviewed.

Non-negotiable

The principal retains the ability to change the policy and revoke the authority. Delegation changes who exercises discretion; it should not silently transfer ultimate control.

WHEN PARTICIPATION BECOMES CHEAP

12. Delegation after AI

AI is automating exactly the visible outputs that many delegate programs learned to reward.

CHEAPER NOW

Proposal summaries, issue classification, forum monitoring, first-pass research, draft rationales, voting-history comparison.

STILL HARD

Mandate interpretation, adversarial judgment, conflict handling, stakeholder consultation, responsibility for failure and willingness to be replaced.

This changes the evidentiary value of governance artifacts. A long rationale once implied substantial reading and synthesis effort. In 2026 it may imply only that an agent can generate prose.

Design shift: delegate programs should move from proof of activity toward evidence of judgment, escalation quality and accountable decision process.

SHIFTED AGENCY PROBLEM

AI moves the boundary; it does not remove the principal-agent problem.

As analysis becomes cheaper, the important questions move upward and downward in the stack: who wrote the policy, which sources the model can trust, how exceptions are escalated, who controls the key and who bears responsibility for an error. Automation can reduce information costs while increasing the need for explicit limits and logs.

AI DELEGATION

Autonomy is a spectrum, not a binary



The first two stages are already operationally plausible at scale. Agentic voting experiments exist, but evidence is not yet sufficient to claim that autonomous representatives produce superior governance outcomes.


MANDATE DRIFT	Does the model still implement the policy principals chose?
MODEL / INPUT ATTACK	Can adversarial content manipulate analysis?
KEY CONTROL	Who can stop signing or execution?
ACCOUNTABILITY	Who bears responsibility when the agent fails?
REVOCATION	Can authority be removed faster than the agent can act?

The meaningful boundary is authority, not authorship

Whether a rationale was drafted by a person or a model matters less than who owns the policy, who approved the conclusion and who can stop execution. AI makes visible governance artifacts cheaper; it does not make authority less consequential. The control design should therefore follow the level of discretion the system gives the agent.

The AI control stack

Automation can compress information costs without making the mandate, key control or accountability problem disappear.



SAFE TO AUTOMATE EARLIER

- Proposal discovery and classification.
- Summarization and source comparison.
- Policy matching and draft recommendations.
- Routine evidence capture.

REQUIRES STRONGER CONTROL

- Material mandate interpretation.
- Conflict and stakeholder trade-offs.
- Signing / transaction authority.
- Model, prompt or policy changes that alter behavior.

The control surface should become stricter as the agent moves from analysis toward irreversible action. At minimum: constrain sources for material claims, log policy/model changes, impose transaction and scope limits, require human approval for high-materiality actions and maintain an emergency revocation path. Fully autonomous representatives remain experimental; the evidence does not yet support treating them as a mature substitute for accountable human governance.

Design question

Can authority be revoked faster than the agent can create an irreversible action? If not, the automation boundary is too permissive.

THE ANTI-HANDBOOK CHAPTER

13. When not to delegate

A handbook about delegation should be able to recommend no delegation.

Delegation adds an intermediary. That is justified only when the reduction in information, coordination or execution cost exceeds the new agency, concentration and monitoring cost.

KEEP DIRECT

Few high-materiality decisions; principals are informed and available; voting cost is low.

USE OPTIMISTIC

Most actions are routine and safe by default; value lies in exception detection and veto readiness.

AUTOMATE

Rules can be specified ex ante and executed safely within hard limits.

USE SPECIALIST BODY

The decision requires narrow expertise but should not grant general political authority.

DELEGATE TEMPORARILY

A quorum or bootstrapping problem exists but should expire when the condition improves.

ABSTAIN / NO GOVERNANCE

The right is immaterial, the mandate prohibits action, or governance itself is unnecessary.


NEGATIVE DESIGN

No delegation can be the responsible outcome.

A handbook that always recommends another intermediary is badly specified. Direct action is better when the principal has capacity and the decisions are few and material. Automation is better when rules can be safely bounded. Abstention can be rational when the right is immaterial or the mandate forbids action. Good governance includes knowing when not to create a representative layer.

SHOULD THIS POWER BE DELEGATED?

A minimal decision tree



The default should not be maximum delegation. The default should be minimum necessary discretion.

The decision tree is intentionally conservative

The tree starts by trying to remove or narrow discretion before assigning a broad representative. That ordering matters. It prevents delegation from becoming a default response to every governance problem and creates a clean handoff to the final framework: if authority still needs to be delegated, the architecture must specify exactly how.

ONE FRAMEWORK

14. The Delegation Architecture Canvas

The handbook's design framework is deliberately singular. Every delegation architecture should answer the same twelve questions.



How to fill the canvas

Use the canvas in order. First name the problem and the principal. Then define the exact authority object before choosing the agent. Only after that should you specify mandate, discretion, incentives and execution. Finish with monitoring, consequence, revocation and contestability. If the design becomes vague near the end, the early allocation decision is probably carrying more discretion than the system can control.

COPY / ADAPT / TEST

The canvas in one page

WHY DELEGATE?	Problem:	PRINCIPAL	Authority owner:
OBJECT	Vote / veto / domain / execution:	AGENT	Individual / org / body / contract / agent:
ALLOCATION	Organic / election / treasury / default / incentive:	MANDATE	Principal / ecosystem / policy / expert:
DISCRETION	General / bounded / rule-based:	DURATION	Open-ended / term / event / sunset:
INCENTIVES	Who pays? What is purchased?	EXECUTION	Governor / Snapshot / multisig / steward / custodian:
MONITORING	Votes / rationales / logs / incidents:	CONSEQUENCE	Reward / warning / reduction / non-renewal:
REVOCATION	Who can remove power? Time / friction?	CONTESTABILITY	Can a better agent realistically replace the incumbent?

Stress test: if any cell can only be answered with "the community will decide," the design probably lacks a defined control surface.

Start with diagnosis, not delegates

Illustrative case: Meridian Protocol is fictional. The purpose is to show how the Delegation Architecture Canvas changes a real design conversation.

ILLUSTRATIVE - NOT A LIVE PROTOCOL

Meridian has 12-20 governance actions in a typical month. Most are routine risk or operational updates. A small number are high-materiality treasury, constitutional or upgrade decisions. Tokenholder attention is uneven and quorum has occasionally become a concern.

OBSERVED SYMPTOM	LIKELY CAUSE	WRONG FIRST RESPONSE	BETTER DESIGN QUESTION
Low turnout on routine votes	Attention cost exceeds materiality	Pay more general delegates	Can routine actions move to bounded / optimistic execution?
Slow risk parameter updates	Technical specialization + latency	Make every delegate a risk expert	Should a narrow expert/steward role own this function?
Quorum misses on major votes	Idle or concentrated voting power	Permanent treasury delegation	Is temporary activation / bootstrap enough?
Institutional holders abstain	Authorization / policy friction	Assume they do not care	What custody and escalation workflow is missing?

Diagnosis

Meridian does not have one "delegation problem." It has at least three different problems: attention, technical specialization and occasional power activation. They should not be solved with one all-purpose delegate layer.

Decompose the authority before assigning agents

Meridian keeps broad political judgment with tokenholders and delegates only the narrower functions that benefit from specialization or speed.

OBJECT	AGENT	DISCRETION	EXECUTION	CONTROL
CONSTITUTION / TREASURY	Tokenholders / opted-in representatives	Broad discretion	Onchain vote	Redelegation / direct vote
RISK REVIEW	Expert risk body	Narrow domain judgment	Recommendation / veto	Election / charter / removal
ROUTINE PARAMETERS	Bounded steward	Coded discretion	Steward transaction	Caps + cooldown + revoker
EMERGENCY ACTION	Security body	Trigger-based	Pause / veto	Expiry + post-event review

WHAT MERIDIAN DOES NOT DO

- It does not create a universal paid delegate program merely because turnout is uneven.
- It does not ask generalists to approve every technical parameter update.
- It does not give a treasury allocator indefinite bootstrap power.

WHAT IT PRESERVES

- Tokenholders retain residual control over constitutional and capital-allocation decisions.
- Specialist authority is domain-limited.
- Operational authority is bounded in code where possible.
- Temporary interventions carry a review or sunset.

Canvas result

One governance system can legitimately contain several authority objects and several agents. The design question is not "who is the delegate?" but "which discretion belongs where?"

Finish with incentives, consequence and exit

A delegation architecture is incomplete until failure changes something.

INCENTIVES	MONITORING	CONSEQUENCE / EXIT
General representatives: no automatic DAO-wide salary. Risk body: scoped term / review compensation. Steward: service-provider contract tied to operational scope. Temporary bootstrap: sunset, not permanent entitlement.	Voting and execution logs. Conflict and affiliation disclosure. Material incidents and response time. Periodic mandate review. Power provenance by source.	Reward reduction or non-renewal. Narrower scope after incidents. Role or signer removal. Treasury delegation recall. Principal redelegation / direct override where supported.

Stress test the design before launch: What happens if the expert body is wrong? If the steward key is compromised? If a bootstrap delegate becomes dominant? If a professional firm changes team? If principals stop paying attention? Each failure should map to a control surface rather than a hope that "the community" will eventually notice.

Final answer for Meridian

Delegate the minimum discretion necessary. Preserve tokenholder control over the high-materiality residual decisions, and make every delegated role legible, bounded, reviewable and replaceable.

THE FRAMEWORK IN TEN LINES

Ten design principles

- 01 Delegate a defined problem, not a fashionable role.
- 02 Separate voting, representation, veto and execution authority.
- 03 Make power provenance visible.
- 04 Limit discretion before adding monitoring.
- 05 Do not mistake activity for judgment.
- 06 Design consequence, not just disclosure.
- 07 Technical exit is insufficient without practical replacement.
- 08 Use specialization when generalist delegation becomes a bottleneck.
- 09 Give bootstrap authority a review or sunset.
- 10 Delegate the minimum discretion necessary - and preserve a credible path to take it back.

How to use the principles

These are not ten independent preferences. They form a sequence. Start by naming the problem and the authority object; narrow discretion where possible; make the origin of power visible; match monitoring and incentives to the actual role; and design a credible state change when the relationship fails. The final test is whether a principal can understand where its authority went and recover it without a governance crisis.

The appendices that follow are reference material for applying that sequence: a comparative system map, a registry of failure and redesign cases, and the method and primary-source notes behind the handbook.

CURRENT + HISTORICAL REFERENCE

Appendix A - Comparative system map

System	Relevant delegation design	2026 / latest relevant state	Why it matters
Arbitrum	Token delegation + paid active-delegate program	ACTIVE RAD July 2026: 35 enrolled, 21 paid. ^[12]	Compensation tied to proposal cadence and rationale rules.
Optimism	Token delegation + optimistic approvals + specialized bodies	CURRENT Season 8 used more optimistic approvals and no direct delegate incentives. ^[3]	Participation capacity can persist with fewer explicit actions.
Uniswap	Token delegation; historical treasury bootstrap	EXECUTED ~12.5M UNI bootstrap delegation recalled 1 Jun 2026. ^{[8][9]}	Clean intervention lifecycle: bootstrap → maturity → unwind.
ENS	Direct delegation + delegator/delegate incentives	PILOT 90-day pilot launched Jul 2026; 90% delegator / 10% delegate pool. ^{[10][11]}	Moves incentives from supply to demand for representation.
Aave	Voting delegation + bounded steward authority	CURRENT Framework v2: delegates unpaid; paid work in contributor/SP roles. ^[4]	Separates representation from professional execution.
ZKsync	Delegation activation + staking/delegate incentive	PAUSED Season 1 ended May 2026; Season 2 paused for redesign. ^[13]	Activation success exposed concentration and custody constraints.
Reserve	DTF delegates + optimistic-governance watchers	ACTIVE renewed six months and expanded Aug 2026. ^[19]	Delegate role changes from voter to vigilant veto-capable watcher.

SYSTEMS 8-14

Appendix A - Comparative system map

System	Relevant delegation design	Status / evidence	Why it matters
Polkadot	Track-specific multirole delegation	CURRENT per-track delegation / undelegation. [15]	One principal can choose different experts by decision class.
Cardano	DRep representative class + SPO / constitutional bodies	CURRENT post-Chang / Plomin governance architecture. [16][17]	Representation is explicit institutional infrastructure, not a UI layer.
Gnosis	Treasury-bootstrapped delegate program	REJECTED 2025 renewal GIP-137 rejected. [20][21]	Operational improvement does not guarantee political renewal.
Lido	Public delegate platform + quarterly incentives	2026 EVIDENCE Q1 2026: 7 eligible delegates, 100% participation, \$75K pool. [23]	Strong lifecycle evidence when delegated VP fell sharply for some participants.
Safe	Token delegation / delegate registry	HISTORICAL Jan 2025 analytics: 40 active, 66 inactive, 115 ghost delegates. [25]	Registry size can radically overstate active representation.
RARI	veRARI delegation + fixed-per-activity incentives	HISTORICAL Q4 2025 report published Jan 2026. [31][32]	Shows simplification when scoring overhead exceeds governance scale.
Snapshot spaces	Space-specific / global delegation, custom registries	CURRENT direct delegation and delegate registries. [2]	Delegation semantics and discovery can vary by space/strategy.

SYSTEMS 15-20

Appendix A - Comparative system map

System / layer	Relevant design	Status	Research use
Cosmos-style validator governance	Staking-linked default representation with delegator override	ARCHETYPE	Political representation can emerge from a non-political economic relation.
Compound-style Votes	Direct ERC20-style delegation	ARCHETYPE technical baseline	Canonical model for separating balance from activated voting power.
Starknet Foundation delegation	Institutional allocator distributes voting power	HISTORICAL DESIGN	Useful for studying bootstrap legitimacy and allocator dependence.
Scroll treasury / reward programs	Treasury VP + multidimensional delegate scoring	HISTORICAL DESIGN	Shows the difference between measuring observable behavior and decision quality.
Institutional custody	Beneficial owner retains asset; separate governance workflow	CURRENT Coinbase Prime / Anchorage. [26][28]	Governance can be technically exercised without moving custody.
Agentic delegation	Software agent assists or executes governance under policy	EXPERIMENTAL	Future boundary: machine-readable mandates, key control and accountability.

The appendix is a comparative design map, not a claim that all programs are currently active. Historical rows are retained only when the redesign or failure pattern remains analytically useful.

FAILURE IS DATA

Appendix B - Failure and redesign registry

Case	Friction / redesign	Lesson
Gnosis GIP-137	Delegate-program renewal rejected after a trial that had improved quorum reliability. [20][21]	Operational metrics do not create budget legitimacy.
ZKsync Season 1	Activation rose, but top recipients concentrated pilot VP; featured list changed; Season 2 paused. [13][14]	Incentives create second-order allocation problems.
Uniswap Franchiser	Bootstrap treasury delegation recalled once original rationale weakened. [8][9]	Temporary interventions need an exit rule.
Aave Orbit → Framework v2	Paid delegate program ended; delegates defined as unpaid while scoped work lives elsewhere. [4][5]	Representation and paid service are separable roles.
Optimism Season 8	Lower visible participation alongside more optimistic approvals and stable capacity. [3]	Do not optimize raw activity when architecture reduces required voting.
Lido delegate attrition	Several delegates lost almost all VP; participation responses differed. [24]	Governance labor and political capital can decouple.
Safe ghost delegates	Large inactive / ghost population in registry. [25]	Profile count is a weak decentralization metric.
Reserve expansion	Program moved into optimistic-governance DTFs where watchfulness matters more than vote count. [18][19]	The job must change with the decision mechanism.

RESEARCH NOTE

Appendix C - Method, currentness and limitations

Publication cutoff

Current-state claims were rechecked through **25 August 2026**. Governance changes quickly; a program described as active may be amended after publication. The handbook therefore treats parameters and program status as point-in-time evidence, not evergreen facts.

Case selection

Cases were selected for analytical value rather than market-cap ranking. Priority went to systems that reveal a distinct mechanism, redesign, failure mode or current frontier: delegation incentives, treasury bootstrapping, optimistic oversight, domain delegation, operational stewards, institutional custody and agentic experimentation.

Measurement limits

No universal Delegation Health Score is published. Cross-DAO concentration requires synchronized onchain snapshots, a consistent denominator and entity resolution. Where those conditions were absent, the handbook uses bounded case-specific observations rather than fabricated precision.

Provider / organization status

Provider relationships are treated as point-in-time. Historical transitions are retained only where they clarify continuity or mechanism, and should not be read as claims that a former provider remains active. The StableLab-to-Anode transition is treated on that basis. [33]

Standard: separate mechanism → stated objective → observed behavior → observed outcome. Do not infer causality from program self-reports alone.

PRIMARY SOURCES / 1-17

Source notes

[1] OpenZeppelin Contracts - Governance / Votes

<https://docs.openzeppelin.com/contracts/5.x/api/governance>

[2] Snapshot - Delegation documentation

<https://docs.snapshot.box/user-guides/delegation>

[3] Optimism - Token House participation and incentives: Season 8

<https://gov.optimism.io/t/token-house-participation-and-incentives-season-8-cycle-39a-45/10537>

[4] Aave - Governance Framework v2

<https://governance.aave.com/t/arfc-governance-framework-v2/25348/1>

[5] Aave - Orbit Program Renewal Q1 and Q2 2026

<https://governance.aave.com/t/arfc-orbit-program-renewal-q1-and-q2-2026/23695>

[6] Aave - Risk Stewards cooldown / pauser role

<https://governance.aave.com/t/arfc-risk-stewards-cooldown-reduction-umbrella-pauser-role-reassignment/25068>

[7] Uniswap - Return 12.5M Delegated Tokens (forum)

<https://gov.uniswap.org/t/temp-check-return-12-5m-delegated-tokens-to-the-governance-timelock/26099>

[8] Uniswap - Return 12.5M Delegated Tokens (executed proposal)

<https://vote.uniswapfoundation.org/proposals/97>

[9] Uniswap / DUNI - Q2 2026 financial statements

<https://governance.cowrie.pro/docs/q2-2026>

[10] ENS EP 6.31 - Delegation Incentives Program

<https://docs.ens.domains/dao/proposals/6.31>

[11] ENS EP 6.48 - Delegation Incentives Program funding transfer

<https://docs.ens.domains/dao/proposals/6.48/>

[12] Arbitrum - Rewarding Active Delegates, July 2026 results

<https://forum.arbitrum.foundation/t/rewarding-active-delegates-july-2026-results/31187>

[13] ZKsync - ZKnomics Staking Pilot Season 1 Review

<https://forum.zknation.io/t/season-1-review-zknomics-staking-pilot-program/992>

[14] ZKsync - Season 1 concentration / featured-list follow-up

<https://forum.zknation.io/t/season-1-review-zknomics-staking-pilot-program/992/3>

[15] Polkadot Wiki - OpenGov

<https://wiki.polkadot.com/learn/learn-polkadot-opengov/>

[16] Cardano Docs - Governance overview

<https://docs.cardano.org/about-cardano/governance-overview>

[17] Cardano Docs - Chang / governance transition

<https://docs.cardano.org/about-cardano/evolution/upgrades/chang>

PRIMARY SOURCES / 18-34

Source notes

[18] Reserve - Delegate Program Trial Run Review

<https://forum.reserve.org/t/reserve-delegate-program-trial-run-review/1584>

[19] Reserve - Delegate Program: the next six months

<https://forum.reserve.org/t/delegate-program-the-next-six-month/1607>

[20] Gnosis - GIP-137 Delegate Program renewal

<https://forum.gnosis.io/t/gip-137-should-gnosis-dao-renew-the-delegate-program/11658>

[21] Gnosis - October 2025 governance summary (GIP-137 rejected)

<https://forum.gnosis.io/t/gnosisdao-governance-summary-october-2025/11765>

[22] Lido - Delegate Incentivization Pilot: lessons learned

<https://research.lido.fi/t/establish-a-public-delegate-platform-and-delegate-incentivization-program/7858/49>

[23] Lido - Delegate Oversight Committee Q1 2026

<https://research.lido.fi/t/establish-a-public-delegate-platform-and-delegate-incentivization-program/7858/94>

[24] Lido - Delegate Oversight Committee, Mar 2025

<https://research.lido.fi/t/establish-a-public-delegate-platform-and-delegate-incentivization-program/7858/48>

[25] Safe - SafeDAO Governance Analytics Report Thread

<https://forum.safe.global/t/safedao-governance-analytics-report-thread>

[26] Coinbase Prime - Governance overview

<https://help.coinbase.com/en/prime/governance/overview>

[27] Coinbase Prime - Custody and governance

<https://www.coinbase.com/en-gb/prime/custody>

[28] Anchorage Digital - Governance

<https://www.anchorage.com/platform/governance>

[29] Anchorage Digital - On-chain Aave governance

<https://www.anchorage.com/insights/introducing-on-chain-aave-governance-at-anchorage-digital>

[30] ScopeLift - Taking over operation of Tally's governance platform

<https://scopelift.co/blog/scopelift-tally-operation>

[31] RARI - Q4 2025 Delegate Incentive Program Report

<https://forum.rari.foundation/t/q4-2025-delegate-incentive-program-report/2568>

[32] RARI - Delegate Incentive Program v2.0

<https://forum.rari.foundation/t/rrc-xx-delegate-incentive-program-v2-0-for-rari-dao-governance/2398>

[33] StableLab governance operations transition to Anode (historical continuity note)

<https://forum.reserve.org/t/stablelab-transitions-to-anode/1440>

[34] Reserve - original Delegate Program Trial Run RFC

<https://forum.reserve.org/t/rfc-reserve-delegate-program-trial-run/1465>

THE CLOSING PRINCIPLE

DELEGATE THE MINIMUM DISCRETION NECESSARY.

AND PRESERVE A CREDIBLE PATH

to take it back.

The goal of a delegation system is not maximum delegation. It is to reduce the cost of collective decision-making without losing control of the underlying authority.

WE HELP GROUPS DECIDE - *then act.*

Decision systems, incentive design and coordination research.

THE DAO DELEGATION HANDBOOK

Research & writing / Enrique Alles Fernandez

START A CONVERSATION -> [ANODE.GG](https://anode.gg)