PRINTABLE EDITION

Sovereign Location — Design Space

Published by the Scintilla Foundation
6 April 2026

Contents

DESIGN SPACE

Why Type 6 Systems Matter

The broader taxonomy of Presence Adjudication Systems is useful because it shows that societies have always needed ways to convert observations about physical presence into judgments that others can rely upon. Witnesses, affidavits, inspectors, courts, databases, and platforms all belong to that larger field.

Once that landscape is visible, however, a more specific question emerges.

Which kinds of Presence Adjudication System are best suited to a world in which coordination is increasingly digital, multi-party, privacy-sensitive, and adversarial?

This section begins from the view that **Type 6 systems — decentralized economic systems — deserve particular attention**.

This is not because every question of presence should be handled by a Type 6 architecture. Many should not. Informal, institutional, and centralized systems will continue to play important roles, and often remain the right tools for their own contexts.

Type 6 systems matter for a different reason. They address a specific and increasingly important problem: how to make consequential claims of physical presence legible across organizational boundaries without requiring either blind trust in a single intermediary or broad disclosure of raw location history.

That problem is becoming more central, not less.

The Problem Type 6 Systems Address

Many traditional systems for adjudicating presence work tolerably well inside a single institutional boundary.

A company can rely on its own workflow system.
A court can hear testimony and review evidence.
An inspector can certify that a visit occurred.
A regulator can treat an official record as authoritative.

These are all real and often necessary forms of adjudication.

But they become less satisfactory when the environment changes. Digital society increasingly involves counterparties who do not fully trust one another, who may not share a common system of record, and who nevertheless need to coordinate around facts of physical presence.

- A logistics network may involve multiple firms.
- An event credential may need to be portable across systems.
- A location-conditioned payment may need to settle across institutional boundaries.
- A compliance-relevant presence claim may need to be contestable by parties outside the operator that recorded it.

In such settings, older PAS types begin to show their limits.

- Some are too local.
- Some are too slow.
- Some are too dependent on centralized authority.
- Some require excessive disclosure.
- Some produce records, but not judgments that are easily replayable or portable.

Type 6 systems matter because they are designed for this harder setting.

What Makes Type 6 Different

A Type 6 Presence Adjudication System does not treat presence as something established once and for all by a single authority. Nor does it treat platform telemetry as self-authenticating.

Instead, it typically introduces a different structure:

- multiple adjudicating actors rather than one exclusive operator
- explicit incentives rather than purely assumed honesty
- challenge or dispute mechanisms rather than silent finality
- durable publication of outcomes rather than private internal logs
- compatibility with bounded claims and privacy-preserving proofs rather than default overexposure

These are not implementation details. They are architectural commitments.

They change the shape of the problem from:

> “Which institution’s record do we trust?”

to something closer to:

> “Under what rules, incentives, and evidentiary constraints can a presence claim become reliable enough for others to use?”

That is a profound difference.

Why the Digital World Pushes in This Direction

The more digitally mediated coordination becomes, the more pressure there is to find evidentiary forms that are not tied entirely to local administrative control.

This is true for several reasons.

First, digital coordination scales faster than institutional trust. Systems can move value, permissions, or decisions across many parties almost instantly, but trust still tends to remain bounded by organizations, jurisdictions, and proprietary infrastructures.

Second, digitally native environments are often adversarial by default. It is no longer safe to assume that all relevant actors share incentives, share context, or will accept a private platform’s internal state as authoritative.

Third, privacy becomes harder to protect if adjudication depends on raw record disclosure. Once systems operate at scale and across institutions, the temptation to over-collect and over-share becomes structurally embedded.

Fourth, consequential digital coordination increasingly requires some form of finality. It is not enough to observe that a platform believes something happened. Other parties may need an outcome that is durable, replayable, and capable of supporting later scrutiny.

Type 6 systems matter in part because they are among the few PAS types built with this entire combination of pressures in view.

Why Not Just Use Type 3 Systems?

This is the most obvious challenge.

Why not simply use centralized digital systems more carefully? Why not improve enterprise logs, secure devices, or trusted platforms rather than introducing economic incentives, committees, disputes, and finalization layers?

The answer is that Type 3 systems remain useful, but they solve a different problem.

A centralized digital system is often excellent when one operator is entitled to define the outcome for its own workflow. A delivery platform can manage its own deliveries. An employer can manage its own attendance system. A telecom operator can maintain its own records. These systems may be entirely appropriate where the relevant relationships are vertical, closed, and institutionally bounded.

But they are weaker where:

- multiple parties need to rely on the same claim
- those parties do not all trust the same operator
- the claim may be contested later
- the outcome may carry financial or legal consequence
- overexposure of underlying traces is undesirable

In such settings, the problem is not simply record-keeping. It is adjudication under conditions of partial trust.

That is where Type 6 begins to justify its additional complexity.

Type 6 Is Not “More Decentralized Therefore Better”

The argument for Type 6 systems is not a general ideological preference for decentralization. It is not the claim that decentralized architectures are always superior, nor that traditional institutions become obsolete. It is certainly not the claim that every presence question should be solved on-chain or in a cryptoeconomic market.

The argument is narrower and stronger.

Type 6 systems are especially well suited to environments where presence claims must become:

- legible across institutional boundaries
- resistant to unilateral control
- contestable under explicit rules
- compatible with privacy-preserving proof structures
- durable enough to support downstream reliance

In other words, they are not better because they are “decentralized” in the abstract. They are better where the problem itself demands a more neutral, adversarially robust, and replayable evidentiary architecture.

The Cost of Type 6

Type 6 systems are not free improvements.

- They introduce design complexity.
- They require incentive engineering.
- They depend on parameter choices that can be subtle and brittle.
- They raise governance questions.
- They can fail if their capital structure is weak, if their dispute model is poorly designed, or if their finality surface is not credible.

This section does not ignore those difficulties. On the contrary, it focuses on them.

The point is not to romanticize Type 6 systems. It is to take them seriously enough to analyze the conditions under which they are actually good.

Conclusion

Type 6 systems matter because they confront a problem that older PAS types often handle poorly: how to adjudicate consequential claims of physical presence in digital environments where no single intermediary should be trusted to define reality for everyone else.

They are not the universal answer to presence adjudication. But they are the most important architectural family for anyone concerned with neutral, privacy-respecting, replayable, and economically accountable forms of digital coordination.

That is why this section turns toward them.

Not because the rest of the taxonomy no longer matters, but because understanding the wider field makes it possible to see where the deepest design challenge now lies.

Trust Models for Type 6 Presence Adjudication Systems

A Type 6 Presence Adjudication System exists because trust is a problem, not because trust disappears.

That point is easy to miss. Systems of this kind are often described as if they “remove trust,” or as though cryptography, staking, and dispute mechanisms somehow make institutional confidence unnecessary. In practice, the situation is more demanding. A serious Type 6 system does not eliminate trust. It restructures it.

That restructuring is the real subject of this page.

The question is not whether a Type 6 Presence Adjudication System relies on trust. It does. The question is **what kind of trust remains, where it resides, how it is constrained, and what happens when it fails**.

This matters because the promise of a Type 6 system is not that no assumptions are required. It is that the assumptions can be made more explicit, more distributed, more contestable, and less dependent on the unilateral authority of one platform or institution.

Trust Does Not Vanish

Every Presence Adjudication System depends on some combination of measurement, evidence, adjudication, and finalization. At each stage, some assumptions remain.

- Sensors may be honest or dishonest.
- Participants may behave sincerely or strategically.
- Verifiers may act independently or collusively.
- Dispute actors may be alert or inactive.
- Publication layers may be durable or weak.
- Governance may be neutral or captured.

A Type 6 architecture does not make these questions disappear. What it does is refuse to concentrate them entirely inside one institution’s internal record. It treats trust as something that should be shaped by explicit rules, adversarial incentives, bounded authority, and challengeable outcomes.

That is why trust models matter so much here. They are not a secondary implementation detail. They are part of the architecture itself.

From Trusted Authorities to Structured Trust

Traditional systems often rely on what might be called **authority trust**. A court, regulator, inspector, platform operator, telecom provider, or enterprise workflow owner is treated as the locus of reliable judgment.

Type 6 systems arise when this model is no longer sufficient.

The relevant problem is usually one in which:

- multiple parties need to rely on the same claim
- those parties do not all trust the same intermediary
- the claim may have financial or institutional consequence
- privacy matters
- the outcome may need to be replayed or contested later

In such settings, the question becomes not “who is the trusted authority?” but “how should trust be distributed, constrained, and exposed to challenge?”

A Type 6 system answers by replacing simple authority trust with a more structured arrangement involving some combination of:

- cryptographic integrity
- economic incentives
- distributed verification
- bounded authority
- challenge rights
- durable publication
- governance constraints

The resulting trust model is usually more complicated than the older one. But that complexity reflects the difficulty of the problem.

The Main Trust Surfaces

Measurement Trust

At the bottom of the system lies the question of observation. How does the claim enter the system at all?

Measurements may come from GNSS signals, radio environments, sensor reports, signed devices, witness observations, or hybrid sources. Even where proofs are used later, the system still depends on some relationship to the physical world.

Measurement trust therefore asks:

- are the observations authentic?
- are they fresh?
- are they resistant to spoofing or fabrication?
- do they reflect the relevant physical event rather than merely some device output?

This is one reason why cryptography alone is never the whole story. A proof can prove something about the inputs it was given. It cannot, by itself, guarantee that those inputs were generated honestly.

Prover Trust

The prover is the party attempting to establish the presence claim.

A good Type 6 system should not need to trust the prover’s word as such, but it will still need to reason about prover incentives and possible attack strategies. A prover may withhold information, attempt to fabricate evidence, exploit measurement weaknesses, or coordinate with corrupt adjudicators.

The trust question here is not “is the prover honest?” but:

- what can the prover gain from dishonesty?
- what evidence can the prover manufacture?
- what constraints make false claims hard or costly?
- what disclosure powers does the prover retain?

In a mature design, the prover should be assumed to be strategic, not saintly.

Verifier Trust

Verifier trust is often the most visible part of a Type 6 PAS.

A verifier may check proofs, evaluate evidence, apply protocol rules, sign outcomes, or participate in committee decisions. But verifiers are not simply abstract validators. They are economic and institutional actors. They may collude, free-ride, disengage, or become captured.

The relevant trust question is therefore not “do we trust the verifiers?” but:

- how many verifiers must fail for the system to fail?
- what incentives do they face?
- how visible is their misconduct?
- can dishonest verification be disputed?
- how replaceable are they?
- how concentrated can the verifier market become?

Type 6 systems matter because they attempt to answer these questions through structured design rather than leaving them implicit.

Watcher or Challenger Trust

Many Type 6 systems do not rely solely on verifiers. They also rely on parties who monitor outcomes and raise disputes when they see something wrong.

This introduces a distinct trust model: not trust in primary judgment alone, but trust in the availability of adversarial correction.

A watcher model assumes that not all errors or dishonest acts need to be prevented in advance, so long as they can be detected and challenged before finalization becomes irreversible.

This raises further questions:

- are watchers economically motivated to act?
- do they have access to enough information?
- how long do they have to respond?
- can they be censored or discouraged?
- what happens if no one watches?

A system that relies heavily on disputes but has no credible watcher economy is only weakly protected.

Finalization Trust

Even if a claim is properly evaluated, the system still needs a way for the outcome to become durable enough that others can rely upon it.

This introduces trust questions around finalization:

- where is the result published?

- when is it considered settled?
- what is the rollback risk?
- what later evidence can reopen the matter, if any?
- who controls the transition from pending to final?

A Type 6 PAS that has strong verification but weak finalization may produce technically elegant results that remain institutionally fragile.

Governance Trust

No sufficiently serious Type 6 system is free of governance.

Thresholds must be set.
Challenge windows must be chosen.
Slash conditions must be defined.
Committee rules must be established.
Upgrade paths must be controlled.

This means every Type 6 PAS contains some governance trust, whether acknowledged or not.

The important question is whether governance is narrow, explicit, reviewable, and institutionally legible — or whether it silently reintroduces the very unilateral authority the system claimed to overcome.

Common Type 6 Trust Models

Committee-Based Trust

In this model, a subset of verifiers is selected to adjudicate a claim or batch of claims. The core trust assumption is that the committee is sufficiently independent and sufficiently honest for the result to be credible.

This model can work well when:

- committee selection is hard to manipulate
- committee size is appropriate to the stakes
- collusion risk is bounded
- challenge rights remain available

Its weakness is that it can silently become oligarchic if the verifier set is too concentrated or the committee formation process is too predictable.

Stake-Weighted Trust

Here, trust is mediated through bonded economic exposure. Adjudicators are trusted not because they are presumed virtuous, but because they stand to lose something meaningful if they behave dishonestly.

This model is powerful because it links system security to capital at risk. But it also raises further questions:

- how much stake is really exposed?
- how quickly can dishonest gains be realized relative to slashing?
- can actors externalize losses?
- how concentrated is stake ownership?

Challenge-Based Trust

In challenge-based models, initial outcomes may be produced relatively cheaply, with the understanding that disputes can correct them before durable finalization.

This often improves scalability and efficiency, but it depends heavily on watcher incentives, evidence availability, and challenge timing. It works best where mistakes or fraud are likely to be observable and economically worth contesting.

Its danger is passive failure: a bad outcome may survive not because it was strong, but because nobody found it worthwhile to challenge.

Hybrid Cryptographic Trust

Some systems combine economic adjudication with stronger cryptographic subsystems.

In such systems, part of the trust burden is shifted away from human or institutional judgment and into formal proof systems. This can be highly valuable, but it should not be overstated. The system may still trust measurement inputs, rule design, challenge structures, or governance even if the proof layer itself is mathematically strong.

Federated Institutional Trust Inside Type 6 Systems

Some systems that look broadly Type 6 may still incorporate known institutional actors — licensed providers, approved measurement sources, regulated attestors, or designated publishers.

This can be sensible. It may improve operational quality, legal legibility, or onboarding. But it also changes the trust model. If too much authority flows back to designated institutional actors, the system may gradually drift toward Type 4 or Type 3 behavior even while preserving Type 6 language.

That is not always wrong. But it should be recognized clearly.

Trust Minimization Is Not the Same as Trust Distribution

A more distributed system is not automatically a less trust-dependent system.

A design may distribute roles across many actors and still leave critical assumptions untouched. It may even make trust harder to reason about if those assumptions become diffuse rather than explicit.

The real goal should not be trust minimization in the abstract. It should be **trust discipline**.

A disciplined trust model is one in which:

- the major assumptions are identifiable
- authority is bounded
- misconduct is visible
- incentives are aligned with honest behavior
- failure modes are understood
- correction paths exist
- governance does not silently dominate the system

That is a more useful standard than the vague claim that the system is “trustless.”

Evaluating a Type 6 Trust Model

A good trust model for a Type 6 PAS should be judged by questions such as these:

Dimension	Question
Explicitness	Are the key trust assumptions clearly visible?
Distribution	Is authority spread across actors, or quietly concentrated?
Incentive alignment	Do key actors lose meaningfully from dishonest behavior?
Contestability	Can bad outcomes be challenged by others?
Observability	Is misconduct visible enough to trigger correction?
Replayability	Can later parties understand how a judgment was reached?
Capture resistance	How hard is it for the system to be dominated by one interest?
Governance boundedness	Are governance powers narrow and legible?
Privacy compatibility	Can the trust model function without default overexposure?
Institutional portability	Can the result be used beyond one operator’s own system?

Conclusion

Type 6 Presence Adjudication Systems do not remove trust. They reorganize it.

Their significance lies in the attempt to move away from unilateral authority and toward a more explicit, distributed, challengeable, and economically disciplined evidentiary architecture. Whether they succeed depends on the quality of their trust model.

A system is not mature because it calls itself decentralized. It is mature when its remaining trust assumptions are visible, bounded, contestable, and proportionate to the role the system is meant to play.

Everything else in this section depends on that.

Privacy / Verifiability Tradeoffs

One of the most persistent assumptions in the design of presence systems is that privacy and verifiability stand in direct opposition.

The thought is simple and familiar. If a relying party wants stronger confidence, it must be shown more of the underlying data. If less is disclosed, confidence must fall. Privacy is purchased at the cost of weaker evidence; verifiability is purchased at the cost of greater exposure.

This assumption is understandable. In many systems, it is also true.

But it is not true in every system, and it is not a law of nature.

For Type 6 Presence Adjudication Systems, this distinction matters enormously. These systems exist precisely because the old architecture — broad data collection combined with unilateral institutional interpretation — is no longer adequate for many digitally mediated forms of coordination.

The real design question is therefore not whether privacy and verifiability are ever in tension. They are. The real question is **what kind of tension this is, where it arises, and how system design can change its shape**.

Why the Tradeoff Appears So Natural

The traditional logic of location systems makes the privacy / verifiability tradeoff appear obvious.

A location reading, route trace, check-in log, or timestamped record is treated as raw material from which confidence is later derived. If a dispute arises, the intuitive response is to ask for more of the record: more coordinates, more timestamps, more metadata, more retained history.

This is how many existing systems work. Confidence is increased by widening visibility.

But this architecture has a cost. The more verifiability depends on access to raw traces, the more every consequential presence claim tends to drag a larger behavioral record behind it. A narrow question becomes linked to broad exposure of movement patterns, timing, context, and often device-level metadata that far exceed the original claim.

That is why the tradeoff feels natural. The system is designed so that confidence grows through overexposure.

The Deeper Problem

This is not just a privacy problem. It is a problem of evidentiary form.

Most current architectures assume that the natural unit of evidence is the underlying data trace. Presence is then treated as an inference drawn from that trace. If the relying party doubts the inference, it asks to see more of the trace.

But that is only one way of structuring the problem.

In many contexts, the relevant question is not “show me everything from which I might infer what happened.” It is “show me that this bounded claim is valid under rules I can rely upon.”

That is a different evidentiary posture.

Once claims are framed in bounded form, the privacy / verifiability relationship changes. The task is no longer simply to hide data while preserving confidence. It is to build systems in which confidence attaches to the claim itself rather than to unrestricted inspection of the full underlying record.

Not All Verifiability Is the Same

Part of the confusion in this area comes from treating verifiability as though it were one thing.

It is not.

A relying party may want confidence in several different senses:

- confidence that the evidence has not been tampered with
- confidence that the evidence corresponds to a real event
- confidence that the claim satisfies a formal rule
- confidence that dishonest adjudicators can be challenged
- confidence that the outcome will remain durable over time

These are all forms of verifiability, but they do not all require the same kind of disclosure.

- Some can be improved through cryptographic integrity.
- Some depend on stronger measurement assumptions.
- Some depend on challenge mechanisms.
- Some depend on durable finalization.
- Some may genuinely require additional disclosure in edge cases.

A mature Type 6 PAS should therefore avoid speaking of verifiability as if it were a single scalar that increases only when privacy decreases.

The Wrong Frontier

Many weak systems accept what might be called the **old frontier**:

- high privacy means low confidence
- high confidence means broad disclosure

This frontier is real in badly designed systems. But it is not the only frontier available.

One of the central ambitions of a serious Type 6 PAS is to move to a different frontier, where confidence is improved not by exposing everything, but by changing the architecture of evidence and adjudication.

This may involve:

- expressing claims in bounded propositional form

- using proof systems that verify predicates rather than expose traces
- separating measurement from disclosure
- allowing disputes to operate on targeted evidence rather than default overexposure
- making finalization and auditability depend on explicit rules rather than institutional black boxes

The goal is not to deny the tradeoff. It is to redesign the system so that the tradeoff becomes less crude.

Bounded Claims Change the Landscape

A crucial move in this redesign is the shift from telemetry to claims.

A telemetry-oriented system asks: what do the coordinates say?

A claim-oriented system asks: what proposition needs to be established?

That proposition may be quite narrow:

- the prover was within a region during an interval
- the asset did not leave a controlled zone
- the participant crossed an event boundary
- the device was present at the site before a deadline

Once the claim is stated at that level, the system can ask a more disciplined question:

What is the minimum information necessary to make this claim usable?

That question is the real turning point.

Selective Disclosure Is Part of the Answer, Not the Whole Answer

Selective disclosure is important, but it is not sufficient on its own.

A system may reveal only a small amount of information and still be weakly verifiable if:

- the underlying measurement is doubtful
- the proof system is poorly designed
- verifiers are unaccountable
- disputes are impractical
- governance is overly discretionary
- finality is fragile

The real objective is therefore not selective disclosure in isolation, but **selective disclosure inside a credible evidentiary and adjudicative architecture**.

When More Disclosure Is Justified

A robust survey of this topic should not pretend that more privacy is always better in every case.

There are situations in which broader disclosure may be justified:

- when stakes are unusually high
- when a claim is challenged credibly

- when fraud patterns require deeper examination
- when legal or institutional due process demands more detailed evidence
- when the bounded claim itself is too coarse to resolve the dispute

The important point is not that broader disclosure never occurs. It is that broader disclosure should be **exceptional, justified, and structured**, not the default evidentiary baseline for every presence claim.

A mature Type 6 PAS should therefore distinguish between:

- ordinary proof mode
- challenge mode
- escalation mode
- exceptional or legal disclosure mode

Evaluating a Privacy / Verifiability Design

A good Type 6 PAS should be judged by questions such as:

Dimension	Question
Claim boundedness	Is the system designed around narrow propositions or broad telemetry?
Disclosure discipline	Does ordinary use reveal only what is necessary?
Proof adequacy	Can the disclosed evidence actually support the intended claim?
Challenge design	Can disputed claims be examined more deeply when needed?
Escalation control	Is stronger disclosure exceptional and rule-bound?
Privacy asymmetry	Who learns what, and is that distribution justified?
Measurement dependence	How much confidence still rests on hidden or trusted inputs?
Adjudicator accountability	Can verifiers or committees hide behind privacy claims of their own?
Finality compatibility	Can the system remain auditable without permanent overexposure?
Institutional usability	Can the resulting claim be relied upon by real counterparties and institutions?

Conclusion

Privacy and verifiability are often in tension, but they are not doomed to remain in the crude form inherited from older location architectures.

The old model ties confidence to broad visibility and treats overexposure as the ordinary cost of evidence. Type 6 systems matter because they make another possibility available: confidence attached to bounded claims, structured proofs, challengeable outcomes, and disciplined disclosure.

That does not abolish tradeoffs. It civilizes them.

And that is the real design goal: not to pretend that privacy and verifiability can always be maximized together, but to build systems in which their relationship is explicit, proportionate, and architecturally well governed.

Proof Architectures for Presence Adjudication

Not every Presence Adjudication System contains a sophisticated proof architecture.

A witness statement is a form of evidence, but it is not usually described as a proof architecture. A regulator's record may settle an issue, but the evidentiary form remains largely institutional rather than formally structured. A centralized digital platform may log location events, but often leaves the interpretation of those logs to internal systems and operator discretion.

Type 6 Presence Adjudication Systems are different. They operate in environments where multiple parties may need to rely on consequential claims of presence without sharing the same institution, trusting the same operator, or accepting broad disclosure of raw traces. In such settings, the architecture of proof becomes much more important.

This page is concerned with that architecture.

Its purpose is not to argue that proof systems alone solve the larger problem of presence adjudication. They do not. Proof is one part of a broader evidentiary and institutional structure. But it is an increasingly important part, because the shape of the proof architecture influences what kinds of claims can be expressed, what kinds of privacy can be preserved, what can be verified formally, and what remains dependent on trust, governance, or dispute.

PAS and PPS

A useful distinction is needed at the outset.

A **Presence Adjudication System (PAS)** is the broader system by which presence claims become socially and institutionally usable. It includes not only evidence formation and verification, but also adjudication, dispute, finalization, and the rules by which outcomes become durable enough for others to rely upon.

A **Presence Proof System (PPS)** is narrower. It is the evidentiary subsystem that transforms measurements or observations into a verifiable claim about presence.

In simplified form:

- **PPS** answers: how is the claim proven?
- **PAS** answers: how does the claim become relied upon?

This distinction matters because a system may have a strong PPS and still have a weak PAS. It may produce elegant proofs, yet remain fragile in dispute, poorly governed, or vulnerable to finality failure.

Why Proof Architecture Matters

A presence claim is not self-interpreting.

To say that a person, device, or asset was within a region during an interval is already to move from the physical world into a structured proposition. The question is how that proposition is supported.

In weak systems, support often takes the form of raw traces, operator logs, screenshots, manually assembled records, or private database entries. These may be usable, but they do not provide much formal discipline. They usually reveal more than necessary, depend heavily on institutional trust, and become awkward under dispute.

A proof architecture matters because it changes the evidentiary form of the claim.

Instead of asking a relying party to inspect broad telemetry and infer what happened, the system can ask a narrower question:

> can this proposition be shown to hold under explicit rules, using evidence that is verifiable at the level the claim requires?

That shift is foundational.

The General Shape of a Presence Proof System

A Presence Proof System can be understood as a chain with several stages.

Stage	Function
Observation	Physical signals or observations enter the system
Claim formation	The relevant proposition is defined
Evidence construction	Observations are transformed into a provable evidentiary form
Proof generation	A proof or structured evidence object is produced
Proof verification	Another party checks that the claim follows under the relevant rules

Not every PPS makes these stages equally explicit, and not every system draws the boundaries in the same way. But this progression is useful because it shows where architectural choices arise.

A system may differ in:

- how observations are trusted
- how claims are bounded
- how much data must be revealed
- whether the proof is public or private
- whether the proof is purely formal or partly institutional
- whether verification is deterministic or judgment-laden

These differences define the proof architecture.

Major Architectural Families

Trace-Based Proof

In trace-based architectures, the proof is effectively the trace itself, or a large portion of it.

A prover discloses raw location data, timestamps, movement history, or device logs, and the relying party inspects these materials to infer whether the claim is satisfied.

This is still the dominant pattern in many real-world systems because it is operationally straightforward and easy to understand. But it has serious weaknesses. It over-discloses by default, scales poorly under dispute, and often leaves interpretation inseparable from institutional judgment.

It is best thought of as a low-maturity evidentiary architecture: sometimes useful, often unavoidable, but structurally weak as a long-term foundation.

Attestation-Based Proof

In attestation-based architectures, a trusted party or trusted device asserts that the claim is true or that relevant measurements were observed.

Examples may include:

- signed device attestations
- secure hardware reports
- witness attestations
- inspector certifications
- trusted measurement providers

This architecture can improve integrity and reduce the need to expose raw traces directly. But it shifts the trust burden onto the issuer of the attestation. The proof becomes strong only to the extent that the attesting party or hardware root is itself trusted.

Attestation-based systems are therefore often useful components, but rarely complete answers on their own.

Predicate-Based Cryptographic Proof

In predicate-based architectures, the system proves not the underlying data itself, but that the data satisfies a defined condition.

This is where cryptographic proof systems become especially important. Instead of revealing all coordinates or measurements, the prover may establish that:

- the measurements are consistent with being within a region
- the claim falls within a time window
- a path crossed a threshold
- a device did not leave a zone during a bounded interval

This architecture is central to the privacy-preserving ambitions of Type 6 PAS because it allows the evidentiary object to be the claim rather than the trace.

Commitment-and-Reveal Architectures

In these architectures, underlying measurements or claim data are first committed to in a binding form and only selectively revealed later if needed.

This can be useful where:

- the system wants to preserve evidence without exposing it immediately
- disputes may require later escalation
- challengers need confidence that hidden data existed at the relevant time
- finalization depends on proving that records were not altered after the fact

Commitments can strengthen integrity and support future dispute resolution, but they do not by themselves prove that the committed data is true.

Hybrid Architectures

Most serious systems are likely to be hybrid.

A mature PPS may combine:

- signed measurement sources
- commitments to raw observations
- predicate proofs for ordinary adjudication
- selective reveal under dispute
- verifier checks for rule compliance

The important question is not whether the architecture is pure. It is whether the different layers fit together coherently.

The Central Design Tension

Proof architectures are shaped by a central tension:

Should the system prove the data, or prove the claim?

A system that proves the data tends to maximize inspectability at the cost of privacy and boundedness.

A system that proves the claim tends to support better disclosure discipline, but also demands more sophistication in proof construction, clearer claim semantics, and stronger confidence that the underlying measurements were handled correctly.

Type 6 systems matter in part because they are among the few PAS types capable of making the second path viable at scale.

Public Proofs, Private Proofs, and Layered Disclosure

Not all proofs are disclosed in the same way.

A proof architecture may be:

- **public**, meaning the relevant proof object can be checked by any party with access to it
- **private**, meaning only designated parties can verify or interpret the proof
- **layered**, meaning ordinary adjudication uses bounded proof, while deeper evidence remains available only under challenge or escalation

Layered disclosure is especially important for Type 6 systems because it avoids a false choice between “show everything” and “show nothing.”

Proof Architecture Does Not Remove Measurement Risk

A strong proof architecture does not make a presence claim trustworthy all the way down.

A proof system may perfectly establish that a claim follows from a set of inputs. But if those inputs were fabricated, spoofed, or poorly grounded in physical reality, the overall evidentiary result remains weak.

This is why proof architecture must always be understood in relation to:

- measurement trust
- prover incentives
- verifier incentives
- dispute design
- governance
- finality

A PPS is not the whole PAS.

Evaluating a Presence Proof Architecture

A PPS should be judged by questions such as these:

Dimension	Question
Claim boundedness	Does the system prove a narrow proposition or require broad data exposure?
Measurement dependence	How much trust is still placed in hidden inputs or attestors?
Privacy discipline	How much is revealed in ordinary use?
Proof soundness	Does the proof actually establish the intended claim?
Verification cost	How expensive is it to check the proof?
Dispute compatibility	Can the architecture support challenges and escalation?
Replayability	Can later parties understand what was proven and under what rules?
Portability	Can the proof travel across institutional boundaries?
Adjudicative fit	Is the proof usable inside a broader PAS, not just elegant in isolation?
Composability	Can the proof architecture coexist with other evidentiary layers?

Conclusion

A Presence Proof System is not the whole of a Presence Adjudication System. But it is one of the places where a system’s evidentiary philosophy becomes visible.

Weak architectures treat presence as something to be inferred from broad records after the fact. Stronger architectures aim to make bounded claims directly provable under explicit rules, with disclosure disciplined to the needs of the claim rather than the appetites of the observer.

That is the deeper significance of proof architecture in this field.

It is not only about formal correctness. It is about deciding what kind of evidentiary object a presence claim should become.

And for Type 6 systems, that choice is fundamental.

Finality Surfaces

A Presence Adjudication System is not complete when it produces an opinion.

It becomes complete, in the stronger sense, when it produces an outcome that others can rely upon.

That distinction is fundamental. A verifier may believe a claim is valid. A committee may sign a result. A challenge window may expire. A record may be published. A blockchain may confirm an entry. But none of these acts is identical to the question that ultimately matters:

> when, where, and in what sense does this presence claim become final?

This page is about that question.

In a Type 6 Presence Adjudication System, finality is not a single event. It is usually the product of several layers: evidentiary acceptance, dispute closure, publication, and durable reliance. The concept of a **finality surface** is useful because it helps describe the point at which a claim passes from being merely evaluated to being sufficiently settled that other systems, institutions, or counterparties can build on it.

Why Finality Matters

Presence claims matter because they affect other decisions.

- A payment may depend on them.
- A credential may depend on them.
- A compliance process may depend on them.
- An access right, a penalty, or a downstream state transition may depend on them.

In each case, another system needs to know not merely that a claim has been examined, but whether it can now be treated as settled enough to act upon.

If finality is weak, then everything built on top of the claim remains fragile. A settlement may need to be reversed. A dispute may reopen. A relying party may discover too late that the supposedly durable result was only provisional.

What a Finality Surface Is

A **finality surface** is the point in the system at which a presence adjudication outcome becomes durable enough for some defined class of reliance.

The phrase is useful because finality is rarely absolute. Different parties may rely on a result at different thresholds. An application may treat a claim as actionable before a court would regard it as conclusively settled. A protocol may treat a challenge window as closed while governance still retains a narrow emergency override.

So the finality surface is not just “the moment it is final.” It is the surface at which the system itself says:

> beyond this point, this outcome is durable enough for these purposes.

Finality Is Not the Same as Verification

A claim may be correctly verified and still not be final.

Verification answers a question like:

> does this evidence satisfy the relevant rule?

Finality answers a different question:

> has the system reached a state in which this result can now be relied upon without ordinary expectation of reversal?

These questions are connected, but they are not the same.

- A system may verify quickly and finalize slowly.
- It may produce a provisional result pending dispute.
- It may publish a signed committee outcome that is still challengeable.
- It may record a result durably in one layer while leaving open another path of contestation.

This is why finality must be treated as its own architectural problem.

The Main Components of Finality

Evidentiary Finality

This is the point at which the system considers the claim sufficiently established at the level of evidence.

For example, a verifier committee may have reached threshold agreement, or a proof may have been accepted as valid under the protocol rules.

Dispute Finality

This is the point at which the ordinary window for challenges has closed, or at which the dispute process has otherwise been exhausted.

Dispute finality is especially important in Type 6 systems because many of them rely on watcher or challenger models.

Publication Finality

This is the point at which the result has been recorded in a durable medium from which later parties can retrieve and inspect it.

In some systems, this may involve a blockchain. In others, it may involve a signed public record, content-addressed publication, or another form of durable shared state.

Reliance Finality

This is the point at which downstream actors are justified in building on the result.

A claim may be evidentially accepted and even durably published, yet some downstream systems may still choose to wait for stronger assurance before acting.

Why “Surface” Is Better Than “Point”

The word *point* can be misleading because it suggests a sharp instant at which uncertainty vanishes.

That is not how most real systems behave.

Finality is often layered, threshold-based, and purpose-relative. It may arrive in stages:

- first the claim is accepted
- then the dispute window closes
- then the result is published durably
- then downstream systems treat it as settled

Calling this a **surface** rather than a **point** helps readers see that finality is a boundary condition in a larger architecture, not merely a timestamp.

Common Finality Models in Type 6 Systems

Immediate Internal Finality

A result is treated as final as soon as the designated adjudicators have accepted it.

This model is simple and fast, but often weak. It leaves little room for challenge and may provide limited confidence when the adjudicators themselves are the main source of risk.

Challenge-Window Finality

A claim becomes final only after a defined challenge period has passed without successful dispute.

This is one of the most natural models for Type 6 PAS because it aligns well with watcher-based security.

Layered Publication Finality

Internal adjudication may occur first, but stronger finality is attached to later publication in a durable shared medium.

This model is often attractive because it separates adjudication from finalization.

Economically Conditioned Finality

Some systems tie finality not merely to elapsed time or publication, but to economic exposure.

A result may be treated as final once:

- the challenge window has closed
- the adjudicators’ bonded risk has expired or settled
- the slashing conditions are no longer live
- and downstream actors know what capital backed the outcome

Exceptional Override Finality

Some systems preserve a narrow exceptional path by which even apparently final outcomes can be revisited under emergency or governance conditions.

This may be prudent in some institutional settings, but it weakens the purity of finality and must be handled with care. If override powers are too broad, finality becomes mostly rhetorical. If they are too hidden, the system appears stronger than it really is.

The existence of override paths should always be treated as part of the finality model, not as an afterthought.

Finality and Trust

Finality surfaces are deeply connected to trust models.

A system's finality is only as credible as the assumptions on which it rests. If finality depends on:

- Watcher vigilance, then watcher incentives matter.
- A publication layer, then that layer's durability and neutrality matter.
- Governance restraint, then governance boundedness matters.
- Stake-backed deterrence, then the real economic exposure matters.

This is why finality cannot be discussed in isolation.

Finality and Privacy

Finality also interacts with privacy in subtle ways.

A badly designed system may achieve durable finality only by durably publishing too much information. A result becomes replayable, but only because the underlying claim has been overexposed.

A better system aims for something more disciplined:

- the outcome is durable
- the relevant proof or adjudication path is replayable
- but raw behavioral traces do not become permanently exposed unless truly necessary

The key question is:

> what exactly becomes durable, and at what level of revelation?

Evaluating a Finality Surface

A Type 6 PAS should be judged by questions such as these:

Dimension	Question
Clarity	Is it clear when a result becomes final, and for what purposes?
Layering	Are the stages of evidentiary, dispute, publication, and reliance finality distinguishable?
Replayability	Can later parties understand what became final and why?

Dimension	Question
Challenge compatibility	Does the finality model preserve meaningful dispute rights before closure?
Durability	Is the finality surface anchored in a medium others can rely upon over time?
Privacy discipline	Does finality require permanent overexposure of underlying traces?
Economic credibility	If finality relies on stake or slashing, is the backing meaningful?
Governance boundedness	Can final results be reopened too easily by discretionary power?
Institutional usability	Can real counterparties and institutions rely on the finality threshold?
Latency	How long must parties wait before reliance is prudent?

Conclusion

A Type 6 Presence Adjudication System does not become important merely because it can evaluate a claim. It becomes important when it can bring that claim to a form of closure that others can rely upon.

That is the role of a finality surface.

It marks the boundary at which presence adjudication becomes durable enough to support downstream coordination — not in the abstract, but in a specific, legible, and institutionally usable sense.

A mature system does not blur this boundary. It defines it carefully.

Security-Capital Surfaces

A Type 6 Presence Adjudication System is never secured by procedure alone.

It may have committees, proofs, challenge windows, durable publication, and explicit rules. All of these matter. But if the system ultimately relies on economically exposed actors to adjudicate claims, defend outcomes, or challenge misconduct, then its security is shaped not only by formal design, but by capital.

That is the subject of this page.

The phrase **security-capital surface** refers to the boundary between what a system can safely adjudicate and what its economically exposed structure can actually defend. It is a way of asking, in concrete terms:

> how much consequence can this system safely bear before dishonest behavior becomes rational?

This is one of the central questions for Type 6 PAS design. A system may appear orderly, cryptographically sophisticated, and procedurally complete while remaining economically weak. If the rewards from corrupting an outcome exceed the credible losses for those who would need to collude, then the security of the system is thinner than its surface presentation suggests.

That does not mean capital is everything. But it does mean that for Type 6 systems, capital exposure is part of the architecture of trust, finality, and adjudication.

Why Capital Matters

A Type 6 PAS differs from older forms of presence adjudication because it does not rely solely on a recognized authority to define the result. Instead, it often relies on verifiers, challengers, publishers, or related actors whose credibility comes partly from the fact that they are exposed to incentives and penalties.

This is an important shift.

- In a court, legitimacy may derive from institutional authority.
- In a platform, control may derive from operational ownership.
- In a Type 6 system, credibility often depends on whether economically relevant actors have enough to lose from misbehavior.

That means security cannot be discussed only in terms of formal process. It must also be discussed in terms of exposed downside.

Security Is Always Relative to Stakes

No adjudication system is secure in the abstract.

It is only secure relative to:

- the value of the outcomes it governs
- the incentives faced by potential attackers
- the cost of collusion
- the chance of detection
- the speed and effectiveness of challenge
- the severity and credibility of penalty

This is especially important for presence systems because the direct on-system value may not reflect the real-world value riding on the claim.

A small on-chain fee may be attached to a presence claim that determines a much larger insurance outcome, logistics release, regulatory consequence, access right, or contractual settlement.

So the relevant security question is not:

> how much value sits inside the protocol?

It is:

> how much value depends on this adjudication outcome, and what would it take to corrupt it?

That wider consequence surface is what makes security-capital analysis essential.

What a Security-Capital Surface Is

A **security-capital surface** is the effective frontier at which a system's economically exposed structure remains credible relative to the value and attack incentives associated with the claims it adjudicates.

Put more simply, it describes the range within which the system's capital-backed deterrence is still stronger than the gains from corruption.

This surface is shaped by several factors:

- how much stake or bonded capital is genuinely slashable
- how quickly dishonest gains can be realized
- whether the relevant actors can externalize losses
- whether watchers are incentivized strongly enough to challenge
- how visible misconduct is
- how long the dispute window remains open
- whether governance can quietly neutralize penalties
- whether real-world value exceeds protocol-visible value by a large margin

A system with a shallow security-capital surface can still function. It simply cannot safely govern high-consequence outcomes.

The Difference Between Nominal and Effective Security

One of the most important distinctions here is between **nominal security** and **effective security**.

Nominal security is what the system appears to have on paper:

- total stake bonded
- advertised slashing rules
- committee thresholds
- published dispute processes
- formal challenge rights

Effective security is what remains after realistic attack conditions are considered.

For example:

- not all nominal stake may be meaningfully slashable
- some participants may be closely aligned or controlled by the same actor
- stake may be borrowed, insured, or externally hedged
- governance may be able to soften penalties
- watcher participation may be sparse
- hidden off-system value may make corruption more attractive than it appears
- adjudicators may gain more from collusion than they fear from punishment

This means visible bonded capital is only the beginning of the analysis. The deeper question is what portion of that capital is truly exposed to credible loss under realistic conditions.

Security Depends on Detection, Not Just Penalty

Capital-backed deterrence only works if misconduct can be detected in time and proven in a way that triggers the relevant penalties.

This is why security-capital surfaces are inseparable from dispute architecture.

A system may advertise large slashable stakes, but if:

- dishonest claims are difficult to observe
- challengers are poorly incentivized
- evidence is too hidden to support timely dispute
- challenge windows are too short
- adjudicator misconduct is hard to attribute

then the capital may remain mostly decorative.

The real security of the system is therefore a product of both:

- **penalty magnitude**
- **penalty realizability**

An uncollectable penalty is not strong deterrence.

The Basic Security Envelope

A useful design intuition is that every Type 6 PAS has a practical **security envelope**.

This envelope describes the class of claims and consequences the system can adjudicate without making profitable corruption too easy.

Inside the envelope:

- honest behavior is more attractive than dishonest collusion
- disputes can realistically correct bad outcomes
- the relevant capital is large enough and exposed enough to deter attack
- finality remains credible

Outside the envelope:

- the rewards from corruption may exceed the credible downside
- watchers may not be sufficiently motivated
- collusion may become rational
- finality may be only performative

Not every PAS must aim for the largest possible envelope. But every serious PAS should know roughly where its envelope lies.

What Expands or Shrinks the Surface

Several design choices affect the strength of a system's security-capital surface.

More Meaningfully Exposed Capital

The most obvious factor is the amount of capital that can genuinely be lost through dishonest participation.

But what matters is not gross bonded capital. It is the portion that is:

- actually at risk
- rapidly slashable
- not easily shielded or externalized
- distributed across actors whose failure is not perfectly correlated

Meaningful capital expands the security surface. Decorative capital does not.

Better Challenger Economics

A system with strong watcher incentives is often much stronger than a system with nominally larger bonded capital but weak dispute incentives.

This is because challengeable systems do not need to prevent every bad act in advance. They need bad acts to be visible, contestable, and punishable often enough that corruption remains unattractive.

Better Claim Boundedness

Bounded claims can sometimes improve security because they make adjudication and challenge more legible. If claims are too vague, too broad, or too context-dependent, then dishonest outcomes become harder to detect and punish.

A more disciplined claim model can therefore expand the effective security surface even without increasing nominal stake.

Better Finality Design

If finality arrives too quickly, dishonest outcomes may become hard to reverse before dispute can operate.

If finality arrives too slowly, honest actors may not find participation worthwhile.

The finality model therefore affects the security-capital surface by shaping how long the system has to detect and punish bad outcomes before reliance hardens.

Better Governance Constraint

A system may appear well-capitalized and well-slashed on paper, yet remain weak if governance can quietly alter challenge rules, soften penalties, exempt favored actors, or otherwise interfere with enforcement.

Governance constraint therefore strengthens the security-capital surface by making penalties more credible.

Hidden Value and Off-Protocol Incentives

One of the hardest problems in this area is that the value secured by a PAS may be much larger than the value visible inside it.

This is especially true for presence systems.

A corrupt presence adjudication may release goods, trigger insurance, unlock a milestone payment, alter legal posture, or satisfy a compliance condition. The on-system fees associated with that claim may be tiny compared with the off-system value at stake.

This means attack incentives can come from outside the protocol entirely.

A verifier may collude not for protocol-native rewards, but for some external payment. A challenger may stay silent not because it is irrational on-chain, but because it is compromised off-chain. A committee may accept a bad claim because the real economic stake is elsewhere.

This is why Type 6 PAS cannot be evaluated purely in token-internal terms.

Security-Capital Surfaces and Use-Case Discipline

A mature system should not treat all use cases as equivalent.

Different presence claims expose the system to different incentive environments. A low-stakes event attendance proof is not the same as a high-value logistics release. A soft reputation signal is not the same as a legally consequential compliance outcome.

This implies that a good Type 6 PAS may need:

- claim classes
- risk tiers
- differentiated security requirements
- different finality thresholds
- caps on consequence exposure
- parameter scaling by claim type

This is not a weakness. It is evidence of design maturity.

Evaluating a Security-Capital Surface

A Type 6 PAS should be judged by questions such as these:

Dimension	Question
Exposed capital	How much capital is genuinely slashable or otherwise at risk?
Concentration	How correlated is control over that capital?
Detection quality	How likely is dishonest adjudication to be observed in time?
Challenge economics	Are challengers paid enough and empowered enough to act?
Attack latency	Can dishonest gains be realized before penalties land?

Dimension	Question
Off-system value	How large are the external incentives to corrupt outcomes?
Governance interference	Can governance weaken penalties or shield actors?
Claim boundedness	Are claims legible enough that bad outcomes can be challenged clearly?
Finality discipline	Does the finality model preserve time for economically meaningful correction?
Use-case fit	Is the system being asked to secure more consequence than its capital can defend?

Conclusion

A Type 6 Presence Adjudication System is secure only within the range of consequences its economically exposed structure can credibly discipline.

That range is its security-capital surface.

To design such a system seriously is therefore to ask not only whether claims can be proven, verified, and finalized, but whether dishonest adjudication remains irrational at the stakes that matter.

That is a harder question than most systems first admit.

But it is also the right question.

Dispute Models

A Type 6 Presence Adjudication System is not secured only by correct initial judgment.

It is secured by the possibility of correction.

That is the starting point of this aspect of design. In systems where presence claims may carry economic, legal, or institutional consequence, it is never enough to say that verifiers will simply evaluate claims honestly the first time. A serious system must also ask what happens when they do not. It must ask what happens when claims are mistaken, fraudulent, collusive, ambiguous, adversarial, or strategically manipulated.

That is the role of dispute.

A dispute model is the part of a Presence Adjudication System that defines how challenged claims are reopened, who may challenge them, on what grounds, under what timing, with what evidence, at what cost, and with what consequences.

Without dispute, a system may still verify. It may still publish. It may still finalize. But it cannot convincingly claim to be adversarially robust.

Why Dispute Matters

Disputes matter because verification is not omniscience.

Even a system with strong proof architecture, bounded claims, and economically exposed verifiers can still fail. Measurement inputs may be fabricated. Claims may be badly formed. Proofs may be valid with respect to dishonest inputs. Verifiers may collude. Watchers may observe inconsistencies only after initial adjudication. Real-world context may reveal that a formally accepted claim was substantively unsound.

A system designed only for undisputed cases is not yet a serious adjudication system.

The true test comes when:

- incentives diverge
- facts are contested
- counterparties disagree
- the stakes are high enough that dishonesty becomes attractive

In such settings, dispute is part of the architecture by which the system remains credible.

What a Dispute Model Does

A dispute model answers several interlocking questions.

It defines:

- **who may challenge** a claim or outcome

- **what may be challenged**
- **when** a challenge is still admissible
- **what burden** the challenger must meet
- **what evidence** can be introduced on challenge
- **who adjudicates the dispute**
- **what penalties** attach to bad behavior
- **what effect** the dispute has on finality

A system with a nominal challenge function but no economically viable path to use it does not really have a dispute model. It has a symbolic gesture toward one.

Dispute Is Not the Same as Governance

A dispute model concerns the adjudication of particular claims or outcomes under existing rules.

Governance concerns the power to alter the rules themselves.

If this distinction collapses, the system becomes unstable in an important way. Instead of saying, “this claim is disputed under the current protocol,” the system begins to say, “this outcome may be changed if powerful actors decide to intervene.”

That is not dispute in the architectural sense. It is discretionary override.

A mature Type 6 PAS should therefore treat disputes as rule-bound processes internal to the adjudication architecture, not as occasions for ad hoc governance rescue.

What Can Be Disputed?

Different systems may allow challenge to different layers of the adjudication stack.

For example, a challenger may dispute:

- the validity of a proof
- the integrity of an attestation
- the admissibility of a claim
- the conduct of verifiers
- the composition of a committee
- the factual basis of an observation
- the interpretation of protocol rules
- the finalization of an outcome

These are not all the same kind of dispute.

Some are **formal disputes**, where the issue is whether the evidence satisfies explicit technical predicates.

Some are **procedural disputes**, where the issue is whether the system’s own rules were followed correctly.

Some are **substantive disputes**, where the issue is whether the accepted claim corresponds to the real-world event it purports to establish.

A mature PAS should be clear about which of these it supports directly, which it supports only partially, and which it leaves to external institutions.

Timing: When Must a Challenge Be Raised?

Disputes are inseparable from time.

A challenge that can be raised forever may destroy usability. A challenge window that is too short may destroy meaningful security. A challenge process that exists only in theory, but expires before a watcher could realistically gather evidence, does not provide much real correction.

A dispute model must determine:

- when a claim becomes challengeable
- how long the challenge window remains open
- whether different classes of claims have different windows
- whether stronger evidence can justify later reopening
- when finality hardens beyond ordinary dispute

These choices affect not only security, but also capital efficiency, operational latency, and downstream reliance.

Who Gets to Challenge?

A dispute model may allow challenge by:

- any participant
- designated watchers
- economically bonded challengers
- affected counterparties
- verifiers or committee minorities
- trusted institutional actors
- some combination of the above

Each model has implications.

An open challenge model may increase adversarial robustness, but can invite spam or griefing unless challenge costs are well designed.

A designated watcher model may improve coordination, but creates dependence on a narrow monitoring set.

A counterparty-only model may reduce noise, but can miss fraud that no directly affected party has sufficient incentive or information to challenge.

Burden of Challenge

Not every challenge should be equally easy.

If challenges are costless and unconstrained, dispute can become a denial-of-service vector. If challenges are too expensive or too burdensome, correction becomes mostly theoretical.

A serious dispute model therefore needs a carefully chosen burden of challenge.

This may include:

- a stake or bond posted by the challenger
- a minimum evidentiary threshold
- a requirement to identify a specific procedural or substantive defect
- penalties for frivolous or malicious disputes
- differentiated thresholds by claim class

The aim is not to discourage challenge in general. It is to make challenge **serious**.

Evidence Under Dispute

A system's dispute model reveals a great deal about what it truly believes counts as evidence.

Some systems permit only the original proof object to be re-examined.

Others allow:

- selective reveal of committed data
- counter-evidence from challengers
- new attestations
- procedural audit records
- committee transcript or signature review
- measurements previously hidden in ordinary adjudication mode

A good dispute architecture distinguishes between:

- **ordinary evidentiary mode**
- **challenge mode**
- **escalation mode**

Without that layered structure, privacy and dispute often end up working against one another.

Who Adjudicates the Dispute?

Several models are possible.

Internal Re-Adjudication

The original verifier set, or a subset of it, re-examines the claim.

This is simple, but may be weak if the original concern is verifier misconduct or collusion.

Expanded Committee Review

A new or larger committee reviews the challenged claim.

This can improve independence, but may increase latency and cost.

Challenger / Defender Structured Contest

The dispute becomes an adversarial proceeding in which the original outcome is defended and the challenger must prove a defect.

This is often attractive in systems where formal challenge grounds can be expressed clearly.

Escalation to Specialized Adjudicators

Some systems may have a distinct dispute layer or specialized dispute committees.

This can improve expertise, but also risks creating a second authority center whose own incentives and capture risks must be examined directly.

External Institutional Escalation

At some point, some disputes may leave the system entirely and enter courts, regulators, or contractual processes.

A Type 6 PAS should be honest about where that boundary lies.

Dispute and Economic Discipline

A dispute model is only strong if dispute outcomes matter.

That usually means some combination of:

- slashing dishonest verifiers
- penalizing false attestations
- rewarding successful challengers
- penalizing frivolous disputes
- reversing or voiding bad outcomes
- delaying or invalidating downstream reliance

A mature dispute model must therefore be economically balanced enough that:

- honest challenge is worth bringing
- dishonest adjudication is worth avoiding
- frivolous challenge is worth discouraging

Common Dispute Failure Modes

Several failure modes recur in immature systems.

Symbolic Challenge Rights

The protocol allows dispute in principle, but watchers lack the information, time, or incentives needed to use it.

Excessive Friction

Challenges are so expensive or procedurally complex that only obvious fraud is ever contested.

Frivolous Griefing

Challenge is so cheap or weakly filtered that it becomes a source of delay, harassment, or denial of service.

Collusive Closure

The parties meant to adjudicate disputes are too aligned with the original decision-makers to provide real correction.

Governance Substitution

Instead of rule-bound dispute, the system relies on discretionary intervention whenever important cases arise.

Privacy Collapse Under Challenge

The system is “privacy-preserving” only until the first meaningful dispute, at which point routine challenge effectively requires full behavioral exposure.

Evaluating a Dispute Model

A Type 6 PAS should be judged by questions such as these:

Dimension	Question
Accessibility	Can legitimate challengers actually use the dispute process?
Timing adequacy	Is the challenge window long enough to permit real review?
Evidentiary depth	Can the system examine enough material under dispute to correct bad outcomes?
Economic balance	Are honest challenges rewarded and frivolous ones discouraged?
Independence	Are dispute adjudicators meaningfully distinct from those whose decisions are under challenge?
Privacy discipline	Does dispute preserve bounded disclosure as far as possible?
Finality compatibility	Does dispute fit coherently with the system’s finality model?
Governance boundedness	Are disputes handled under rules rather than ad hoc override?
Procedural clarity	Is it clear what can be challenged, by whom, and on what grounds?
Correction power	Can a successful dispute actually alter the outcome in a meaningful way?

Conclusion

A Type 6 Presence Adjudication System is credible not only because it can adjudicate, but because it can be challenged.

That is the role of a dispute model.

It converts the possibility of error, fraud, or collusion from a fatal weakness into a design problem: who may object, how, when, with what evidence, and with what consequences.

A mature system answers those questions clearly. It does not hide them in procedure, defer them to discretion, or pretend they are edge cases.

Because where presence claims matter, dispute is not an afterthought. It is one of the conditions under which finality, security, and trust become believable.

Governance and Parameter Control

No serious Type 6 Presence Adjudication System is free of governance.

- Thresholds must be chosen.
- Challenge windows must be set.
- Committee rules must be defined.
- Admissible claim types must be specified.
- Reward and slashing parameters must be calibrated.
- Upgrade paths must be controlled.
- Emergency powers, if any, must be bounded.

This is not a flaw. It is a fact of system design.

The mistake is not that governance exists. The mistake is to pretend that it does not, or to treat it as though it were separate from the trust model, dispute model, and finality model of the system.

The central question is not whether governance is present. It is **what governance is allowed to control, how that control is exercised, how visible it is, and whether it quietly undermines the neutrality the system claims to provide**.

Why Governance Matters

A Type 6 PAS exists because no single operator should be able to define presence claims for everyone else by unilateral fiat.

But that ambition can be weakened in a quieter way.

A system may distribute verification, introduce disputes, require stake, publish outcomes durably, and still leave decisive power concentrated in a governance layer that can:

- alter rules opportunistically
- change thresholds after the fact
- weaken penalties
- protect favored actors
- reopen final outcomes selectively
- narrow or widen admissible evidence at will

When that happens, decentralization at the adjudication layer may be mostly superficial. The system may look distributed operationally while remaining centralized constitutionally.

That is why governance must be examined directly. A mature PAS should not only ask how claims are adjudicated, but also who governs the machinery of adjudication itself.

Governance Is Not Adjudication

Adjudication concerns the handling of particular claims under existing rules.

Governance concerns the power to define, modify, interpret, or suspend those rules.

This distinction matters because systems often appear more neutral than they really are if governance powers are kept offstage. A protocol may say that claims are evaluated objectively, while leaving it implicit that the claim format, dispute timing, verifier set, risk thresholds, slashing conditions, or publication rules can all be changed by some smaller authority center.

A mature Type 6 PAS should therefore make governance powers explicit and bounded.

Parameter Control Is a Form of Power

Much of governance in real systems takes the form of parameter control.

This can sound technical and innocuous, but it is often one of the most important forms of authority in the entire architecture.

Parameters determine things such as:

- how much stake is required
- how committees are formed
- how large committees must be
- how long challenge windows remain open
- what counts as valid evidence
- when finality is reached
- what penalties apply to dishonest actors
- what rewards are paid to challengers
- which claim classes are permitted
- which risk tiers exist
- whether emergency overrides are available

To control these parameters is not merely to tune performance. It is to shape the trust, privacy, finality, and security properties of the system itself.

Some Parameters Are More Dangerous Than Others

Not every parameter creates the same constitutional risk.

Some parameters are operational. Others are structural.

Operational parameters may affect convenience, efficiency, or throughput without altering the deeper nature of the system very much.

Structural parameters, by contrast, affect:

- who holds effective authority
- what outcomes can be challenged
- how much disclosure is normalized
- what level of capital discipline exists
- whether finality is meaningful
- whether a verifier market can become captured
- whether governance can exempt itself from ordinary discipline

A mature PAS should distinguish between these categories.

The Core Governance Tension

A Type 6 PAS needs enough adaptability to:

- correct mistakes
- improve parameters
- respond to adversarial learning
- manage unforeseen edge cases
- evolve as the surrounding environment changes

But too much flexibility undermines exactly the qualities the system is meant to provide:

- neutrality
- replayability
- institutional legibility
- bounded authority
- credible finality

The goal is therefore not “no governance.”

The goal is **disciplined governance**: governance that is real, explicit, and capable of responsible maintenance, but also narrow enough and constrained enough that it does not silently dominate the adjudication layer.

Common Governance Models

Founder or Operator Governance

In early-stage systems, governance often sits with the founding team, operator, or a closely held organization.

This is common, and in some cases unavoidable at first. It can permit rapid iteration and coherent stewardship.

But it is also the weakest model from the standpoint of neutral adjudication. It concentrates constitutional power in a small center and makes trust in the system inseparable from trust in its operator.

For a system aspiring to mature Type 6 status, this model is usually transitional at best.

Tokenholder or Stakeholder Governance

Some systems assign governance rights to tokenholders, stakers, or economically bonded participants.

This can broaden participation, but it does not automatically solve the problem. Stakeholder governance may still be highly concentrated, capture-prone, or poorly aligned with the interests of those most affected by adjudication outcomes.

Council or Committee Governance

Some systems rely on elected or appointed councils, protocol committees, or constitutional chambers to manage upgrades and structural parameters.

This can improve deliberation and institutional legibility, especially if membership, powers, and procedures are explicit. But it also introduces a governance class whose incentives and accountability must be examined directly.

A council can stabilize governance, or simply formalize concentrated power. The difference lies in its mandate and its constraints.

Layered Governance

More mature systems may adopt layered governance, in which:

- some parameters are fixed or very hard to change
- some are adjustable under ordinary process
- some require supermajority, delay, or multi-body approval
- some can be changed only prospectively, never retroactively
- some emergency powers exist but are narrow and auditable

This is often one of the healthiest approaches because it recognizes that not all parameters should be governed in the same way.

Governance and Time

Governance power is not defined only by what can be changed. It is also defined by **when** change can take effect.

Time matters because a system's legitimacy often depends on whether participants can know the rules under which they are acting.

A governance system that can alter critical parameters instantly may make the system highly agile, but it also weakens predictability and institutional reliance.

A mature Type 6 PAS should therefore think carefully about:

- notice periods
- delayed activation
- prospective-only changes
- non-retroactivity principles
- explicit treatment of in-flight claims

Without these protections, parameter control becomes a hidden form of adjudicative power.

Governance and Exceptional Powers

One of the hardest governance questions concerns exceptional powers.

- Should a system retain an emergency pause?
- A catastrophic invalidation power?
- A path for legal intervention?
- A constitutional override in cases of obvious systemic failure?

There is no universally correct answer. But there are clearly bad answers.

Bad systems pretend such powers do not exist when they effectively do. Others define them so broadly that finality becomes mostly rhetorical.

A serious Type 6 PAS should be explicit if exceptional powers exist, narrow in how they are defined, and legible in when and how they can be exercised.

Governance and Neutrality

A PAS becomes valuable when parties with different interests can rely on it without assuming that one side effectively owns the rules.

Governance therefore matters not only because it changes parameters, but because it determines whether the system can plausibly present itself as a neutral evidentiary architecture rather than as a tool of one institution, one coalition, or one economic bloc.

Neutrality here does not mean politics disappears. It means that the constitutional structure of the system is disciplined enough that no actor can easily convert governance power into quiet adjudicative dominance.

Governance and Specification

As systems mature, governance should ideally move closer to specification and further away from informal discretion.

That means:

- important parameters are named and classified
- powers are enumerated rather than implied
- procedures are explicit
- change thresholds are clear
- activation timing is legible
- audit trails are durable
- constitutional assumptions are documented

This is one of the reasons specification families matter. A serious PAS should not merely have governance. It should be able to describe its governance in a form that others can study, critique, compare, and eventually assess for compliance.

Evaluating Governance and Parameter Control

A Type 6 PAS should be judged by questions such as these:

Dimension	Question
Explicitness	Are governance powers clearly stated, or partly implicit?
Scope	Which parameters are governable, and which are intentionally fixed?
Structural sensitivity	Can governance alter trust, dispute, privacy, or finality properties directly?
Constraint	Are important powers bounded by delay, supermajority, or layered approval?
Non-retroactivity	Can rule changes affect already-submitted or already-finalized claims?
Override risk	Do exceptional powers exist, and if so, how narrow are they?
Transparency	Are parameter changes durable, legible, and auditable?
Capture resistance	How hard is it for one actor or coalition to dominate governance?
Institutional legibility	Can external parties understand what governance can and cannot do?
Neutrality compatibility	Does governance preserve the system’s claim to neutral adjudication?

Conclusion

A Type 6 Presence Adjudication System is not defined only by how it verifies claims. It is also defined by how it governs the rules, parameters, and powers through which verification becomes adjudication.

A system does not become neutral merely because it distributes verification. It becomes more plausibly neutral when its governance powers are explicit, bounded, delayed where necessary, resistant to capture, and unable to quietly hollow out the evidentiary structure below them.

That is a demanding standard.

But if a Type 6 PAS is meant to support serious cross-institutional reliance, it is also the right one.

Design Principles for Type 6 Presence Adjudication Systems

A Type 6 Presence Adjudication System should not be judged only by whether it is novel, decentralized, or cryptographically sophisticated.

Those qualities may matter, but they do not by themselves make a system good.

A good Type 6 PAS is one that solves the right problem in the right way. It must make consequential claims of physical presence usable across organizational boundaries without collapsing into unilateral authority, indiscriminate surveillance, weak finality, or decorative staking. It must be able to support reliance while remaining privacy-disciplined, adversarially robust, and institutionally legible.

That is a demanding standard. It is also the right one.

The purpose of this page is to state the main design principles that follow from the preceding analysis. These are not implementation instructions. They are architectural commitments: principles that distinguish serious Type 6 systems from weaker systems that merely borrow the language of decentralization or cryptographic proof.

Begin From the Claim, Not the Trace

A mature presence system should begin by asking what proposition needs to be established, not how much telemetry can be collected.

This is one of the most important design shifts in the whole field. Weak systems treat presence as something to be inferred from raw traces after the fact. Stronger systems begin by defining the claim itself: a person was within a region during an interval; an asset remained inside a controlled zone; an event boundary was crossed under valid conditions.

If the system begins from the trace, overexposure tends to become normal. If it begins from the claim, proportionate evidence becomes possible.

Treat Presence as an Evidentiary Problem

Presence is not just a data problem and not just a sensing problem. It is an evidentiary problem.

That means the system must be designed not merely to observe, but to support judgments that others can rely upon. It must be able to answer:

- what exactly is being claimed
- what evidence is admissible
- how that evidence is verified
- how challenges work
- when outcomes become durable
- what remains inspectable later

A system that collects impressive measurements but cannot bring them into a legible adjudicative form is not yet a serious PAS.

Minimize Unilateral Authority

The defining promise of Type 6 systems is not that authority disappears. It is that authority is no longer silently concentrated in one operator’s internal record.

A good Type 6 PAS should minimize the degree to which any single platform, verifier, committee, data provider, or governance body can define reality for everyone else without meaningful constraint.

Where unilateral authority remains necessary, it should be explicit and narrow.

Make Trust Explicit and Disciplined

A serious system should never pretend to be trustless.

Trust does not vanish in Type 6 architectures. It is redistributed, formalized, and exposed to challenge. The design goal should therefore be trust discipline.

The major assumptions should be visible. Their scope should be bounded. Their failure modes should be understandable. Their correction paths should be real.

Use Capital as Discipline, Not Decoration

Stake and bonded capital should be treated as instruments of discipline, not symbols of seriousness.

A system with large nominal stake but weak detection, weak slashing, weak dispute rights, or highly correlated control may be far less secure than it appears.

A good Type 6 PAS should therefore link economic exposure to actual adjudicative risk. Capital should be meaningfully slashable. Misconduct should be attributable. Collusion should be costly. Hidden escape routes from penalty should be minimized.

Design for Challenge, Not Just Initial Judgment

A system is not serious because it can produce an answer. It is serious because it can survive disagreement.

This means dispute architecture is not a secondary add-on. It is part of the core design. A good Type 6 PAS should assume that some claims will be wrong, fraudulent, collusive, or misleading, and should provide a real path by which such outcomes can be challenged and corrected.

Separate Ordinary Proof From Escalation

A well-designed system should not force every ordinary claim to carry the full evidentiary burden of the hardest possible dispute.

Ordinary adjudication should operate on bounded claims and proportionate evidence. More invasive or more detailed evidentiary access, where it exists, should belong to structured challenge or escalation modes rather than being normalized for everyday use.

A good Type 6 PAS should therefore distinguish clearly between:

- ordinary proof mode

- challenge mode
- escalation mode
- exceptional override or external review, if any

Make Finality Legible

A claim is not useful merely because it has been evaluated. It becomes useful when it reaches a form of closure that others can rely upon.

A good Type 6 PAS should therefore define its finality surfaces clearly. It should be possible to understand:

- when a claim is evidentially accepted
- when ordinary disputes have closed
- where the durable record of the outcome exists
- what kind of downstream reliance is justified and when
- whether exceptional reopening paths exist

Finality should not be vague, hidden, or purely rhetorical. The system should state what becomes final, at what threshold, for what purposes.

This is especially important where claims support payments, credentials, access rights, compliance outcomes, or later institutional scrutiny.

Publish Durable but Disciplined Outcomes

A mature Type 6 PAS should produce outcomes that can be referenced later, inspected by relevant parties, and relied upon across institutional boundaries. But it should do so without turning durable publication into permanent overexposure of raw behavioral traces.

This means the system should think carefully about what becomes durable:

- the claim
- the proof object
- the adjudication result
- the challenge history
- the finalization record
- or some combination of these

Not every layer of evidence needs to become publicly exposed forever in order for the outcome to be replayable. A good system should publish enough to support durable reliance without making indiscriminate transparency its default evidentiary model.

Keep Governance Real but Bounded

Governance is unavoidable. The relevant question is not whether governance exists, but whether it is disciplined enough to preserve the neutrality and replayability the system claims to provide.

A good Type 6 PAS should distinguish between parameters that are operationally adjustable and parameters that are structurally constitutive. It should be much easier to tune routine settings than to alter core trust, dispute, privacy, or finality assumptions.

Match Security to Consequence

Not every presence claim carries the same stakes, and a mature system should not pretend otherwise.

A good Type 6 PAS should recognize that different claim classes may require different security thresholds, challenge windows, finality conditions, or evidentiary burdens. Low-stakes attendance proofs, high-value logistics releases, and legally sensitive compliance claims should not necessarily be treated as though one security model fits all.

This implies use-case discipline. The system should know what its security-capital surface can actually defend, and it should avoid inviting reliance beyond that envelope without stronger protections.

A serious PAS does not merely ask what is technically possible. It asks what is responsibly supportable.

Preserve Institutional Legibility

A presence adjudication system does not become useful merely because it is internally coherent. It becomes useful when other parties can understand how to rely on it.

A good Type 6 PAS should therefore be legible not only to protocol designers, but also to institutions, counterparties, regulators, auditors, and technically informed outsiders. Its claims should be understandable. Its trust assumptions should be articulable. Its dispute process should be explainable. Its finality conditions should be clear. Its governance powers should be documented.

This is not a concession to older institutions. It is part of what makes the system portable beyond its own internal culture.

Prefer Specification Over Informal Doctrine

As systems mature, principles should increasingly become specifications.

A strong Type 6 PAS should be able to describe:

- its claim semantics
- its proof architecture
- its dispute rights
- its finality thresholds
- its capital discipline
- its governance powers

in forms that are stable enough to study, compare, critique, and eventually assess for conformance.

Conclusion

The purpose of design principles is not to dictate one implementation. It is to make clear what kind of architecture deserves to be taken seriously.

A Type 6 Presence Adjudication System is not good because it uses cryptography, staking, or committees. It is good when those elements are arranged in ways that make consequential claims of presence neutral enough, challengeable enough, durable enough, and privacy-disciplined enough for real counterparties to rely upon.

That is the design task.

The next page asks the stronger question that follows from these principles: what would an ideal Type 6 Presence Adjudication System actually look like?

Properties of an Ideal Type 6 Presence Adjudication System

A taxonomy can describe the field. A design space can clarify the tradeoffs. But eventually a more demanding question has to be asked:

What should a good system actually look like?

This page is an attempt to answer that question for **Type 6 Presence Adjudication Systems**: systems that seek to adjudicate consequential claims of physical presence through distributed verification, explicit incentives, challengeability, and durable finalization rather than through unilateral institutional authority alone.

The claim of this page is not that one perfect implementation already exists. Nor is it that every domain needs the same architecture. The claim is narrower and more important:

if a Type 6 Presence Adjudication System is to deserve serious reliance, it must exhibit a recognizable set of properties.

These properties are not cosmetic. They are what distinguish a mature presence adjudication architecture from a system that merely borrows the language of decentralization, cryptography, or privacy while remaining structurally weak.

The System Should Begin From Bounded Claims

An ideal Type 6 PAS should begin from bounded, adjudicable claims rather than raw telemetry.

Its central evidentiary object should not be an unstructured stream of coordinates, device events, or operator logs. It should be a proposition that another party can understand and rely upon:

- that a person was within a defined region during a defined interval
- that an asset remained inside a controlled zone
- that a device crossed a threshold under stated conditions
- that an attendance condition was satisfied

This matters because claim shape determines the entire downstream architecture. If the system begins from traces, then overexposure and discretionary interpretation tend to become normal. If it begins from claims, then proportionate evidence, selective disclosure, and clearer dispute become possible.

The System Should Be Evidentiary, Not Merely Observational

A serious PAS is not a sensor network with branding. It is not a map, a tracking dashboard, or a telemetry warehouse.

It is an evidentiary architecture.

That means the system must do more than observe or estimate location. It must support a structured path from observation to evidence, from evidence to adjudication, from adjudication to dispute, and from dispute to durable reliance.

In an ideal system, it is always possible to answer:

- what exactly was claimed
- what kind of evidence supported it
- under what rules it was assessed
- how it could have been challenged
- when it became final enough for reliance

The System Should Minimize Unilateral Power

No single operator, verifier, data source, platform, or governance body should be able to define reality for everyone else without meaningful constraint.

This does not mean that all roles must be symmetrically distributed, nor that every function must be maximally decentralized. It means that consequential authority should be bounded, visible, and challengeable.

An ideal Type 6 PAS does not rely on hidden sovereignty.

The System Should Make Trust Explicit

An ideal Type 6 PAS should never claim to remove trust altogether.

Instead, it should make trust assumptions visible enough to inspect, bounded enough to reason about, and disciplined enough to challenge.

The system should be able to state clearly:

- what it assumes about measurement integrity
- what it assumes about prover incentives
- what it assumes about verifier independence
- what it assumes about watcher participation
- what it assumes about finalization layers
- what it assumes about governance powers

The System Should Support Selective Disclosure by Default

An ideal Type 6 PAS should not require full behavioral exposure in order to establish ordinary claims.

Its ordinary mode of operation should be one in which:

- the claim is narrow
- the evidence is proportionate
- disclosure is bounded
- confidence does not depend on revealing entire movement histories

This does not mean that richer evidence is never needed. Disputes, escalations, and legal processes may justify deeper inspection in some cases. But the ordinary evidentiary burden should remain disciplined.

A system that claims to protect privacy but requires overexposure for routine use is not yet well designed. It has merely deferred surveillance to the point of adjudication.

The System Should Prove the Claim, Not Merely Expose the Data

An ideal Type 6 PAS should be designed, wherever possible, around proving that a bounded proposition holds rather than dumping the underlying data for others to interpret.

A weak system says: here is the trace; you decide what it means.

A stronger system says: here is a claim, here is evidence that it satisfies the relevant rule, and here is how that evidence can be checked or challenged.

This shift is what makes privacy-compatible verifiability possible.

The System Should Be Challengeable in Practice, Not Only in Principle

An ideal Type 6 PAS should be secured not only by its initial adjudication layer, but by a credible possibility of correction.

This means disputes must be real.

Challenges should be:

- economically viable to bring
- procedurally clear
- temporally possible
- evidentially meaningful
- capable of producing actual correction

A mature Type 6 PAS assumes that some claims will be wrong or adversarial and is designed accordingly.

The System Should Make Dishonesty Costly in a Real Sense

Economic discipline matters only when the downside is real.

An ideal Type 6 PAS should expose adjudicating actors to meaningful loss if they behave dishonestly, collusively, or recklessly. This requires more than nominal stake. It requires:

- capital that is genuinely slashable
- misconduct that is attributable
- challenge processes that can trigger enforcement
- governance that cannot casually neutralize penalties
- sufficient distribution that security is not mostly performative

The system should know the limits of its own security-capital surface. It should not imply that any level of consequence can safely rest on its outputs if the economically exposed structure cannot actually defend that reliance.

The System Should Know Its Security Envelope

Not every presence claim has the same stakes, and not every PAS should pretend otherwise.

An ideal Type 6 system should understand which classes of consequence it can responsibly support, and under what conditions. It should recognize that low-stakes attendance proofs, high-value logistics releases, and legally sensitive compliance claims may require different evidentiary burdens, challenge windows, capital requirements, or finality thresholds.

A mature system is not one that claims universal applicability. It is one that knows where its security envelope lies and designs accordingly.

The System Should Make Finality Legible

A presence claim becomes important when other parties can rely on it.

That means an ideal Type 6 PAS must make finality clear.

It should be possible to understand:

- when a claim has been accepted evidentially
- when ordinary challenges are no longer admissible
- where the durable outcome is recorded
- what kind of reliance is justified and when
- whether exceptional reopening paths exist

Finality should not be vague, operator-dependent, or purely rhetorical.

The System Should Publish Durable Outcomes Without Normalizing Surveillance

Durability matters. But permanent overexposure is not the only way to achieve it.

An ideal Type 6 PAS should make outcomes replayable and inspectable without treating the indefinite publication of raw location traces as the normal cost of institutional memory. It should think carefully about what becomes durable:

- the claim
- the proof object
- the adjudication result
- the challenge record
- the finalization marker

Not every layer must be equally public forever in order for the system to remain auditable.

A mature PAS should therefore aim for durable legibility rather than indiscriminate transparency.

The System Should Distinguish Adjudication From Governance

No serious PAS is free of governance. But governance should not quietly swallow adjudication.

An ideal Type 6 system should make clear:

- what is decided by protocol rules
- what can be challenged by participants
- what can be changed by governance
- what powers are exceptional rather than ordinary
- which changes are prospective only
- which structural features are intentionally hard to alter

A system in which governance can casually rewrite the trust model, dispute model, or finality model is not yet constitutionally mature.

The System Should Be Legible to Parties Outside Itself

An ideal Type 6 PAS should not require total immersion in its internal culture to be understood.

Its trust assumptions, claim semantics, dispute rights, finality conditions, and governance powers should be intelligible to:

- counterparties
- auditors
- regulators
- technically informed outsiders
- institutions deciding whether to rely on its outputs

Legibility is part of neutrality. A system that cannot explain itself cannot easily ask others to rely on it.

The System Should Be Specifiable

A mature Type 6 PAS should be capable of specification.

That means its key properties should be describable in structured, stable, and auditable form:

- what counts as a valid claim
- what the proof architecture guarantees
- what dispute rights exist
- how finality is reached
- how economic discipline works
- what governance can and cannot do

This is not a bureaucratic add-on. It is one of the marks of maturity.

The System Should Be Normatively Honest

An ideal Type 6 PAS should be honest about what it is for, what it can defend, and what it cannot solve.

It should not imply:

- that cryptography removes all trust
- that decentralization automatically produces neutrality
- that privacy eliminates dispute
- that finality is the same as truth
- that every real-world conflict can remain internal to the protocol
- that all use cases are equally suitable

Normative honesty is not modesty for its own sake. It is one of the conditions of credibility.

Conclusion

The purpose of an ideal is not to pretend that implementation is easy. It is to make clear what maturity would look like.

In the case of Type 6 Presence Adjudication Systems, maturity does not mean eliminating all trust, all dispute, all governance, or all ambiguity. It means arranging these unavoidable realities in a way that makes consequential claims of physical presence more neutral, more disciplined, more contestable, more durable, and less surveillance-dependent than older architectures allow.

That is the promise of this design space.

It is also the standard by which systems in this category should be assessed.