PRINTABLE EDITION

Sovereign Location — Frameworks

Published by the Scintilla Foundation
6 April 2026

Contents

FRAMEWORKS

Concepts and Terminology 1
Taxonomy of Presence Adjudication Systems 4
Selective Disclosure and Presence 10

Concepts and Terminology

This site uses a number of terms in a more precise sense than ordinary usage might suggest. Some are familiar words given sharper meaning. Others are conceptual terms introduced in order to make an emerging design space more legible.

This page is not intended to be an exhaustive glossary. It is a compact guide to the site's core working vocabulary: the terms that recur often enough, and carry enough conceptual weight, that it is worth stating clearly how they are being used here.

Presence

Presence refers to the fact that a person, device, or asset was within a defined place during a defined interval, under conditions relevant to some claim.

In ordinary language, presence can sound trivial: simply being somewhere. In the context of this site, however, presence matters because it often functions as a condition of coordination. A delivery may depend on it. An inspection may depend on it. A credential, payment, compliance outcome, or access decision may depend on it.

Presence is therefore not treated here as a mere descriptive detail. It is treated as something that may carry evidentiary, institutional, and economic consequences.

Adjudication

Adjudication refers to the process by which a claim is assessed under explicit rules and brought to an outcome that others can rely upon.

This is a broader idea than simple verification. A system may verify that a signature is valid or that a proof satisfies a predicate, but adjudication concerns the larger question of how a claim becomes accepted, rejected, disputed, or finalized in a socially or institutionally meaningful way.

The use of this term is deliberate. Presence claims often matter not because they are interesting to observe, but because they affect decisions, rights, obligations, or settlement. That makes adjudication central.

Verification

Verification refers to the process of checking whether some evidence or claim satisfies the relevant rules.

In this site's vocabulary, verification is an important component of a larger evidentiary structure, but it is not the whole structure. Verification may establish that a proof is valid, that a condition has been met, or that a formal claim follows from available evidence. But verification alone does not settle every institutional question.

This distinction matters because many systems are described as “verifying presence” when in fact they are only asserting it, or only producing data from which another party is expected to infer it.

Selective Disclosure

Selective disclosure refers to the principle that a presence claim should reveal no more information than is necessary for the claim to become usable.

This is one of the site’s most important framing concepts. In many situations, the relevant question is not where someone was at every moment, but whether they satisfied a bounded condition: being within a region during an interval, crossing a threshold, remaining inside a zone, or attending a defined event.

Selective disclosure therefore stands against the assumption that full behavioral traces are the normal price of evidentiary confidence. It does not mean secrecy for its own sake. It means proportionate revelation: enough to support the claim, but not so much that the evidentiary act becomes indistinguishable from surveillance.

Sovereign Location

Sovereign Location is the design principle that presence claims should be adjudicable under neutral, privacy-respecting, replayable rules, without requiring blind trust in a single authority or unnecessary disclosure of raw location history.

The term *sovereign* does not imply isolation or total independence from institutions. It refers more specifically to the idea that no single platform, database, or intermediary should hold exclusive authority over the truth of presence claims.

Sovereign Location is therefore not simply about privacy, nor simply about decentralization. It is about the structure of evidentiary power: who can produce, interpret, contest, and rely upon claims of presence, and under what rules.

Proof of Presence

Proof of presence refers to evidence that a person, device, or asset satisfied a defined presence condition.

The exact form of such proof may vary. In some systems it may take the form of signed attestations, system records, or institutional certifications. In others it may involve cryptographic commitments, zero-knowledge proofs, or multi-party adjudication.

What matters conceptually is that the proof is tied to a claim. It is not merely raw telemetry or a coordinate trace. It is evidence presented in relation to a proposition that another party may need to accept, reject, or dispute.

Coordination Primitive

A **coordination primitive** is a recurring condition that many different systems need to reference in order to make consequential decisions.

Identity is a coordination primitive because many systems need to know who acted. Time is a coordination primitive because many systems need to know when something happened. Presence increasingly belongs in the same category because many systems need to know whether someone or something was in a relevant place during a relevant interval.

To call presence a coordination primitive is to argue that it should not be treated as an incidental feature buried inside isolated products. It is becoming a shared problem across many domains, and therefore a candidate for more general infrastructure.

A Note on Usage

These terms are meant to work together.

- **Presence** is the underlying fact pattern.
- **Proof of presence** is the evidentiary form in which that fact may be presented.
- **Verification** checks whether the evidentiary claim satisfies relevant rules.
- **Adjudication** determines how the claim is ultimately handled in a broader institutional or system context.
- **Selective disclosure** constrains how much information must be revealed along the way.
- **Sovereign Location** names the design principle that these processes should not collapse into surveillance or blind trust in a single intermediary.
- **Coordination primitive** explains why all of this matters beyond one narrow application domain.

Readers do not need to memorize these terms before reading the rest of the site. But keeping their intended meanings in view will make the surrounding arguments much easier to follow.

Taxonomy of Presence Adjudication Systems

Human coordination depends on the ability to establish facts about physical reality. One of the most important of these facts is **presence**: whether a person, device, or asset was within a relevant place during a relevant interval, under conditions that matter to someone else.

Questions of presence are everywhere, even if they are not always described in those terms.

- Was a courier present at a delivery address?
- Did a worker attend a job site?
- Was a vehicle within a restricted zone?
- Did a participant attend an event?
- Did an inspection occur at a facility?

Historically, societies have developed many different ways of answering such questions. Some rely on memory and social recognition. Others rely on witnesses, signatures, inspectors, courts, databases, platforms, sensors, or cryptographic systems. These mechanisms differ not only in technology, but in authority structure, trust model, privacy posture, cost, portability, and durability.

This page proposes a taxonomy of **Presence Adjudication Systems (PAS)**: systems that convert observations about physical presence into judgments that other actors can rely upon.

The purpose of the taxonomy is not to force every real-world system into a perfectly clean box. It is to give readers a usable vocabulary for comparing the ways societies and institutions establish presence, and for understanding why newer digital systems are emerging.

What Is a Presence Adjudication System?

A **Presence Adjudication System** is any mechanism that transforms observations about the physical world into a socially actionable determination of presence.

In simplified form:

```
PAS = Observation → Evidence → Adjudication → Finalization
```

Where:

- **Observation** is how the underlying fact enters the system
- **Evidence** is how that fact is represented
- **Adjudication** is how a judgment is reached
- **Finalization** is how the judgment becomes durable enough for others to rely upon

This definition is intentionally broad. It includes not only advanced cryptographic protocols, but also witnesses, affidavits, inspectors, courts, and centralized databases. That breadth is important. Presence adjudication is not a new invention. What is new is the growing demand for systems that can perform it at digital scale, across institutional boundaries, under stronger privacy and trust requirements.

Why a Taxonomy Helps

A taxonomy helps in three ways.

First, it places digital systems in historical context. A modern protocol is not trying to invent the idea of presence adjudication from nothing. It is entering a field of long-standing institutional forms.

Second, it clarifies tradeoffs. Different types of PAS solve different problems well. Some are cheap but local. Some are authoritative but opaque. Some are privacy-invasive but operationally efficient. Some are portable but governance-heavy.

Third, it allows newer systems to be understood as alternatives within a broader design space rather than as isolated technical curiosities.

Major Types of Presence Adjudication Systems

The following taxonomy is best read as a set of recurring types. Real systems may combine features of several types, and some systems may evolve from one type toward another over time.

Type	Name	Core adjudication model	Typical examples	Main strengths	Main limitations
0	Informal Social Recognition	Presence accepted through shared social knowledge	village memory, community recognition, local familiarity	low cost, socially natural	weak durability, low portability, low auditability
1	Testimonial Systems	Human statements treated as evidence	witnesses, affidavits, signatures, declarations	flexible, familiar, legally legible	interpretive, reputation-dependent, contestable
2	Institutional Systems	Formal authority determines outcome	courts, customs authorities, inspectors, regulators	strong legitimacy, enforceability	expensive, slow, centralized
3	Centralized Digital Systems	One operator records and interprets machine data	delivery apps, fleet tracking systems, mobile phone logs	scalable, operationally efficient	opaque, privacy-invasive, trust-heavy

Type	Name	Core adjudication model	Typical examples	Main strengths	Main limitations
4	Federated Systems	Multiple institutions jointly attest or validate	consortium networks, regulated data exchanges	broader trust base, shared governance	complex coordination, governance-dependent
5	Cryptographically Anchored Systems	Digitally signed evidence with integrity protection	signed sensor reports, secure hardware attestation	stronger integrity, tamper resistance	still trust-anchored to issuers or hardware roots
6	Decentralized Economic Systems	Open or semi-open adjudication via incentives and disputes	stake-secured verifier markets, challenge systems	capture resistance, replayability, cross-party legibility	design complexity, incentive sensitivity
7	Strong Privacy-Preserving Systems	Presence established with minimized disclosure	zero-knowledge presence proofs, private region membership proofs	better privacy discipline, bounded revelation	technical complexity, still needs governance/adjudication context
8	Hypothetical or Future Systems	Not yet mature, but conceptually possible	ubiquitous trusted sensing, advanced multi-party sensing fabrics	potentially powerful new capabilities	unresolved feasibility, governance, and civil implications

Reading the Types

These types can be understood as a historical and structural progression, but not as a simple ladder of improvement.

Type 0 and Type 1 systems are ancient and still important. Much of everyday human life continues to rely on social recognition, testimony, and local trust. They are often fragile by modern digital standards, yet they remain cheap, flexible, and institutionally familiar.

Type 2 systems introduce formal authority. Courts, customs offices, inspectors, and regulators do not merely observe presence; they render judgments that can carry force beyond the immediate moment. These systems are often stronger in legitimacy and enforceability, but they are also costly, centralized, and difficult to scale gracefully.

Type 3 systems represent the dominant contemporary digital pattern. Platforms and enterprises collect location data, store logs, and use their own internal systems to determine what happened. These systems are operationally powerful, but their evidentiary structure is often weak from an external point of view. They scale well, but they usually require either trust in the operator or overexposure of raw data.

Type 4 and Type 5 systems begin to address some of these weaknesses. Federation broadens the base of authority. Cryptographic anchoring can improve integrity and tamper resistance. But both types often remain dependent on trusted issuers, hardware roots, institutional governance, or closed networks.

Type 6 and Type 7 become especially important for the digital world because they address a deeper problem: how to establish consequential presence claims without simply collapsing into unilateral platform control or indiscriminate surveillance. They are not automatically better in every context, but they are better aligned with the needs of open, networked, multi-party coordination.

Evaluation Dimensions

No taxonomy is useful unless it supports comparison. Presence Adjudication Systems can be compared along several recurring dimensions.

Dimension	Question
Authority structure	Who has the power to determine the outcome?
Trust model	What assumptions must be trusted for the result to be meaningful?
Privacy posture	How much information becomes visible, and to whom?
Evidence granularity	What kinds of claims can the system express?
Auditability	Can later parties reconstruct how a result was reached?
Portability	Can the judgment travel across institutions or contexts?
Economic accountability	Are adjudicators exposed to meaningful incentives or penalties?
Finality	When does a decision become durable enough to rely upon?
Capture resistance	How difficult is it for the system to be corrupted or dominated?
Cost and scalability	How expensive is the system to operate, and how broadly can it be used?

These dimensions make it easier to see why different PAS types suit different environments. Courts score differently from mobile platforms. Signed sensor systems score differently from witness testimony. A decentralized protocol may outperform a centralized platform on capture resistance or replayability while performing worse on simplicity or institutional familiarity.

Why Type 6 Matters

For the purposes of this site, **Type 6 — Decentralized Economic Systems** deserves particular attention.

This is not because every presence question should be handled by a decentralized economic protocol. Many should not. Informal, institutional, and centralized systems will continue to exist, and often remain appropriate in their own domains.

Type 6 matters because it offers one especially important answer to a distinctly modern problem: how to adjudicate consequential presence claims in digital environments where no single intermediary should be trusted to define reality for everyone else.

Its characteristic features include:

- independent or semi-independent verifiers
- explicit incentives and penalties
- dispute or challenge mechanisms
- durable publication of outcomes
- compatibility with cryptographic proofs and bounded claims

This makes Type 6 especially relevant where presence must be made legible across organizational boundaries, where privacy matters, where outcomes may carry financial or institutional consequence, and where unilateral control by one operator is undesirable.

The deeper design questions raised by Type 6 systems belong in the Design Space section. Here, the point is simply to mark why this type deserves special attention within the broader taxonomy.

The Role of Blockchain

Blockchain does not measure physical reality, and it should not be described as though it does.

Its role within some PAS types is narrower and more important than that. It can provide:

- durable publication
- neutral coordination
- economic settlement
- replayable finalization

In other words, blockchain is not itself a presence adjudication system. It is one possible component in the finalization and incentive structure of certain PAS types, especially Type 6 systems.

What the Taxonomy Makes Visible

The value of this taxonomy is not only classificatory. It changes how the problem appears.

It makes clear that presence adjudication is a longstanding civilizational function rather than a niche problem invented by modern protocols.

It allows readers to recognize systems they already know — affidavits, inspectors, courts, platform logs, signed reports — as members of a broader family.

And it creates the conceptual bridge needed to understand why newer systems are emerging. If presence has become a coordination primitive for digital society, then older PAS types will increasingly show their limitations. Some are too local. Some

are too trust-heavy. Some are too privacy-invasive. Some are too slow or too institutionally bounded.

The question is not whether older PAS types disappear. They will not. The question is which types are best suited to the demands of increasingly digital, multi-party, privacy-sensitive coordination.

That is the larger design problem this site is concerned with.

Conclusion

Presence Adjudication Systems are a foundational but often overlooked part of social and institutional life. They are the mechanisms by which societies turn observations about the physical world into judgments that others can rely upon.

Understanding them requires more than technical description. It requires attention to authority, trust, privacy, incentives, portability, and finality.

This taxonomy is intended as a framework for that understanding. It gives readers a vocabulary for interpreting familiar systems, comparing historical and digital alternatives, and seeing why newer forms of presence adjudication are emerging.

Everything that follows in the site’s later design-space discussions depends on that comparative foundation.

Selective Disclosure and Presence

Presence should be understood as a selective disclosure problem, not merely a tracking problem.

That distinction matters because most contemporary location systems begin from the wrong side of the question. They assume that the basic task is to collect, retain, and interpret as much spatial data as possible. Presence then appears as a downstream inference drawn from a much larger behavioral record.

But in many important settings, that is not the real problem at all.

The real problem is usually narrower: how can a person, device, or asset demonstrate a bounded fact about presence without disclosing more than the situation requires? Not “where was this entity at all times?” but “can it establish that it was within this region, during this interval, under these conditions?”

Once the problem is stated that way, the design priorities begin to change.

From Tracking to Claims

Tracking systems are built to observe trajectories. They gather coordinates, timestamps, device identifiers, routes, and surrounding metadata over time. Their logic is cumulative: the more information collected, the richer the picture that can later be reconstructed.

A presence system, by contrast, does not necessarily need a rich picture. It needs a usable claim.

That claim may be quite modest. A courier was at the delivery point during the agreed window. An inspector was present at a site. An attendee crossed an event boundary. An asset remained within a controlled zone during a relevant interval.

In each case, the underlying evidentiary question is limited. Yet current systems often answer it by collecting and exposing far more than the claim itself requires.

Selective disclosure begins by refusing that default.

It asks whether the claim can be supported at the level of what matters, rather than at the level of everything that happened around it.

Why This Is a Framework Question

Selective disclosure is sometimes treated as a privacy feature added after the fact, as though a fully formed location system first collects whatever it likes and only later decides to reveal less of it.

That is too shallow.

In the context of presence, selective disclosure is not merely a user preference or interface setting. It is a design principle that shapes how the problem itself is understood. It affects what kinds of claims are representable, what counts as relevant evidence, how much information counterparties should receive, and what kind of adjudication becomes possible.

Framed this way, selective disclosure is not about concealment for its own sake. It is about discipline. It is the idea that evidentiary systems should disclose what is necessary for the claim and no more than that.

That principle becomes especially important once presence carries real consequences. When payment, access, compliance, credentialing, or liability depend on a presence claim, the temptation is often to demand maximal visibility. Selective disclosure challenges that instinct.

The Shape of the Narrower Claim

One reason modern location systems over-disclose is that they are often built around the wrong unit of meaning.

The natural output of a location stack is usually a coordinate stream, or something close to it. But the natural unit of social and institutional use is often not a coordinate. It is a proposition.

A proposition might be:

- this person was within a defined region
- this device entered a site during a valid interval
- this asset did not leave a restricted zone
- this participant satisfied an attendance condition

These are not just smaller pieces of data. They are differently structured claims. They are bounded, contextual, and tied to some decision that another system needs to make.

Selective disclosure becomes possible when systems are designed around such propositions rather than around maximal telemetry. The point is not merely to hide data. It is to express the relevant fact at the right level of abstraction.

Presence Without Behavioral Exposure

This matters because location data is rarely neutral.

A movement history can reveal habits, relationships, routines, political participation, commercial activity, vulnerabilities, and patterns of life that far exceed the original question being asked. A system that demands raw traces in order to establish a narrow presence claim effectively forces the subject to disclose a much broader behavioral record than the situation warrants.

That is not simply inefficient. It alters the balance of power between the party making the claim and the party demanding evidence.

Selective disclosure offers a different model. Instead of treating presence as something to be reconstructed from a pool of retained traces, it treats presence as something that can be established through bounded revelation. The subject need not surrender the whole map of their movements merely to prove one limited fact.

This is one of the key conceptual shifts in Sovereign Location. The relevant question is no longer how much data can be collected, but how little must be revealed for the claim to become usable.

Selective Disclosure Is Not Secrecy

It is important, however, not to confuse selective disclosure with refusal or opacity.

A system that reveals nothing useful may protect privacy, but it does not solve the coordination problem. Presence claims often matter precisely because another party needs to rely on them. A delivery receiver, regulator, insurer, employer, venue, or counterparty may need some basis for confidence that a condition has been met.

So the aim is not to avoid evidence. It is to make evidence proportionate.

Selective disclosure therefore sits between two bad extremes. On one side is surveillance: reveal everything and let institutions sort it out later. On the other side is unusable opacity: reveal so little that the claim cannot meaningfully be relied upon.

A mature presence system must find a better balance.

Adjudication, Not Just Presentation

Selective disclosure is often misunderstood as a matter of presentation: what a user interface shows to a viewer, or what fields are hidden in a report.

That is too superficial.

In a serious presence system, selective disclosure has to reach deeper than presentation. It has to shape the evidentiary and adjudicative architecture itself. The question is not merely what a viewer sees, but what the system treats as necessary to establish the claim in the first place.

This is why the topic belongs in the Frameworks section. It is part of how the field should think.

Institutional Consequences

Once presence is framed this way, a number of design consequences follow.

Systems should be designed around specific claim types rather than raw data exhaust.

Verification should focus on whether a defined proposition has been established, not on whether an institution has accumulated enough telemetry to feel comfortable.

Evidence should be proportionate to the consequence at stake.

And the subject of the claim should, wherever possible, have more meaningful control over how the underlying informational record is disclosed.

These are not merely user-experience preferences. They are architectural commitments. They shape who holds evidentiary power, how disputes are conducted, and whether presence becomes compatible with privacy or permanently tied to surveillance.

Conclusion

To treat presence as a selective disclosure problem is to move toward a more mature model of digital evidence.

It means recognizing that the relevant social act is usually not tracking but proving. Not accumulation but adjudication. Not maximal visibility but bounded legibility.

This does not eliminate hard questions. Systems still need to decide what kinds of claims are valid, what level of confidence is sufficient, how disputes should work, and when stronger disclosure is justified. But it does clarify the direction of travel.

A presence system should not begin by asking how much location data can be collected. It should begin by asking what kind of claim needs to be established, and what the minimum necessary disclosure is for that claim to become usable.

That is the core framework idea.

Presence is not merely something to observe. It is something to disclose, selectively and under rules, when the situation requires it.