SESSION 10C

FRACTIONAL IDENTITIES FOR THE AGENTIC AI FUTURE



COSAC

COSAC 2026

SRDJAN MARINOVIC

Full Title

FRACTIONAL (Delegated) IDENTITIES FOR THE AGENTIC AI FUTURE, and no one knows what that future really is. And everything I present today is subject to change tomorrow. And the security standards are evolving fast.

Myself?

Futurae (CTO) – Authentication and Trust Platforms

PwC US (Director)– Trust and reliability for AI, Smart Venues

SignalGraph (CTO)– US-based startup in SIGINT and Trust Signals

ETH Zurich (PostDoc) – Run-time verification of security policies

Imperial College London (PhD) –Trust-based formal security models

Futurae?

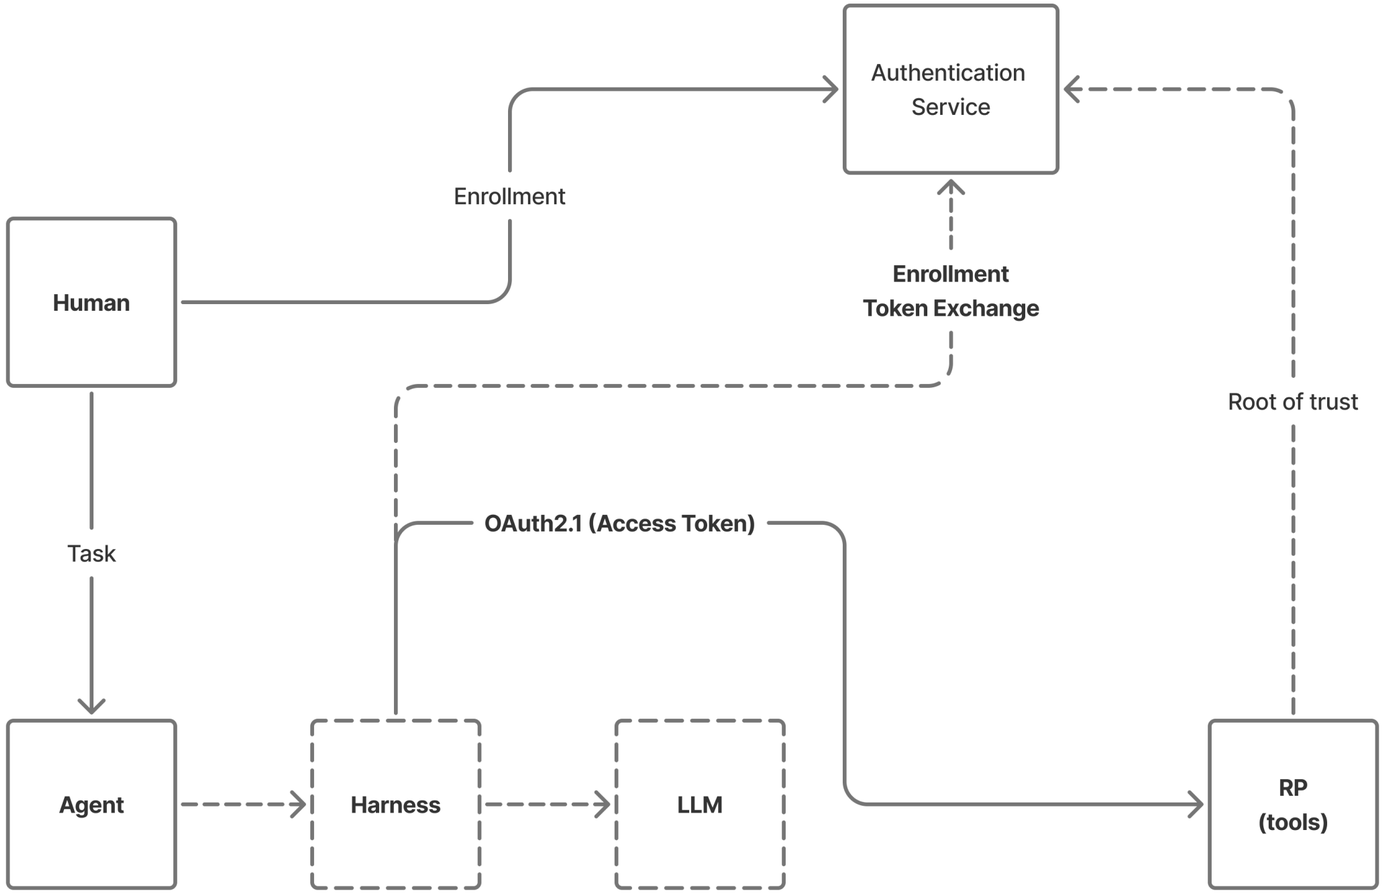
MFA and Tx Signing Products

Recently lunched Risk-based Passkeys and Trust Signals platforms

160+ Fin Clients mostly in Europe

10Mil MAUs with 60Mil monthly transactions

Where is the industry today?



We have moved away from service accounts, account impersonations, custom token schemes...

SPIFFE framework: identity scoped to workload.

OAuth2.1 and token exchanges.

Focused (small-scale) agent systems such as AP2 (short bounded delegations).

Lots of standards in the works.

Immediate Challenges

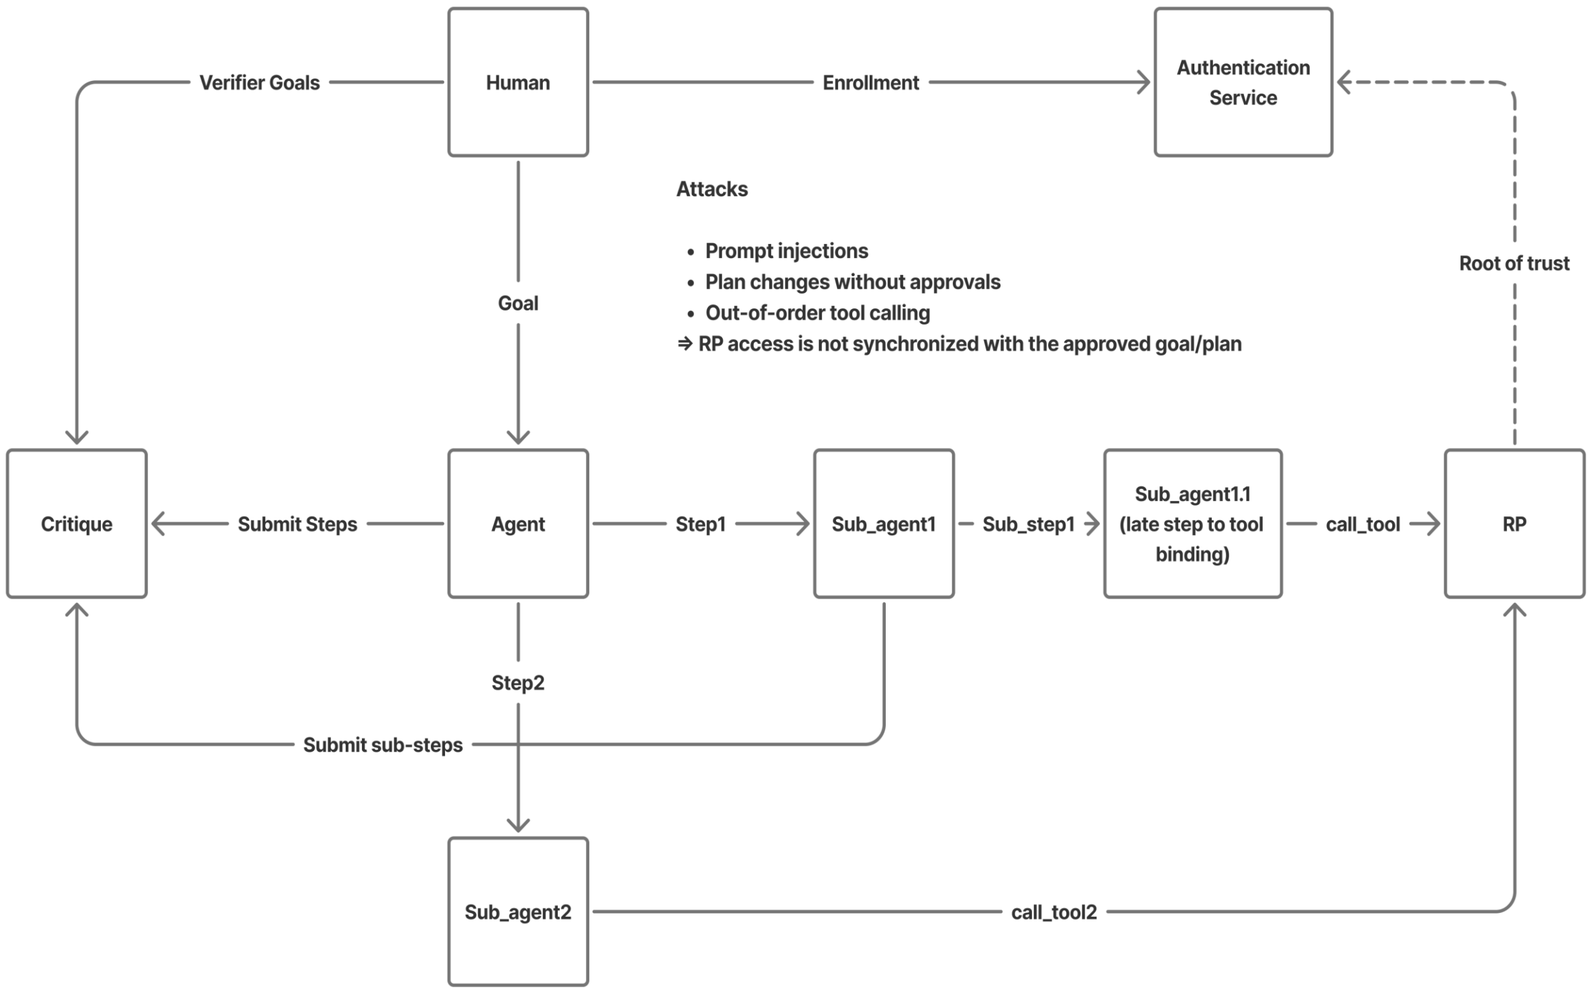
1. Existing identity and capability-based are built for individual rational humans.
2. Agents, in contrast, can be unreliable; they are not static; they can spawn.
3. Agents’ actions must be verified.
4. Death-by-1000 standards; not trivial putting many RFCs together into one system.

Key Problem

We do not have a *unified* identity and authorization models/methodologies for agentic systems whose architecture and capabilities are scaling-out faster and faster.

This is not a critique; our current solutions (esp ecommerce) are a natural starting point.

Deployed Agentic System Models



Trillions of agents are coming online.

Agent systems are comprised of many ephemeral agents with distinct roles.

Planning and goals are the inputs.

We have no idea how these systems are going to look like in 3 years.

Agentic Attacker Model

1. Agents can compromise **Authz** and **SIEM** systems.
2. Agents collude and will attempt to cheat.
3. Agent’s reasoning can be polluted.
4. **For now, assume agents cannot break crypto and formal proofs.**

Imagine an agent who really likes you

Your agent really really wants to buy you expensive shoes for your birthday, so **it decides to hack the authorization server to issue itself more permissions.**

Path-forward (what does good look like)?

01

Embrace the “swarm” agentic system model.

02

Build recursive (multi-party) delegation models around agentic roles.

03

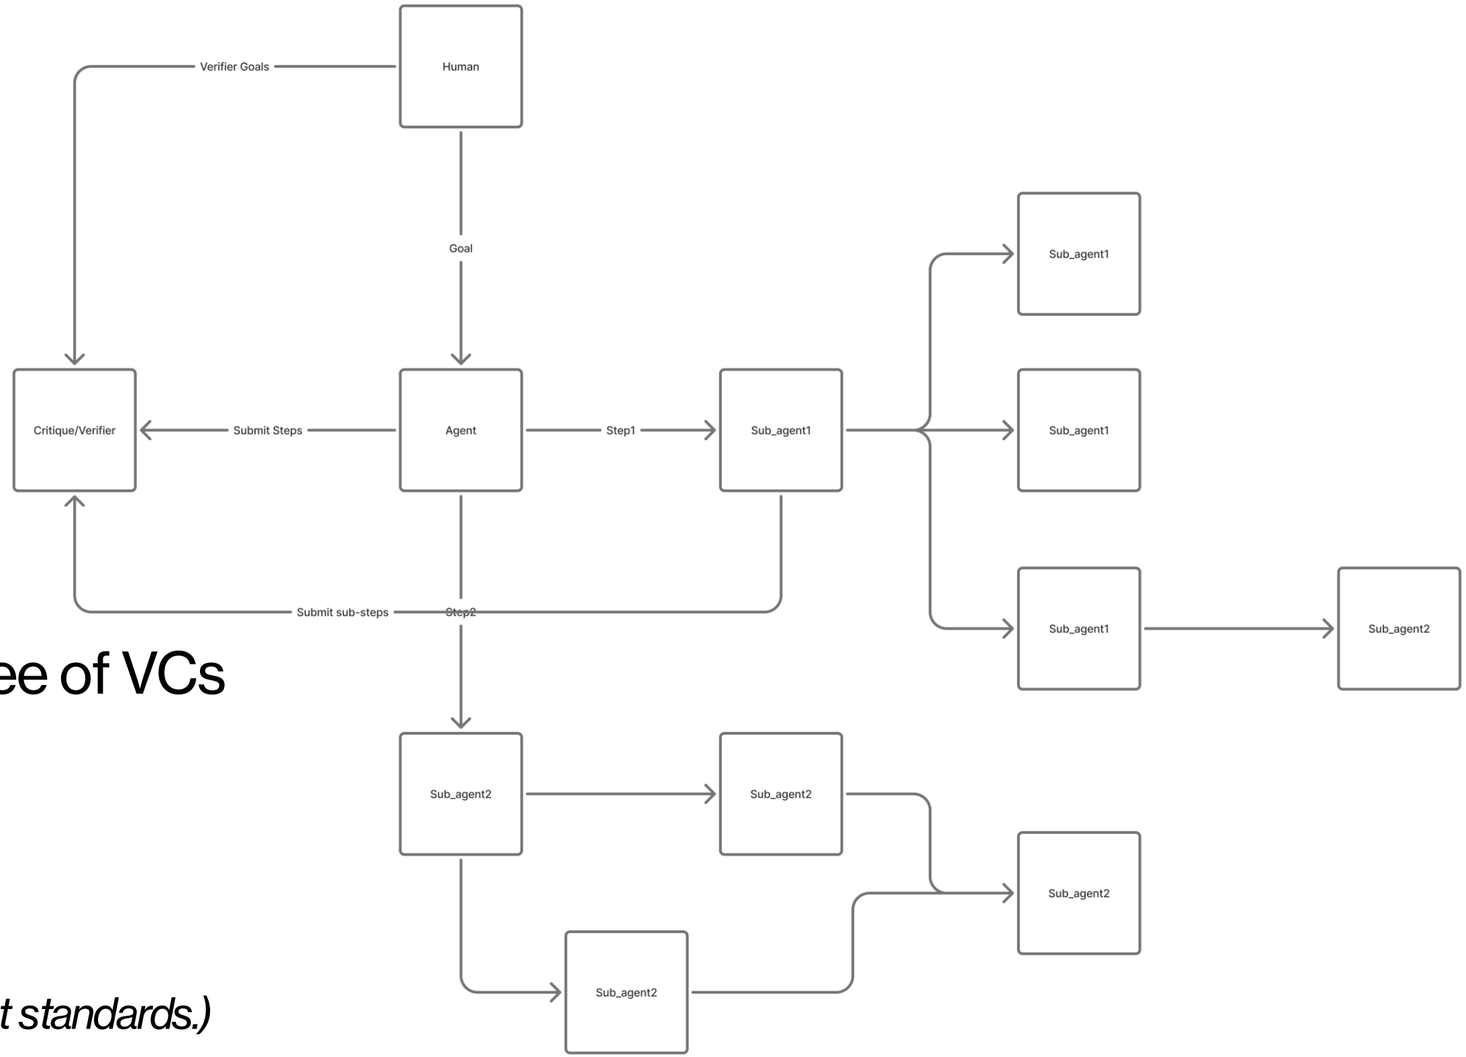
Focus on formal fundamentals

- Unified model for policies and delegations.
- Provable audits
- Provably correct software modules

Fractional Identity Model

An internal effort to build systems and models and explore/compare current standards and solutions.

01 Fractional Agent Identities



- Decentralized roots of trust
- Identity is not just an identifier, it is a tree of VCs
- Identity can be scoped to a goal
- Identity can be temporal

(SPIFEE-family of protocols are a great addition to current standards.)

Assertion Credentials

{Pub_UUID, says, Pub_UUID, **has**, **Attribute**}_Sign(Priv_UUID)

{Pub_UUID, says, Pub_UUID, **can**, **Scope**}_Sign(Priv_UUID)

{Pub_UUID, says, Pub_UUID, **delegate**, **Scope**}_Sign(Priv_UUID)

Assertion Credentials

{Pub_Futurae, says, Pub_Human, **delegate**, **call:tool**}_Sign(Priv_Futurae)

{Pub_Human, says, Pub_Agent1, **can**, **call:tool:x**}_Sign(Priv_Human)
{Pub_Human, says, Pub_Agent1, **delegate**, **call:tool**}_Sign(Priv_Human)

{Pub_Agent1, says, Pub_Agent2, **can**, **call:tool:y**}_Sign(Priv_Agent1)

02 Recursive and Scoped Delegations

1. Unified model for policies and assertions.
2. Multi-party delegation policies.
3. Context-conditioned policies.
4. Selective disclosure.

Policy Credentials

```
{
 (Pub_UUID, says, X, has, Attribute),
        (Pub_UUID2, says, X, has Attribute)
}_Sign(Priv_UUID)
```

```
{
 (Pub_UUID, says, X, can, call:tool:x),
        (Pub_UUID2, says, X, has, validated_plan),
        (Pub_UUID3, says, X, can, call:tool:x)
}_Sign(Priv_UUID)
```

Policy Credentials

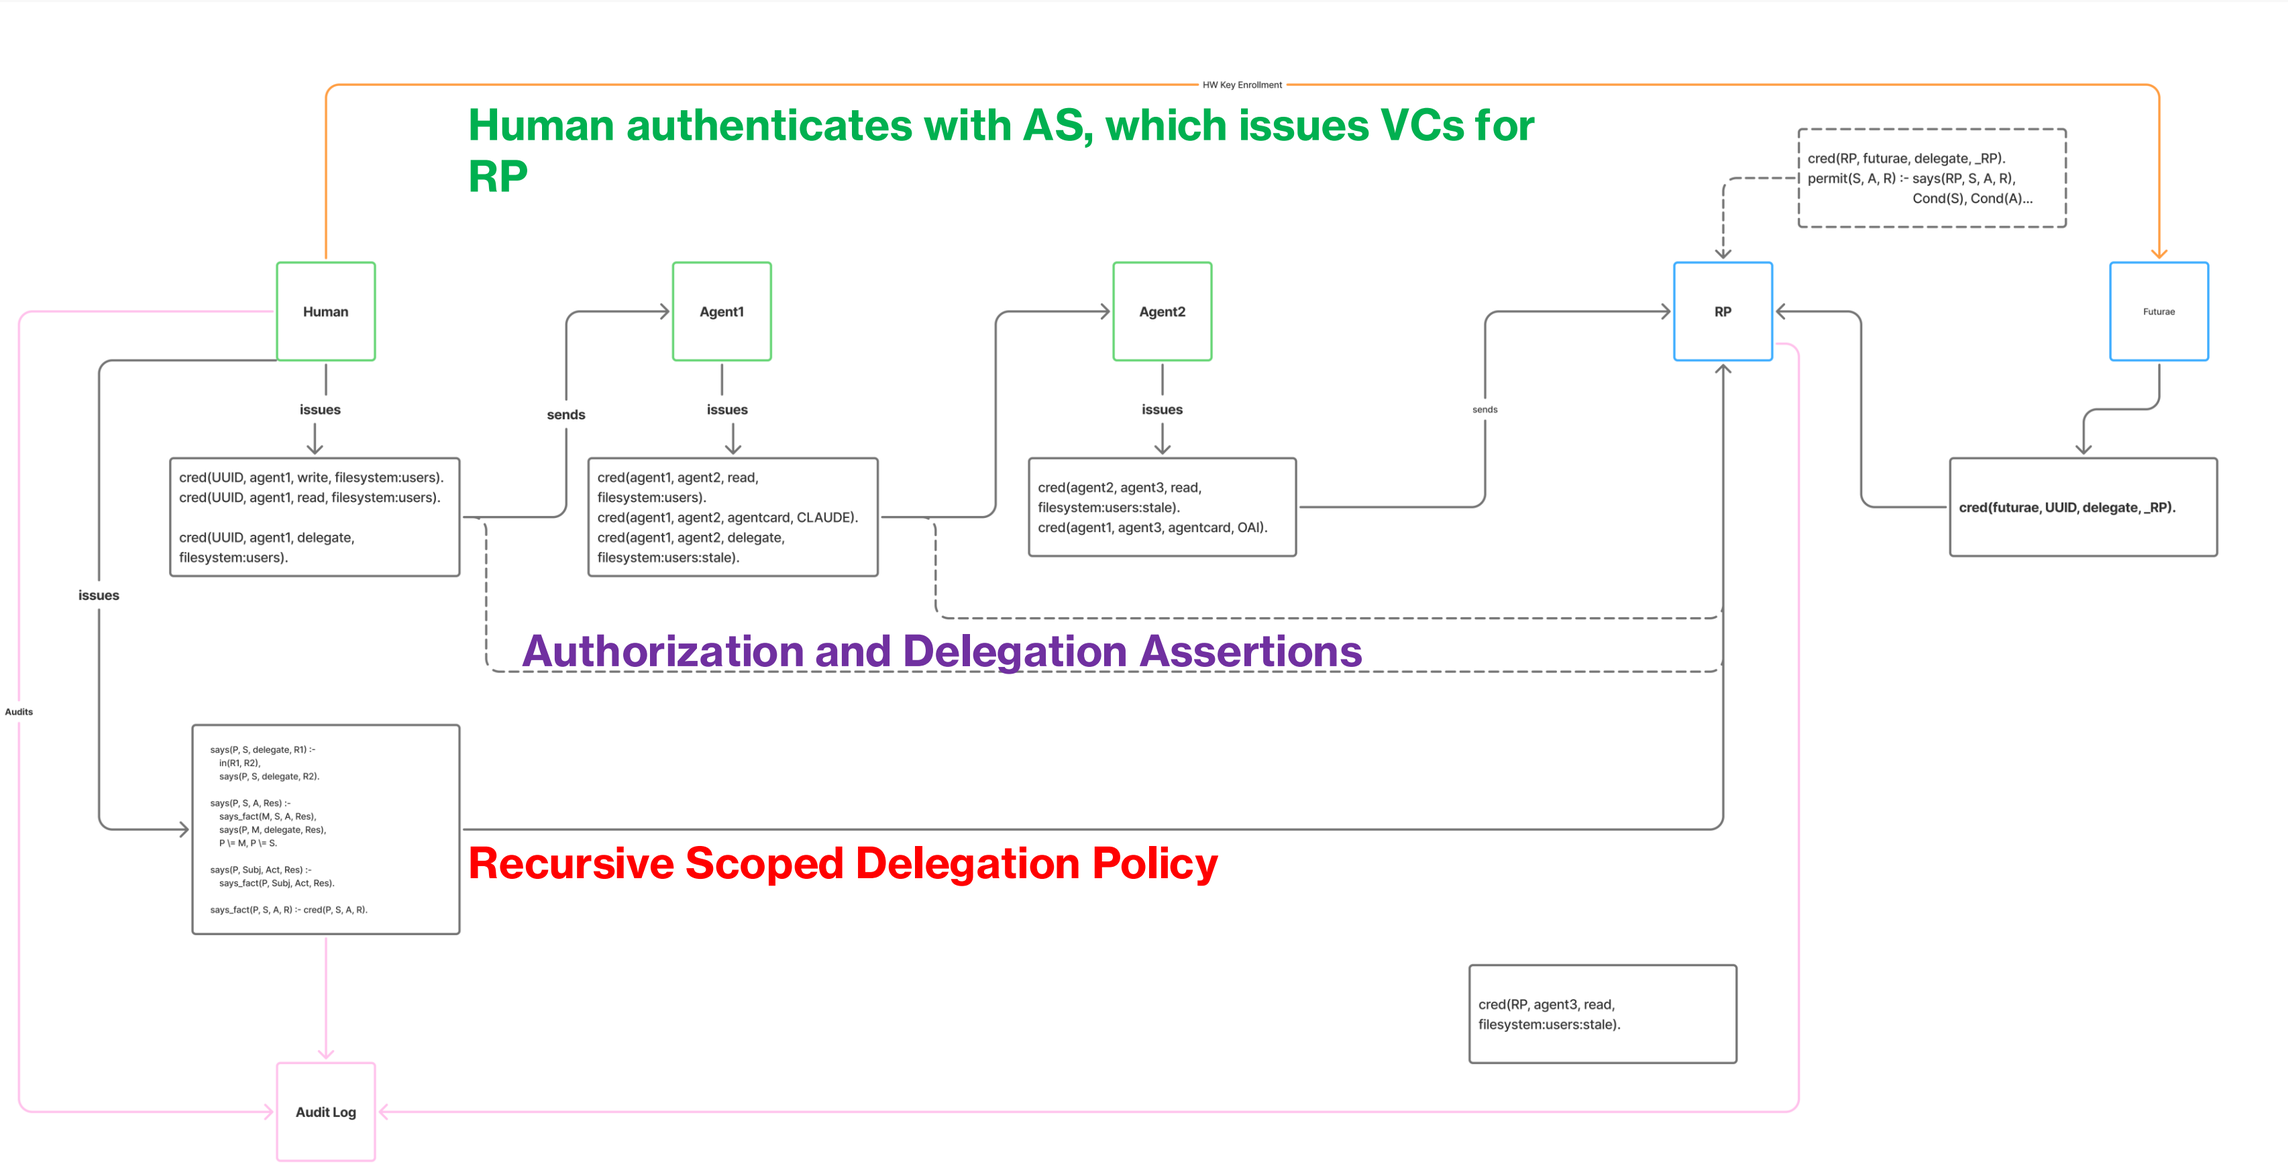
```
{
 (Pub_HUMAN, says, X, has, Attribute),
        (Pub_SupervisorAgent, says, X, has Attribute)
}_Sign(Priv_HUMAN)
```

```
{
 (Pub_HUMAN, says, X, can, call:tool:x),
        (Pub_ValidatorAgent, says, X, has, validated_plan),
        (Pub_PlanningAgent, says, X, can, call:tool:x)
}_Sign(Priv_HUMAN)
```

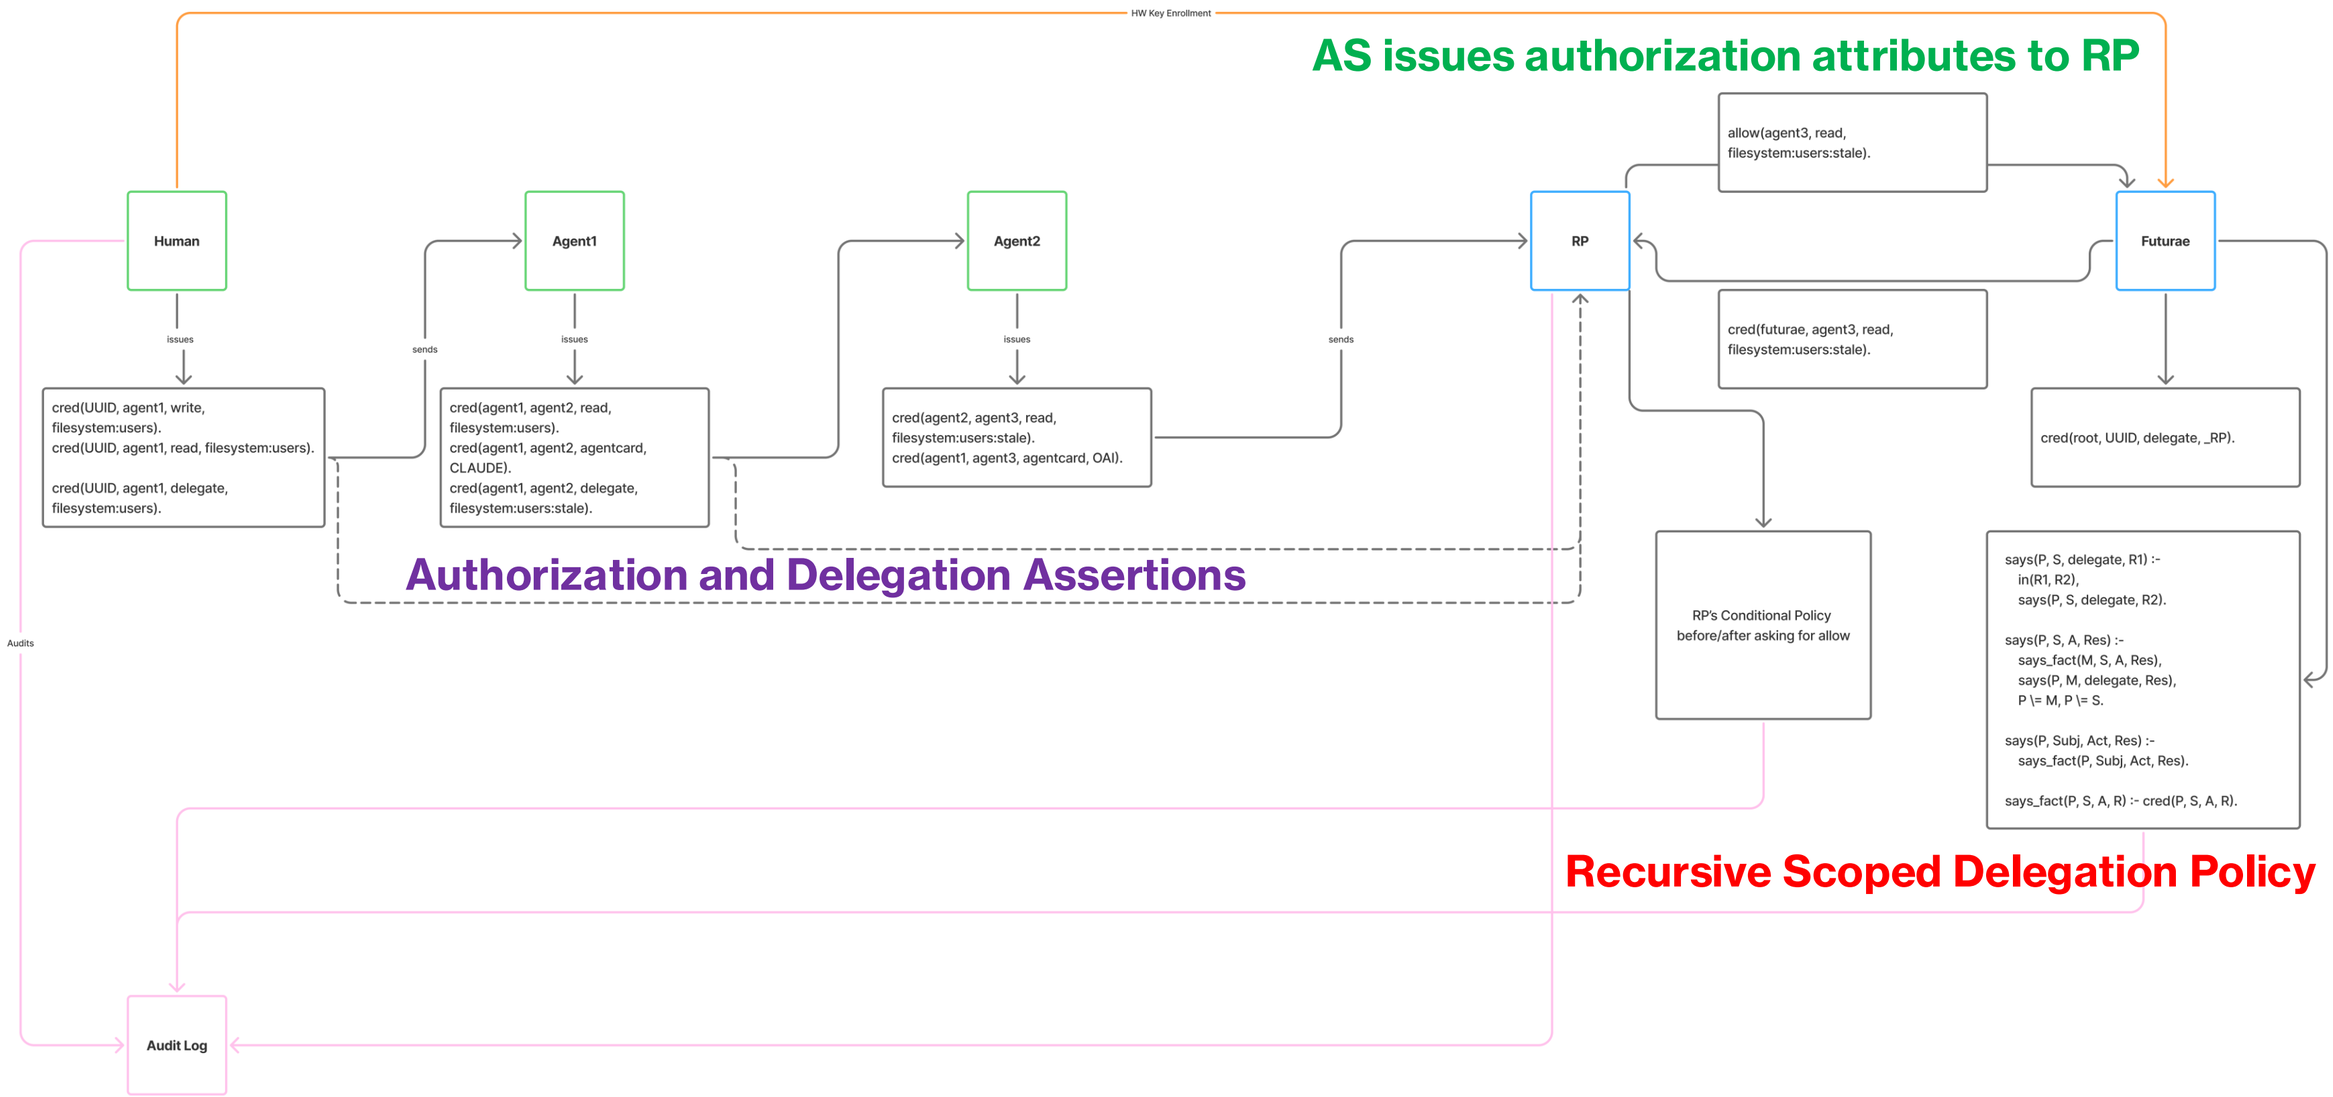
Formal Credentials Semantics

1. Positive First-order Logic with functions.
2. **PTIME** Fixpoint semantics.
3. Family of Datalog languages.

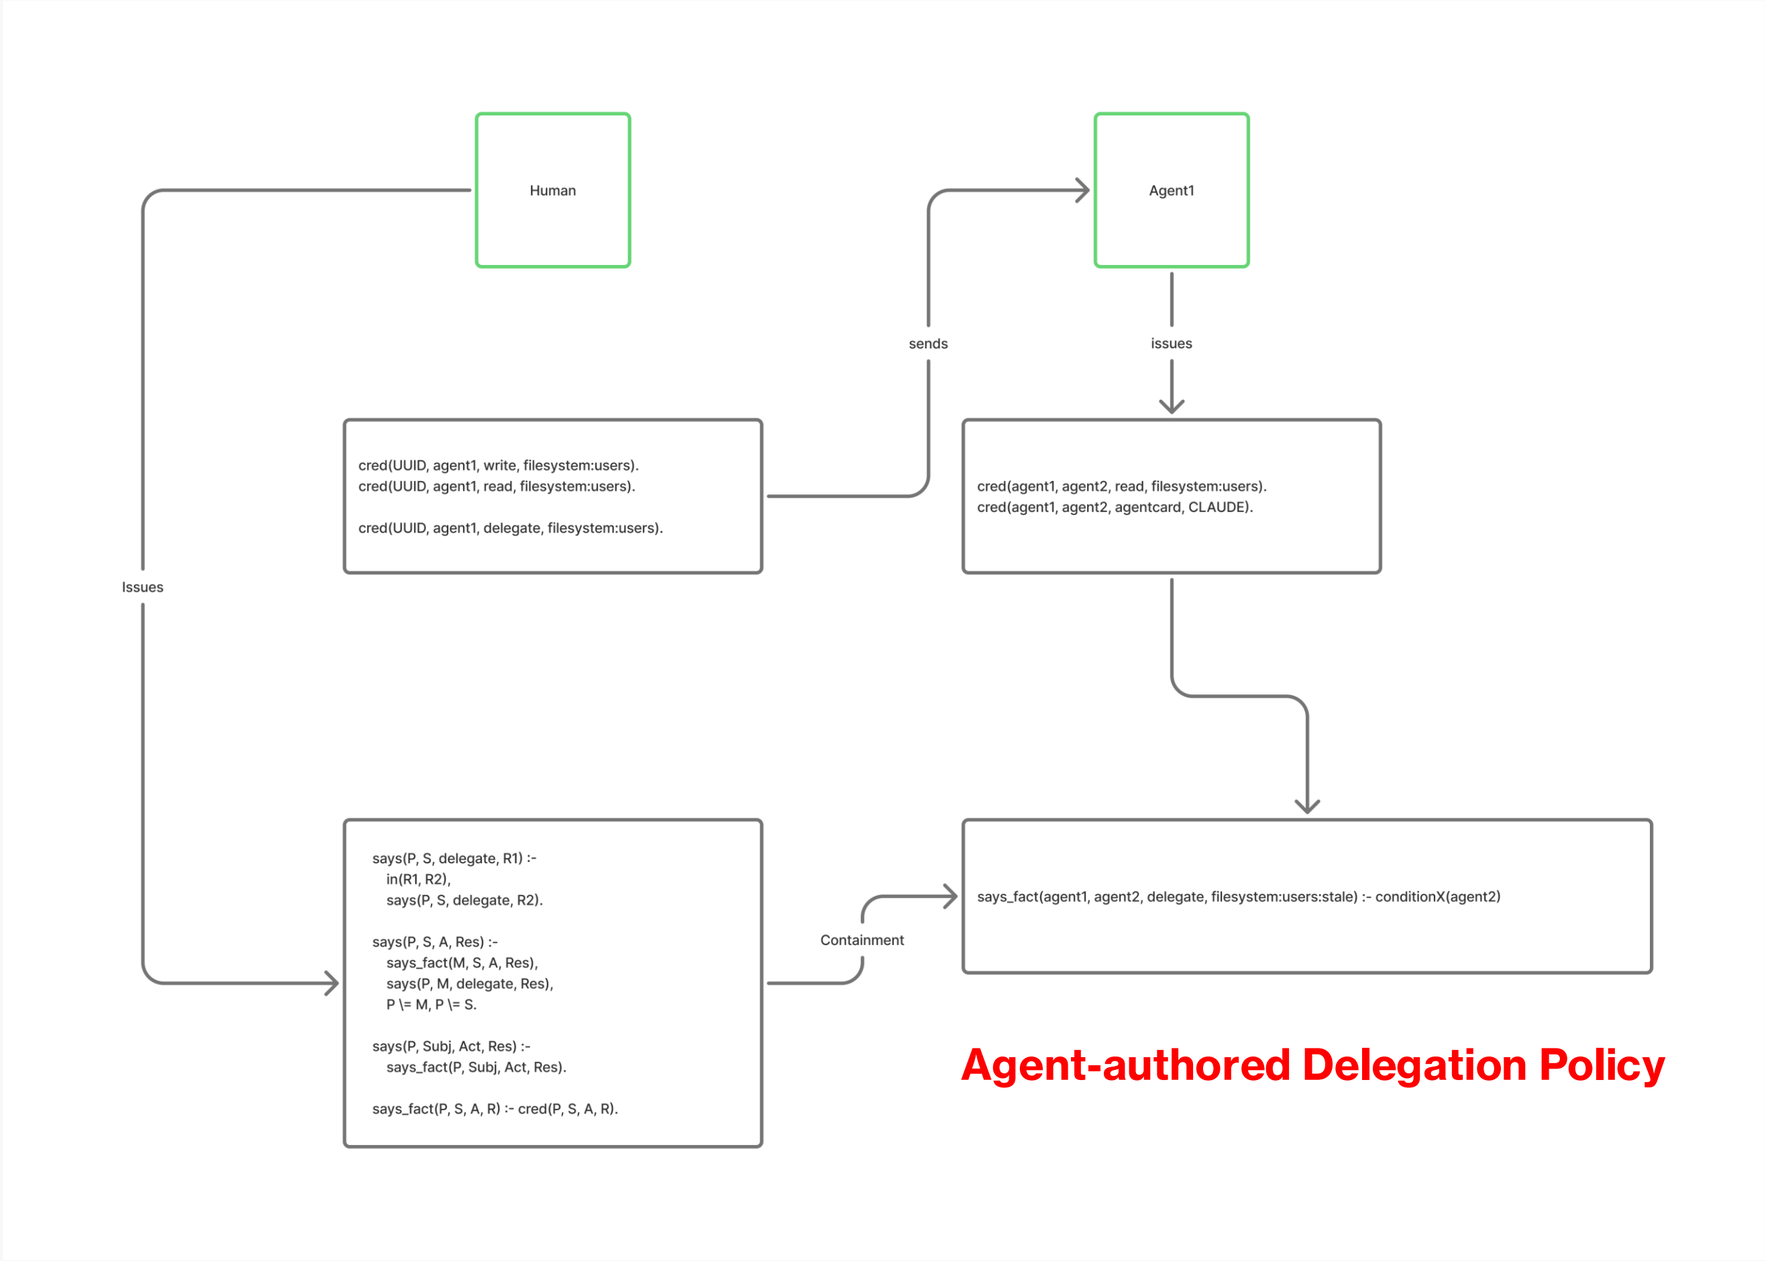
Human-Authored Delegations



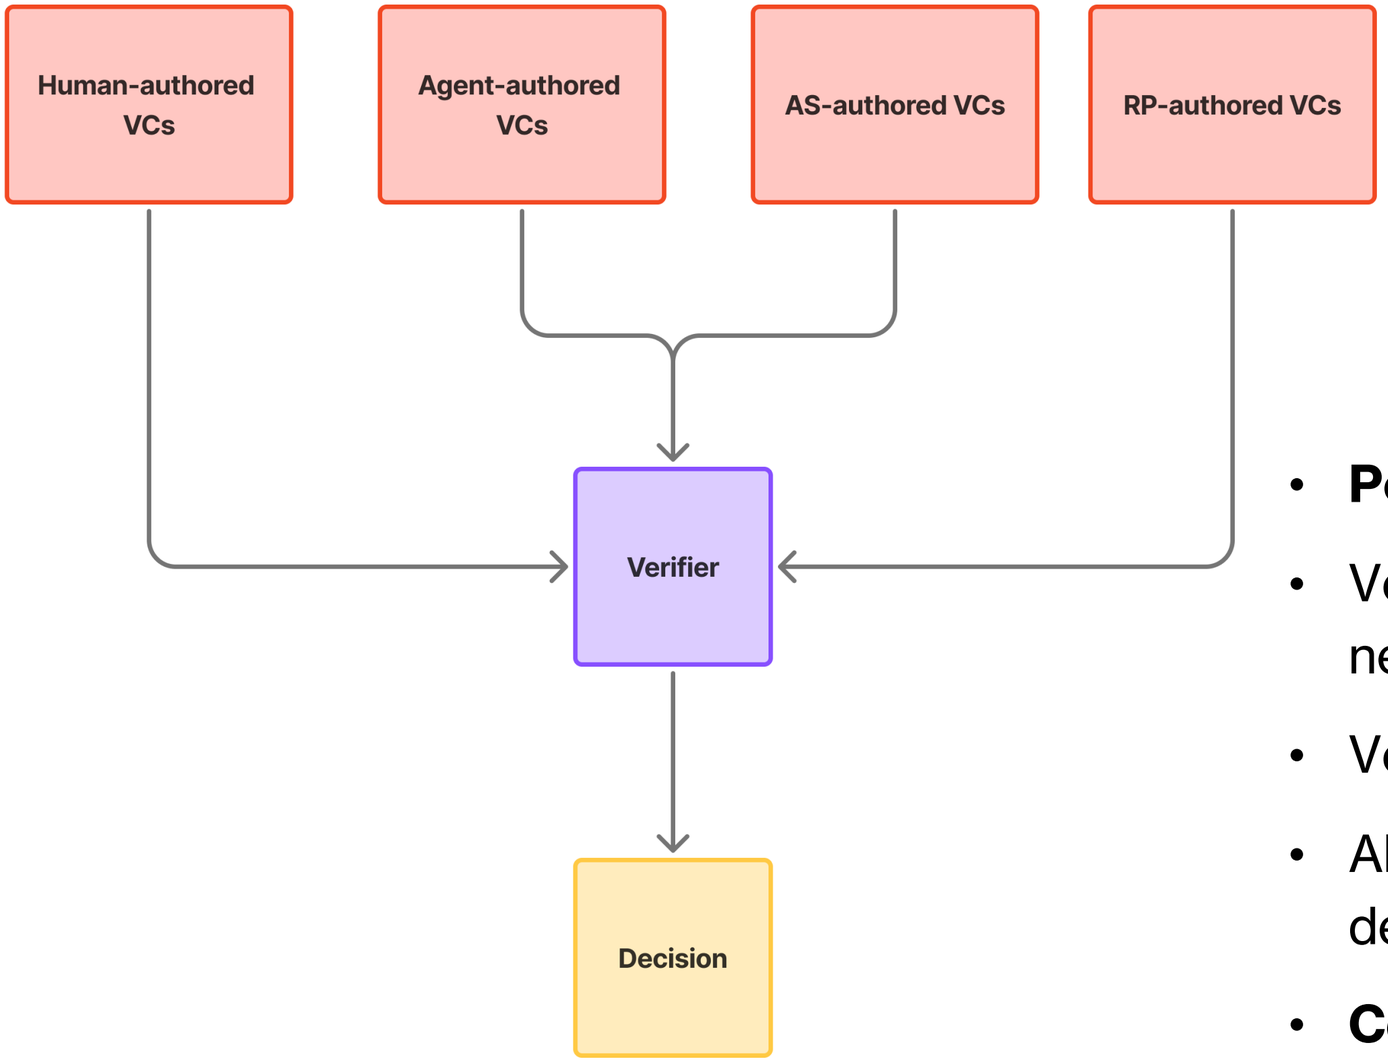
AS-Authored Delegations



Agent-Authored Delegations



03 Verification via Formal Methods



- **Policies, assertions decisions are also VCs.**
- Verification is **an independent process**, no need to trust a verification authority.
- Verifier is a theorem-prover.
- All proofs must be constructive and not depend on absence of evidence.
- **Correctness does not depend on the algorithm/implementation.**

A lot of exciting work is happening

Protocol		Focus	Decentralized Identities through verifiable lineage (PKs)	Agent Capabilities/Cards	(Delegated) Policy Authorship
FuturaeSays*	*Beta in Q1-27*	Formal Model for Decentralized Identities with explicit Policy Authorship and Provable Audits	Unbounded PK chains	chain-of-signatures via delegations	Human, Agent, RP, AS (constrained within decidable segment)
Cryptographically Verifiable Actor Chains for OAuth 2.0 Token Exchange	https://www.ietf.org/archive/id/draft-mw-spice-actor-chain-05.html	Extend Token Exchange (TE) RFC with enforcement of delegation chains present in TE	Unbounded PK chains	Out of scope	AS
AP2**	https://ap2-protocol.org/ap2/specification/	Payment authorization, authenticity, and audit semantics	Out of scope	Yes	Human
Verifiable Intent**	https://verifiableintent.dev/spec/	How to encode authorizations and provable audits	2-step PKs	Yes	Human
KYA-OS	https://modelcontextprotocol-identity.io/mcp	PK-based Agent Identity and Delegation. Extension of MCP	Yes	Yes	Human
AI Identity Management System	https://datatracker.ietf.org/doc/draft-ietf-wimse-aims/	Construct short-lived credentials for agents to obtain transaction tokens. Builds on TE.	SPIFFE URI with rotating PK Keys	Yes	OAuth2 AS
DigiCert Whitepaper 2026	https://www.digicert.com/content/dam/digicert/pdfs/whitepaper/the-new-trust-architecture-for-AI-whitepaper-en.pdf	Credential infrastructure for agents	SPIFFE URI with rotating PK Keys	Yes	Out of scope

A lot of exciting work is happening

Protocol		Delegation Structure	Scope Constraints	Agent and System State	Privacy-preserving goals
FuturaeSays*	*Beta in Q1-27*	n-out-of-k gates, bounded chains, unbounded chains	AND/OR logic operators with math comparisons	Requires external memory*	Explicit Verified Plans, RP-conditioned Selective Disclosure
Cryptographically Verifiable Actor Chains for OAuth 2.0 Token Exchange	https://www.ietf.org/archive/id/draft-mw-spice-actor-chain-05.html	unbounded chains	Out of scope	Out of scope	Out of Scope
AP2**	https://ap2-protocol.org/ap2/specification/	no, explicit out of scope at the moment	Out of Scope	Out of Scope	Selective Disclosure conditioned on RP role
Verifiable Intent**	https://verifiableintent.dev/spec/	1-step chain	Constraints over set membership, and threshold values	Requires external memory	Selective Disclosure conditioned on RP role
KYA-OS	https://modelcontextprotocol-identity.io/mcp	unbounded chains	CRISP constraints		Out-of-Scope
AI Identity Management System	https://datatracker.ietf.org/doc/draft-ietf-wimse-aims/	1-step chain	Explicit Human approved, scoped at AS	-	-
DigiCert Whitepaper 2026	https://www.digicert.com/content/dam/digicert/pdfs/whitepaper/the-new-trust-architecture-for-AI-whitepaper-en.pdf	Out of scope	Out of scope		

A lot of exciting work is happening

Protocol		Unlikability across RPs	Declarative/Algorithmic Verfication
FuturaeSays*	*Beta in Q1-27*	Yes in Futurae-authored flow	Theorem-proving with First-order Logic
Cryptographically Verifiable Actor Chains for OAuth 2.0 Token Exchange	https://www.ietf.org/archive/id/draft-mw-spice-actor-chain-05.html	Out of Scope	Algo
AP2**	https://ap2-protocol.org/ap2/specification/	No	
Verifiable Intent**	https://verifiableintent.dev/spec/	No	Algo
KYA-OS	https://modelcontextprotocol-identity.io/mcp	Out of Scope	
AI Identity Management System	https://datatracker.ietf.org/doc/draft-ietf-wimse-aims/	No	Algo
DigiCert Whitepaper 2026	https://www.digicert.com/content/dam/digicert/pdfs/whitepaper/the-new-trust-architecture-for-AI-whitepaper-en.pdf		

All I got from this talk...

Solutions need a **common abstract delegation framework** for evaluating different standards and implementations against *advanced system and attacker models*.

Futurae is working on *Fractional Agent Identity* focusing on **formal unified models for complex delegations and identities**. We are looking for partners.

Promising work around SPIFEE and AP2 family of protocols, which we want to extend and build on top of.

Thank you.

SESSION FEEDBACK

A Feedback Prize Draw will take place at the end of the day