

INTEROPERABILITY & STANDARDS · BREAKOUT · OCT 23, 12:05 · LL20C

When Agents Spawn Agents

Securing recursive delegation in multi-agent systems

Anishma Mavuram

EmpowerID · RDCP, the Recursive Delegation Capability Profile

WHAT CHANGED


The unit of risk is no longer the tool call.
It's the **delegation edge**.

YESTERDAY



One principal, one token, one call. Classic OAuth answers "may *this client* call *this API*?"


TODAY



Agents expose sub-agents to their LLM *as tools*. Every spawn is an **edge**, and every edge asks: is this descendant acting under a **valid, revocable, least-privilege** authority that traces back to a consenting human?

Without a zero-trust chain across those edges, you get: confused deputies · delegation laundering · privilege creep · forged lineage · audit trails that cannot prove who authorized what.

Anatomy of a trustworthy delegation edge



1 · Who issued it?

A root principal — a human or a service owner — gave a durable, versioned **consent**. Every edge traces back to it.

2 · Who holds it?

A specific agent with a **cryptographic identity**; the credential is bound to that identity's key, so it can't be lent or stolen as a bearer token.

3 · What exactly does it allow?

An explicit **scope** — tools, constraints, expiry — that is never wider than the parent's. Least privilege by construction.

4 · Where does it sit in the chain?

Lineage: the path from root to this edge, carried in the credential and anchored in a grant store the verifier can check.

5 · Can it be taken back?

Revocation at any point in the chain kills everything beneath it — within a bound that is published, not assumed.

6 · Can we prove what happened?

Evidence: tamper-evident receipts linking spawn → grant → token → tool call → outcome, written before the side effect.

Four ways a delegation chain goes wrong

1 · Confused deputy

A downstream agent holds broad authority and is talked into using it for someone else's purpose. The tool server trusts the badge, not the intent.

2 · Delegation laundering

An agent that is *denied* an action spawns a child that gets a *fresh* credential from somewhere else — and the denied action happens one hop down.

3 · Privilege creep

No rule says a child must be narrower than its parent, so scopes accumulate hop by hop. Level five holds more than level one ever consented to.


4 · Forged or unprovable lineage

The log says `svc-account-7` called `pay_invoice`. It cannot say which human, which consent, which chain — or whether the chain was ever real.

All four share one root cause: authority gets **re-minted** at each hop instead of being **derived, narrowed, and traceable**.

FAILURE 1

Confused deputy: a legitimate token, used for an illegitimate request



The ERP can only check what the token carries — and the token says nothing about **who asked C to do this**.


Why it happens

- C's authority is *standing*: it belongs to C, whoever is asking
- Nothing links C's call to the request that caused it (B's)
- B's limit (read-only) is invisible to the ERP
- The decision to act was made by an LLM reading untrusted text

What would have stopped it

When C acts *for B*, it holds a token **derived from B's**. B's authority has no `pay_invoice`, so C's derived token has none either — the ERP checks the token's scope and denies. No judgement about intent needed.

Delegation laundering: denied as A, done as B



B never received A's token. B has its own client ID and secret in its config, and the IdP issues it the admin scope it was **registered** with — for any caller. The IdP was never told B is working for A, under the user's read-only consent.

Why it happens


- **Spawning is unguarded:** B is just a tool in A's code; A invokes it and nobody asks "may A spawn B?"
- B logs in with *its own* service-account credentials, not ones derived from A's
- The IdP decides on B's identity alone; it has no idea A is behind the request
- The log says "ops-bot deleted the ledger" — A and the user are nowhere in it

What would have stopped it

Guard the spawn: the PDP decides "may A spawn B, and with what ceiling?" **Derive the token:** B gets A's ceiling by exchange — no `delete_ledger`, because A never had it — with the chain `user → A → B` inside.

Privilege creep, and an audit trail that can't testify

When least privilege isn't enforced per hop



- "Add the scope the child needs" becomes union, not intersection
- Long-lived child refresh tokens outlive the parent's task
- Agent types approved with the full catalog, so every edge is already maximal
- Result: an agent five levels deep holds authority no human consented to

When the trail can't prove who authorized what

```
2026-10-21T14:03:11Z svc-finance-bot
  tool=pay_invoice amount=48,200 USD
  result=ok
  authorized_by=? consent=? chain=?
```

- Logs record *which identity* acted, not *under whose authority*
- Nested "act" claims are informational — any hop can write anything
- Receipts written *after* the side effect vanish with the crash
- The evidence may exist somewhere — it cannot be produced, intact, when it matters

Five rules, then an architecture that enforces them

1 Grant before token

Authority is a row in the grant store, in a tree. A token is only a pointer to a row that exists.

2 Authority never widens at any hop

child = parent $\cap$ approved-edge ceiling $\cap$ optional narrowing. Design-time edges must be strictly narrower than the parent type.

3 No LLM in the security path

The model asks for a child; brokers, the authorization server and the policy engine decide.

4 A child's credential never transits the parent

The parent gets an opaque delivery reference; the child fetches its own secret, as itself.


5 Re-verify on every call; revoke with one write

Short tokens, composite introspection, a kill switch that takes the whole subtree down within a published bound.


IN ONE LINE

Every spawn becomes a **recorded grant**: issued by the root principal, scoped no wider than its parent, bound to the child's key, written to the grant store before the child runs, and revocable all the way down.

The architecture, in one picture



What happens when Agent A says "delegate to B"



What the LLM did

Called a tool. Nothing else. It never saw a credential and never chose a scope.


What exists afterwards

A grant row for B in the grant store:
parent = A's grant, path from root,
scope $\subseteq$ A's, expiry $\leq$ A's. A token that points at it.

What never existed

A token without a grant. A child wider than its parent. A credential the parent could reuse as B.

Authority is computed by intersection — never by addition



Two guarantees

- **Design time:** each approved edge must be a proper subset of the parent type's catalog. A shorter TTL alone doesn't count.
- **Run time:** child $\sqsubseteq$ parent $\cap$ ceiling at every hop, or the exchange is refused.

A real algebra, not string matching

- Tools: set containment over the tenant's finite catalog
- Typed constraints (min, max_floor, set $\sqsubseteq$, bool): a child may only tighten
- Unknown constraint $\rightarrow$ reject. Missing parent constraint $\rightarrow$ reject.
- Canonical form (RFC 8785) so every language hashes the same bytes

The token carries the authority and the lineage

```
{
  "iss": "https://as.example", "sub": "alice",
  "client_id": "agent-b", "agent_id": "payment-
  executor",
  "exp": now + 5 min, "jti": "tok-7c1",
  "cnf": { "jkt": "<B's key thumbprint>" },
  "authorization_details": [{ "type": "mcp_tool",
    "tool_arns": ["tool:erp:pay_invoice"],
    "constraints": { "currency": ["USD"],
  "approval_above": 1500 } }],
  "grant_id": "7d3c...",
  "grant_path": ["a1f0...", "9b2e...", "7d3c..."]
}
```

RFC 9068 JWT access token · RFC 9396 authorization_details · RFC 9449 /
8705 key binding · gold = RDCP extension claims

authorization_details · what it may do

The scope travels in the token. The PEP checks “is this tool in the list, within these constraints?” locally, on every call — no round-trip to the AS on the normal path.


grant_id + grant_path · where it came from

grant_id names this token's row in the grant store.
grant_path is every ancestor, root → self: the user's consent, Agent A, then this grant. If any id on the path is revoked, this token is dead.

Checked at the PEP, every call

Signature · key binding · tool ∈ authorization_details · no id on grant_path in the revoked list. The AS is asked only for irreversible tools or once a cache bound expires (next slide).

One write at the root — the whole subtree goes dark



One update in the grant store marks the whole subtree inactive — every row whose path starts with the revoked id. No token is touched; the tokens simply stop pointing at anything live.


How the PEP finds out

- **Signal** — the AS sends a *grant revoked* event to every PEP (Shared Signals Framework / CAEP). Each PEP keeps a **deny list of revoked grant ids**.
- **Path check** — on the next call, the PEP compares the token's `grant_path` to its deny list. Any ancestor on the list → deny. The whole subtree dies with one entry.
- **Introspection** — for irreversible tools (pay, delete, send) the PEP asks the AS live, every time: zero staleness.
- **Backstop** — tokens live ≤ 5 min and children get no refresh tokens.

What else pulls the switch

- Disabling an agent type or an edge in the registry revokes every grant on it
- Owner leaves · consent expires or is narrowed · signing-key compromise

Receipts that can testify



each receipt carries `grant_id · jti · agent_id · versions · trace id` · and the previous receipt's hash

Written before the side effect

Intent → attempt → completion → reconciliation. If the durable write fails, the call is refused. A crash can lose a result, never the record that it was attempted.

Answers the question the log couldn't

Who consented, under which version, through which edges, with which key, approved by whom. Human approvals are single-use artifacts consumed atomically at the call.

Tamper-evident, not just stored

Each receipt carries the hash of the one before it, so the chain can be independently checked — not merely trusted.

Compose what exists; extend only where the standard stops

NEED	STANDARD WE REUSE
Derive a child token from a parent token	OAuth Token Exchange · RFC 8693
Express fine-grained, structured authority	Rich Authorization Requests · RFC 9396
Bind the token to the holder's key	DPoP · RFC 9449 · mTLS · RFC 8705
Workload identity for every agent	SPIFFE / SPIRE X.509-SVID
Ask "may this edge be delegated?"	OpenID AuthZEN Authorization API 1.0
Verify a token is still live	Token Introspection · RFC 7662 · JWT AT · RFC 9068
Hash the same bytes in every language	JSON Canonicalization · RFC 8785
Tell every PEP a grant is gone	Shared Signals Framework / CAEP (OIDF)

Where RDCP extends

grant_id

grant_path

narrowing algebra

subtree revoke + published bound

Private claims and HTTPS-namespaced identifiers, documented for registration. Nothing is invented where an RFC already answers.

Why standards-first

So the gateway, IdP, policy engine and agent framework can come from four vendors and still agree on what a delegation *is*.

The basics any recursive-delegation solution must have

- 1 Cryptographic identity for every agent process**
Workload identity, not shared service accounts. In-process children share the harness's blast radius — know that.
- 2 Grant before token, in a durable tree**
Parent id, path, expiry $\leq$ parent, one root per consent.
- 3 Monotonic attenuation with a defined algebra**
Intersection over a finite catalog; typed constraints; reject what you can't compare.
- 4 Sender-constrained, short-lived tokens**
No bearer tokens; no child refresh tokens; continuation re-checks the tree.
- 5 Credentials never transit the parent**
Delivery references; the child fetches as itself.
- 6 Re-verify on every call**
Check the grant and the path; irreversible tools always introspect live.
- 7 Subtree revocation with a published bound**
One write at the root; facts to every verifier; say out loud how stale you allow yourself to be.
- 8 Lineage in the token, truth in the grant store**
Nested claims are hints; the authorization server's grant row is the source of truth.
- 9 Evidence before side effects**
Hash-chained receipts; approvals as single-use artifacts.
- 10 No LLM in the security path**
The model requests; deterministic components decide.

Crossing the trust domain: when the child agent lives in another organization

Today · one trust domain

- One Delegation AS holds every grant
- The token's `grant_id` points at a row that AS owns
- One PDP decides whether a parent may spawn a child, and with what
- One grant store to revoke from; one set of PEPs to signal

Everything works because the issuer, the policy and the store are the same party.

Open · the child is a remote agent in another company

- **Least privilege:** their AS doesn't know our tool catalog. Whose rules decide that the remote token is "narrower"? There is no standard comparison of `authorization_details` (RFC 9396 §6.1).
- **Lineage:** our `grant_path` is meaningless to their PEP. A cross-domain path needs issuer-qualified grant ids and a way to verify it without trusting the hop that carried it.
- **Revocation:** our kill switch is a row in our store. How does it reach their PEP, and how do we know it did?
- **Evidence:** receipts from two ledgers and two clocks — who adjudicates a dispute?

Purpose & run binding

A valid token works for any permitted tool during its life. Binding authority to *this task* is unsolved.

Injection inside the ceiling

Attenuation bounds the damage; it doesn't stop a tricked agent doing a permitted bad thing.

Standards gaps

No registered delegation claims; no interop test suite for attenuation.

Undo

Revocation stops the *next* call. Nothing un-pays an invoice.

TAKE-AWAYS

Secure the edge. Prove the chain. Keep the kill switch.

Secure the edge, not just the call

The delegation edge is the unit of risk. Each one must be issued, scoped, bound, recorded and revocable.

Compose standards

Token Exchange, RAR, DPoP/mTLS, SPIFFE, AuthZEN, introspection, JCS — plus a small, documented set of extensions.

Prove it, then be able to pull it

Lineage in the token, truth in the grant store, receipts before effects, and a kill switch that reaches every level on a published clock.

Anishma Mavuram

EmpowerID · anishma.mavuram@empowerid.com · RDCP — Recursive Delegation Capability Profile

THANK YOU

Questions?

Anishma Mavuram

EmpowerID · anishma.mavuram@empowerid.com

RDCP — Recursive Delegation Capability Profile