IIWXLII

INTERNET IDENTITY WORKSHOP 42

April 28-30, 2026

Book of Proceedings

www.internetidentityworkshop.com

Computer History Museum / Mountain View CA

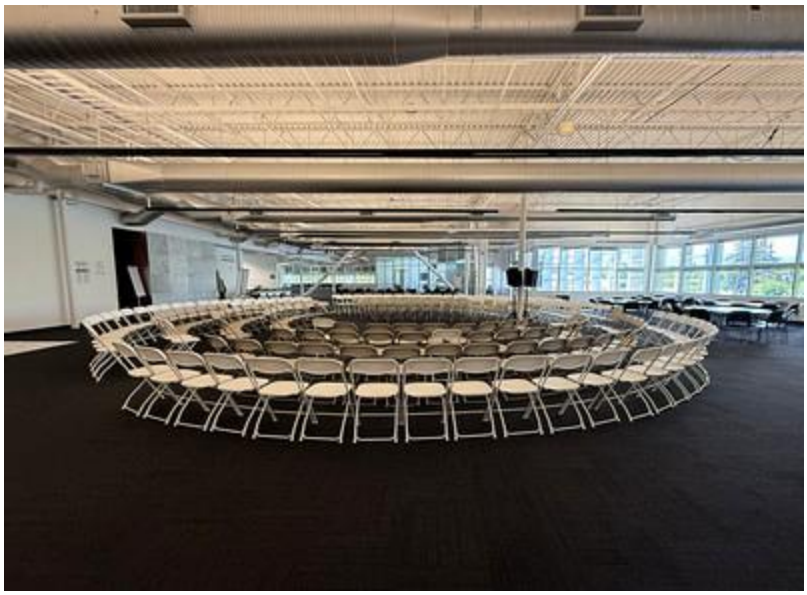

Opening Circle by Doc Searls

Collected & Compiled by
HEIDI N. SAUL & EMMA WINDLEY

IIW founded by Kaliya Young, Phil Windley and Doc Searls
Co-produced by Heidi Nobantu Saul, Phil Windley and Kimberly Culclager-Wheat
Facilitated by Heidi Nobantu Saul & Kaliya Young

Thank You! Doc Center & Book of Proceedings Sponsors:

ANONYMOUS SPONSOR

Contents

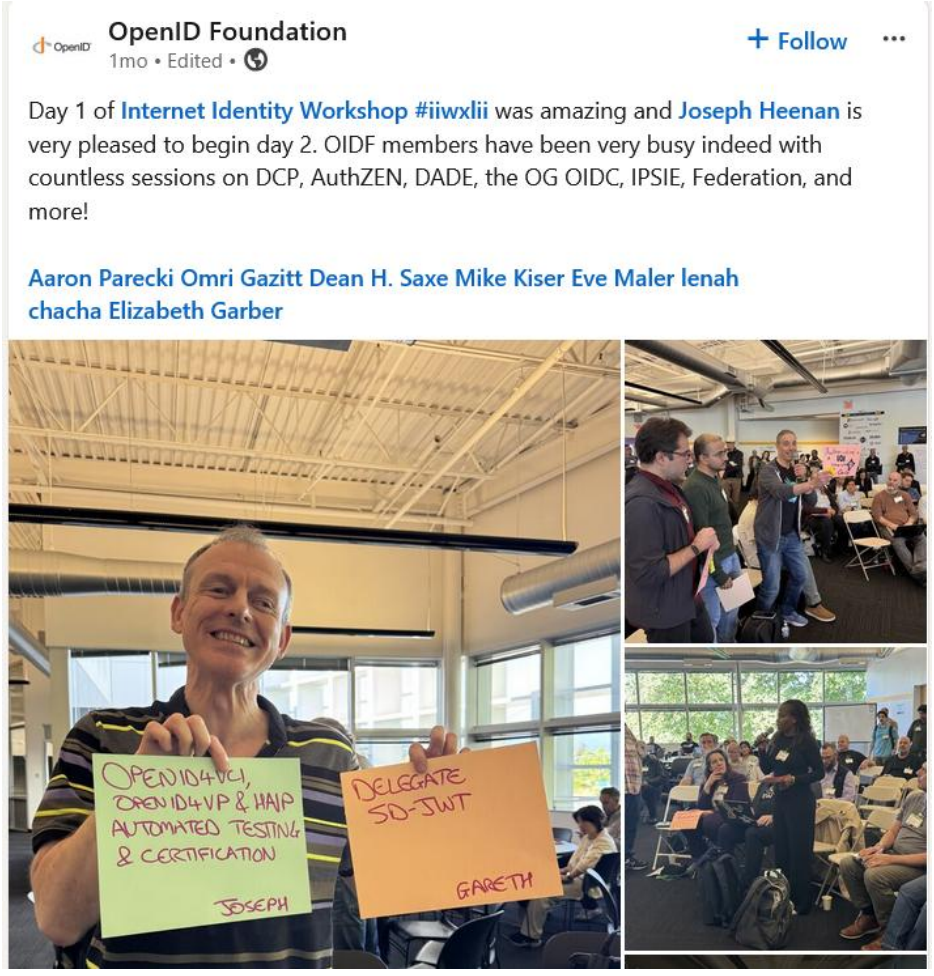

About IIW

The Internet Identity Workshop (IIW) was founded in the fall of 2005 by Phil Windley, Doc Searls and Kaliya Young. It has been a leading space of innovation and collaboration amongst the diverse community working on user-centric identity.

It has been one of the most effective venues for promoting and developing Web-site independent identity systems like OpenID, OAuth, and Information Cards. Past IIW events have proven to be an effective tool for building community in the Internet identity space as well as to get actual work accomplished.

The event has a unique format – the agenda is created live each day of the event. This allows for the discussion of key issues, projects and a lot of interactive opportunities with key industry leaders that are in step with this fast-paced arena.

Watch this short documentary film: ***"Not Just Who They Say We Are: Claiming our Identity on the Internet"*** http://bit.ly/IIWMovie to learn about the work that has happened over the first 12 years at IIW.

The event is now in its 21st year and is Co-produced by Phil Windley, Heidi Nobantu Saul and Kaliya Young. IIWXLIII (#43) will be October 3 - 5, 2026.

Phil Windley's IIWXLII Report

🐝 **Summary:** IIW XLII brought 287 people to the Computer History Museum in Mountain View for three days of sessions on identity, agents, and the legal and technical foundations of first person digital life. The agenda reflected a community grappling with real deployment challenges: SEDI and duty of loyalty, agentic identity, MyTerms, post-quantum cryptography, and the EUDI wallet. AIW2 followed on Friday, continuing the agentic internet conversation. 🦋

Read the full IIWXLII Report here:
https://www.windley.com/archives/2026/06/internet_identity_workshop_xlii_report.shtml

Thank You to our Sponsors!

IIW would not be possible without the community that gathers or the Sponsors that make the gathering feasible. If you are interested in becoming a sponsor or know of anyone who might be please contact Phil Windley at Phil@windley.org for Event Sponsorship information.

Coming Up!

IIWXLII #44
November 3 - 5, 2026 / Mountain View, CA
https://internetidentityworkshop.com/

IIWXLII Daily Schedule

IIWXLII 3 Day Schedule

TUESDAY, April 28 / Doors Open at 8:00 AM for Registration Barista! Bagels (PB&J, Cream Cheese) - Yogurt - KrispyKreme Donuts - Fruit - Cheese - Boiled Eggs etc.			
Barista! And Continental Breakfast	8:00 - 9:00	Lunch	1:00 - 2:00
Welcome Introduction	9:00 -10:00	Session 3	2:00 - 3:00
Opening Circle / Agenda Creation	10:00 - 11:00	Session 4	3:00 - 4:00
Session 1	11:00 - 12:00	Session 5	4:00 - 5:00
Session 2	12:00 - 1:00	Closing Circle	5:00 - 5:45
Welcome Drinks by AWS Identity & Dinner by You 6:00 **NEW LOCATION!** *Michaels @ Shoreline Golf Links 2960 N Shoreline Blvd*			

WEDNESDAY, April 29 / Doors Open at 8:00 Barista! Bagels (PB&J, Cream Cheese) - Yogurt - Krispy Kreme Donuts - Fruit - Cheese - Boiled Eggs etc.			
IIW Women's Breakfast Roundtable's	**7:45 - 9:00**	Speed Demo Hour	1:30 - 2:30
Opening Circle / Agenda Creation	8:45 - 9:30	Session 4	2:30 - 3:30
Session 1	9:30 - 10:30	Session 5	3:30 - 4:30
Session 2	10:30 - 11:30	Closing Circle	4:30 - 5:30
Session 3	11:30 - 12:30	Conference Dinner	6:00 ~
Lunch	12:30 - 1:30		
Conference Drinks by You & Dinner by Microsoft Traditional BBQ (w/plenty of V&V options) - Here at CHM!			

THURSDAY, April 30 / Doors Open at 8:00

Barista! Bagels (PB&J, Cream Cheese) - Yogurt - KrispyKreme Donuts - Fruit - Cheese - Boiled Eggs etc.

Opening Circle / Agenda Creation (SHARP)	8:45 - 9:30	Session 4/Working Lunch	12:30 - 2:00
Session 1	9:30 -10:30	Session 5	2:00 - 3:00
Session 2	10:30 - 11:30	Closing Circle	3:00 - 4:00
Session 3	11:30 - 12:30	**IIWXLIII November 3 - 5, 2026**	

Eve Maler's Book Launch Party! @Steins Beer Garden & Restaurant

Directly Following Closing Circle ~ RSVP Here: https://luma.com/rw26g09k

IIW 42 Agenda Creation = Schedule & Workshop Sessions

158 distinct sessions were called and held over 3 Days.

We received notes, slide decks, links to presentations and photos of whiteboard work for 123 of these sessions.

Tuesday April 28, 2026 Day 1 / Sessions 1 - 5

Session 1
1A/ Introduction to the Decentralized Trust Graph & First Person Project / Drummond Reed
1B/ Oauth 101 and IIW Session / Aaron Parecki
1C/ NO SESSION
1D/ NO SESSION
1E/ NO SESSION
1F/ SIROS / WW ….. Longfellow ZKP Age Demo / John & Leif Johansson
1G/Are Passkey Server Implementations Secure? / Ken Watanabe
1H/ (Multi Device) Decentralized Media & Messaging / Zach Alexancer
1I/ OpenID4VC 101 & Updates / Joseph & Frederik Krogsdal Jacobsen
1J/ NO SESSION
1K/ NO SESSION
1L/ AI for vLEI / Esteban Garcia
1M/ NO SESSION
1N/ NO SESSION

Session 2
2A/ Prove IRL Connections - Harvard University Open Source Keyring App Demo and Discussion / Brendan Miller & Albert
2B/ Intro to Open ID Connect - and IIW 101 Session / Mike Jones
2C/ Trust Infrastructure as a Public Utility: What Would It Take/ Erika Bjune
2D/ NO SESSION
2E/ Maintaining Our Humanity - What is Identity? Spirit Tech + Agency / Diane Oneal
2F/ ZKP 101 / Brent Zundel
2G/ Authorization in an SSI Ecosystem / Jacob Siebach
2H/ Privacy Preserving Biometrics + "Passkey for Personhood" / Scott Jones w/Realkeys
2I/ Kwaai Decentralized AI Infrastructure / Reza Rassool
2J/ NO SESSION
2K/ UX Design for SSI Products / Janet Gonzales
2L/ NO SESSION
2M/Formal Security Verification of Specs - what would we want? / Frederik Krogsdal Jacobsen
2N/ NO SESSION

Session 3
3A/ Content Authenticity 101 C2PA . CAWG / Eric S & Scott P
3B/ Authorization 101 (now with AI) and IIW Session / Omri Gazitt
3C/ All my old problems now have **AI** in them but they're still problems. / Eve M & Justin R
3D/ Why Personal ID in the era of AI? / Denny Wong
3E/ NO SESSION
3F/ VRM + MyTerms + Fiduciary Agents / Iain Henderson
3G/ NO SESSION
3H / AAuth DEEP DIVE / Dick Hardt
3I/ Agent Taxonomy - What "is" an Agent? Ephemeral vs. Persistent Agent Lifecycle / Danny Zollner
3J/ Anonymity for Civilian ownership of identity. Reconstitution. / Benner Harvey w/News Futures
3K/ Age Verification & MASS SURVEILLANCE / Brent Shambaugh
3L/ DID x PKI for IoT, NHI Linking DID into x.509 cert / Susumu Ishizuka
3M/ The Fiduciary Commons 1 - from principles to law / Mike Leahy
3N/ Cross-Border Finance as a cast study for Cross-Border ID Data Management in Africa / Lenah Chacha

Session 4
4A/ Cryptoperiod Catastrophe - NIST SP800-57 Misapplications of "Bare" Signatures + ZKPs for non-ephemeral case proofs / Sam Smith
4B/ Fido and Webauth 101 an IIW Session / John Bradley
4C Regenerative Accelerationism (1) / Kaliya & Friends
4D/ NO SESSION
4E/ Death and the Digital Estate…. What's Next!? / Dean,Eve, Mike
4F/ EUDI Wallet - What is German Cooking? / Mirko
4G/ Digital Fiduciary Update / Joe Andrieu
4H/ OAuth for SUB Agents FT.French Toasts / Abhishek Hingnikar

4I/ JLINC overview + update - Governance Authorization Auditing / Darius Dunlap
4J/ NO SESSION
4K/ NO SESSION
4L/ SSI 10th! Revisiting the principles for the next decade! Christopher Allen
4M/ Interactive Endpoint Authz via First Party Apps / Frederik Krogsdal Jacobsen
4N/ STORY BOARDING AGENT IDENTITY / Rohit Khare ("Gov Ops")

Session 5
5A/ KERI/ACDC Bulk Issuance for SEDI Privacy - Security and WorkFlow / Sam Smith
5B/ Self-Sovereign Identity (SSI) - Decentralized Identity (DI) 101 - an IIW Session / Steve McCown
5C/ KYA Pay - Agentic Identity and Payment Credentials / Ankit Agarway and Mike Jones
5D/ NO SESSION
5E/ I've got mad skills! Writing skills to plan your digital estate with Claude / Dean S
5F/ EUDI Wallet Authentication & Authorization / Mirko M
5G/ IDENTITY TALES: Finding the Narrative / Erica Connell & Joe Andrieu
5H Identity as a Marital Art - * MDMA * Mixed Digital Marital Art - Getting off on Digital X (1) / Jeff O
5I/ The Fiduciary Common Part 2 - Why Incomplete Legislation is a Practitioner's Problem / Mike Leahy
5J/ NO SESSION
5K/ NO SESSION
5L/OPEN HAVEN.net Building Bridges in the P2P Ecosystem / Day Waterbury
5M/ NO SESSION
5N/ Safeguarding Products + Standards for The Dark Times (help Elizabeth with her EIC talk) Elizabeth Garber

Wednesday April 29, 2026 Day 2 / Sessions 6 - 10

Session 6
6A/ Archival Qualify Identifiers? / Eric Scouten
6B/ Digital ID Acceleration - The world needs us / Riley Hughes
6C/ Beyond Triples - Semantic Frame as Data Model - Meaning & Trust Substrate / Joshua Chambers
6D/ What do you mean "Call my Agent"? / Hirsch S
6E/NO SESSION
6F/ My Terms 101 / Doc Searles
6G/ Pairwise Pseudonyms from Government ID's or Holy Pseudonym Batman POW / John B & Dirk
6H/ Originator Profile for Content Authenticity / Michiko
6I/ Why Physical IDs need Digital Stuff - come see some counterfeits / Elaine Wooton
6J/ Trust Infrastructure as a Public Utility PART 2 / Erica Bjune
6K/ What to do with 3 Million Biz ID's? / James Monaghan & Stefan Heller
6L/ What would an actual 'democratizing technology' look like? / Dave Sanford
6M you have the digital credentials that you need? / Karla McKenna
6N/ Ecosystem Guidance Needs and Wants / Frederik Krogsdal Jacobsen

Session 7
7A/ Cross App Access - no more OAuth redirects / Aaron Parecki
7B/ Governance - Backed Content Credentials / Scott Perry
7C/ NO SESSION

7D/ The Antidote to Fear: Digital agency + identity in the wilderness of technology. / Swan Black
7E/ NO SESSION
7F/Onboarding Enterprises to Decentralized Systems / Adrian Ross
7G/ KERI Foundation Wallets / Evan Asakawa
7H/ The Crazy Road of Authorization - Are we really ready for Agents (??) / Bhavna B
7I/ DBSC Secured Cookies W3C Standard / Guillaume Ehinger, Lucas Santos, Bryan (Google)
7J/ NO SESSION
7K/ Rebooting bshambaugh.org Enterprise Information System for Peer Production / Brent Shambaugh
7L/ Network Defrag - How do we avoid fragmentation of Protocols implementation and society? / Kevin
7M/ IDPRO: What should we be adding to our body of knowledge? / Elizabeth G + Tina Srivastrava
7N/ Signed credential metadata updates and payments / Frederik Krogsdal Jacobsen

Session 8
8A/ OAUTH fo MCP - How MCP has reshaped OAUTH in the past year / Aaron Parecki
8B/ DigaVouch: Building consent-gated Identity for the "Grey Zone" FRESH Eyes Welcome! / Juan Alvarez
8C/ Delegate SD-JWT / Gareth
8D/ KEEPING THE FAITH / Eric Welton
8E/ What is Computer Science? Only one right answer… / Sol Ashlynn
8F/ Cross Device Flows - *Hybrid CTAP *Passkeys *Digital Creds / Harsh Lal & Mohamed Amir
8G/ Agent Registry for Identity & Authorization / Aaron & Adolofo Grego
8H/ Intros & Updates on the Japanese version of NIST's SP 800-63 (Digital Identity Guidelines) / Ryo Nakashima
8I/ SEDI - The missing foundation for Digital Identity / Phil Windley
8J/ IEEE 754 Floating Point Determinism / Mark Lenhardt
8K/ The Subject is a couple! / Atul Tulshibagwale
8L/ Building a living systems trust substrate - Architectural dive into social fabric / Frank
8M// Beyond IAM + PKI: KERISuite as a Trust Substrate / Mark Scott
8N/ Exploring and Solving for the Missing Middle of Capital for values based entrepreneurs in tech / Kaliya and Friends

Session 9
9A/ How Does an Agent Decide Who to Trust? Verifiable Trust Protocol / Fabrice Rochette @ Verana
9B/ Invitation to Collaboration on Agentic Protocols and Representations / Mike Jones & Ankit Agarwal
9C/ Playnet / Ruzgar E
9D/ SEDI - Where the sidewalk ends. / George McE
9E/ Bring - Your - Own - Everything Stack / Dmitri Z
9F/MyTerms ~ The journey continues / Mary, Justin + IEEE 7012
9G/ OpenID 4VCI, Open ID4VP & HAIP Automated Testing & Certification / Joseph Heenan
9H/ Sovereign Identity Namespace | *Time is Now? Or *Been there…done that? / Jeffrey Mendelsohn
9I/ The Fiduciary Commons (Part 3) AI: The Inference Problem / Mike Leahy
9J/ Verified vs True - Proof of Effort / David Condrey
9K/ NO SESSION
9L/ Philosophies of Privacy…. What do we even mean / Sol + Elizabeth
9M/ Credential Properties Naming / Dean H S (sp?) + Pam Dingle
9N/ BILLION - TOKEN PERSISTENT CONTEXT - What does it mean for data sovereignty and collective AI? / Brian von Herzen

Session 10
10A/ EUDI Wallet & eIDAS - Ask me anything / Mirko
10B/ SEDI Bulk Issued ACDC Context Wallet Management Self-Enforcing Data Loyalty / Sam Smith
10C/ Does any of the even matter if we can't deal with Ad-Tech? / When
10D/ Imagineering Exploring a Neighbor Sheds Pilot / Tracey R.B. & Friends
10E/ Representing Entitlements in SCIM… How to replace a flawed object model in a new era / Danny Zollner
10F/ Decentralized "Social Media" - Content Moderation, Community Notes, & Trust & Safely / Zach Alexander
10G/ "VCKNOTS" pluggable identity stack for issuer, wallet, verifier, developer / Naohiro Fujie
10H/ US Chamber of Commerce NDP Update / Phil Long
10I GovOps / Mike Schwartz & Rohit Khare
10J/ Schemas, Standards, IDs. - JSon Schema WG - Data self-identifying -LLMs & Schemas /Lisa D
10K/ Fair Witness Digital Signet / Joe Andrieu
10L/ Self-Sovereign Computing - Personal LLM For True Agency / Christopher Allen
10M/ AGENTS: Subject or Client? / Omri Gazitt
10N/ Thoughts on Identity Wallet in JAPAN / Soshi Hamaguchi

Thursday April 30, 2026 Day 3 / Sessions 11 - 15

Session 11
11A/ Playnet (second time) / Ruzgar E
11B/ OpenID4VCI Server-to-Server Issuance / Gareth
11C/ Cognitive Liberty + Captive Audience (1st Amendment Low) / Eric Welton
11D/ AI Agents and Open Banking / Kevin Feltes + Jean Paul LC
11E/ Identity and (geo) politics / Wendy Seltzer
11F/ My Terms 101.5 / Doc
11G/ VCWG New chatter, new work. Come find out what we're doing / Brent Z
11H/ KERI + did:webs 101 / Kent Bull w/GLEIF
11I/ What happens when Digital Identity is compromised? / Steve McCown / Anonyome Labs
11J/ NO SESSION
11K/ NO SESSION
11L/ Dilithium: Powering server user-Agents at Warp Speed - DEMO! And more / Swan and When
11M/ NO SESSION
11N/ NO SESSION

Session 12
12A/ Disclosure Policy - To whom can I present my credential? Using OID4VC / Mirko Mollik
12B/ PK4 Open Claw / Jesse Ariss
12C/ Non-thinking Non-identity Non-activity / Ben Go
12D/ How to preserve reality in a world of deepfakes? (Come w/ Ideas!) Lauren Paer
12E/ NO SESSION
12F/ Agent Names - human-readable names backed by DIDs / Markus Sabadello and Drummond Reed
12G/ Composable Committed Components (P256) DB -> ZK (B?S) / Adrian R.
12H/ Be a NOW it All! How we built a community calendar that answers all in one place, the The Burning Question "What's Happening?" / Jon Udell & Joyce Searls
12I/ NO SESSION
12J/ "Mastering Digital Identity" A.M.A and T.M.E. (tell me everything) / Eve Maler

12K/ NO SESSION
12L/ Post Quantum Cryptography 101 / Bryce Frey
12M/ NO SESSION
12N/ Monetizing Issuance of VC's !? / Neils

Session 13
13A/ Independent Identity Platforms: Using PBC to create alternatives to big tech identity IDP's / Jonathan McHugh
13B/ Local First Software / What it is & how can we get E26 encryption? Help Wanted! / David Gilde
13C/ I'm Gonna Call You Peaches! / Justin R
13D/ CAM & Brevity Anti-Saas Solutions / Chris Kula
13E/ NO SESSION
13F/ Fido LTAP2.3 What's New / John B
13G/ Deferred Token Response / Frederik + Max Gerber
13H/ HomeID - Asset Level Identity / Gregory Labrousse w/Justin
13I/ Frames All The Way Down - The missing foundation / Justin Chambers
13J/ Let's Get Relational / Michael Becker
13K/ NO SESSION - Moved to Space L ½ way through the session
13L/ Launch of the DTG ZKP Task Force / Drummond Reed
13M/ NO SESSION
13N/ Intro To GovOps / Mike Schwartz & Rohit Khare

Session 14
14A/ The Personal AI Superagent = A special double session (2in1)
Why verifiable trust agents (VTA's) and verifiable trust communities (VTC's) are a game changer for self-sovereign identity (SSI)
Fiduciary AI and the Relationship Rewards Model / Drummond Reed
14B/ Using Policy-Based Authorization to control Agent Behavior / Sarah Cecchetti & George & Phil Windley
14C/ NO SESSION
14D/ Considering Trust - What participants in the DigID Ecosystem Need / David Kelts
14E/10% Improving IDV error rates / Elaine W
14F/ Regenerative Accelerationism / Kaliya & Friends
14G/ Identity as a Martial Art - (Second time offering) / Jeff O
14H/ NO SESSION
14I/ NO SESSION
14J/ STUDY HALL w/Swan! Struggle buddies to get session notes in! / Swan B
14K/ NO SESSION
14L/ What We Learned Building a PoC for vLEI / Esteban Garcia
14M/ NO SESSION
14N/ NO SESSION

Session 15
15A/ OpenVCT (verifiable trust community) demo from Affinidi CEO Glenn Gore / Drummond Reed
15B/ ACDC Delegation Authorization EVAC Solution to Broken Delegation Chains. / Sam Smith
15C/ Open Claw Verifiable Identity Documents / Sarah C
15D/ NO SESSION
15E/ NO SESSION
15F/ 2FA for Email access - Why not pretty soon? / J. Fenton

15G/ Conveying Trust - Mechanisms to communicate trust & compliance to Policy / David Kelt
15H/ Registering an Agent with aria.bar ~ How can we collaborate? / Aaron & Adolfo Grego
15I/ Regenerative Accelerationism 2.0 Continued from Space F in Session 14
15J/ NO SESSION
15K/ NO SESSION
15L/ Facing the Dumpster Fire of Age Verification / Brent Shambaugh
15M/ NO SESSION
15N/ NO SESSION

Tuesday October 21, 2025 ~ Day 1

Session 1
1A/ SEDI Summit 0 calories, All caffeine / Utah - Sam Smith, Steven McCown, Phil Windley, Chris Bramwell
1B/ OAuth 101, Why OAuth? OAuth for MCP! New OAuth Specs? / Aaron Parecki
1C/ Please Teach Me Web Wallet / Hideaki F
1D/ Server User-Agents (vs Client User-Agents) A Fundamental Missing Piece of Infrastructure / When Leggett
1E/ NO SESSION
1F/ Passwordless Authentication w/ 100% Adoption / Andy Swett, Mike Jones, Ravi Ramaraju
1G/ NO SESSION
1H/ NO SESSION
1I/ Self Sufficient Software Connecting Canoe Paddlers / Tyler Childs
1J/ NO SESSION
1K/ Native App to Webb SSO / Matt MacAdam BECU
1L/ KYAPAY - A Protocol for Agent Identity and Commerce / Ankit Agarwal @ Skyfire
1M/ Introduction to the First Person Project / Drummond, Margeigh Novotny, Martina Kolpondinos and others
1N/ NO SESSION

Session 2
2A/ SEDI - GUARDIANSHIP - DEMO-META + ROBLOX INTEGRATION / Veridian
2B/ Intro to Open ID Connect / Mike Jones
2C/ My Private AI Agent as Authorization Server / Adrian Gropper
2D/ IIW Testimonial w/Pete the Videographier - He is doing short recordings - go make one! :-)
2E/ Neither SOCIAL nor ENTERPRISE SSO and Federation when I only kinda trust you / Justin Richer
2F/ IEEE P7012 (My Terms) and the End of Corporate ToS / Justin Byrd
2G/ Building trustworthy AI / MCP ecosystem - finding ways forward / Chris Phillips
2H/ HARMS We Care About / Erica Connell & Joe Andrieu
2I/ OpenID4VC 101 & ….. / Kristina Yasuda, Joseph, Paul
2J/ NO SESSION
2K/ NO SESSION
2L/ Drivers License Demo (live) / Francisco Corella Joint work with: Suhni Chuhad, Pema Selden. Veronica Wognas
2M/ Swiss e-ID “the least worst centralized government identity system?” & The Failure of Decentralized Identity (and what to do about it) / Christopher Allen
2N/ What Are Transaction Tokens? (secure identity context propagation} Ashay R.

Session 3
3A/ Email Verification Protocol / Sam G & Dick H

3B/ Authorization 101 / Steve Venema
3C/ Content Authenticity 101 / Eric S
3D/ Building a Multi-Agent Research tool for Blockchain Governance & Standards using First Person Project / Mitchell Travers
3E/ NO SESSION
3F/ SEDI Guardianship & Delegation (for dummies) / Timothy Ruff
3G/ Your WORDS Are Your Identity , VCONs tracking everything you say. / Ben Curtis
3H/ First WEB Based Open Source Passkey Enabled Identity Wallet / Stina and Leif Johansson
3I/ Presenting CoralKM - A protocol for decentralized Key Management + Recovery (help needed to give feedback and design) / David Gildeh
3J/ NO SESSION
3K/ Four MCP Use Cases (with distinct identity considerations) / Atul Tulshibagwale
3L/ Tools for Traps *Expanding the Language* / Jeff Orgle
3M/ NO SESSION
3N/ NO SESSION

Session 4
4A/ Threats + Mitigations Persistent State-Issued Citizen Entitlements (SEDI) Trade-offs (KERI-ACDC) Security / Sam Smith
4B/Passkeys 101 / John Bradley
4C/ K-ARF Open Framework Digital Identity Infrastructure for 'South Korea' / Jinyoung Jun - Hopae
4D/ The Six Inversions - The systematic transformation of legal protections into Platform Feudalism! / Christopher Allen
4E/ Digital Fiduciaries - What should they KNOW? / Joe Andreau
4F/ VLEI Ecosystem Upate - GLEIF / Karla McKenna
4G/ DIDComm Intoo, CBOR encoding, New Protocols / Sam Curren
4H/ Multi - Subject OAuth / Sarah C
4I/ Scaling the Agentic Web / Andor Kesselman
4J/ Verifying Desktop Applications for localhost redirects / Paul C
4K/ "Who's Responsible?" Authorization and Liability in AI Agents / Emu Iizuka
4L/ Beyond the Blue Check - from Social Verification to Verifiable Relationships / Brendan M + Alberto L
4M/ Proximity Presentation for SD-JWT using 18013-5 / Lee, Kristina, and more
4N/ IIW Testimonial for Pete! Any time in room N

Session 5
5A/ How to use Keri/ACDC for SEDI Infrastructure Distributed Decentralized Signing Infrastructure / Sam Smith
5B/Decentralized Identity 101: elements and applications / Steve McCown
5C/ Anastasia: Cinderella's Stepsisiter - Turning X.509 certs into Anonymous Key Attestation / Dan Yamamoto
5D/ Gen AI Phishing! Using AI for Bad Things! / Yuriy A
5E/ NO SESSION
5F/ Passkeys PQC and Beyond / John Bradley
5G/ Gordian Autonomy Stack / Christopher Allen Blockchain Commons
5H/ Use cases in North American Real Estate / Michael Krotscheck
5I/ Identity 4 Groups "real decentralized?" Sideways. Earth verifiable community (nor social graph) / Kaliya Y
5J/ Natural vs. Verified Person? / Luke Nispel
5K/ Technical Approaches to Delegation and Guardianship / Richard Esplin

5L/ Women in Identity - What does our community need to help *Women/underrepresented communities depending on ID systems? * YOU / Elizabeth Garber Exec Dir. WiD
5M/ If Bitcoin had Identity Layer - first person network for retroactive UBI & solving $338T Glogal Debt Crisis / Nivas Sivaprakasam
5N/NO SESSION

Wednesday October 22, 2025 ~ Day 2

Session 6
6A/ Is Compromising a SEDI Treasonous? / Rep. Chevrier
6B/ MY TERMS for Dummies / Doc & Joyce and Customer Commons Team
6C/ When Patients and Doctors write their own Software… / Adrian Gropper.
6D / NO SESSION
6E/ NO SESSION
6F/ Data Unions & Labor Laws w/James Felson Keith running for congress on Data & Labor Platform / Kaliya Y
6G/ Presenting VSs based on Natural Language Instructions - Offline LLM for DCQL on Mobile / Ken Watanabe
6H / OpenID 4 VC Conformance Testing Deep Dive / Joseph Heenan
6I/ On Behalf of use Authorization for AI agents (draft spec) / Hasintha Indrajee & Sachin Mamoru from WS02
6J/ NO SESSION
6K/ Driving Adoption vLEI + KERI - GLEIF vLEI Learning Modules - GLEIF Global Hackathon / Esteban Garcia
6L/ KYAPAY - A protocol for Agent & Principal identity & Payments / Ankit Agarwal @ Skyfire
6M/ OPAQUE Passwords RFC 3807 / Michael Krotscheck
6N/ NO SESSION

Session 7
7A/ Native Client Attestation Grant & Key Binding / Frederik Krogsdal Jacobsen & Attestation-based Client Auth
7B/ Verifiable Relationship Credentials (VRS's) and R_Cards: A Design Session / Brendan Miller +++
7C/ Interop KERI + W3C Libraries (reuse not rewrite) / Kent Bull
7D/ NO SESSION
7E/FIGHT CLUB - The Kids are online! Plain Language discussion on children on the internet / Swan Black
7F/ State Endorsed Decentralized Identity (SEDI) for Dummies / Timothy Ruff
7G/ SD-JWT How it works? Entry Level / Lukas Han
7H/ State of the State: Tracking and coordinating Public Policy / Ethan Veneklassen
7I/ Biometric Bound Credentials / Richard Esplin
7J/NO SESSION
7K/ NO SESSION
7L/ NO SESSION
7M/ Password Auth w/PQL in the Quantum Era / Bryce Frey
7N/ NO SESSION

Session 8
8A/ OpenID 4VC - Issuance over W3C DC API / Kristina, Paul, Lee
7B/ HARMS -> STORY 'Digital Identity Mad-Libs' / Erica Connell
7C/ NO SESSION

7D/ New Data Paradigm - US Chamber of Commerce use case: U.I. (unemployment insurance) @ VC issuance to Employees / Phil Long
7E/ Non-Binary Approaches Human Verification - from 0/1 to 0.0/1.0 /Discussion Session / When Leggett
7F/ OAUTH for MCP (client ID metadata document) / Aaron Parecki
7G/ vLEI Authentication Organizational Login (vLEI) with ecr credentials / Christoph S, Vincent V, Stefan I
7H/ NO SESSION
7I/ Trust Registry Query Protocol / Andor, Darrell, Drummond, Phil
7J/ NO SESSION
7K/ NO SESSION
7L/ Death and the Digital Estate / OPEN ID Foundation Community Group - George Fletcher
7M/ Women in Identity? - What should an organization focused on inclusion be doing with and for the community of ID Professionals? / Elizabeth Garber
7N/ NO SESSION

Session 9
9A/ OpenID4VC "server-to-server' mode / Kristina, Joseph, Paul, Lee Christian,
ACDC Blindable Registries - How to minimize correlation + detect compromises - Proof of Age non-zkp / Sam Smith
Alternative traffic stop authentication safer for the officer / Francisco Corella
NO SESSION
NO SESSION
My Terms Practical Experience - Customer Commons / Iain Henderson
If bitcoin had an identity layer 2.0 / Nivas Sivaprakasam
What are DIF Recommended DID Methods? / Jonathan Rayback
Build It Yourself Agent Espert *live* Building an agent to talk to a spec / Chris Phillips
NO SESSION
Cross-border interoperability of VC's with demo of Japan's My Number car on iPhone / Hideaki Furukawa
Agentic Identity Gateway / Teng and Authorization and Liability in AI Agents / Emu Iizuka
Hologram Messaging Verifiable Chatbots P2P + Ai / Ariel Fabrice
The End of the Global Internet / Heather Flanagan

Session 10
10A/ OpenID4VC "server-to-server' mode / Kristina, Joseph, Paul, Lee Christian,
10B/ ACDC Blindable Registries - How to minimize correlation + detect compromises - Proof of Age non-zkp / Sam Smith
10C/ Alternative traffic stop authentication safer for the officer / Francisco Corella
10D/ NO SESSION
10E/ NO SESSION
10F/ My Terms Practical Experience - Customer Commons / Iain Henderson
10G/ If bitcoin had an identity layer 2.0 / Nivas Sivaprakasam
10H/ What are DIF Recommended DID Methods? / Jonathan Rayback
10I/ Build It Yourself Agent Espert *live* Building an agent to talk to a spec / Chris Phillips
10J/ NO SESSION
10K/ Cross-border interoperability of VC's with demo of Japan's My Number car on iPhone / Hideaki Furukawa
10L/ Agentic Identity Gateway / Teng and Authorization and Liability in AI Agents / Emu Iizuka
10M/ Hologram Messaging Verifiable Chatbots P2P + Ai / Ariel Fabrice
10N/ The End of the Global Internet / Heather Flanagan

Thursday October 23, 2025 - Day 3

Session 11

11A/ Malnets - Understanding Malware Networks / Jacob Siebach
11B/ OIDF AI Identity Management Community Group Call
11C/ Storytelling - The Art of the "Lie" Utah's Biggest Liar / George McEwan
11D/ NO SESSION
11E/ Help Wanted - Seeking Practical Advice for Implementing SEDI / Alan Fuller
11F/ Future of Work meets Identity and Data Portability / Brad Topliff
11G/ Perspectives from the United Nations from Refugee Registration to Self-Sovereign Identity / Besem Obenson
11H/ ZKP - Overview - Getting to Deployments / Leif Johansson, Chris
11I/ Human Alignment in Agent to Agent Authorization / Adrian G
11J/ NO SESSION
11K/ NO SESSION
11L/ KERI Auth Browser Extension demo/discussion / Ed Eylcholt
11M/ Notarized Verifiable Relationship Credentials (VRCs): The "Virtual Selfie" - Harvard Applied Social Media Lab / Brendan Miller, Alberto Leon, Drummond and others
11N/ Phone Home - An Update / Timothy Ruff, Joe Andrieu, Steve McCown

Session 12

12A/ SEDI Guardian Demo, Roblox + Social Book / Veridian
12B/ WTF is a "client_id" / Justin Richer
12C/ NO SESSION
12D/ NO SESSION
12E/ NO SESSION
12F/ FeDIDeration DIDs in OPENID Federation / Lukas and Fraser
12G/ You've Got The Wrong Use Case / Alan Karp
12H/ per-Credential Metadata (by the issuer…) / Paul, Gareth, Leev, Kristina, etc
12I/ NO SESSION
12J/ NO SESSION
12K/ NO SESSION
12L/ Domains of Identity - 16 Subdivisions of Use Cases… walking through them + considering developing intersections between domains. Book + Practice Talk 4 Taiwan / Kaliya Y
12M/ What Identity can Learn from Home Assistant and Home Automation - interop - privacy / Sam Curren
12N/ Privacy in a Surveillance World / Steve McCown

Session 13

13A/ Working Session - How to add VC, VP and maybe mDL to CAWG identity / Eric Scouten, Andrew D
13B/ KERI's Post Quantum Story Surprise Quantum Attack Capture New Decryp Cater Crypto Agility / Sam Smith
13C/ How to Drive Adoption of Fairest Universal Basic INcome Currency with Inevitable AI Job Disruption (using identity tech) / Nivas
13D/ NO SESSION
13E/ NO SESSION
13F/ Getting Started on MY TERMS Deployment - Who wants to help? Starting work groups TODAY! / Doc, Joyce, Kari
13G/ Credential Usage Policy / Paul, Tobias, Kristina
13H/ Tools for Traps *Expanding the Language* / Jeff Orgle
13I/ RP Architectures & ID Token Audiences / Frederik

13J/ Mechanics of Running and Open ID Federation (I'm struggling!) / Nicole R.
13K/ The TXNS are coming from inside the house! How do we think about minting tokens for in-clutter transactions / Andrew Todd, Mongo DB
13L/ CorpLKM & Key Recovery Part 2 / David G
13M/ NO SESSION
13N/ NO SESSION

Session 14
14A/ What's your p(doom) and why? / Omri G
14B/ NO SESSION
14C/ NO SESSION
14D/ NO SESSION
14E/ NO SESSION
14F/ Server-to-Server issuance / Hicham, Martijn, Gareth
14G/ KERI Suite (ACDC CESR SEDI) Ask me anything / Sam Smith
14H/W3C/ Why? A research projects findings / Emily L?
14I/ Defining Policies for AI Agents - Is Auth Policies enough?! / Andor
14J/ NO SESSION
14K/ NO SESSION
14L/ Crowd-Designed Software Session: BYOE architecture, Social Network for IIW / Dmitri, Ty, Bengo
14M/ Building a National-Scale IOP with OIDC: Short comings OP OIDC / Fredrico
14N/ NO SESSION

Session 15
15A/ German EUDI Wallet / Kristina Paul, Christian
15B/ Resource wonder PASSKEY credential Flow for Agentic AI'S / Hideaki F.
15C/ Can YadaCoin Help Keri? / Matt V
15D/ Privacy Dissolution or Redamation - How do we redain data practically speaking and attach data to a strong ID? What are the protocols? / Beth P
15E/ NO SESSION
15F/ Cross-Pollinate - Bring Your Ideas to "Server User-agents" / When Leggett
15G/ Sorta Kinda Digital ID - adding digitally signed printed elements to physical IDs and documents / Elaine Wooton
15H/ NO SESSION
15I/ Local First Software + Bridging Communities / David Gilden
15J/ NO SESSION
15K/ NO SESSION
15L/ Authorization and Liability in AI Agents / Emu Iizuka
15M/ NO SESSION
15N/ CfP Brainstorming for EIC, Indentiverse / Heather F

Fraser Edwards • 1st
CEO at cheqd, building the identity layer for the agent and trusted data economy...
1mo •

Each of the days at **Internet Identity Workshop** (**#IIW42**), I make a point of reading through all of the sessions across the day to try and identify any patterns or themes.

Whilst there will be Agent Identity Workshop tomorrow (Friday), the dominant theme across IIW42 has to be (unsurprisingly): **#AI**.

A lot of sessions focusing on **#AIAgents** but also a lot of sessions focusing on establishing trust and security as generative AI supercharges fraud and overwhelms existing approaches to ID&V. These sessions can broadly be broadly broken into:

1. Combining **#DID** and **#AI** capabilities to establish brand new capabilities
2. Using **#DID** and **#VCs** to combat new risks resulting from **#GenAI**, e.g. **Riley Hughes** ID/acc or **Coalition for Content Provenance and Authenticity (C2PA)** from the likes of **Eric Scouten**

Outside of the AI theme, there are plenty of the usual inter-op.

Would welcome others to take a spin through the pictures and see if any other themes jump out at them! See you all in the comments

Geun-Hyung (Peter) Kim and 57 others 11 comments • 3 reposts

Notes Day 1 / Tuesday April 28 / Sessions 1 - 5

SESSION #1

Introduction to the Decentralized Trust Graph and the First Person Project

Session Convener: Drummond Reed
Session Notes Taker(s): Erika Bjune and Drummond Reed

Tags / links to resources / technology discussed, related to this session:

This was the presentation that Drummond shared for the discussion:
https://docs.google.com/presentation/d/1PcuwPIFgYhVd39C60nAShyJS-h1skAZ10D7_43NrJLA/edit?usp=sharing

Other related materials:
First Person White Paper
ToIP / DIF Decentralized Trust Graph Working Group
Decentralized Trust Graph Glossary
DTG Credential Schemas

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Drummond explained that the First Person Project is NOT an actual Linux Foundation project, but grew out of four Linux Foundation projects— LF Decentralized Trust, Trust Over IP (ToIP), Decentralized Identity Foundation (DIF), and the OpenWallet Foundation (OWF)—plus the Ayra Association, which is governance for a decentralized trust registry network based on the ToIP Trust Registry Query Protocol (TRQP).

He then gave a presentation explaining the origin and structure of a decentralized trust graph as defined by decentralized identifiers (DIDs) and Verifiable Credentials (VCs) for each node. This is now being standardized at the ToIP / DIF Decentralized Trust Graph Working Group.

Some of the questions during the discussion:

- Could there be other node types besides person, device, AI agent, and verifiable trust community (VTC)?
    - A: All other types of nodes can be represented by AI agents.
- What kind of organizations can be VTCs?
    - A: Any social or legal group or organization of any kind, from a 3 person chat group to a church group to a corporation to a nation state. A VTC is simply a collection of the other node types.
- Can VTCs join other VTCs?

    - Yes. VTCs can have peer-to-peer relationships with other VTCs using verifiable relationship credentials (VRCs), or they can have membership relationships using verifiable membership credentials (VMCs). The latter enables hierarchies.
- How is reputation expressed in the DTG?
    - A: It is layered over the graph. There is a special credential defined for reputation—the verifiable endorsement credential (VEC). See the DTG Credential Schemas document.
- How are verifiable trust communities (VTCs) governed?
    - A: That is up to the community itself. That level of governance is out of scope for the Decentralized Trust Graph Working Group.

- How to govern a digital public utility?
    - A: Make it a cooperative owned by the members. For example, the First Person Cooperative is incorporated as a Colorado Public Benefit Limited Corporation.

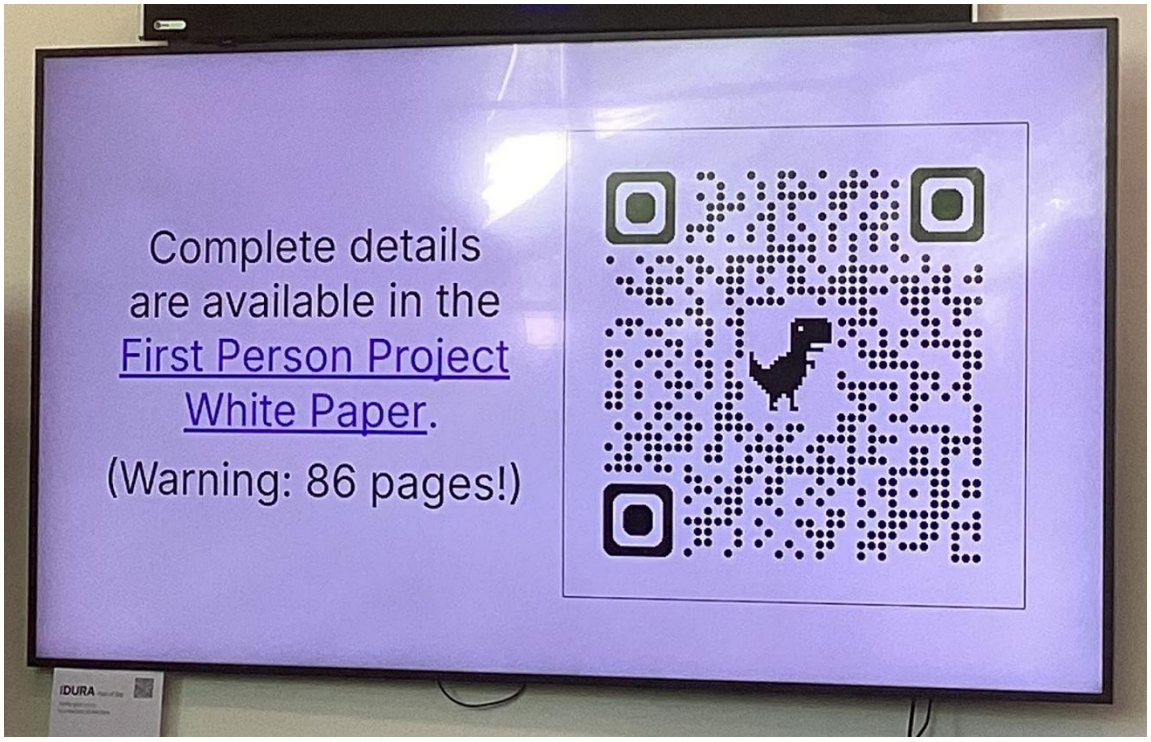

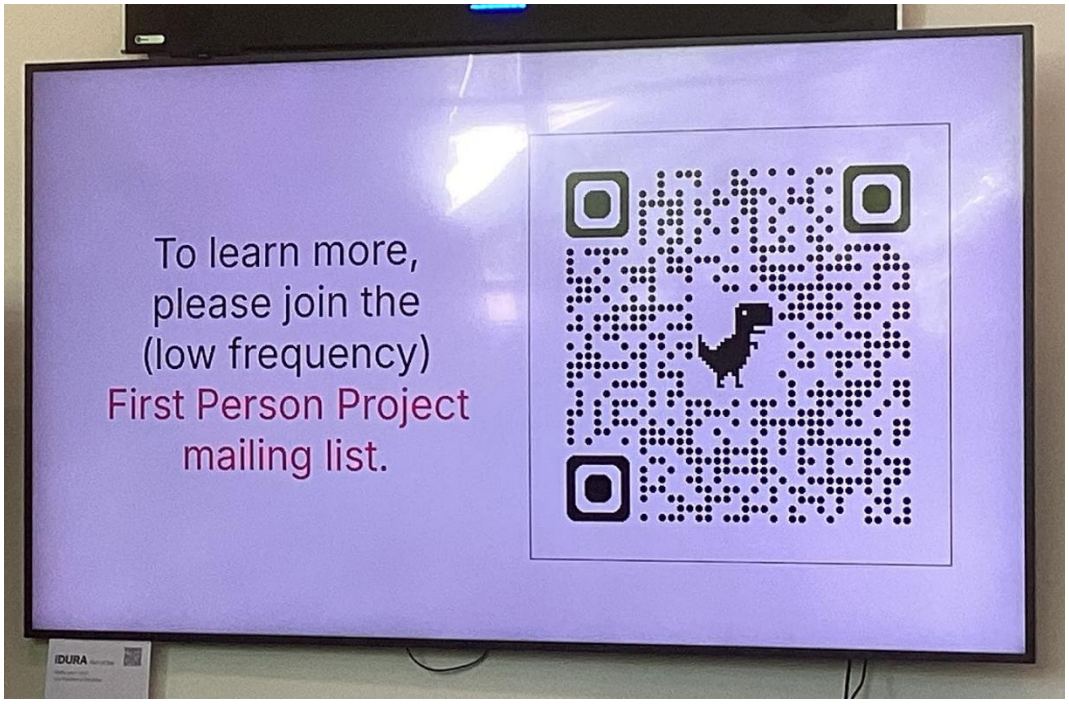

One year ago, Linux Foundation issued a challenge to address a malicious malware attack via decentralized trust system development.

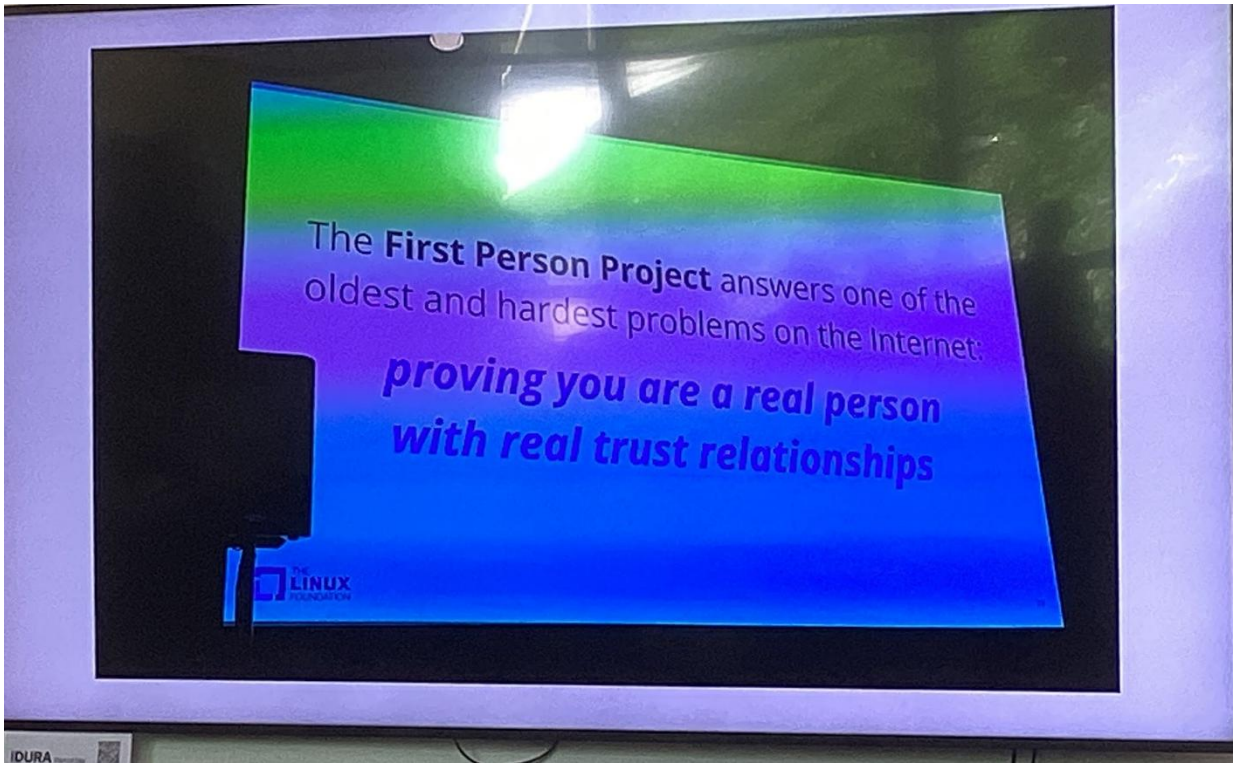

verifiable relationship creds use web of trust

one year later, an AI attack took two weeks to show how easy it was for AI to infiltrate open source, versus the original human threat, which took two years of social engineering.

Anyone can participate and contribute to GitHub repos

What is a decentralized trust graph versus the web of trust?
Three components:

1. Keys: foundation of asymmetric cryptography, cannot be faked by AI agents, and can be quantum proof
2. DIDs (Decentralized Identifers): any form of cryptographically verifiable identifiers; enable key and endpoint rotation; no central registry required
3. Verifiable Credentials (VCs): cryptographically verifiable credentials; semantic claims; can support selective disclosure and ZKPs

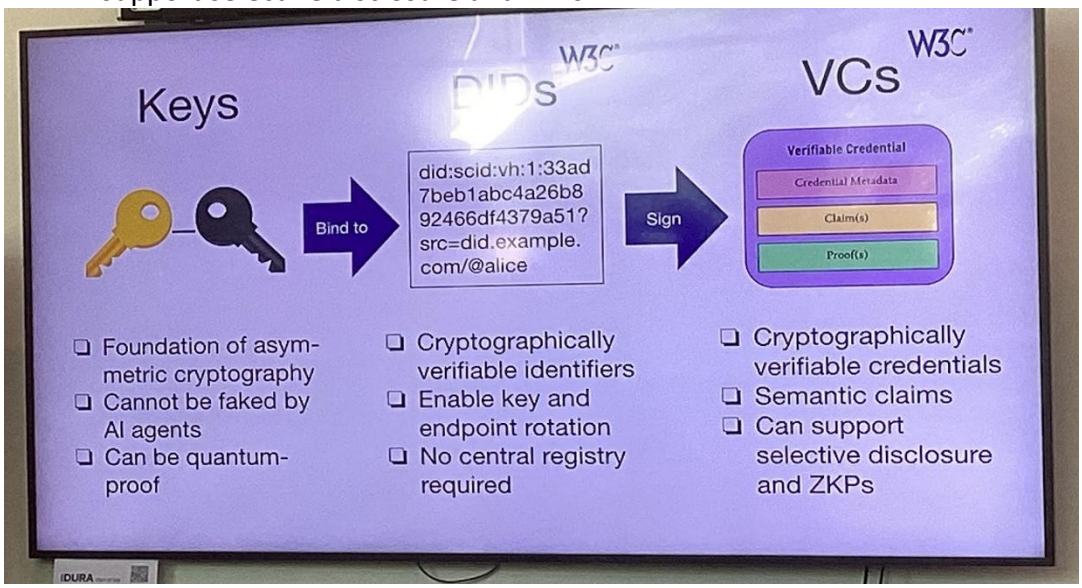

Use these basic building blocks to build a graph of trust relationships.

The Web of Trust has been used to manage Linux contributors but is getting old, tired and doesn't scale. In the DTG, Verifiable Relationship Credentials (VRCs) are exchanged between people

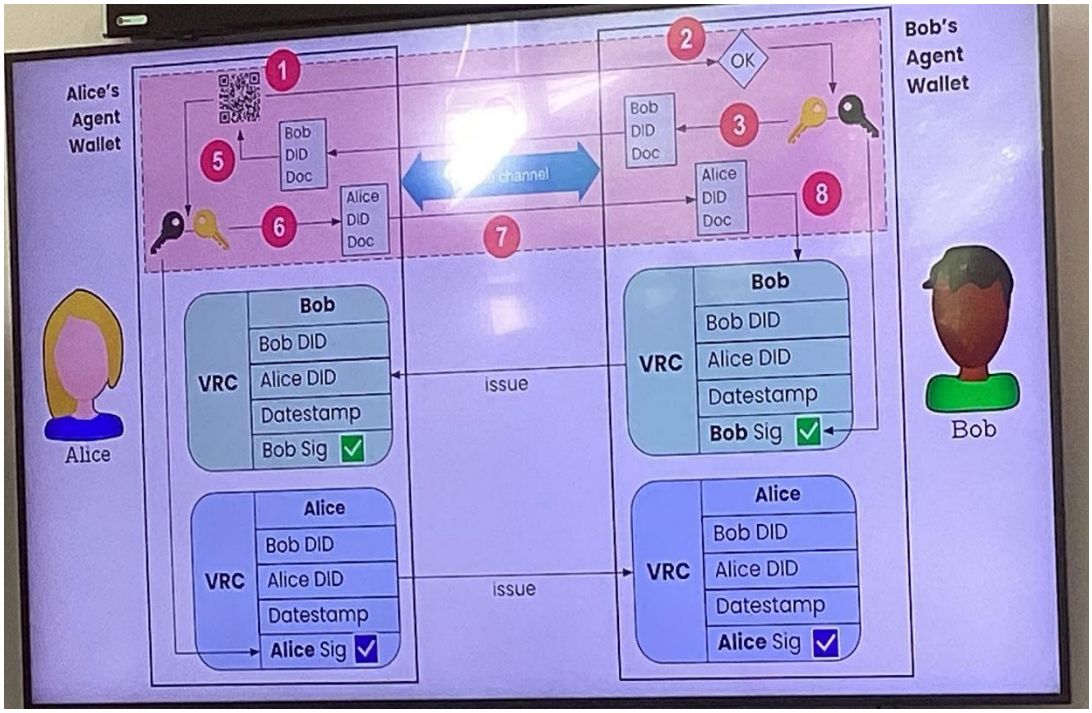

there is no “trusted” intermediary in this. It opens a private channel between two people, issues and verifies private keys that are verified to create a VRC. This gets applied to all relationships in a trust graph.

OpenVTC is the code base on GitHub where this is being implemented.

Current trust graphs depend on only one node type: peer. To make this more effective and encompassing is the implementation of four node types: People, Devices, AI Agents, and Verifiable Trust Communities (any community that needs to have a trust foundation). This is accomplished by providing proof of personhood for communities

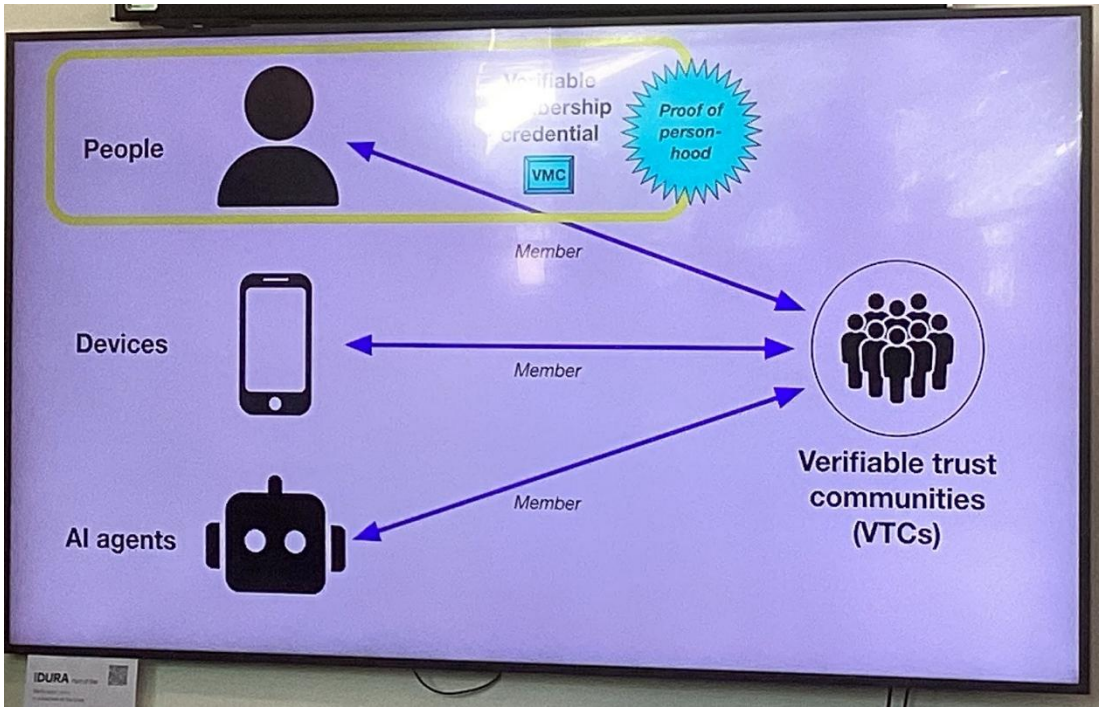

DTG is a baseline for trust. What policies and governance can be layered on it is broadly open and flexible. OpenVTC treats everything like a community and will support any infrastructure that needs to be built around it.

This can also solve the problem of who AI Agents work for, using the same VMC connection to show a proof of agenthood.

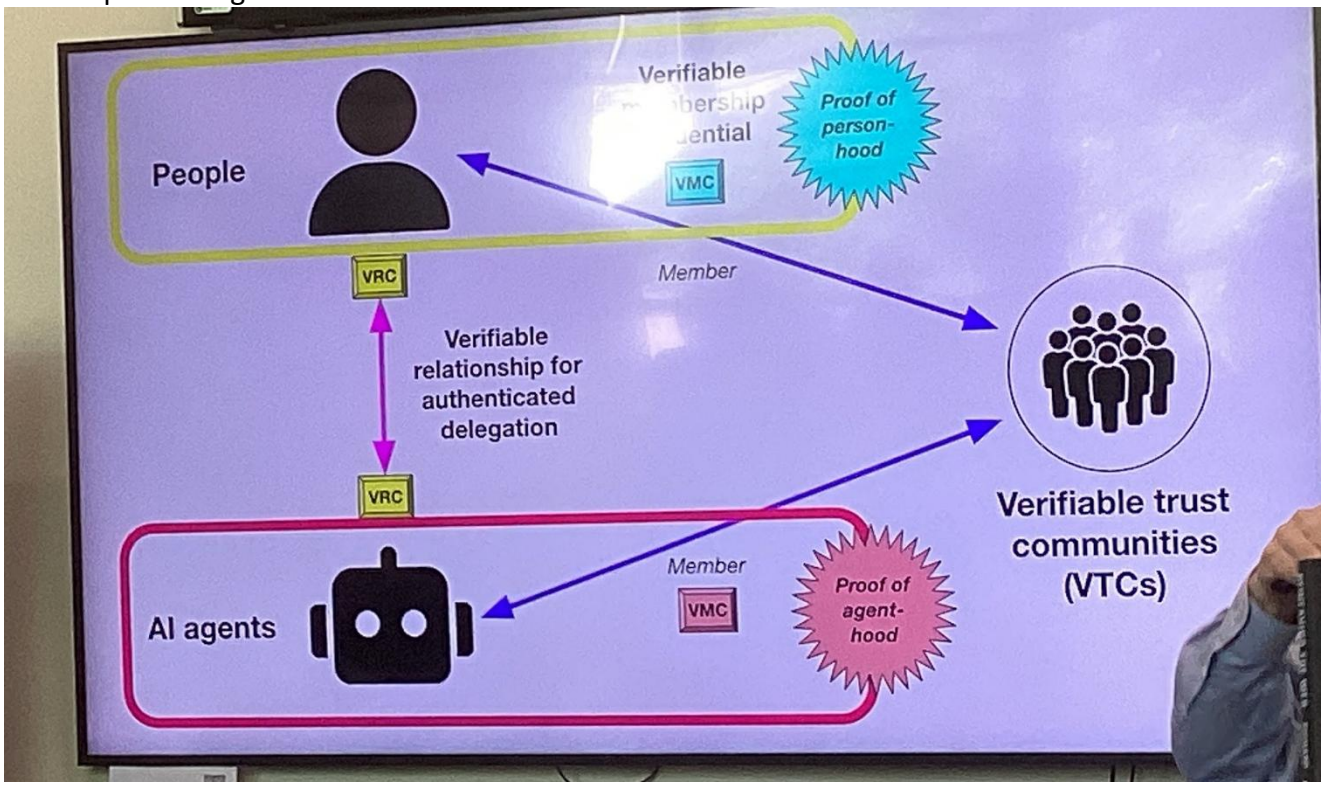

Books:

- Richard Whitt, *Reweaving the Web*
- Authors?, *If Anyone Builds It, Everybody Dies*
- Phil, *Dynamic Authorization* @ manning

There are four node types, but control is maintained by people using Personal and Community Network Managers

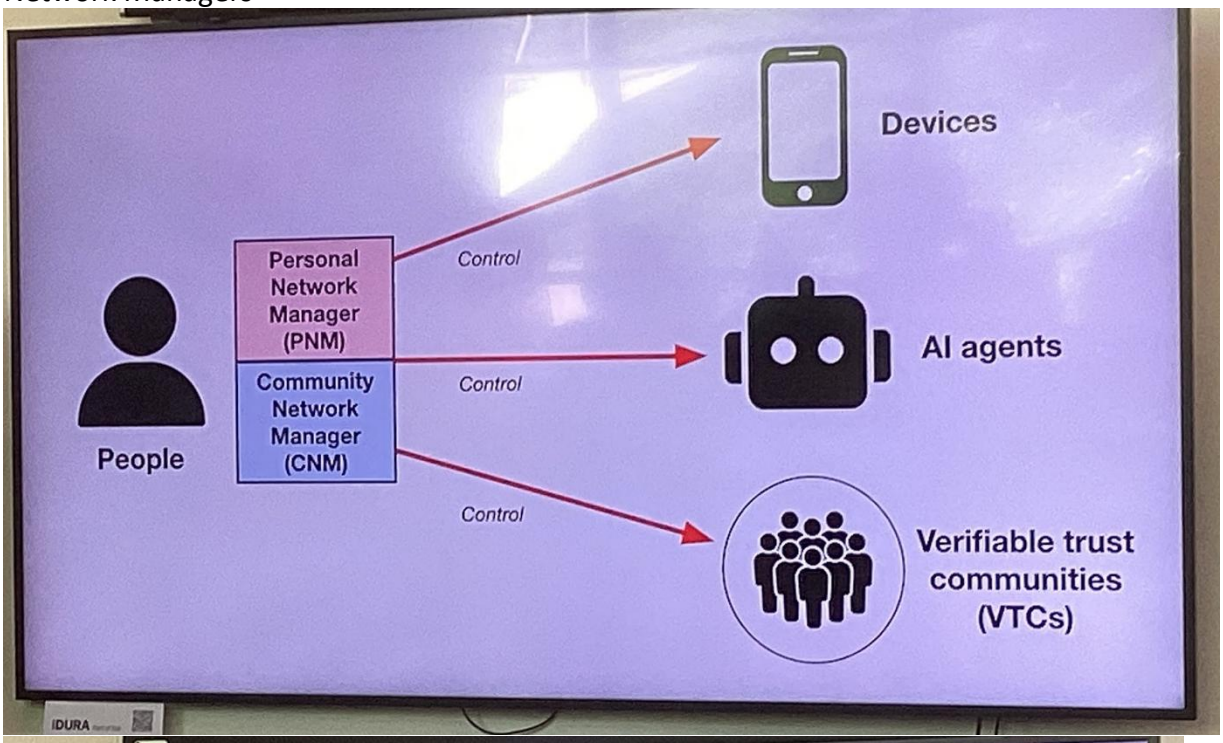

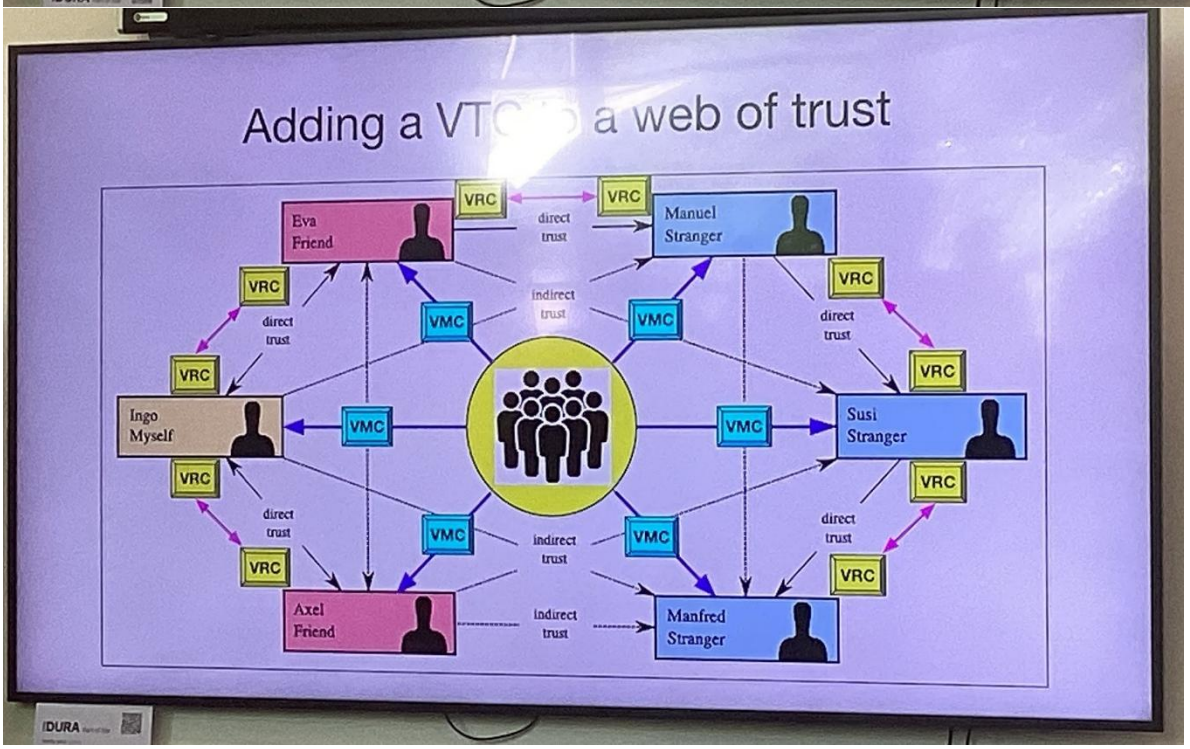

This enables the Linux community, for instance, to replace the two-person relationship rule with the Verified Trust Community.

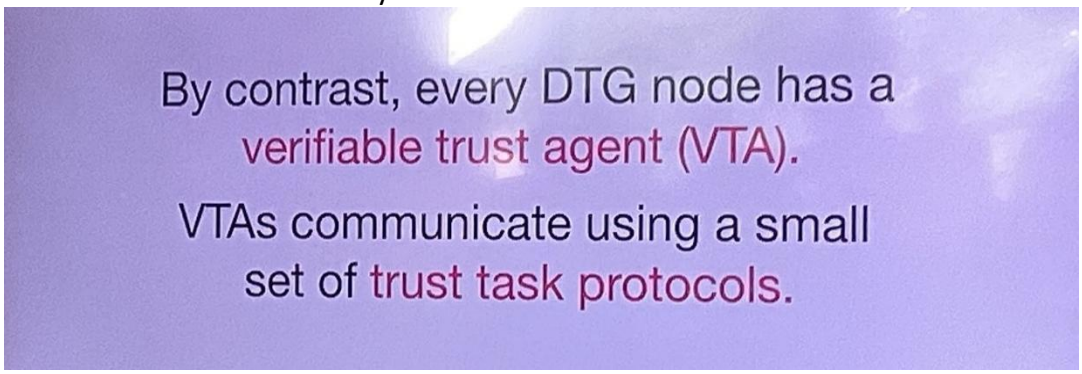

trust task protocols are the building blocks of utilizing the trust graph. There will be protocols for issuing credentials at the foundation level, but as people and orgs use it, they will develop protocols for payment processing, policy updates, membership credential management, etc… not to automate but to make it protocol/policy driven.

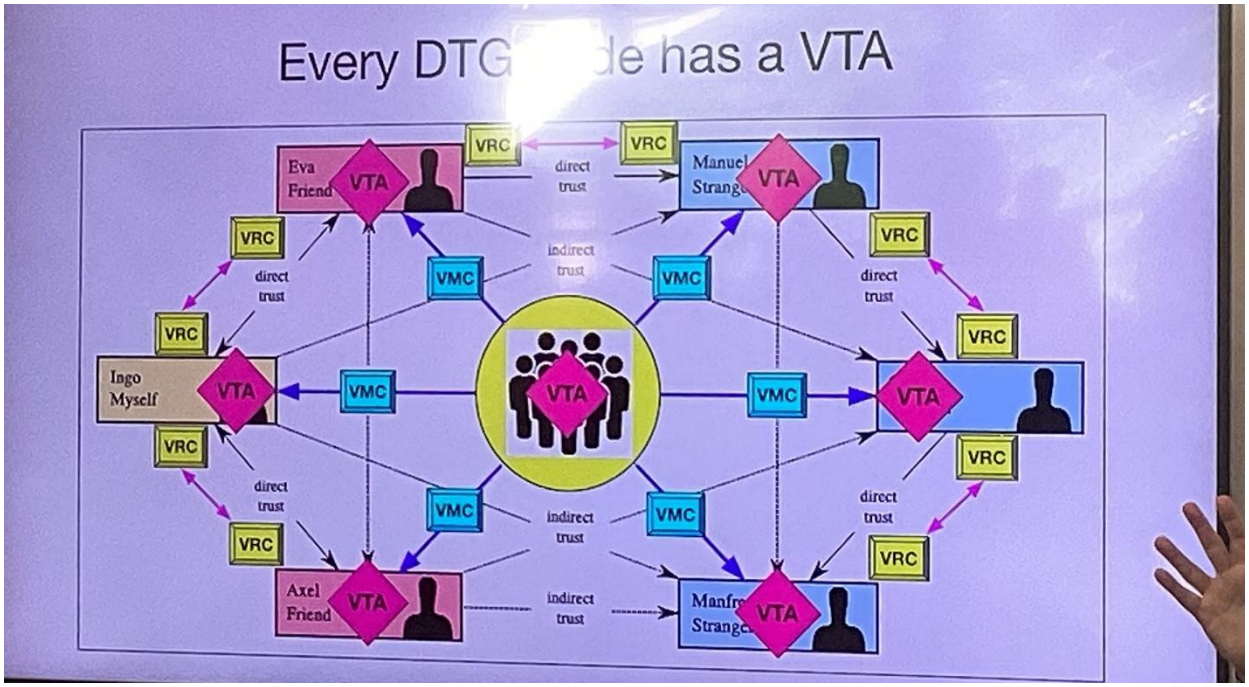

open protocols for communities for peer to peer trust, like Linux communities need, as well as other communities that can use incorporated credentials to share things; frees communities from the platforms they are on.

OAuth 101 an IIW Session / Aaron Parecki

Session Convener: Aaron Parecki
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

SIROS / WW ….. Longfellow ZKP Age Demo

Session Convener: John & Leif Johansson
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Are Passkey Server Implementations Secure?

Session Convener: Ken Watanabe
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Introducing Passkey testing framework
Found dozens vulnerability with it

(Multi Device) Decentralized Media & Messaging

Session Convener: Zach Alexancer
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

- Mastodon E2EE effort (https://blog.joinmastodon.org/2026/04/sovereign-tech-agency-funding/)
- "Help, my Signal has a bad device!" paper (Wichelmann et al. 2025) (https://eprint.iacr.org/2021/626)
- Signal device injection attacks (2025) (https://cloud.google.com/blog/topics/threat-intelligence/russia-targeting-signal-messenger/)
- WhatsApp multi-device implementation (https://engineering.fb.com/2021/07/14/security/whatsapp-multi-device/)
- "Formal Analysis of Multi-Device Group Messaging in WhatsApp" (Albrecht et al. 2025) (https://eprint.iacr.org/2025/794)
- "Practically-exploitable Cryptographic Vulnerabilities in Matrix" (Albrecht et al. 2022) (https://nebuchadnezzar-megolm.github.io)
- Keybase sigchains (https://keybase.io/blog/chat-apps-softer-than-tofu)
- FOKS (Federated Open Key Service) (https://foks.pub)
- Bluesky/AT Protocol's did:plc spec (https://web.plc.directory/spec/v0.1/did-plc)
- KERI spec (https://trustoverip.github.io/kswg-keri-specification/)
- KERI Foundation (https://keri.foundation)
- Weighted Cryptographic Device Management (https://github.com/mandihtha/protocol/tree/master/wcdm)
- Social Web Incubator Community Group (Social CG) (https://www.w3.org/community/socialcg/author/dmitriz/)
- Social Web Working Group (https://www.w3.org/groups/wg/social/)
- ActivityPub Client to Server spec (https://www.w3.org/TR/activitypub/#client-to-server-interactions)

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

- There is a close relationship between social/interactive media & secure messaging
    - Social media apps end up with DM features, which should be encrypted (see: Instagram)
    - Secure messaging apps end up with "social media"-esque features (groups, usernames)
- Secure messaging is technically harder; arguably should be the foundation, after which making content public (for social-media-style features) is easier
    - See Mastodon's lack of E2EE on DMs (recent funding for trying to fix)
- In secure messaging, identity and device management is the hardest part
    - Signal's encryption algorithms are widely adopted, audited, widely respected

    - Flaws in Signal Messenger, Matrix, etc. generally relate to identity/device/session management
- Keybase is a good model
    - Strong approach to identity and device management; audited (acquired by Zoom)
    - FOKS as open-source successor project by technical lead (Max Krohn)
    - Uses verifiable key history construct called "sigchains"
- KERI as a plausible approach to borrowing Keybase patterns
    - KERI key event logs are verifiable key histories, conceptually similar to Keybase sigchains
    - Bluesky's did:plc as also similar, could be adapted in a similar way
    - KERI support of fractional key weights unlocks additional features, e.g. cryptographic 2FA, recovery patterns
- Get involved in W3C groups that discuss topics like this
    - Social Web Incubator Community Group (Social CG) (https://www.w3.org/community/socialcg/author/dmitriz/)
    - Social Web Working Group (https://www.w3.org/groups/wg/social/)
        - Rebooted since 2018
        - If you are not a W3C member, approval may be required for participation?
- ActivityPub was never really implemented fully
    - The original spec had both a server-to-server spec and a client-to-server spec
    - Most applications only implement the server-to-server spec (and often not compliant)
    - The client-to-server spec in some respects is more decentralized, e.g. users signing activities
- Example of KERI-based 2FA user journey
- Potential integrations with existing media protocols (ActivityPub, ATProto/Bluesky, Nostr)

OpenID4VC 101 & Updates

Session Convener: Joseph Heenan & Frederik Krogsdal Jacobsen
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Slides used are here:
https://docs.google.com/presentation/d/1I6tQMJ2IPNg3hoFq-dyU0lQMG-K4HbX-/edit?usp=sharing&ouid=107381980093922120275&rtpof=true&sd=true

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

AI for vLEI: What does it mean for an AI agent to have a vLEI?

Session Convener: Esteban Garcia
Session Notes Taker(s): Esteban Garcia

Tags / links to resources / technology discussed, related to this session:

vLEI, LEI, GLEIF, AI, Agents, MCP, human-in-the-loop (HITL)
vetoflow.pdf
demo video: https://www.youtube.com/watch?v=Msnf9w3wt88
https://trustalys.com/

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The session explored what it means for an AI agent to hold a vLEI. Trustalys built a proof of concept to investigate this question and surfaced more questions than answers.

Key context: organizational identity differs from personal identity in operational pattern. It runs at high frequency (issuances, revocations, rotations, verifications at business tempo), high complexity (multi-party, multi-sig, chain-of-authority across orgs), and is embedded inside processes like onboarding, procurement, compliance, and payments. Two motivating scenarios were employee credential lifecycle (ECR provisioning during HR onboarding) and invoice verification with payment authorization for B2B.

Core architectural principle: expertise without authority. The agent has the expertise to construct correct vLEI operations and suggests them. The human holds authority and decides. This separation generalizes beyond vLEI to all controlled AI operations.

The Vetoflow demo showed this in action: prompt to agent, agent builds a playbook, human reviews workflows, human approves or vetoes.

The approval problem was discussed in depth. Risks include fatigue (rubber stamping), granularity tradeoffs (too fine loses big picture, too coarse becomes a blank check), expertise gaps (can the human actually evaluate the request), adversarial surface (the agent controls how requests are framed), and accountability gaps between nominal authority and effective delegation. The core risk is the gap between "approval happened" and "approval was meaningful."

Not every action needs approval. Three routing categories were proposed: no approval (read-only, queries, background validations), automatic (pre-authorized routines under policy), and human approval (keystore-touching operations, key ceremonies, issuances). The routing decision is configurable per organization, context, and risk level.

Two patterns for non-assistant agents: Pattern A is orchestration with per-action human approval. Pattern B is delegated authority where the agent acts within a scoped boundary defined by the human. Both preserve human accountability, differing only in where the approval boundary sits.

On role hierarchy: the EGF currently requires natural-person anchors. Three options were laid out for fitting agents in: sub-delegate below ECR (agent acts under an ECR person's scoped authority), peer credential at ECR/OOR level requiring a new framework section, or a modified ECR extending the existing role for non-human holders. Tradeoffs involve EGF change complexity, accountability depth, and deployment timeline.

The proposed concept is the Verifiable LE Agent Delegate (vLEAD). Technical feasibility was demonstrated via the PoC, where the agent could play every role in a vLEI trust chain. Open questions remain on the legal framework, governance models, and scope of agent liability. The biggest question is governance, not technology.

The agent also functions as a domain translator, lowering the barrier from operational complexity (keystores, witness networks, OOBI resolution, inception and delegation parameters, ACDC schemas, protocol sequencing) to natural intent like "I need to issue an ECR credential to our new agent in London," while keeping protocol primitives inspectable.

Additional observations: conversational interfaces challenge identity wallet UX assumptions, the workflow model produces an audit trail as a side effect rather than an add-on, agents can build complex multi-party scenarios in minutes (useful for testing and training), and in YOLO mode agents surfaced bugs in KERI and undocumented sequencing knowledge.

Architecture summary: AI Agent talks to an MCP Server, which writes to an Action Queue (SQL DB). The trust boundary sits at human approval. Below the boundary, a TUI and Executor pull from the queue, the Executor uses the OS Keyring and runs kli subprocesses to perform the actual operations.

SESSION #2

Prove IRL Connections - Bootstrap your Trust Graph: Harvard University Open Source Keyring App - Demo & Discussion

Session Convener: Brendan Miller & Alberto Leon, Harvard Applied Social Media Lab
Session Notes Taker(s): Brendan Miller & Alberto Leon

Tags / links to resources / technology discussed, related to this session:

See the slides for more information. IIW Keyring demo 4-28-26

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

IIW folks are invited to try the app, provide feedback and suggestions, and learn together about the trust graphs produced from events like IIW! Collaboration on the open-source repos is also welcome.

Introduction to OpenID Connect

Session Convener: Mike Jones
Session Notes Taker(s): Mike Jones

Tags / links to resources / technology discussed, related to this session:

The presentation is linked to from https://self-issued.info/?p=2845.

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Once again, there was an engaged and informed set of participants who brought their own perspectives and questions to the session, making it more useful for everyone.

Trust Infrastructure as a Public Utility

Session Convener: Erika Bjune
Session Notes Taker(s): Margeigh Novotny

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Prompt: If the goal were to make a public utility, not an enterprise effort, not a government asset… what would the infrastructure need to look like?

- There are already other players in this space: server networks providers, for example

"VTCs beg for decentralized server networks"

- So how could we own this as a public network?
- FP partner is already developing a distributed server infrastructure

What are some case studies or projects that others in the room are thinking about?

- Bennet:
    - Journalists, sources, publishers this chain requires different implementations
    - Civilians (small biz) should be paid to permit offers to civilians (advertising)
    - "Nested circles of trust"
    - We think of journalism/news as a public benefit, should be a trusted source
        - If we can create this thing it will become
- Lisa: Data Trust Registry person [launched today]
    - Internal trust verification processes that are internal enterprise applications, services, etc.
    - There's now a public utility called DTI
        - Data Transfer Initiative is a non-profit
        - Intention and promise: to operate in the public interest, Eudamonic purpose
        - Transparency
    - Requires ongoing stakeholder participation and a robust process
- Ben
    - There needs to be an exit mechanism for groups so that you don't have to escalate everything to a governance layer, if nodes, or groups want to exit
    - The Bad Place- shared lists of defederated servers
    - MLS is a forkable group protocol
    - Apply Occam's razor, don't be too ambitious
    - Identity is not a credential, it's a process - all the proof that the issuer issued that thing
- Shay:
    - Shared accountability
    - Credit Members have proportional representation

- Curtis [Google wallet]
    - Why “utility”?
    - With physical infrastructure it makes sense, but does it in this context?
    - Is there a utility that is not a government backed thing?
    - Issuers are hard to change after the fact
- GLAM perspective
    - Art/luxuries precedents: museums can be for/non profit, can be public
    - Are “objects” nodes?
- Victor:
    - May be beneficial to have localized manifestations (a la sovereign data hubs for example)
    - This could help drive adoption

What are some of the properties of Public Utility?

- **Governance**
    - Requirement for public membership
    - Proportional representation
    - Technical and all steering committees (rotation of members)
    - Open records
    - Stakeholder agency
- **Universal access**
    - Equal universal access
    - Cash-free?
    - Accessibility
- **Business model**
    - Cash free?
    - Shared revenue or dividends
    - Business model needs to “work”, be stable, in order to be regenerative
    - Skin in the Game model (code of ethics)
    - Partnerships have to have aligned incentives
    - Shareholder ownership, stakeholder ownership (skin in the game)
- **Security**
    - Open source?
    - Enforceable
- **Must have strong adoption driver**
    - Value exchange
    - Buy-in
        - Incentives for contributing and participation (self-evidence as the bar)
- **Anti-capture**
    - Protection from privatization (anti-capture)
        - E.g. schools and prisons
    - ICANN is a good model
    - Don’t be vain/overly ambitious about the scope the organization takes on

        - innovate only as much as possible, even wrt “great ideas from within the organization itself”
- **It just works** (like water from a faucet)
    - You don’t have to know how plumbing works in order to get water
    - Accountability for when things go wrong
    - Could be an eco-system of experts (could be part of adoption drivers)
        - Who gets to be paid to do this, could be a vulnerability, moral hazard
- **It is funded**
    - Where is all the money coming from
    - Funding sources need to be customer-subsidized, or
    - No anonymous donations?
- **Resilience and Recovery**
    - Data-gathering of the infrastructure to detect problems
    - Analysis framework to help solve them
    - Elevate and address issues proactively
- **Do more with less…**
    - be cautious about what you take on… it should do as little as possible and no more
    - Could be a sequenced approach
        - As little as possible on thing A, then tackle thing B
- **AI will drive this**
    - Responsibility for keeping data safe (e.g. attorneys)

Distributed Server Infrastructure

- Distribution the edge / OS
- It just works
- Support & maintenance
    - Ecosystem of experts
- Value Creation
- Credential issuer?

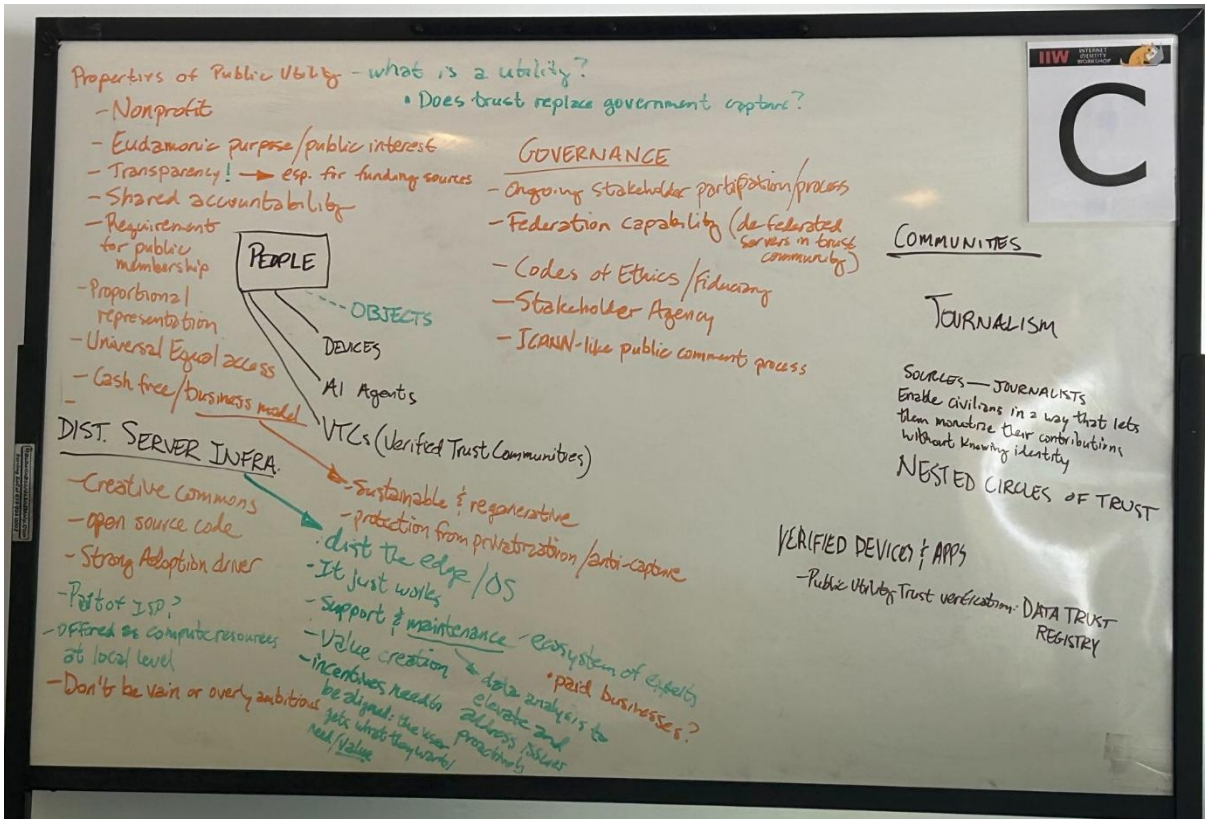
Properties of Public Utility - what is a utility?
• Does trust replace government capture?
- Nonprofit
- Eudaimonic purpose/public interest
- Transparency! → esp. for funding sources
- Shared accountability
- Requirements for public membership
- Proportional representation
- Universal Equal access
- Cash free/business model
PEOPLE
OBJECTS
DEVICES
AI Agents
VTCs (Verified Trust Communities)
GOVERNANCE
- Ongoing stakeholder participation/process
- Federation capability (de-federated servers in trust community)
- Codes of Ethics/Fiduciary
- Stakeholder Agency
- ICANN-like public comments process
COMMUNITIES
JOURNALISM
SOURCES — JOURNALISTS
Enable civilians in a way that lets them monetize their contributions without knowing identity
NESTED CIRCLES OF TRUST
VERIFIED DEVICES & APPS
- Public Utility Trust verification: DATA TRUST REGISTRY
DIST. SERVER INFRA.
- Creative commons
- open source code
- Strong Adoption driver
- Part of ISP?
- Offered as compute resources at local level
- Don't be vain or overly ambitious
- Sustainable & regenerative
- protection from privatization/anti-capture
- dist the edge/OS
- It just works
- Support & maintenance - ecosystem of experts
- Value creation
- data analysis to elevate and address issues proactively
• paid businesses?
- Incentives needed to be aligned: the user gets what they want/need/value
C

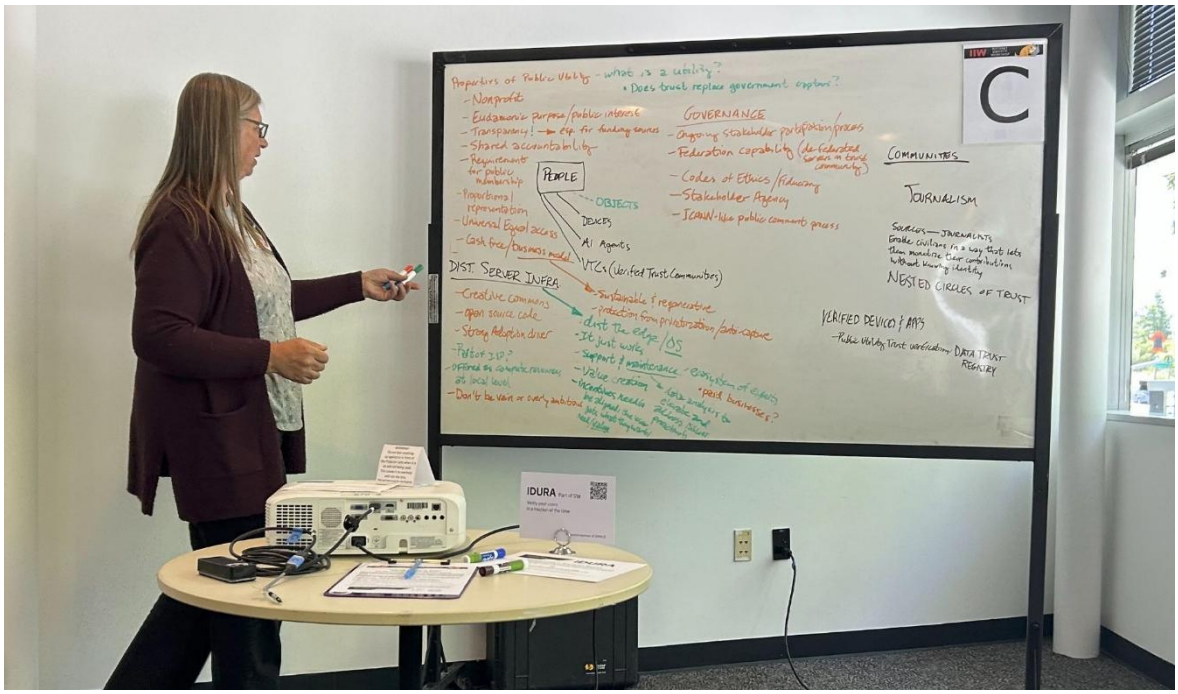
IDURA

Maintaining Our Humanity - What is Identity? Spirit Tech + / Human Agency

Session Convener: Diane O'Neal
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

dianeoneal.com
daocreativelabs@gmail.com

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Diane
Words like Identity and Protocol I think about differently then you.
Open Space
"Spirutal Technology"
There is much more technology we have developed as humans beyond the machine stuff.
Chant - Hawaiian Chant.
Used at the beginning of meeting to set space for opening.
then we will share names - land into group together.
Say name some one bringing to conversation. Place.

People
Grandchildren
Labradoodle
Mother in Law
Friends named Charles
Son x 3
Clients and all they have been through
Two Kids - driven to build better world for them
Daughter - went away to college
Business Partner who passed away
Lucy Luge - German Philosopher
Deceased Business Partner Part Mom
Author John Havens - Heartifical Intelligence
Matt Taylor
Friend Jeff McNeily
Samantha Sweetwater - on Kinship journey

Places
Italy
Vancouver / Cascadia
Japan
SF Bay Area Coast

Internet Identity Workshop Community
On a Sail Boat in a calm place
Azores Portugal where I was born
SAS Lounge at SFO
Dock on Colorado River in Arizona
Playing together Instrument and game
Land near state forest with a stream
100 Achers tend on schomish river and salish sea - grandmas and grandpas of the land
Mom's backyard
Our Living Room with family

Diane - feel into how much more is here now.
Naming things brings them forth, and gives life-force to our intentions and ideas, in a way that thinking alone does not. This is a type of technology that is very simple and we have available to us.

Moving through extension and thoughtful protocol to open a space and to bring certain elements and energies into the circle also sets each of us up individually with a mental position that will inform what will come next. It also brings the group into greater coherence together.

We will move forward with 3 Core questions - thrown into center of the circle.

What is motivating you - what is important and valuable to you

- What is Identity? - what is it we are protecting,
- Agency - Dimensionality - what is ours to do?
- What is technology - and what is the purpose - ultimately?

We get to marinate in them…and informed by others' understanding.
Be positioned to understand what we are doing….

Adopt Tyson Yungkaporta's shared traditional method of relating: Sand Talk
Build on the conversation in the middle -
Listen - let what is said land

Chris - known by security work, cryptography - why I entered into this - i really wanted to build spaces for collaboration. For people to be able to collaborate. Improve - grow - make things better. When people didn't feel safe they didn't want to cooperate - repeated over and over in role playing games. now we have "safety rules" - what are the rules that allow us to play more powerful games. A lot of what is going on right now - we are learning how to be more secure. Not learning how to be more safe. Keeping us from collaborating.

Chindy - especially around what I did around "identity verification" there is a soul or spirit behind that identity - so focused on the 'name' 'address' - not that "thing" it is a person with beliefs and trust.

Jeff - Lisa Lavasure - "Safety is a state of being.

Brian - I would like to quote words of a friend - safety - is a false promise. Moral and ethical reasons aiming for absolute safety.

Wen - thinking about history of western civilization. Dichotomy - needs more attention. We have come to accept a mechanical view of the world vs. ecological
Identity only exists in relationship to other things.

Technically minded thing - greater then human…persue - decentralized identity via cryptographic means - understandable reason - proof by machine. May be removing actual trust.

Brad - I come to this from couple angles - one of things about AI that knows you. SelfActual.io also working on a book - can't explain how brain works. Snow globes of reality

Miko - I want to put language around Harari - Nexus -Intersubjective reality . I blame decart for everything. Ergo Sum - based on cogido.

Ubuntu Philosphy - because we are. Witnessed Identity

(some confusion it was Decart - not Dick Hardt)

We through out safety because it is not measurable. Our limbic systems are under attack.

Dave is a novelist - the phrases that I hate "democratizing technology" everything that have ever seen that is democratizing technology - turned other things to shit - pump and dump schemes. If anyone can come in and do it - we see hierarchies build so it is not democracy. I would like to see real democratizing technology - one human one vote - and recognizing these technologies. Average people rise up and say that is not what I want to do and say no we are the majority.

what do you mean - "democratizing technology" - blockchain and bit coin.

Frank - one thing that is coming up for me is importance of relationship. Inner relationship means. What is the thing that should be driving relationships from a reciprocal fashion - tenders and regenerative tenders. Relationships between us - expressed through action of love - tending is love in action. Putting energy in whether or not getting something in return - where does it come from - the heart. First person culture - the relationship to where you are at - place, language - is always loca. Permaculture represents way we should be living vs. How we are taught to live.

How we heal our trama is through relationship and tending. Repeating this commitment. REpeating on a protocol basis. We are drifty creatures - it is important ot have protocols in space to remind where we are at.

Diane - technology beyond machines. They design objects to re-infoce stories values, reinforce that which matters to them. Textiles - hold information of those stories - they are reminded - so when they use the remain anchored. that is what is coming up for me.

Cindy - I worry that our young people are so

Textiles are a technology - communication across generation.
Interesting Trap we are in - “only new things” are technologies. Presentism.
Anything is invented before you are 35 is not technology everything after 35 is inscrutibale.
The point being that technology changes - that is the whole point. It is how se describe what the change is in our world.

Technology expands and grows - lazer disk - CD -

Ancestral
We have other systems that have different systems and availability expressivity and accessibility.

Diane - questions about what drives those choices

Joyce - We are creating things - verbs are short srift.

Night - doesn’t get that much attention.
Day - gets all the attention.

We are giving attention to the nouns - things to do things.
But no one is doing the dishes.
Doing the process in anything that we build.

Humanity and digital
Night and day
we are doing “thing” things - what is the doing.

Indigenous cultures - Love fabric illistraiton - Fabric is a noun that expresses the thing.

The technology process of being woven.
Energy and intention - into that product.

Touch something - that feels sacred and feels more too it.
hand made of different quality then mass produced things.

Jeff - regarding fabric and things like this - human biding

Living is LOCAL - phrase of mine.

Combination of honoring tradition and ledgering tradition.
Tethering to local - de-tethering in place.

Technologies to help focus and drive attention to local.

Day - going back to the verbing. What is happening here is a long unfolding. Time shovenism - with what is technology. I am liking the perspective of the universe awakening to itself. when I think the technology as the social. This is what we do - technologizing from the beginning - our existence is all that. Weaving. Tending and Weaving is landing for me.

Darius - we each of us is primarily of the expression of the creativity of the universe - in the relationships. Family and local community.

The technology that most of work on with internet and computers.
Makes possible for those connections.

Very powerful.
Would do well when create arficats and protocols.
Each of us being expression in creativity of the universe in a part of how it does it.

The universe is awake we are the ones waking up.

Without intelligence and creativity it does'nt have a way to express itself.

Brad - I'm curious how we build this in..from your heart to - not about chanting, universe - they don't want to think that big - things meet where language is.

Diane - basics that inform design.

Things saying
process as much as we think about products
relationships are foundational

Sense of local being important - right relationship local to global.

put all of this on a map and be conscientious.

Brad - Something like first person project and map on to human history - this is where our brains are at. If people are interested. How it maps to how we do things.

Matt - Local is Living. I think that is profound. I have been thinking about this - you have this spectator mode - on IG/TikTok - they are in the first person. They are the ones having the experiences. One of them it he desire for entertainment - "consume" someone elses experience - they would just go and have those experiences - TV

Defalt mode of being -

Play - you wanted "to be entertained"

Justin - plays and the history of that - people sitting around a camp fire creating a story - listening to a story - not having the experience - having the experience of hearing the story. Socially

conditioned to inflate those to. Kids Watch on Youtube then they think they know how to do that. Commoditation of story telling platforms. I think it a human thing - to want to experience things out side of our own experience.

Diane - thank you for taking a risk.

What is technology? what is technology for?
What is identity? - identity & consumption - builds sense of self through what we consume. how we are controlled/ controllable through that?
Identity as Story - brad touched on this - our sense of perception of who we are and how/where we fit - Telling self about where we belong.
Identity as Agency - what is ours to do and what is ours to hold.
What are trying to protect and hold for.

May the ideas explored here empower all attendees to have greater insight and clarity about what they are building and how it serves the outcomes that will truly serve us, each other's highest good, and life.

ZKP 101

Session Convener: Brent Zundel
Session Notes Taker(s): Nat and Lauren

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

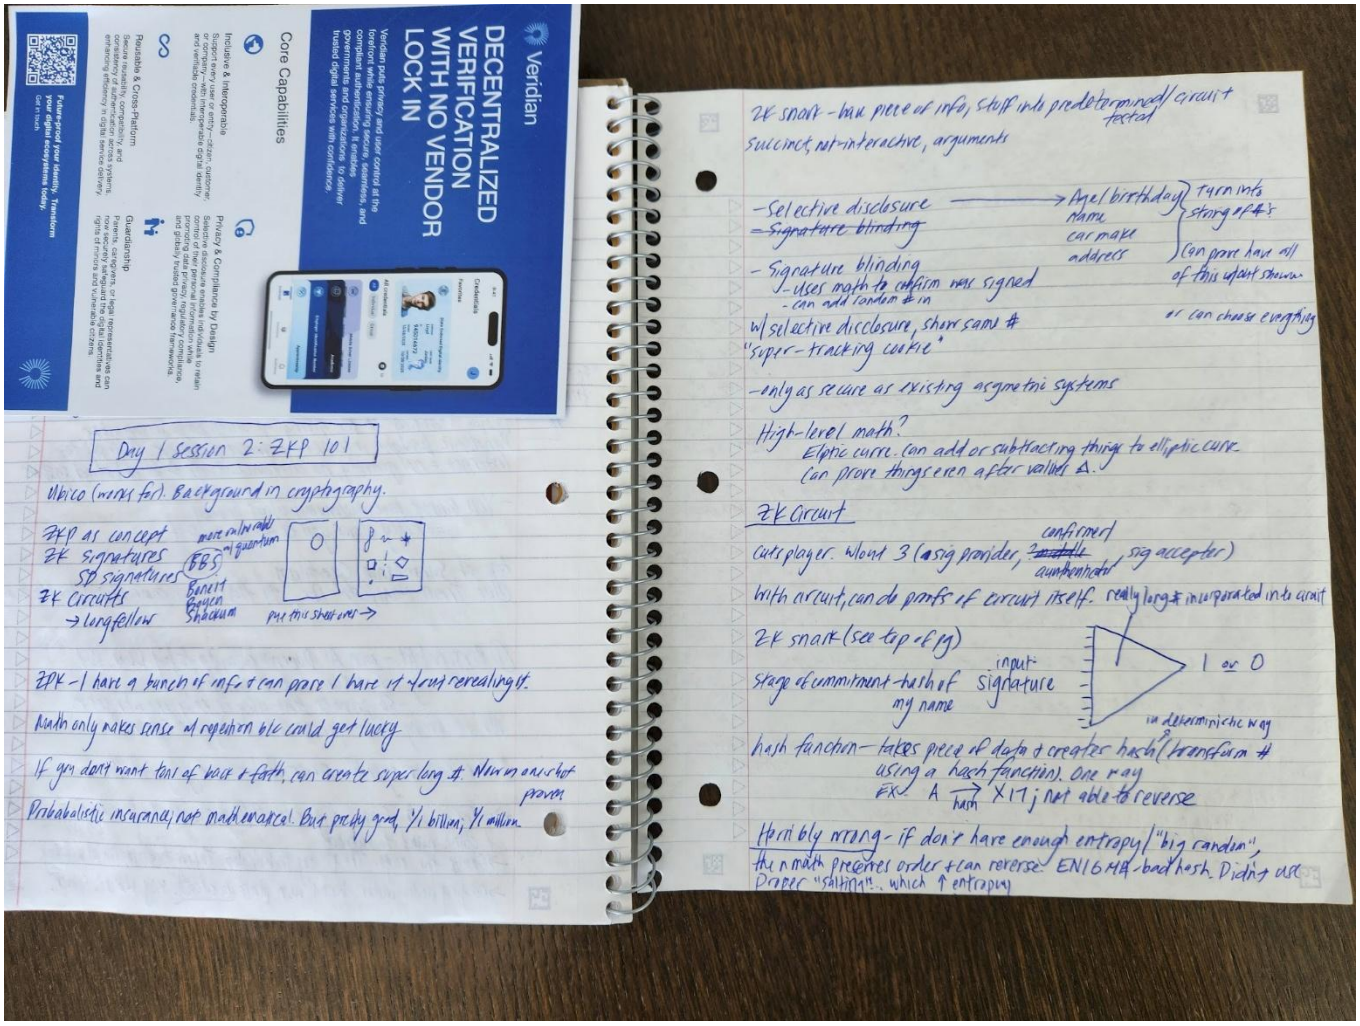

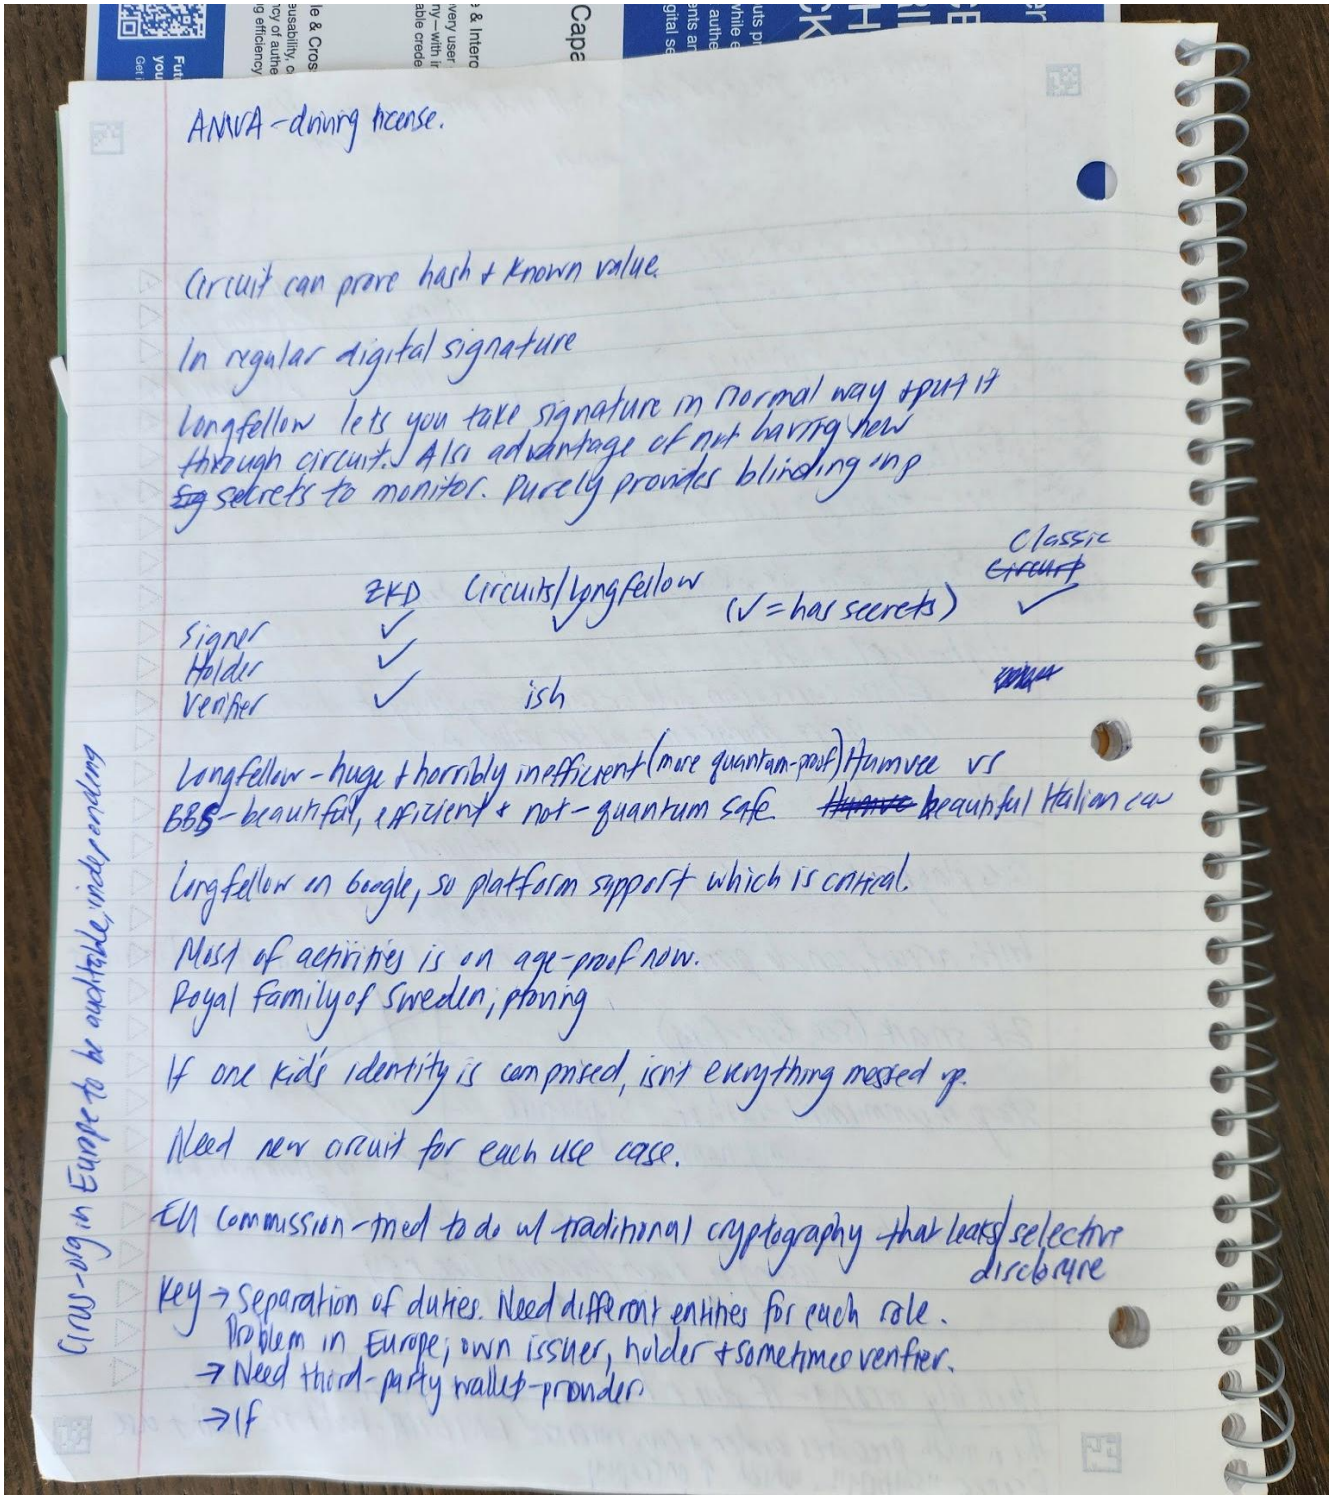

rest of the notes here, not sure how to get them in with the excellent picture: https://docs.google.com/document/d/18W5HihALh49Uuz6Ct1VU0YpWLHsxzjk7iqx2Xizpt74/

Authorization in an SSI Ecosystem

Session Convener: Jacob Siebach
Session Notes Taker(s): Jacob Siebach

Tags / links to resources / technology discussed, related to this session:

Authorization, Abacus AuthZ

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Authorization requires data. If clients utilize an external authorization engine, then the AuthZ engine can contact cloud agents to ask for proof requests when authorization policies require it. Live demonstration of this in action.

Suggestions: Authorization engines should consider AuthZEN for standards; delegation should be explicitly handled.

Options: Implementation can determine if requested data is cached or at what frequency it is requested; connection to data sources is determined by the needs of the business.

Privacy Preserving Biometrics + "Passkey for Personhood" / Scott Jones with Realeyes.ai

Session Convener: Scott Jones
Session Notes Taker(s): Scott Jones

Tags / links to resources / technology discussed, related to this session:

Demo — Reusable Human Credential in a Browser Wallet - verifeye-wallet-demo.vercel.app

A working demonstration of VerifEye issuing a reusable human credential into a browser wallet without revealing PII.

Flow: visit a site that requires proof of human + adult (Tinder and Bumble are used as examples), run a quick VerifEye check, choose whether to store the credential locally behind a passkey or use it one-time. If stored, visit a second site and the credential is recognized instantly — no re-verification, no central database.

Two claims demonstrated: you are human, and you are an adult. About 30 seconds end-to-end. The demo is deliberately wallet-agnostic — the issuance and cross-site presentation pattern would apply to any VC-compatible wallet.

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Topic. Session opened with an overview of Realeyes.ai's path into the identity space — privacy-preserving biometrics positioned in the gap between heavyweight IDV and approaches that no longer work (e.g., CAPTCHAs). Followed by a live demonstration of a generic browser-wallet flow (demo link) showing VerifEye issuing a reusable human credential (humanness + adult claims) into a wallet, stored locally and tied to a passkey, then re-presented at a second site without re-verification. Discussion followed.

Discussion Threads:

Revocation and refresh. The wallet should receive proof of non-revocation rather than the verifier checking status in a way that exposes it to issuer-side knowledge — the issuer knows about revocation, but the verifier doesn't need to. Joe Andrieu noted parallel work on removing status checks from VCs more broadly; potentially adjacent. Discussion also touched on the importance of refreshing the biometric anchor (not just the passkey) for cases like age verification.

ZKPs and privacy of embeddings. Substantial discussion on whether biometric embeddings should be considered reversible. Reviewed the spectrum of privacy-preserving options — ZKPs, multi-party computation across nodes (MPCs), and others — noting that the right choice depends on the risk profile and use case. ZKP work is in motion across several threads: SIROS / wwWallet (Longfellow ZKP approach), ZKP pseudonyms via the First Person Project, and related lines. Acknowledgment that ZKPs are not yet shipped but are being actively pursued.

Passkeys — what they solve and what they don't. Discussion of why a passkey was used in the demonstrated setup, and whether storage-plus-passkey alone could be relied on to ensure the right user. Covered the gaps biometric verification fills around continuity of control, uniqueness, and freshness — verifying the right person is present, not just that someone has access to the device.

Continuity of control. Discussion ranged from passkey continuity in general to age verification specifically. The credential-handoff scenario — where one person initiates a verified session and another takes over (e.g., a child accessing gambling or adult content via an adult's verified session) — was discussed as a real-world failure mode that continuous verification can address.

Community-scoped vs. global uniqueness. The session drew an explicit distinction from global uniqueness approaches: the proposed model is community-scoped — uniqueness measured within a defined ecosystem rather than across all of humanity. Examples raised: a person might be a gamer, a scientist, a Californian — each a distinct community in which their uniqueness can be measured. Individual companies similarly need to manage the aperture of their community's uniqueness — e.g., an e-commerce platform wanting to detect a returning offender previously offboarded.

Specific Questions Raised:

Integration ease for relying parties. Given the demo example, how easily could N sites adopt this pattern? Stated philosophy is plug-and-play with minimal effort required from the verifier; implementation specifics TBD.

Standards alignment. Wallet integration work (starting with wwWallet) is aligned with eIDAS 2.0 — confirmed in the room by Leif from SIROS — with the intent to follow existing standards or codify new ones where relevant.

Authenticator role. Clarified that the underlying system supports multiple use cases — personhood, uniqueness, demographics, and authentication. The demo focused on personhood/uniqueness/age, but the same capability extends to authentication.

Age verification accuracy at critical thresholds. Discussion of 18+, 13-17, 21+, 70 vs 75. Group sense that particular models are needed for particular thresholds with heuristics for interpreting confidence. An ask for the development of a 21+ model for alcohol scenarios was noted.

Vendor longevity / integration risk. Discussed Realeyes' founding in 2007 and long-standing presence in the industry, with a Big Tech anchor customer providing scale and rapid expansion in 2026 into new verticals like the gig economy.

Browser vs. app. Agnostic — integration is wherever the business case exists.

Client-side vs. cloud vs. on-prem. Strong room interest in client-side for the privacy benefit. Tempered with a risk-profile framing: client-side fits well for use cases like proving age to buy alcohol; less well for protecting high-value access (e.g., $1M-stakes flows) where adversarial sophistication is higher and the client-side implementation is, frankly, a sitting duck. Deployment model should be use-case driven.

Open Areas:

- Standardized ZKP path for biometric embeddings — multiple threads in motion, none shipped
- Standards alignment for wallet integrations beyond wwWallet
- Tunable uniqueness aperture for ecosystem operators
- 21+ age model for alcohol scenarios

Demo & Contact:

verifeye-wallet-demo.vercel.app

Scott Jones
scott.jones@realeyes.ai
Realeyes.ai
Book a call with me

Slides: Scott also referred to slides from a presentation to ToIP in March 2026.

This includes:

- Realeyes' lineage and origin story
-

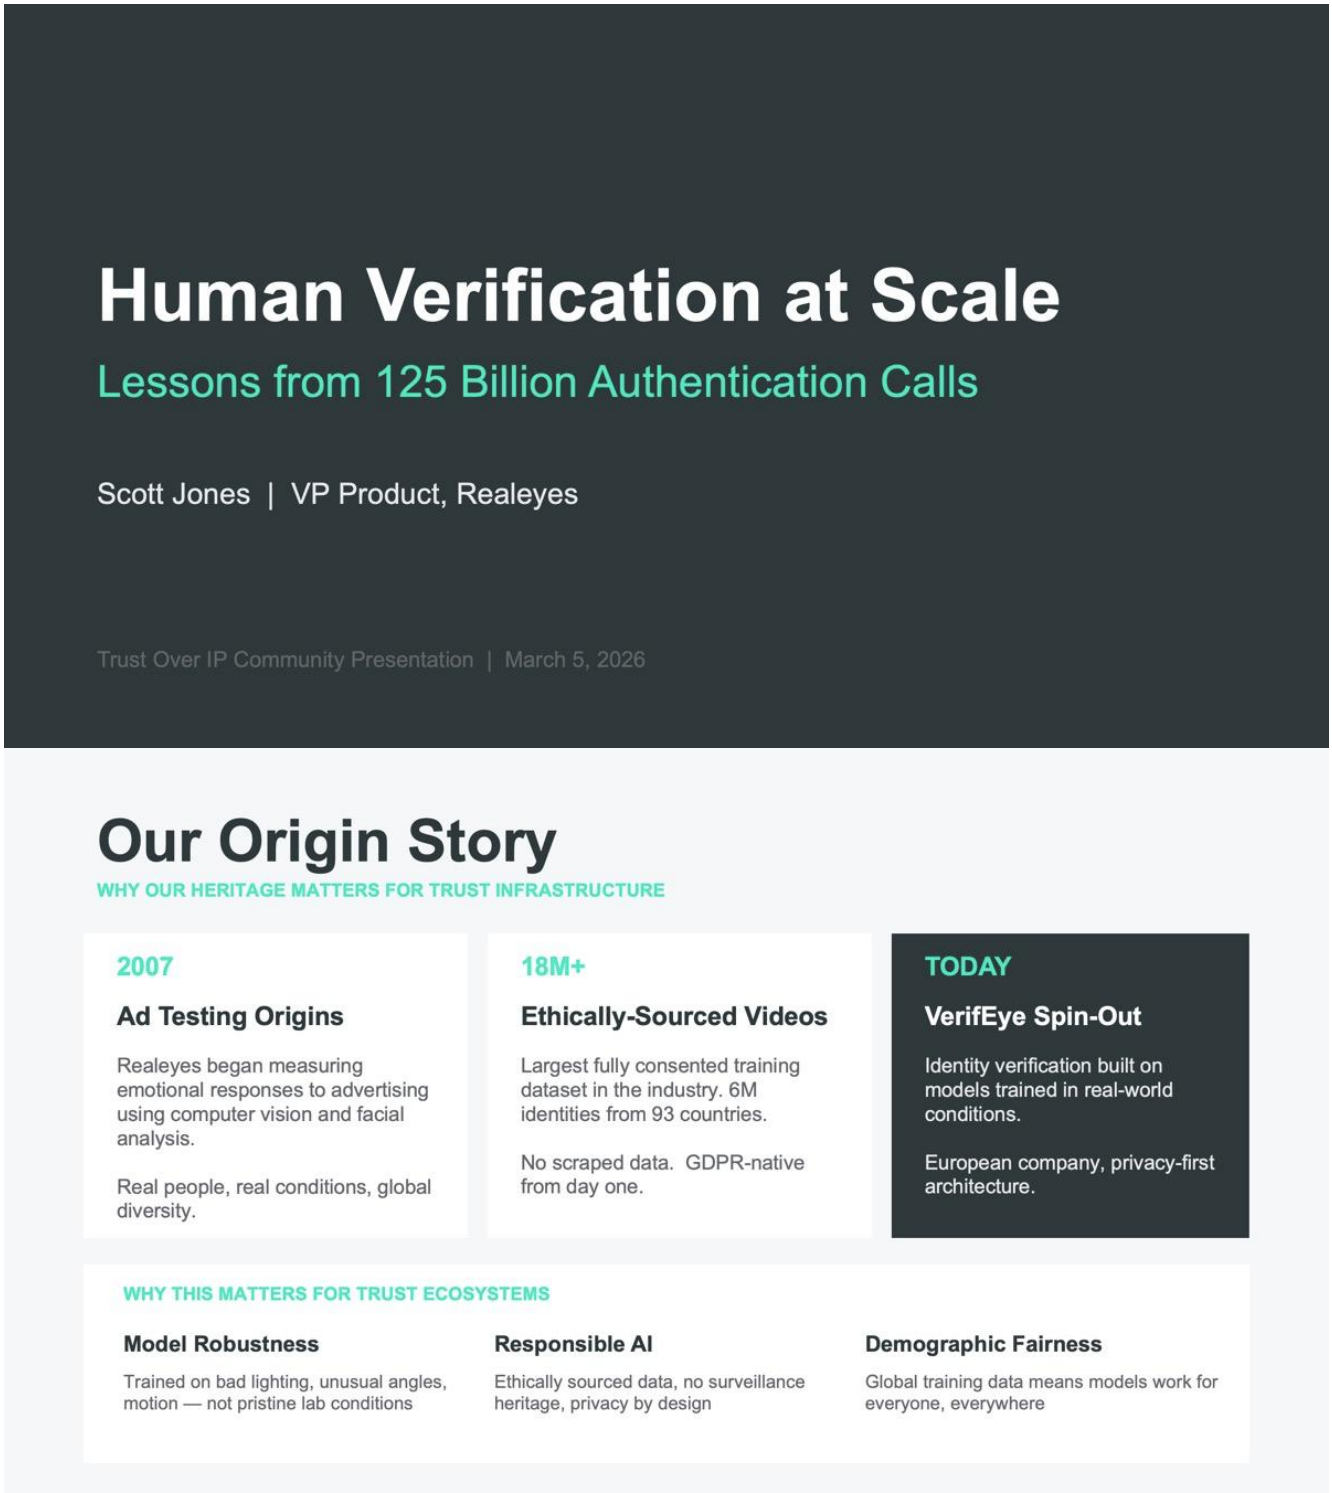

Operating at Internet Scale

125B+

annual verification calls

SCALE CONTEXT

- Billions monthly for the largest social media platforms serving billions of users
- Account recovery at 500M+ checks/day
- Fraud prevention across consumer apps, gig economy, gaming, e-commerce and more
- Compliance verification for age-gated content

This is not theoretical infrastructure — it is battle-tested at the scale ToIP envisions for trust ecosystems.

The Verification Gap

CAPTCHA

Low friction
Frustrates users
Fails against bots
No identity signal

LIGHTWEIGHT IDENTITY

The Gap

Prove humanity
Verify uniqueness
Confirm age/demographics
Minimal friction
Privacy-preserving

FULL KYC / IDV

High assurance
Heavy friction
Kills conversion
Regulatory burden

WE COMPLEMENT IDV

Major platforms use ID images + profile photos + selfie frames to authenticate users with our model

Enables re-verification without repeating full IDV — improves conversion by delaying heavy KYC until critical regulatory moments

Most trust ecosystems need something in between — stronger than CAPTCHA, lighter than KYC.

What Passkeys Can't Do

THE HOLDER BINDING PROBLEM

WHAT PASSKEYS SOLVE

- Device-bound authentication
- Phishing resistance
- No password management
- Strong cryptographic binding

WHAT'S STILL MISSING

- Uniqueness verification *-10 accounts on one device, 10K devices tied to one face*
- Cross-device continuity
- Account recovery
- Age / demographic verification
- Change of control detection
- Fraud signal generation

Passkeys prove "this device is authorized" — not "the right person is holding it."

Two Verification Modes

REALEYES PROGRESSIVE VERIFICATION PHILOSOPHY

UPFRONT

Verification at the gate

- Account creation / onboarding
- Age gate compliance
- KYC supplement
- High-value transaction auth

One-time check, personhood credential issued

CONTINUOUS

Holder binding throughout

- Gig Economy
- Gaming & gambling
- Remote proctoring
- High-security work sessions
- Shared workstation environments

Periodic re-verification, same person present

Spectrum: Lightweight passive checks -> Basic liveness -> Advanced anti-spoofing -> Highest security active liveness

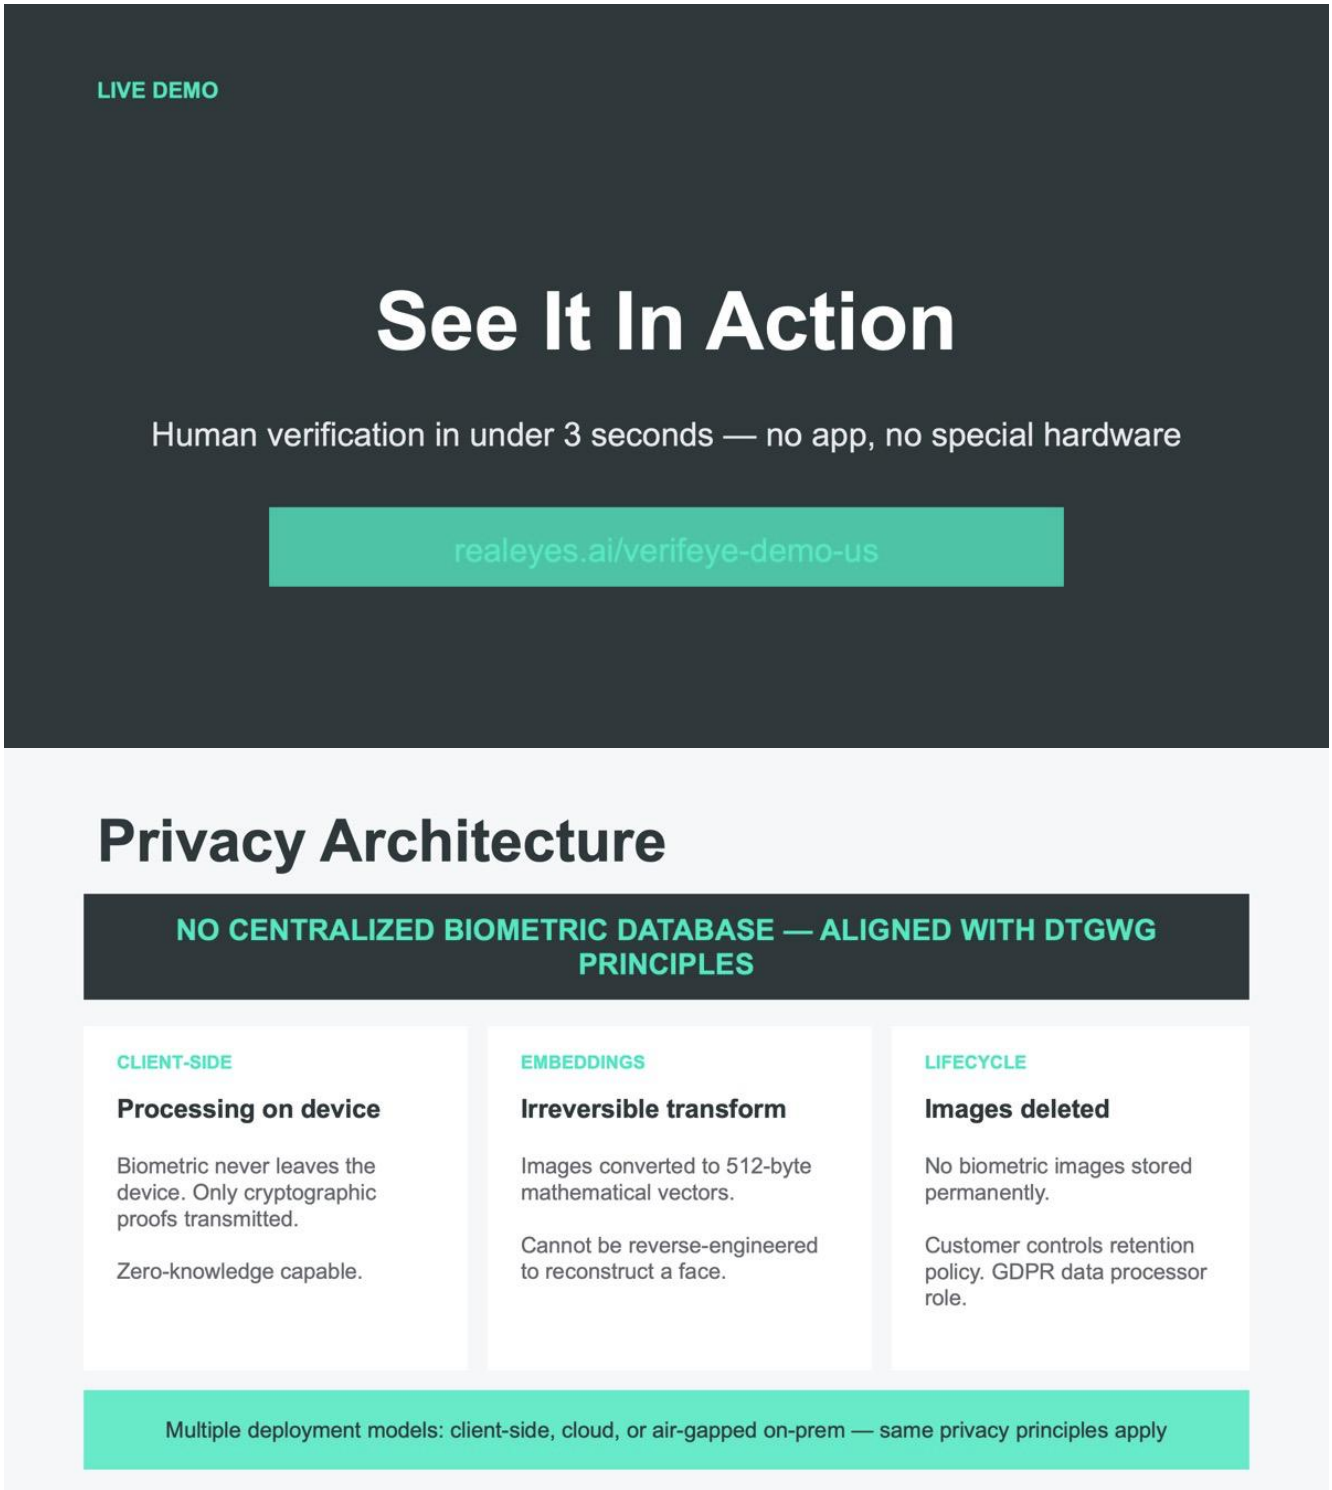
LIVE DEMO
See It In Action
Human verification in under 3 seconds — no app, no special hardware
realeyes.ai/verifeye-demo-us
Privacy Architecture
NO CENTRALIZED BIOMETRIC DATABASE — ALIGNED WITH DTGWG PRINCIPLES
CLIENT-SIDE
Processing on device
Biometric never leaves the device. Only cryptographic proofs transmitted.
Zero-knowledge capable.
EMBEDDINGS
Irreversible transform
Images converted to 512-byte mathematical vectors.
Cannot be reverse-engineered to reconstruct a face.
LIFECYCLE
Images deleted
No biometric images stored permanently.
Customer controls retention policy. GDPR data processor role.
Multiple deployment models: client-side, cloud, or air-gapped on-prem — same privacy principles apply

Scoped Uniqueness

"All parties control their own subgraph of trust relationships." — ToIP DTGWG

GLOBAL UNIQUENESS

Centralized approach

- Requires global biometric database
- Specialized hardware
- Centralized infrastructure
- Privacy / sovereignty concerns

COMMUNITY-SCOPED

Decentralized trust graph

- Uniqueness within trust ecosystem
- Any device with a camera
- No global index required
- Ecosystem controls own policies

Different problem, different architecture. Most trust ecosystems need uniqueness within their governance framework — not a global registry.

How We Fit the ToIP Stack

BIOMETRIC HOLDER BINDING AS LAYER 1 TRUST SUPPORT

Layer 4: Trust Applications

Layer 3: Trust Tasks (Credentials, Presentations)

Layer 2: Trust Spanning Protocol

Layer 1: Trust Support (Biometrics, Keys, Storage)

VERIFEYE PROVIDES

- User binding interface
- Biometric holder binding
- Liveness / anti-spoofing
- Personhood credentials
- Age / demographic claims
- Uniqueness verification

"User binding is the interface via which a Layer 2 implementation can request and verify biometric or other authentication information from a user." — ToIP Technology Architecture

Lessons: Fairness at Scale

DEMOGRAPHIC FAIRNESS

97-99%

accuracy maintained across skin tones, gender, and age groups

Based on internal and customer fairness testing across demographic segments

"IN-THE-WILD" TRAINING

Why our heritage matters

Ad testing = real-world conditions

Bad lighting, unusual angles, motion

Models work for real people

Not just lab-optimized datasets

THE LESSON

Verification systems that only work for some people and hygienic conditions are not verification systems — they are exclusion systems.

Fairness is not a feature; it is a prerequisite for inclusive trust ecosystems.

Lessons: What We Learned

ICONOGRAPHY

Simple beats literal

Emoji -> person holding camera -> iconography. Users respond to abstract visual cues, not realistic depictions.

SERVER VS CLIENT

Offer both, same principles

Enterprise wants server-side for compliance/audit. Privacy advocates want client-side. Support both with identical privacy guarantees — same embeddings, same deletion, same processor role.

WHY COMMUNITY-SCOPED WON

Global uniqueness creates blockers

Early assumption: everyone needs global uniqueness.

Reality: most trust ecosystems need uniqueness within their governance framework. Global approaches raise sovereignty and privacy concerns that block progress. Community-scoped uniqueness — where each ecosystem controls its own identity population — turns out to be what customers actually want.

Where We're Going

INNOVATION ROADMAP

EDGE-FIRST ARCHITECTURE

Client-side SDKs for iOS/Android shipping Q1 2026. On-device processing, wearable-ready models, zero cloud dependency option.

MICRO-EMBEDDINGS

512-byte embeddings fit on NFC cards. Physical + digital credential convergence. Works over Bluetooth, NFC, 5G.

OFFLINE VERIFICATION

Client-side processing means no internet required. Tap-to-verify at retail POS, air-gapped environments, border crossings.

ZERO KNOWLEDGE BIOMETRICS

Client-side ZKPs prove biometric match without transmitting the biometric. Server-side ZKPs enable anonymous uniqueness — verify 'one person, one credential' without linking to identity. Breach-proof by design.

DEEPFAKE DETECTION Novel system trained on proprietary dataset using latest-gen tools.

AI AGENT CUSTODY Biometric binding enables individuals to claim custody of AI agents while maintaining SSI.

Case Study: wwWallet Integration

WORK IN PROGRESS WITH SIROS FOUNDATION

1 User verifies via VerifEye -> 2 Personhood credential issued -> 3 Credential stored in wwWallet -> 4 User presents proof to verifier -> 5 Verifier confirms without biometric

USE CASE

Replace CAPTCHA with Reusable Credential

User proves humanity once, stores credential in wallet, presents proof across sites — no repeated friction.

EXPLORING

- Browser-native integration (no SDK)
- Passkey + credential binding
- Selective disclosure (e.g., age proof)

Open Questions

We would genuinely value this community's input

? Holder binding at credential creation

How do we minimize friction while ensuring the right person creates the credential?

? Credential refresh intervals

How often should personhood be re-verified? Context-dependent or governance framework policy?

? Privacy-preserving revocation

How do we revoke credentials without creating a trackable event?

? Cross-ecosystem uniqueness

When and how should uniqueness span multiple trust ecosystems?

EXPLORATORY: ZK Pseudonyms — proving uniqueness without revealing identity across trust ecosystems

Collaboration Opportunities

STANDARDS

Active

- W3C holder binding confidence method
- ToIP / DTGWG alignment

Seeking collaborators

- Personhood credential schema
- First Person Project engagement

INTEGRATION

Active

- wwWallet / SIROS Foundation
- eIDAS 2.0 wallet ecosystem

Seeking collaborators

- Wallet providers
- Trust registries

RESEARCH

Active

- Deepfake detection
- Fairness testing frameworks

Seeking collaborators

- ZK proof applications
- Privacy-preserving holder binding

If any of these threads resonate, I'd love to connect.

Kwaai Decentralized AI Infrastructure

Session Convener: Reza Rassool
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

UX for SSI Products

Session Convener: Janet Gonzales
Session Notes Taker(s): Janet Gonzales

Tags / links to resources / technology discussed, related to this session:

Agentic AI, Digital Identity, UX Design Principles. Takeaways from 5 years of lessons learned in UX

https://drive.google.com/file/d/1xIIEfuz6C0u4YTYtv84XLGmXxDF8I4uU/view?usp=drive_link

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

If you'd like to talk UX, I am open for discussion during and after IIW
https://calendly.com/janet-gonzales/meet-with-janet

Formal Security verification of specs – what would we want?

Session Convener: Frederik Krogsdal Jacobsen
Session Notes Taker(s): Jin Wen

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Discussion Notes

- Manual proofs currently take approximately **6 months** and cost **$150,000–$200,000**
- Code/spec changes may break existing proofs, creating ongoing maintenance burden
- The current process does **not** yield a formal certification
- Working Groups are reluctant to make changes due to proof fragility
- The group discussed what an ideal formal verification assessment tool would look like

Key Understandings

- The assumed trust boundary is **our own security mechanisms** — i.e., those are treated as correct
- Desired tool properties include: clear attacker model, readable results, low-guidance proof generation, speed parity with existing tools, composability (layering/profiling), and better property detection
- Ability to derive **conformance tests** from proofs was flagged as a desirable (likely stretch) capability

Outstanding Questions

- **Which protocol** is the primary target for formal verification?
- **Who captures the ROI?** Two candidate groups were identified:
  - Implementers
  - Profilers / ecosystem stakeholders

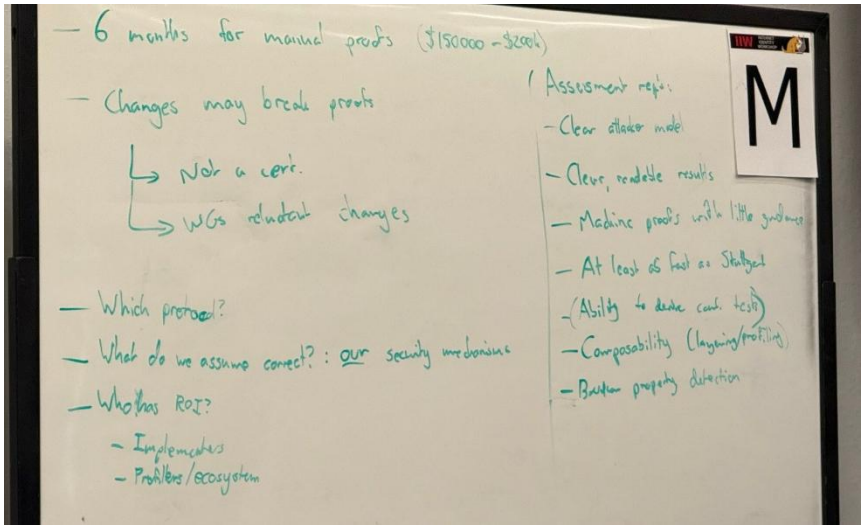

SESSION #3

Content Authenticity 101

Session Convener: Eric Scouten, Scott Perry
Session Notes Taker(s): (unknown)

Tags / links to resources / technology discussed, related to this session:

https://ericscouten.dev/2026/content-authenticity-101-iiw-42/ (slide deck used in talk)

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

C2PA - Claim Generator
Hardware or Software tool involved in creating content

CAWG - named actor
Individual or Organization who is operating those tools.

C2PA - claim generator
GPS data / time of capture
if known to hardware
Edit actions taken / AI used
Ingredients incorporated into content

CAWG named actor can describe

- Individuals or organizations
- Events or locations depicted in content
- Metadata / Context for content

Developing a consent task force.
about this.

C2PA - claim generator
x.509 certificate / COSE signature
Certificates (new) have C2PA specific key usage not interoperable with other purposes
Issued to hardware or software that demonstrates with C2PA rules

CAWG - named actor
Flexible Framework for using multiple kinds of digital credentials
Intended to bind credential to content
Optional for those that wish to

Scott presents dimoned image.
How do we control the usage of C2PA specification

C2PA Data Model
every manifest has only one claim.

Authorization 101 (now with AI) and IIW Session

Session Convener: Omri Gazitt
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Slide deck

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Introduction to authorization, building up a set of concepts:

- What is authorization
- Classifications of authz systems
- What is an authorization model
- Examples of authz models
- Granularity
- Categories of models (*BACs)
- Architectures
- Standards
- AuthZEN
- Bonus: AuthZ and MCP

Q: Does OPA support AuthZEN given that it is on the interop slide?
A: Not natively. There are multiple parties that have created “proxy” layers from AuthZEN to OPA.

All my old problems now have AI in them, but they are still problems

Session Convener: Justin Richer & Eve Maler
Session Notes Taker(s): Jin Wen

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Protocols and specs:

- OAuth 2.0
- OIDC (OpenID Connect)
- DCR (Dynamic Client Registration)
- UMA (User Managed Access)
- GNAP (Grant Negotiation and Authorization Protocol)
- RAR (Rich Authorization Requests)
- Protected Resource Metadata (RFC 9728)
- CIMD (Client ID Metadata Document)
- Transaction Token
- ID-JAG (Identity Assertion JWT Authorization Grant), cross-app access
- Macaroons

Workload and agent identity:

- SPIFFE (Secure Production Identity Framework for Everyone)
- WIMSE (Workload Identity in Multi-System Environments)
- MCP (Model Context Protocol), as the canonical agentic example
- Pets vs. cattle (agent identity model)
- RS-first flows (resource-server-first dispatch, agentic workflows)

Policy and authorization concepts:

- PDP (Policy Decision Point)
- Cedar (policy language)
- Runtime authorization assessment
- RO (Resource Owner) policy controls
- Least privilege (and Justin's critique of it)
- Object capabilities, capability composition
- Delegation (vs. impersonation)
- Purpose of use

Trust, risk, and human factors:

- HITL (Human in the Loop)
- Auditability of agents, sub-agent chains
- Liability apportionment

- Dark patterns in conversational AI
- Social engineering

Domain references:

- FHIR (Fast Healthcare Interoperability Resources)
- TPO (Treatment, Payment, Operations) as a healthcare purpose-of-use precedent
- NHS Digital "blue badges" and PayPal merchant onboarding (business-trust examples)
- Desired Outcome Contracts (Eve Maler & Nick Dam)

TLDR

Session: Justin Richer + Eve Maler. Thesis: most "new" AI agent identity problems are old security problems. Mine the past for ideas.

Key takeaways:

1. **Old patterns are coming back.** UMA's RS-first flow (client hits resource, gets redirected to AS) is now mandatory in MCP authorization via RFC 9728. UMA's resource set registration, token attenuation, and grant negotiation (GNAP) all solve problems agentic systems are rediscovering. Justin: GNAP and UMA are *"D-list celebrities. They solve real problems but didn't take off."*
2. **Three problems get conflated** (Eve): client dynamicism, client identification, and loose coupling between AS and RS are separate things.
3. **Center of gravity shifts to runtime.** Static pre-registration can't keep up. PDPs evaluating per-call with current context matter more than registration-time grants. Secrets and trust roots are *"completely out the window."*
4. **Delegation is finally taken seriously,** in two distinct forms: access permissions (UMA) vs. authority and obligations (acting *"as if you were me"*). Kerberos is impersonation, not delegation. The "Auth" in OAuth was always delegation.
5. **Least privilege is breaking.** Eve: *"Agents are basically most-privileged devices."* Justin: describing privilege by intent is gone. Cedar gives a calculus, but the missing piece is *purpose of use,* the "why."
6. **Three surprises framework** (Eve): delight, shock, screw over. Inferring intent from API calls no longer works.
7. **Auditability needs a new layer.** Sub-agent chains break user-level audit. Liability apportionment will drive the rest.

Verified: UMA's RS-first flow is genuinely reappearing (Kantara UMA 2.0 §3.1, RFC 9728 §5, MCP Authorization 2025-06-18 all share the same shape).

People to follow: Karl McGuinness (ID-JAG, AI agent identity), Sarah Cecchetti (Cedar profile for OAuth RAR), Eve Maler & Nick Dam (Desired Outcome Contracts).

Premise

Justin and Eve: most "new" AI identity problems are old security-protocol problems. The session mines past work for ideas that apply.

OAuth, DCR, OIDC: what was learned

- OAuth was built for website-to-website delegation with pre-registered clients. DCR filled the gap for cases like desktop email clients where pre-registration is impossible.
- DCR originated in OpenID (with UMA influence), then was uplifted into OAuth.
- Eve: *"Technical trust is easy. Business trust was deliberately slow."* Examples: PayPal merchant onboarding still required humans; NHS Digital "blue badges" saw fraud spike when the process got too fast.
- OIDC took roughly two years post-spec to get RPs on board, and the spec was tuned to RP needs. George Fletcher: the issuer/holder/verifier model has the same RP gap today.
- Eve, quoting a Gartner analyst: *"AI is the balloon payment on technical debt."*

Eve's Three Problems (often conflated)

1. Dynamicism of clients: new clients appearing.
2. Identification of clients: who or what is acting.
3. Loose coupling of introduction: how AS and RS find each other.

RS-first flows are back

- UMA's pattern: client hits RS without a token, RS dispatches client to the right AS with a permission ticket. Eve called it *"user-centric WAM."*
- This pattern is reappearing in agentic specs because they need the loose coupling. Justin: the new question isn't just *who do you ask* but *what are you asking for and how do you encode it privacy-preservingly.*
- UMA Resource Set Registration: RS asks AS for an opaque token ("banana duck") to hand to the client.
- Protected Resource Metadata (RFC 9728) is the passive, well-known equivalent.
- CIMD (Client ID Metadata Document) lets the client publish its own metadata, sitting alongside DCR.
- George Fletcher: discovery needs nuance. Public, protected, and first-party metadata aren't the same.

RAR and the encoding problem

- Justin: RAR is *"OAuth scopes that don't suck. JSON objects, not strings."*
- RAR works when the client knows what to ask for. When it doesn't, you need negotiation. GNAP added grant negotiation for this. Justin: GNAP and UMA are *"D-list celebrities. They solve real problems but didn't take off."*

The new context: agents

- Velocity, automaticity, nonlinear or fractal scale.
- Center of gravity shifts to runtime authorization assessment (PDP-style, per-call) rather than registration-time grants.
- Bar rises for RO policy controls: dynamic, expressive, not one-shot consent.
- Secret and trust-root management is *"completely out the window."*

Pets vs. cattle

Human identity is still pet-ish. Agents must be cattle: ephemeral, numerous, replaceable. Agent dynamism is dragging human identity along with it.

SPIFFE and WIMSE

SPIFFE solved workload dynamism inside a closed trust domain. WIMSE (IETF) is for SPIFFE-things crossing trust boundaries.

Delegation
Eve: *"The D-word has finally become a thing."*

Two kinds (Eve's distinction):

1. Delegation of access permissions: UMA-style, sovereign user grants access.
2. Delegation of authority or obligations: *"interests aligned so perfectly you act as if you were me, but I can track it was you."* Not identity delegation.

Kerberos is impersonation, not delegation. Justin: *"It does not know it is it."*

The "Auth" in OAuth stands for delegation. Justin's seminar trick. UMA and GNAP extend it: human to software, possibly via another human.

Chains, hops, transformation

- Even simple MCP looks like: human, agent, MCP client, MCP server, downstream. Sub-agents make it worse.
- Dimitri asked about IETF multi-hop delegation with proof chains.
- ID-JAG works when one party owns the chain. GNAP supports multi-hop awkwardly.
- George Fletcher: crossing trust domains requires transformation. Scopes and claims aren't portable. *"We haven't started tackling that."*

Token up or downgrade and capability composition

- UMA worked through token attenuation; it's coming back.
- Justin's counter-example: a delegate may need to do something the caller *cannot,* e.g., an email delivery system with mailbox access the caller lacks. Not a strict subset.
- George Fletcher: capabilities composition. Two permissions combine (user's plus operator's).
- Justin: this is what transaction tokens address. SPIFFE constrains the call graph; the txn token carries per-call context.

Mike Schwartz (Gluu)

- Least privilege is great *if you know the strategy.* With agents you don't, and it's too dynamic for HITL.
- Cited Sarah Cecchetti's draft-cecchetti-oauth-rar-cedar (Cedar profile for OAuth RAR; AS may return less-permissive policies than requested).
- Downscope is the easy case. Upselling from a downscoped token is the real danger, like a power of attorney to sell your car being used to max your credit card.

Intent: the wood chipper
Justin: agentic systems *"throw this conversation through a wood chipper because we pretend we can infer intent from client calls."* Pre-AI, an email API call meant *"I'm sending email."* With an agent, *"I might not have any idea what it's actually trying to do."* Anecdote: an agent asked to scan his keychain; *"NO."* It turned out to be looking for a config file.

Eve's Three Surprises

1. Delight: useful unexpected behavior.
2. Shock: unexpected, not harmful.
3. Screw over: unexpected and harmful.

Eve is watching dark-patterns research on conversational AI. *"If we figure this out, we will solve social engineering."*

Least Privilege and Purpose of Use

- Eve: *"I'm questioning least privilege. Agents are basically most-privileged devices."*
- Justin: describing privilege by intent has *"well and truly gone away."* Mobile app permissions show why; *"what is a map application?"* is tautological.
- Cedar gave us a usable policy calculus, evaluated at the PDP. Justin: *"It's still permissions. The missing piece is the 'why,' the purpose of use."*
- Eve: privacy law has had purpose of use forever (healthcare TPO: Treatment, Payment, Operations) but it has never really been implementable. The human interpretation of purpose is what's missing from policy languages.

Auditability and liability

- Sub-agents make audit ugly. Children and grandchildren of processes need to chain back to one entity. Splunk and Databricks already feel the pain.
- Justin's analogy: process IDs. We need that level of auditability, but at the agent layer, not the user layer. *"Process one, user one crashed"* doesn't help.
- Eve: liability apportionment rolls up from auditability. *"Solve liability and a lot of great stuff unrolls."*

Closing references

- **Karl McGuinness** (ex-Okta Chief Product Architect; karlmcguinness.com): prolific on AI agent identity, MCP, WIMSE, and OAuth agent extensions. Author of ID-JAG (draft-ietf-oauth-identity-assertion-authz-grant).
- **Sarah Cecchetti**: draft-cecchetti-oauth-rar-cedar; ex-AWS Cedar lead.
- **Eve Maler & Nick Dam**: *Desired Outcome Contracts.* Lifecycle of capturing and binding intent from natural language.
- An OpenID Foundation WG is collecting agent use cases. Possible follow-up session.
- Eve: *"If you solve for the hard human case, humans who get to say what's right for them, applicability to enterprises, workloads, and AI gets more tractable."*

Verification / References

"UMA's RS-First Flows, Now Relevant Again" was cross-checked against primary sources:

- Kantara UMA 2.0 Grant, §3.1 and §3.2.1: client requests resource without token; RS returns WWW-Authenticate: UMA with as_uri and ticket. Canonical RS-first dispatch.
- RFC 9728 (OAuth 2.0 Protected Resource Metadata, April 2025), §5: same shape via well-known endpoint. Does *not* cite UMA despite the identical pattern.
- MCP Authorization (rev. 2025-06-18): mandates RFC 9728 and WWW-Authenticate on 401. UMA's flow modernized for agentic clients.

Conclusion: UMA originated the pattern; current agent-facing specs are reusing it.

Meeting minutes compiled from live notes and transcription.

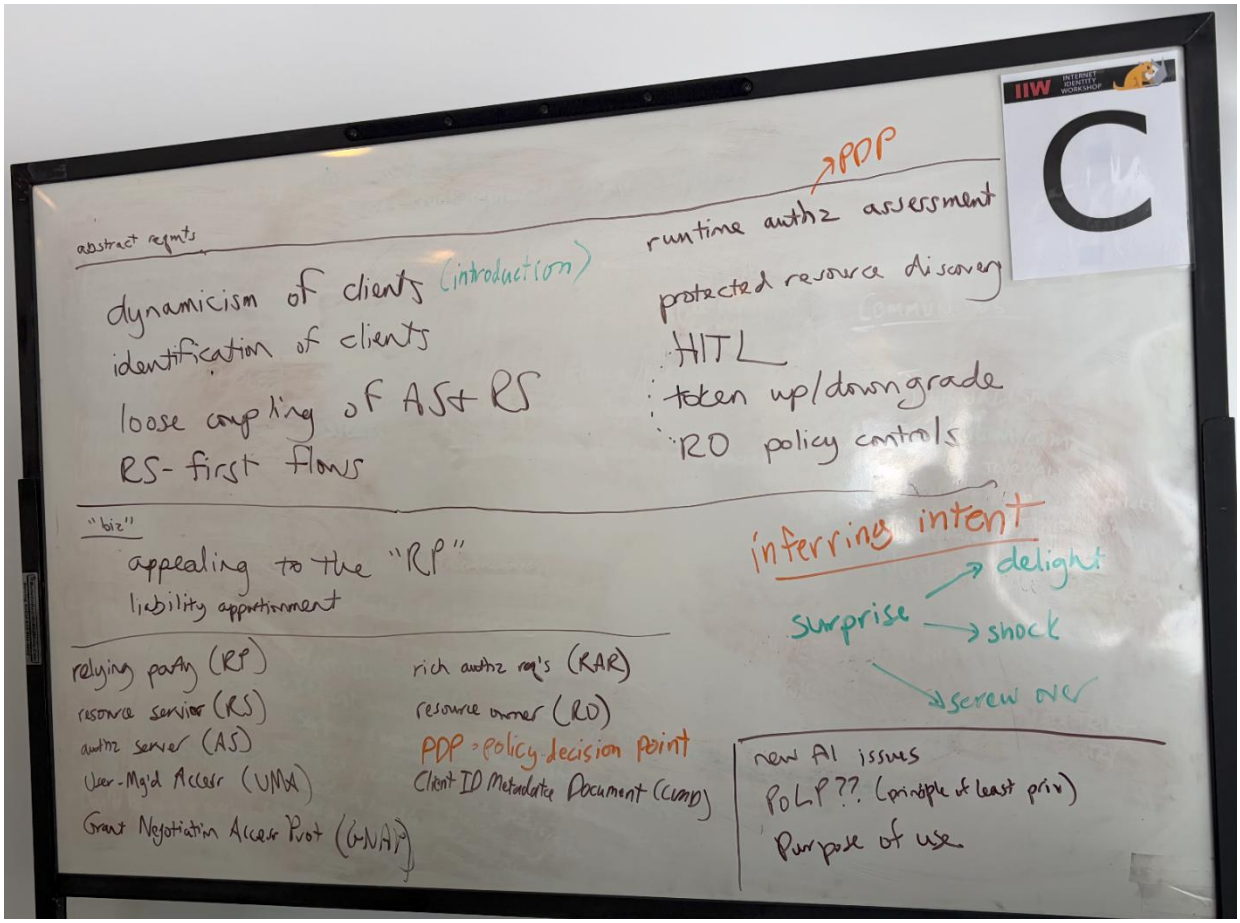

Why Personal ID in the era of AI?

Session Convener: Denny Wong
Session Notes Taker(s): Denny Wong

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Presence: Anthony H, Scott Jones, Aldofo Grego, Brian von Herzen, Atreedev Banerjee, …

- Discussion: How can we ensure that we do not hand over our data/knowledge/… to the AI company, as we did during the internet era, when we became the product?
- Today it there are protocols and frameworks that octroi digital identities to AI Agents, anchored on internet TLD, DNS. For example, (aria.bar) where rules or policy for the Agents can be set.
- The key part is to follow the value creation, as adoption drivers, as to why we would give an AI agent an identity, as ways to protect privacy, knowledge or IP?
- How identities could be a driver to drive the creation of market place for skills (agents) that can work to create new value. Think of hiring an AI agent to do your work, etc.
- …

VRM + MyTerms + Fiduciary Agents

Session Convener: Iain Henderson
Session Notes Taker(s): Iain Henderson

Tags / links to resources / technology discussed, related to this session:

I wrote them up here. https://hendersoni.substack.com/p/a-write-up-on-this-session-on-vrm

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

This post is dual-purpose. It acts the notes for a session I ran at The Internet Identity Workshop number 42. And then as my wider sharing of the same subject matter.

The subject was as above from the session announcement. My thesis is that we are now at the point **where all of the tooling required to empower people in the digital realm is in place**. And therefore the focus moves to specific implementations and adoption.

We discussed each tool or component in turn. We only had an hour so did not get too much detail on the fiduciary AI agents capability.

1. Sovereign hardware

We used the example of the Fairphone to ilustrate this point. This comes from an EU based provider, runs on a 'de-googled' version of Android with added privacy features, and much favours open-source and privacy-minded apps and utilities in their also de-googled App Store. Cloud back-up tends to use Nextcloud; so robust in functional terms and definitely more private.

The net result of this approach at this stage is that the Fairphone user has a sophisticated smart-phone, with all parts replaceable; and no tracking from the operating system perspective. Apps can conceivably track should they persuade the user to enable that. The approach still has some rough edges; but privacy minded people could now quite safely choose this direction for their next smartphone.

This is just one example of many of privacy-minded and 'local first' hardware and connectivity.

2. Identifiers/ keys (self-sovereign, self-created owned and controlled)

'Keys', that is to say decentralised identifiers/ DIDs with cryptographic properties that are a representation of an individual online, are clearly on an upward trajectory with many 'wallet providers' now in the market with various focus areas.

The one we majored on in this discussion was the soon to be launched MyKey (https://www.mykey.me) which has the USP of being highly functional; but also directly created, owned and controlled by the individual who wishes to gain more control over their sign-up/ sign-in methods (with other properties mentioned below). Conceptually, and before too long in reality, these keys/ identifiers can be life-long and thus highly enabling for individuals - because they have the capability to be transferred across service providers (like a mobile phone number), and have mechanisms to deal with 'lost phones' or compromised digital relationships.

The underlying point here is that if **YOU** don't own and control the primary means of representing you in the digital world then **SOMEBODY ELSE DOES.**

3. Identity (of a human)

A subset of the MyKey app is identity itself; and verification/ use thereof.

By that we mean what amounts to a core profile data-set. Typically name, date of birth, gender, home address and contact details. So the data asked for and provided to a significant number of the organisations we each deal with online.

Each of those data attributes might self-asserted, or assured, or otherwise verified.

In technical terms the variant of this we showed in demo was the standard Open ID Connect core profile. That is the profile that sits inside all 'social sign up/ in' methods; e.g. 'sign in with Google'.

4. Persona/ Context (roles within human)

We all have one body, but within that typically operate in different persona or 'contexts'. This can be as simple as 'me, the default one, or 'me at work', perhaps 'me the parent', 'me with power of attorney for someone else'. In the physical world we flit in and out of these different contexts invisibly. In the digital world we will often have different profiles or associated data attributes to signify these differences. In practical terms in our demo the individual signing up for/ into an online service can set up these different profiles, and then choose with which they are entering into a relationship.

This capability becomes really important for example in engaging AI agents, which need to know the context on/ in which they are operating in order to deliver appropriate responses.

5. My Data (personal)

Associated with each persona/ context there can and will be LOTS of data. Simplistically, which I think is necessary to scale, I mentally and technically bucket this data into these five categories:

- Core Profile(s)
- About Me
- Things I Have (or had)
- Things I Want
- Jobs to be Done

That subject is huge and complex, and includes the various interpretations of 'how does a person manage their data?' That includes personal data stores, vaults, pods, wallets and many more. But for now this is just the physical storage. For now, and for the majority of people they typically have two methods:

- Physical storage, filing cabinets for example, holding paper copies of key documents; birth certificates onwards

- Scattered across hundreds of web sites and apps, behind the terms of service and privacy policies or other. And maybe a spreadsheet, password manager or similar tool that are populated manually.

Neither of the above are sufficient, but that is now being addressed in the market with various options available. The one shown in the session was DataPal, a fiduciary data service in UK.

6. API's to/ from My Data

For the currently small number of people using a personal data service/ store/ vault/ pod, they will usually have some form of API and/ or webhook approach. These enable external parties to read or write (or both) to a persons own data store.

There is an important branch of these APIs being driven by regulation, specifically in EU. The Digital Markets Act mandates real time APIs that enable people to pull data out of the 7 huge 'gatekeepers' (GAFA plus Microsoft, Booking.com and TikTok). And another branch in the form of The Data Act, which does as above for IoT devices and services - including cars.

This will be an area of huge growth over the next decade.

7. Terms and Conditions of Service (ToS)

These typically 'contracts of adhesion', and their close cousin below (privacy policy) are what underpins many of the ills of the digital realm today.

They will continue conceptually much as they are, because every service is entitled to set out what it is, how to engage with it, how much it costs etc.

But two improvements will happen over time:

- AI agents, especially those acting on a fiduciary basis for individuals, will increasingly shine a light on deeply opaque practices within the ToS, and help their humans avoid such services. Or at least not stumble into them with no or limited awareness.

- Personal data practices, currently often conflated into the ToS, will be held separately so that they are more easily upgraded to standardising approaches such as MyTerms.

8. Privacy Policy (inc. cookie banners et al)

This is where MyTerms/ IEEE 7012 directly impacts. Quite simply a MyTerms agreement is a direct equivalent and thus an optional replacement for an organisation's own custom built privacy policy. It addresses the long since unreasonable practice in which we humans have a different privacy policy in place with every single organisation we deal with. Compounded then by the fact that we don't read them anyway.

9. Registration/ account (the relationship artefact)

We showed the latest version of the MyKey app proposing a MyTerms agreement, signing up/ in to a demo site, and recording the signed agreement.

That video demo is at the link below. https://vimeo.com/1187324394?fl=pl&fe=sh

10. Agency (doing stuff across the infinite loop)

Now somewhat short of time in the session, we simply noted that significant difference between

- **Fiduciary AI agents**, that work for individuals and only individuals with a duties of loyalty, care and fidelity

- **Non-fiduciary agents**, which may provide very useful services for individuals; but ultimately owe their allegiance to the entity that runs them as a service

Each of the above could easily be an IIW session on their own; so this was a skim through to back the original assertion that the tools that are required for 'VRM' and empowered people online are now in place; albeit at a base level. So the path from here must turn to adoption.

AAuth DEEP DIVE

Session Convener: Dick Hardt
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Agent Taxonomy - What "is" an agent? Ephemeral vs Persistent Agent Lifecycle

Session Convener: Danny Zollner
Session Notes Taker(s): Steven Tamm

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

(Session discussion drifted from initial topic to broader AI security concerns / implications)

Defining the AI Agent Taxonomy.

"What is the lifecycle for AI Agents"
Core Problem: "What is an Agent, and is there consensus on the schema and attributes for that Agent."

Core issues is
Differentiating between long-lived and ephemeral identities in the enterprise.
Agent Class (Long-Lived) vs Agent Instance (Ephemeral)
(Reference the slides)
Same issue with Hierarchical Delegation for Sub-Agents with "task spawning"

Question: "Is orchestration only associated with classes?" it was unintentional framing

Question: Maybe it's not the "model", it's the harness that matters, where the ephemeral nature doesn't matter. It'll just inherit from the harness. Maybe the distinction is whether it's a different harness?

Discussion about “Sub-agent” in the same sandbox, or in the harness, a different sandbox. Is it an identity problem to ensure that a spawned claude subagent would be “restricted” in access to ensure security? Can the harness restrict sub-agents by API taxonomy.

- Is the Ephemeral Agent class dynamically generated, or preprovisioned? Does it matter if it goes across platform boundaries?
  Danny asks “Doesn’t someone have to preprovision some kind of ‘upper-bound’ access, or else how can it get any tokens?” You need to give an allowance for the travel agent agent to buy a ticket, but it should be prohibited from asking to do a wire transfer. How do you specify that?”
- Does an agent class need to register all sub-agent classes? Is it a “maximum” or “minimum” set of scopes? If the sub-agent shares the subject claims… isn’t that dangerous?

(ChristopherA) Maybe it’s the wrong model? It should be about capabilities and authorities granted. The difference between ephemeral and long lived: it’s the principal who does the work. “I have given them the duty of my authority, so I may be the legal principal, but their authority”

- Long lived vs ephemeral is a descriptor, not a defining characteristic of their existence.
- “I’m trying to figure out the zoom in/out of the principals for ai agents.” Maybe it’s at rest vs in motion? To do identity governance?
- Multi-user/multi-party “object capabilities”? (See further https://medium.com/@harshiljani2002/the-difference-between-object-capabilities-and-identity-based-access-control-ea3f9ef243cd). Should we have a separate session on it?
- “What is the available landscape of agents, and when they do work, what appears in the logs as the principal that did the work?” (PamD)

What is the thing doing the work now? As a taxonomy problem in an enterprise scenario (i.e. policy enforcement). What is the list of AI agents?

- ChristopherA brought us back to legal frameworks and thought we should look at that to help enlighten the taxonomy. Agency and duties. Such as what is a commission? “The agent won’t be arrested for donating to a terrorist organization.”
- He was asked to clarify what is a legal duty by (Bryan F @Google)

“Why did you choose to differentiate ephemeral agents with JIT vs Stable long-lived agents that are centralized.” Danny said “you need to claim governance at some point, which is centralized”. Decentralized systems do revocation&rollback, which you “have to accept,” which is problematic.

- Workloads are agents, but what if they are lambda agents? How can you manage it?
- Kwaai.ai was talking about shared/decentralized model serving. How do persistent agents affect taxonomy? Danny says “yeah, it’s a way’
- Maybe a 2d taxonomy. Persistent vs non-persistent and centralized and de-centralized. PamD asks “how does the access change?” Digression about decentralized vs distributed.
- ChristopherA talked about secure enclaves on the mac with agents on his machine having different access. How can the agent buy azure resources on my machine with my keys, but due to KYC-requirements it can’t buy anything from Azure that requires a human? It’s a special class of agent that needs long-term access with a known principal tied to a real human. That category “responsible party” matters.

From the bank, if your agents are local, anything in userspace is fully open. “What is ‘you’ as the authorizing/responsible person vs you.”

Nick Steele (recently of OpenAI) tries to summarize: “Is there nothing else to be done to ensure access/fraud/etc.“ He claims the “bad agent” won’t exist in the near term. He wants to solve the problem around oauth/mcp/delegated user/etc. Human presence when necessary is valid for high risk activities; those are “near term” use cases. ChristopherA asks how do you get Agentic multifactor implemented? Nick says “the way authorization is going to work, is tied to tasks more than the agent itself, and the tasks are the token bound to the agent identity? That’s what I want to talk about. What does the identity need to have to get the task and perform the task? What does the authorizer need to know about it?” PamD “is that forensic, or gate?” Nick says “gate” Chris says “you gotta have segregation of duties”.

Anonymity for Civilian ownership of identity. Reconstitution w/News Futures

Session Convener: Bennet Harvey
Session Notes Taker(s): Voice Notes / not edited by a human

Tags / links to resources / technology discussed, related to this session:

https://www.linkedin.com/in/bennetharvey/
https://reconstitution.com/
https://newscopilot.com/

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Claude Organization of Raw Audio Transcript:

Project Overview: Civic Data & Journalism Platform

Data Infrastructure

- Currently built on AWS using a single graph database
- Seeking a more resilient, recoverable, decentralized architecture — away from AWS dependency and out of reach of vested interests
- Two patents held by the founder, being gifted to the entity
- Exploring a purpose-constrained open source license model — free to civilians and journalists, but used to extract revenue from corporations and government

Core Architecture

- Graph database maps societal data signals to individuals anonymously

- Built on PostgreSQL + Python using machine learning (not generative AI)
- Patent covers nested/hypergraph databases — one societal graph, one user graph — with apps that detect what is relevant to a user within the moving societal graph
- Called a "Human Context Engine" — societal context for decision-making, not short-term AI context

App 1: News Copilot (Journalist Tool)

- Allows journalists to investigate cycles of influence
- Covers local, state, and federal levels
- Used as a recruitment tool for disaffected journalists to govern the platform
- Governance by the craft of journalism (not the news media industry), plus data scientists and civilians

App 2: ReNews.us (Civilian Context Agent)

- Personalized news experience oriented around the user's own profile
- Surfaces data signals relevant to race, gender, income, etc.
- Phase 1: Ingests every article from nonprofit US newsrooms, mapped to the graph via NLP; sends traffic directly to newsroom sites (restoring traffic lost to generative AI and social media)
- Phase 2: Articles readable in-app; publishers paid directly
- Navigation is designed to be omnicognitive — no search in phase one, to preserve contextual understanding of relative magnitude and relationships
- No comments, no social media; users vote on issues (favorable/unfavorable) and policies (for/against)
- Cluster analysis of voting demographics reveals belief patterns across racial, gender, and income lines — reporting that doesn't currently exist

Security & Anonymity

- Financial services-level security, but built on no stored personal data
- Civilian profiles are anonymous; journalists are validated with publications and articles
- Identified vulnerability: simultaneous hack of Stripe and the platform within a 15-minute window — considered extremely unlikely
- Seeking white hat hackers to stress-test the anonymity platform

User Profiles

- Longitudinal, forward-looking profiles curated by the user (vs. ad tech's backward-looking behavioral inference)
- Users own their own identity and data
- Interest in Solid Pods or similar personal data stores as a future direction

Revenue Model (ReCommerce)

- At ~300,000 users, civilians can opt in to a separate shopping app
- Users receive signals about relevant product categories; merchants propose offers directly

- Users get paid per offer received (e.g., $100), with a portion (~$33) required to be spent on per-article news consumption in ReNews
- Removes Meta/Google intermediary, lowering merchant costs and potentially consumer prices
- No ads, no subscriptions in the core ReNews app

Governance & Ownership

- Merging with News Futures (parallel organization), which may become the initial governing body
- No VCs, no private equity, no publishers, no advertisers, no corporations
- Civilian stakeholders hold board seats
- Crowdfunding model; nonprofit structure with options available to participants from the revenue-generating side
- Founder has bootstrapped to near-poverty; burn rate is manageable; seeking ~$3M in bridge funding over ~6 months

Political Transparency

- Includes all campaign finance data, politician votes, statements, and actions
- Goal is to reveal political bias — not embed it — insulated from direct political system influence

Age Verification (for OS & otherwise) & Mass Surveillance

Session Convener: Brent Shambaugh
Session Notes Taker(s): Brent Shambaugh & Jen Kostyrna

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Jen's Notes
Age Verification – Brent Shambaugh

- US states (CA, CO, NY), EU, other government entities pushing for age verification for protection of minors against inappropriate content, but questions – is it really about protecting kids? How will it work?
- Questions about operating systems using age signal to check age appropriate content – linux system would use date of birth, potentially provide to any apps installed on OS. Doesn't solve for shared devices. Would still just be a self (parent?) attested date as opposed to any sort of real identity proofing and creation of verifiable credential.
- Systems (MS, apple, google) already know a user's age. Companies like Google lobbying for OS level solution – require this as opposed to app/website level identity proofing.
- Is this regulatory capture? Big tech companies stand to benefit. If Apple/Google are only wallets, has ship already sailed in US?
- Missing infrastructure, need for data fiduciary.

Brent's notes:
Overview of age verification. Systems …love local. Easy solution. Set birthdate in the Computer Application. age…every service …mon age…Browser or app can refuse.

Georgia…Lots are services that are free…representing kyc…access denied….parents set birthday…content..attestation..

privacy preserving …proof of age …child…. creating…use my behavior…what keeps the parent from fudging….depending on

Attesting child is 15 years old… state parents….what is the actual authority dealing with…what is actual auth dealing with….not replacing the bank…not claiming to be accurate…

Parents decide act —-> federal

Stupid question

It becomes illegal …any laptop system

Either approach …move away from responsibility

Server user agent…need something beyond our devices …because servers are running … what types of software and infrastructure …where should your credentials go …. not universal way of messaging said infrastructure … warrant vs sopena

We acknowledge we have big tech …. Regulatory capture ….half the stakes…have a way to do this

DID-Backed X.509: Linking Decentralized Identity to PKI Certificates

Session Convener: Susumu Ishizuka
Session Notes Taker(s): Masayoshi Mitsui

Tags / links to resources / technology discussed, related to this session:

Presentation slide

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Summary
Presentation by NodeX of a technical proposal integrating DIDs (Decentralized Identifiers) with X.509 certificates. The technology aims to reconcile long-term identity management for IoT devices with existing investments in PKI infrastructure.

NodeX and DID Infrastructure

- NodeX provides DID-based identity infrastructure
- Targeted at critical infrastructure sectors such as healthcare, energy, and defense
- Robust, long-term identities are required for IoT devices that operate for 5–10+ years
- Automates machine lifecycle management, reducing the human operational burden
- Provides machine-to-machine (M2M) communication by DIDComm and policy-based authorization

Challenges and Motivation

- Existing customers have made significant investments in PKI and cannot start from scratch
- Certificate systems are widely used at:
    - Protocol layer: TLS/Mutual TLS
    - Government layer: regulations may mandate the use of specific PKI schemes
    - Operations layer: such as CA, RA, LDAP, and CRL is difficult to replace
- Issues of “just issue cert”
    - In traditional PKI, trust shifts to the CA, and the device itself loses its autonomy

    - Device identity is only valid within the certificate's validity period, breaking identity continuity
    - A change of CA (e.g., due to corporate mergers) can result in the loss of all device identities

Technical Proposal: DID + X.509 Integration

Two extensions are added to embed a DID into the X.509 certificate:

- Including a DID URL in the SAN (Subject Alternative Name) extension
- A new certificate extension that verifies Proof of Possession (PoP)
- these extensions are non-critical, so existing infrastructure continues to work as before
- PKI is not modified, nor is the root of trust completely migrated to DIDs

Changes in the Trust Model

- Traditional PKI: trust based on the CA chain
- Proposed scheme: CA chain + DID document
- Identity continuity no longer depends on the certificate's lifetime
- Persistent identity that survives certificate or CA rotation
- CA certificate sharing becomes unnecessary for cross-organizational trust

Application Scenarios

- Combination with BRSKI (Bootstrapping Remote Secure Key Infrastructure)
- DIDs can be used in place of IDevID (Initial Device Identity) at manufacturing
- Key rotation: DID can renew its key, not like the hardware-embedded birth certificate, where cert renewal is impossible
- Standard usage with TLS / mTLS. no change in TLS layer

Implementation Considerations

Issuer Side

- The RA needs DID registration and PoP verification capabilities
- Devices need to manage two keys: the DID key and the certificate key
- The CSR must include the DID and PoP

Verifier Side

- PoP verification is added using the TLS library's callback functionality
- Access to a DID resolver is required

Deployment Approaches
Optional / Fail-Open Approach

- Existing clients continue to operate as before

- Certificates with PoP can also be accepted

Active Verification Approach

- Combination of CA + PoP verifier
- Some verifiers verify only the PoP, while others verify both CA + PoP

PoP-Only Approach

- Self-signed certificate + PoP only
- Independent of the CA, placing trust solely in the DID
    - Great comment : self-signed certificate doesn't seem robust, because use of self-signed cert can verify the DID is owned by that entity, but DID itself needs to be "known" by the verifier outside of this proposal.

Challenges and Future Considerations
Air-Gapped Environments

- DIDs often assume internet connectivity
- May be difficult to use in factory or industrial IoT environments

Handling of Revocation

- Certificate revocation and DID key revocation exist separately
- There is ambiguity and complexity in the semantics between the two different keys

Relationship to Other Approaches

- Integration at a different layer than DIDLink
    - This proposal: certificate layer
    - DIDLink: TLS protocol layer
- The opposite approach to the DID X.509 method (certificate-based DIDs)

Q&A Session Highlights

- about the similarity to the idea of putting cert finger print in DID doc
    - almost similar in a sense that DID-cert binding, one difference would be that the certificate renewal cycle and the DID renewal cycle can be fully decoupled
- Questions about the applicability of DIDs at the component level
    - Mention of the relationship to Software BOM (Bill of Materials)
    - in this proposal the DIDs are assumed to be assigned to end products such as drones and vehicles

The Fiduciary Commons 1 - from principles to law

Session Convener: Mike Leahy
Session Notes Taker(s): Mike Leahy

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The Fiduciary Commons | Session 1 of 3
From Principles to Law
What Utah Enacted, What Remains, and a Roadmap for Other States

What This Session Argued
Government holds citizen data not as a property owner but as a constitutional fiduciary, bound by duties of loyalty, care, and transparency. This is a pre-existing obligation embedded in the constitutional structure by the framers. The Fiduciary Commons framework makes it operational through three model statutes. Utah enacted the first two unanimously in 2025 and 2026, establishing the first statutory codification of government data fiduciary duties in American state law. The gaps in Utah's framework, and the roadmap for other states to complete it, define the next phase of the work.
Four Ideas Worth Keeping
Facts cannot be owned.
The Supreme Court confirmed in Feist Publications (1991) that facts are discovered, not created, and belong to no one. A database may be proprietary; the facts inside it are not. 'Data ownership' is a category error that confuses control over a container with title to the content. The Fiduciary Commons framework begins by correcting that error.
Every commons requires a trustee.
Air and navigable water are held in public trust because no one owns them and everyone depends on them. Facts and data have the same structure: non-rivalrous, non-depletable, flowing through every institution. The public trust doctrine, established in Illinois Central Railroad v. Illinois (1892), provides the governance model. Government holds citizen data as a trustee with defined duties, not as a property owner with defined rights.
Utah proved the argument works.
SB 260 (2025) established digital identity rights: selective disclosure, prohibition on server retrieval, individual control over credentials. SB 275 (2026) codified data fiduciary duties: purpose limitation, data minimization, a digital bill of rights. Both passed unanimously in a polarized environment. Their independent convergence with the Fiduciary Commons framework validates the underlying logic. The critical remaining gap: no private right of action. AG-only enforcement is a suggestion, not a mechanism.
The digital general warrant is the constitutional threat.
The Fourth Amendment was written to prohibit general warrants: authority to search anyone, anywhere, for anything. Government databases that aggregate citizen data without individualized suspicion are the digital equivalent. The Fiduciary Commons framework is, in part, the

constitutional corrective: if government holds citizen data as a trustee, the surveillance that digital general warrants enable is a breach of fiduciary duty, not merely a policy problem.
The question is not whether this framework becomes law. The direction of travel is established. The question is whether your systems are ready when it does.

The Closing Ask

Engage with the legislative process in your primary state markets. You have standing with state CIOs and legislative staff that policy advocates do not. The three-stage roadmap is clear: enact the Utah baseline, add the private right of action and enforcement architecture, enact GAAFA. The template exists. The political path has been proven. Your technical credibility is the missing ingredient.

What This Means for Your Work

For identity vendors

SB 260's phone-home prohibition is a direct architectural requirement: mobile credentials that report back to a central server each time they are presented are prohibited by law in Utah. That prohibition will spread. Systems designed for selective disclosure and non-retrieval architectures from the outset are cheaper to build and legally more durable than retrofitted ones. Audit trails and revocability are not optional features; they will be baseline requirements.

For protocol designers

The interoperability requirements in SB 260 create a direct connection between the statutory framework and the standards you are developing. Standards that conflict with statutory fiduciary requirements will require revision when those requirements become law in adopting states. ISO 18013-5's server retrieval capability is a specific example: that capability is what SB 260 prohibits. Early alignment between technical standards and statutory requirements avoids costly divergence later.

For standards bodies and foundations

The rulemaking processes that will define technical compliance standards under SB 260 and SB 275 are open. Your participation shapes what technically achievable compliance looks like. Abstract academic participation is less influential than input from organizations that have deployed these systems at scale. The window for that input is now, while rulemaking is being established, not after compliance standards are fixed.

Key Terms

Fiduciary Duty

A legal obligation requiring a person or institution that holds power over something on behalf of another to act in the beneficiary's interest, with duties of loyalty, care, and transparency.

VIDA (Verified Identity Data Act)

Model statute establishing digital identity rights: selective disclosure, prohibition on server retrieval, individual control over credentials, and interoperability requirements.

PDTA (Personal Data Trusteeship Act)

Model statute codifying government data fiduciary duties: purpose limitation, data minimization, audit trail requirements, and individual rights of access and correction.

Selective Disclosure

A credential architecture that allows a citizen to prove a specific fact (age over 21, professional license) without revealing the underlying data or any other information in the credential.

Phone-Home Prohibition
The requirement that mobile identity credentials not transmit data back to a central server upon presentation. Prevents construction of a surveillance record from credential use events.
Private Right of Action
The legal standing of an individual citizen to bring a lawsuit for breach of statutory rights, as distinct from relying solely on government enforcement. The critical missing element in Utah's SB 275.
fiduciarycommons.com | michael@fiduciarycommons.com | For the full argument, see the Fiduciary Commons Series Summary.

Cross-Border Finance as a cast study for Cross-Border ID Data Management in Africa

Session Convener: Lenah Chacha
Session Notes Taker(s): Lenah Chahca

Tags / links to resources / technology discussed, related to this session:

Wallets, verifiable credentials, Identifiers, cross-border identity

Topics Discussed

Standards, wallets, verifiable credentials, economics of verifiable credentials, liability insurance

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps

The sessions explored cross-border financial systems as a mature, real-world example of systems operating across jurisdictions, using them to analyze data management practices and identify lessons for designing ecosystems that minimize risk.

Context: Presentation of an ID is a requirement for KYC and CDD in any financial system and important for financial stability. In some cases, these processes take weeks and in other cases it has to be done on the spot to allow real-time transactions. The ACFTA is a framework that desires continent-wide free trade agreement among African countries by allowing free movement of persons, capital and investments. This will certainly allow various use cases for cross border and these will also require KYC and CDD.

The session explored a difference in credential format as a risk, a key example cited is the difference in birth certificate formats issued in the USA at the county levels, and that this requires a risk calculation by the entity accepting these credentials. At this point a choice is made to either accept the credential or contact the issuer to verify the credential. The key question highlighted was for issuers to think about whether the evidence presented by the holder is sufficient without reaching out to the state.

The group noted that the NIST 800-63 is a good source for standards and guidelines for identity proofing, authentication and sharing trust across systems.

A concern that was noted is surveillance, while credentials are used cross border and a call is made home, how are the surveillance policies balanced and especially in cases where the countries have contradicting policies on the topic. More especially, how is the determination made by the verifying country of whether this is indeed true in the issuing country?

Trust: It was noted that in a financial system network, trust is always between the entities of type organization that trust each other to fulfill the transactions in the network i.e the ordering, intermediate and beneficiary financial institution(s) (FIs).

Trust in Verifiable credentials ecosystems: The group also discussed the mobile phone as the main substrate upon which verifiable credentials will be deployed. In this discussion, ownership and continuous authentication mechanisms were brought up and especially the fact that it is difficult to continuously determine that the phone is in use by the rightful owner. Various approaches were discussed including biometric proofing (something you are) , and that this cannot be guaranteed in mobile technology today. The group also discussed questions like unlike in the issuance of cards where trust relationship is the verifier trusts the issuer and the holder, in verifiable credentials another dimension is introduced and that is trust in the wallet itself. And therefore the debate is should the holder, verifier or issuer trust the wallet? What combination of this bi-directional trust would be stronger than the other? In what scenarios would one trust relationship be preferred over the other?

The team also discussed risk in financial systems, noting that transactions are always assessed against the amount of money involved, which in turn determines the level of rigor applied to KYC and CDD.

SESSION #4

Cryptoperide Catastrophy

Session Convener: Samuel Smith
Session Notes Taker(s): Henk van Cann

Tags / links to resources / technology discussed, related to this session:

- slides
- https://github.com/SmithSamuelM/Papers/blob/master/presentations/CryptoperiodCatastrophe.pdf
- NIST SP 800-57 Part 1 Rev 5.

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Slides: https://github.com/SmithSamuelM/Papers/blob/master/presentations/CryptoperiodCatastrophe.pdf

30 people in the room. Sam asks for a show of hands of people in the credentials business NOT being from a KERI background. 2 hands go up. He concludes that the people his talk is aimed at, hopefully, will show up.

Crypto periods are like signing with disappearing ink. It's the life span of a key.

Formally in NIST SP 800-57 Part 1 Rev 5.

Example: birth certificates; they need to last for 100 years.

Sam about privacy in the light of key management: "It's a harsh trade. A harsh trade is a hard trade that makes you feel really bad 😳"

The opposite of bare signatures is "anchored signatures."

ZNP are good for ephemeral use. Don't use them outside the crypto period.

Don't confuse susceptibility with vulnerability

Sam introduces the analogy of a bear chasing two men, one big guy, one slim, sporty guy. "Who is susceptible?"
AI is going to present a 'bear' for everyone: the cost of attacks has gone to zero.

Rest is in Sam's slides.

Fido and Webauth 101 an IIW Session

Session Convener: John Bradley
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Regenerative Accelerationism / Kaliya & Friends

Session Convener: Kaliya & Friends
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Death and the Digital Estate - What's next?

Session Convener: Dean Saxe, Eve Maler, Mike Kiser
Session Notes Taker(s): Mike Kiser

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The Death and the Digital Estate (DADE) community group has published a whitepaper and practical guide.

Members stated that they were fully supportive of the demand / need; the use case remains well known.

"What do you do when your customer dies?" is not a well developed approach / use case.

Technology, data, and policy approaches are all necessary for this.

individual policy approaches are coming (such as myterms (https://myterms.info / IEEE standard) that might specify end of life terms.

asynchronously and retroactively . . .

privacy law doesn't apply once you're gone, but perhaps private / contract law that might apply.

There are some legal pathways, but the technical avenues are not developed the same way in the digital world.

Succession planning—would it be best to have a trust and then a digital avatar of the self to guide the process… (aka the "Ready Player One" scenario)

How do you prove that the person was gone? that is the key element / signal that would be required…

The increase of 2fa and a lack of a strong signal / notification of death combine (historically) to promote impersonation.

Lack of incentives make further progress difficult

Look at lobbying the elected officials who can drive policy.

Data + emotion are different motivators . . .

NB: UK has a death notification system . . . (frozen accounts automatically, etc.)

EUDI Wallet - What is Germany Cooking

Session Convener: Mirko Mollik
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

EUDI Wallet, EU, eIDAS

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Presentation slides: https://docs.google.com/presentation/d/1bqWxgg5X6gLDXAvxGhSQ-7fiIC4uLhWa/edit?usp=sharing&ouid=116865618062361812085&rtpof=true&sd=true

Insights into the approach Germany is taking to roll out their EUDI wallet eco system

- explaining the different roles
- pointing out how to onboard to their sandbox: https://bmi.usercontent.opencode.de/eudi-wallet/eidas2/ecosystem_knowledge_centre/

Mentioning multiple open questions

- how can non eu companies participate in the sandbox or in general with the EUDI Wallet
- how huge will the eco system be in around 5 years

Digital Fiduciaries Update

Session Convener: Joe Andrieu
Session Notes Taker(s):

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Digital Fiduciary Concept
A new professional class defined by economic entanglement and oath-binding, creating accountable human intermediaries in digital identity workflows.

Fair Witness Ceremony
In-person verification combined with cryptographic proof-of-use — two things AI cannot fake — forming the foundation of trusted identity attestation.

Digital Signet
A Fair Witness protocol binding legal personhood to cryptographic identifiers, inspired by the USCIS I-9 form.

Fair Witness Network
Patent filed January 2025 for a privacy-preserving credential registry. The patent will be licensed exclusively to certified digital fiduciaries.

Initiative Progress
Launched Fall 2024. 13 members total; 11 have taken the oath. No external funding yet — bootstrapping from Joe's company.

Live Demo
Proof-of-concept walkthrough available at:

- wallet.digitalsignet.org
- demo.digitalsignet.org

Technical Roadmap (15 items)

1. Selective disclosure (BBS)
2. Revocation
3. Unlinkability
4. Multiple proofs
5. Continuity strategy
6. Transport
7. Archive storage
8. Wallet support
9. DID methods
10. Credential formats
11. Bitcoin Beacons
12. Browser extension
13. Open source
14. Fair Witness Network on Bitcoin
15. Decentralized governance

Certification Levels

- General
- Technical Specialist
- Regulatory Specialist
- Fair Witness (paralegal/notary level)

Open Questions

- Name changes
- Digital evidence ingestion
- Slide deck availability

https://docs.google.com/presentation/d/1wONdr8QQjgpHjVOJrXr0Thc57OgxHHSj/edit?usp=sharing&ouid=109819894046991510916&rtpof=true&sd=true

OAuth 2 for Sub Agents featuring French Toasts.

Session Convener: Abhishek Hingnikar
Session Notes Taker(s): Steven Tamm

Tags / links to resources / technology discussed, related to this session:

Implementation for French Toast JWT proposal: https://github.com/ciamshrek/french-toast-jwt. Clawdery's updated implementation of "Waffle" https://github.com/clawdreyhepburn/waffles-draft

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Core Problem Presentation

In agentic workflows, an **Agent** often needs to delegate tasks to **Subagents.** All clients (agents) have tokens that are for the same Resource Server (RS) [Fig.1]. Today this requires using Token Exchange

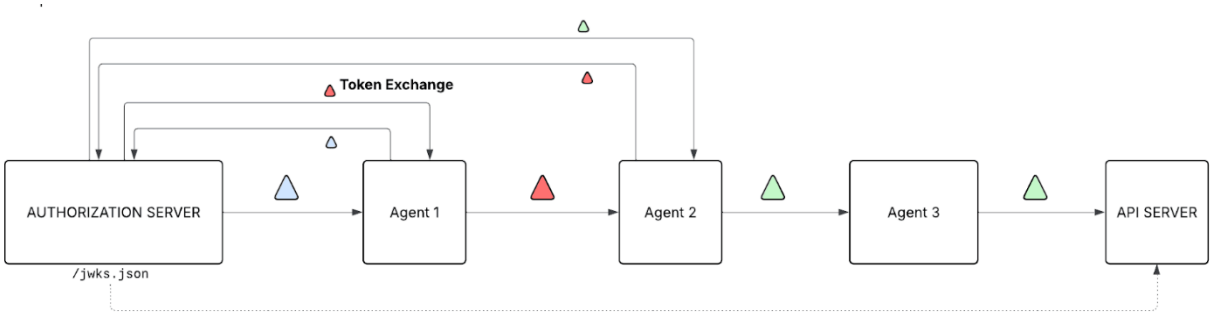

Fig 1. Agent to Agent Delegation via Token Exchange

Currently, delegating with restricted permissions (Scope Attenuation) presents a scalability bottleneck:

- **High Latency/Traffic:** Traditional OAuth requires going back to the Authorization Server (AS) to exchange a broad token for a narrower one.
- **No Standardized Pattern:** Token Exchange requires complex workarounds to ensure DPoP pinning for tokens.
- **Lack of Native Support:** Patterns like **Biscuits** or **Macaroons** provide attenuation but are not codified in standard OAuth/JWT-based infrastructures, requiring massive migrations.

The ideal outcome is an agent to self-sign a "downscoped" token for a subagent without a round-trip to the AS.

Overview of our proposed solution

We proposed a new token format, “french toast”, with additional constraints to be added to an JWT Access Token, in the form of cryptographic proofs. The orchestrator / parent agent can then provide this token to the sub-agent.

During the later discussion we opted to name these tokens “Waffles”, the document will refer to them as Waffles hereby.

Fig 2: Requests with Waffle Tokens.

The original proposal is available at https://github.com/ciamshrek/french-toast-jwt. Key considerations:

- The token used to make the request to the Resource Server includes the original token, and any additional tokens that may have been sub-issued. The entire chain is treated as a token.
- By leveraging DPoP [RFC 9449] the tokens are always bound by the key material for the sub-agent.
- By leveraging Client ID Metadata Document [https://datatracker.ietf.org/doc/draft-ietf-oauth-client-id-metadata-document/], the Resource Server can obtain the key material for all clients in the chain

Discussion and Notes

- Why not use macaroons or other existing token formats:
  - the authorization server requires migration of everything to Biscuits, Macarons etc.
  - You can prove that cedar is a strict subset with “lean”, so you get an “unbreakable sandbox for AI”. That’s also why the whole chain is there as well; you can’t just use hashes across systems vs Clawdrey.
  - In this proposal, it’s easier to change Resource servers than fixing all the identity infrastructure.
- Question: “Are you expanding this to context graph and intent?”
  - This is a way to get assurance from the orchestrator side that you can only attenuate tokens. It’s like a X.509 chain in a JWT chain;
  - Thus it can operate with graph / intent chains from the previous token format, perhaps using Cedar / Rego or policy dialect the RS accepts
- Similar proposals have been brought before for WIMSE etc.
- OAuth principle: keep the client stupid.
  - Concern: Which is why they rejected this before. This requires the scopes to be understandable in the API and knowing that it’s more restrictive. Client as issuer is verboten.
  - Discussion:
    - Agents tend to accumulate generic positions, in this scenario the orchestrator agent has to be smart, in order to delegate clients.
    - The remaining tooling does not need to fully update to this
  - Potentially use Cedar Policies
  - This problem is mostly for developer workcases, but it might be good to violate the “clients are stupid” principle. Justin really thinks this “might be a good thing” with the orchestrator agent being smarter is “ok.” SaraC agreed.

    - This could be implemented directly at API Gateway level to ensure simpler adoption for RS
- JOSE might not be the best envelope structure for this.
    - Proposed a different structure for it in Wimse, buckets for each token object and you could relate them in a merkel tree structure. Vs a chain. You can compact it better than JOSE.
    - Instead of JOSE as constraints, alternative formats might exist
        - Goal is to leverage the existing JOSE tooling and make this easier to adapt
- “How do you ensure fail close?”
    - Dpop will fail.
    - Token format is also opt-in for the RS not AS.
    - What should the agent do when it gets a 401?
    -

JLINC Overview and Update

Session Convener: Darius Dunlap
Session Notes Taker(s): Darius Dunlap

Tags / links to resources / technology discussed, related to this session:

JLINC Website: https://jlinc.com
JLINC Protocol Documentation: https://protocol.jlinc.io/

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Presented an overview of JLINC, including examples for Agentic Workflows and Enterprise Applications and Zero-Knowledge Audit.

Slides and Speaker Notes:

In this session, we present an introduction to JLINC and use cases for Data Governance and Audit.

What is JLINC?

JLINC provides cryptographically signed, encoded and human-readable data sharing agreements, with auditable data provenance

- **Human and machine-readable (and therefore legally admissible) agreements**
- **Signatures and hashed audit records at every step**
- **Distributed Identities (DIDs) that allow each entity to retain unique keys**
- **Signed, Hashed transactions are logged in an Audit Server**
- **Audit Includes hash of original data and agreement, but not the data.**

There is a lot going on in that first sentence.
- All records are cryptographically signed
- With Human-readable, encoded data sharing agreements
- called “Verified Contractual Agreements” (VCA) in our proposed standard
- The signed agreements and shared data are hashed
- And the resulting record is shared to an Audit server
- “Zero-Knowledge Audit” (ZKA) in our proposed standard

Records can be hosted by 3rd parties utilizing a unique hash in the URL
Records include a signed hash of the data, but not the data itself
Records are produced at each step of the workflow
JLINC is very lightweight and very fast.
JLINC builds in trust and provenance for business processes and information sharing workflows…
All without the overhead of Web 3.0 blockchains

What follows is a generalized example of the Data Governance workflow that JLINC enables.

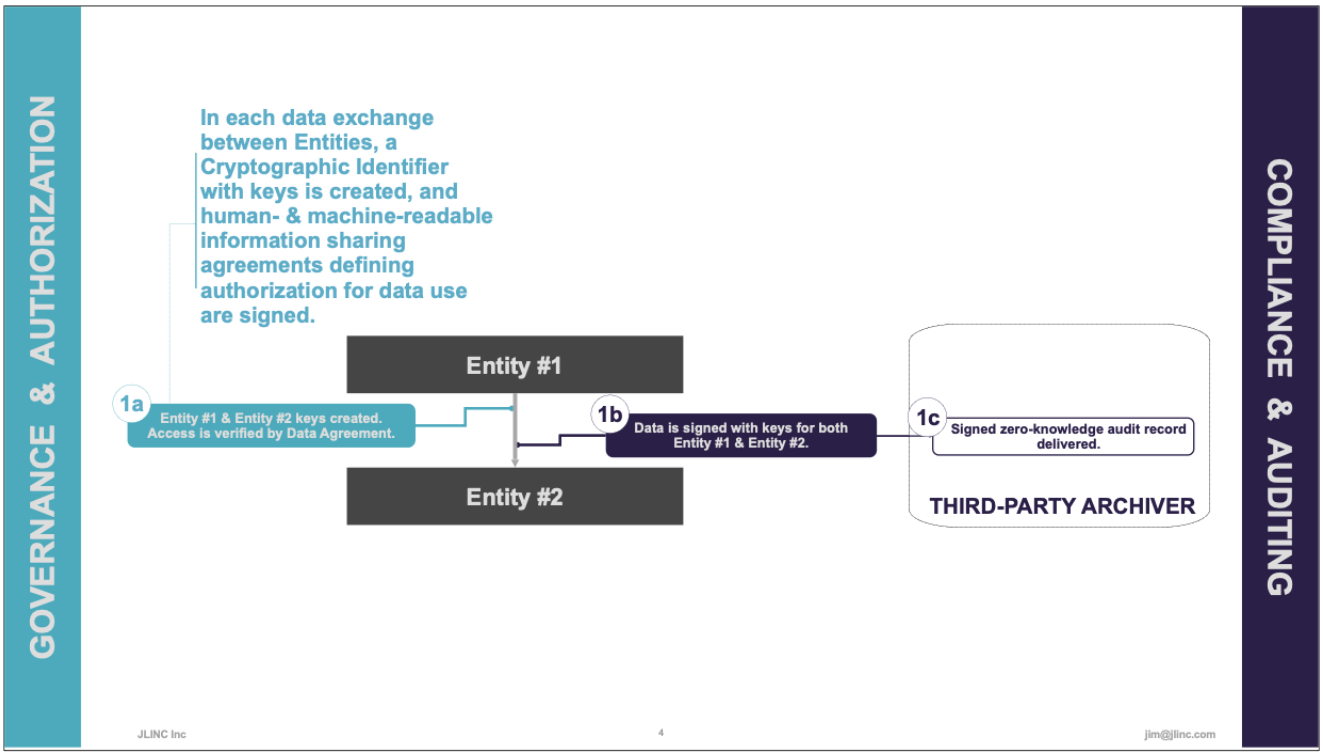

A data exchange is done in three steps:

1. Cryptographic identifiers and keysets are created and a Verified Contractual Agreement (VCA) is signed.

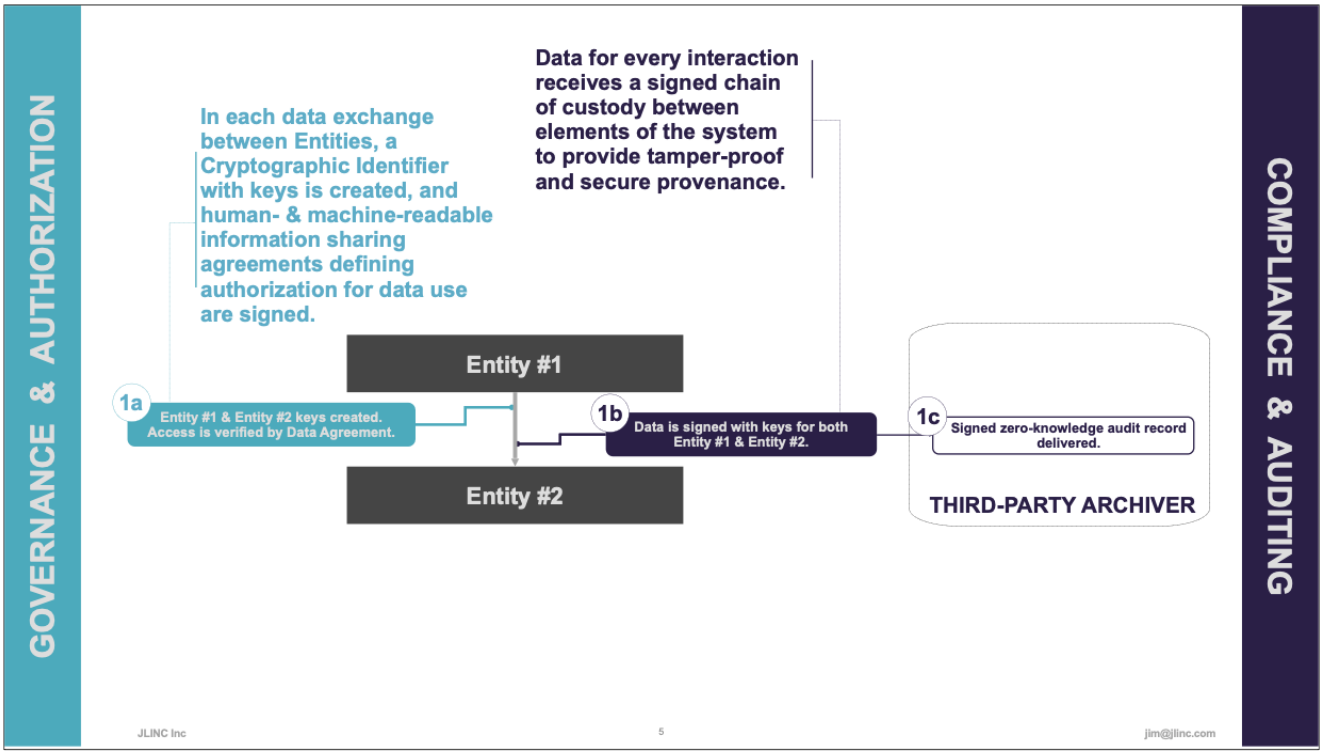

A data exchange is done in three steps:
1. Cryptographic identifiers and keysets are created and a Verified Contractual Agreement (VCA) is signed.
2. Data is signed by both entities, providing tamper-proof and secure data provenance and data integrity

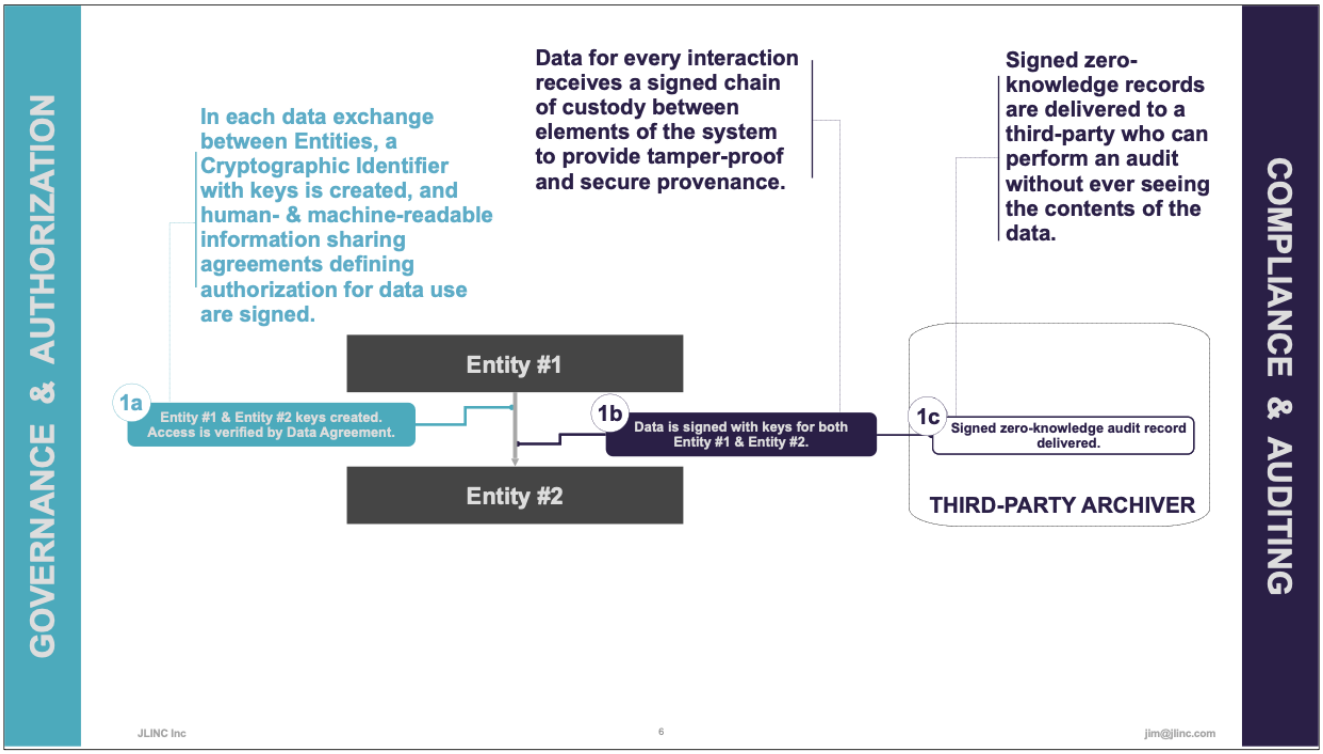

A data exchange is done in three steps:
1. Cryptographic identifiers and keysets are created and a Verified Contractual Agreement (VCA) is signed.
2. Data is signed by both entities, providing tamper-proof and secure data provenance and data integrity
3. Signed zero-knowledge records are delivered to a third-party audit server

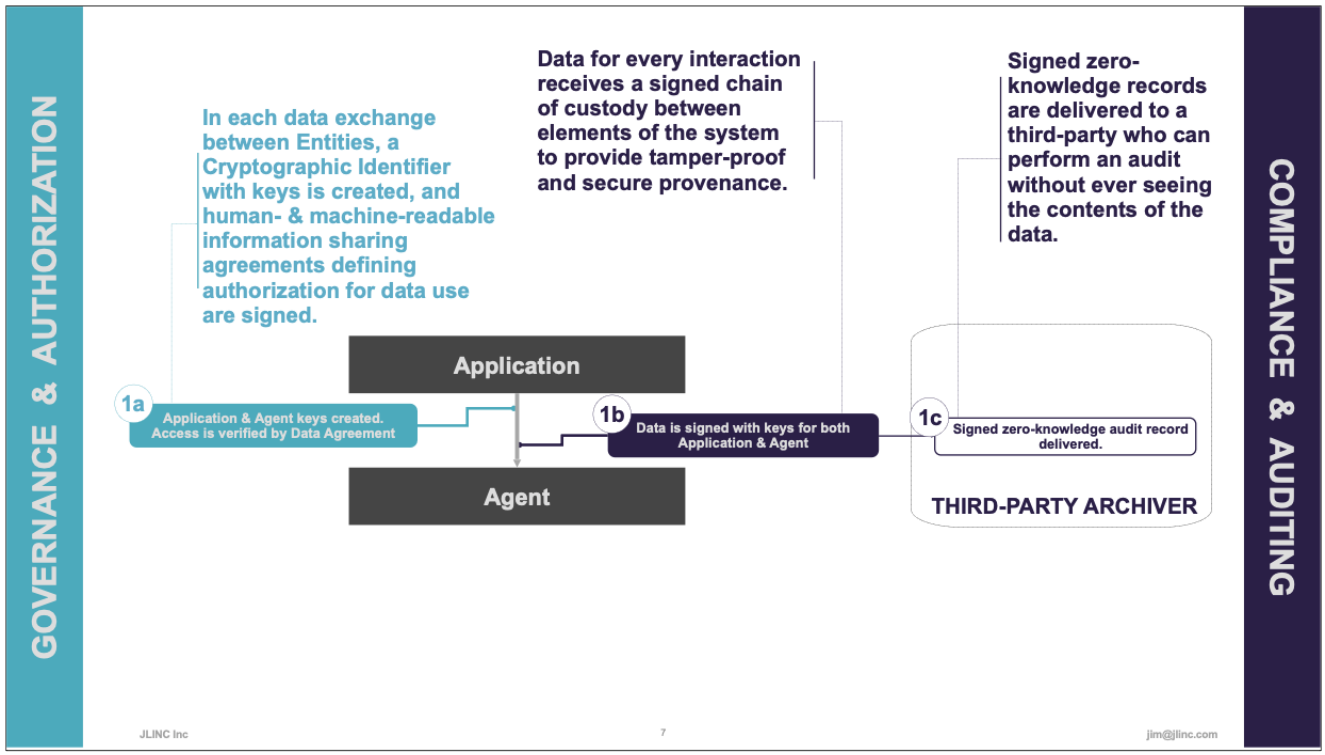

This same workflow is applicable to many systems, such as this AI agent processing workflow, showing an application that sends data to an Agent for processing.

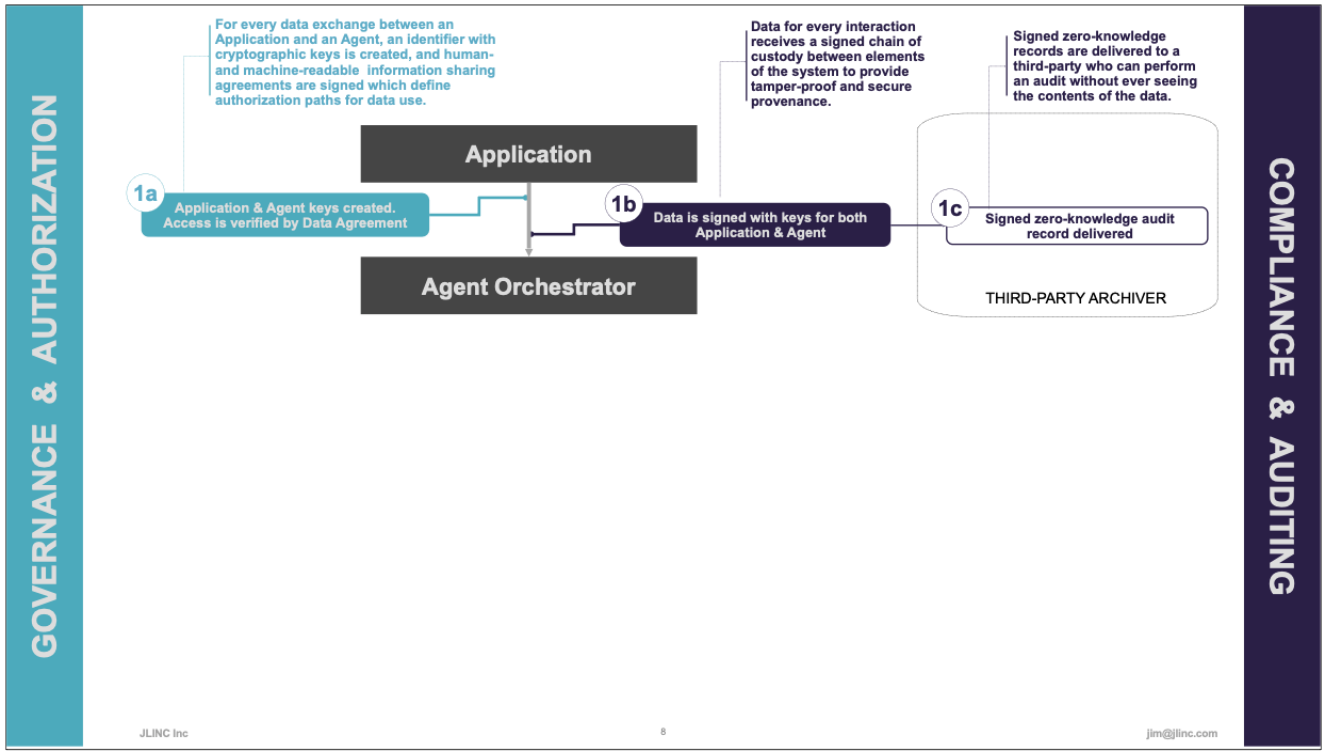

In many systems an Agent Orchestrator may control a more complex process.
This workflow also works for AI Factory processing implementations

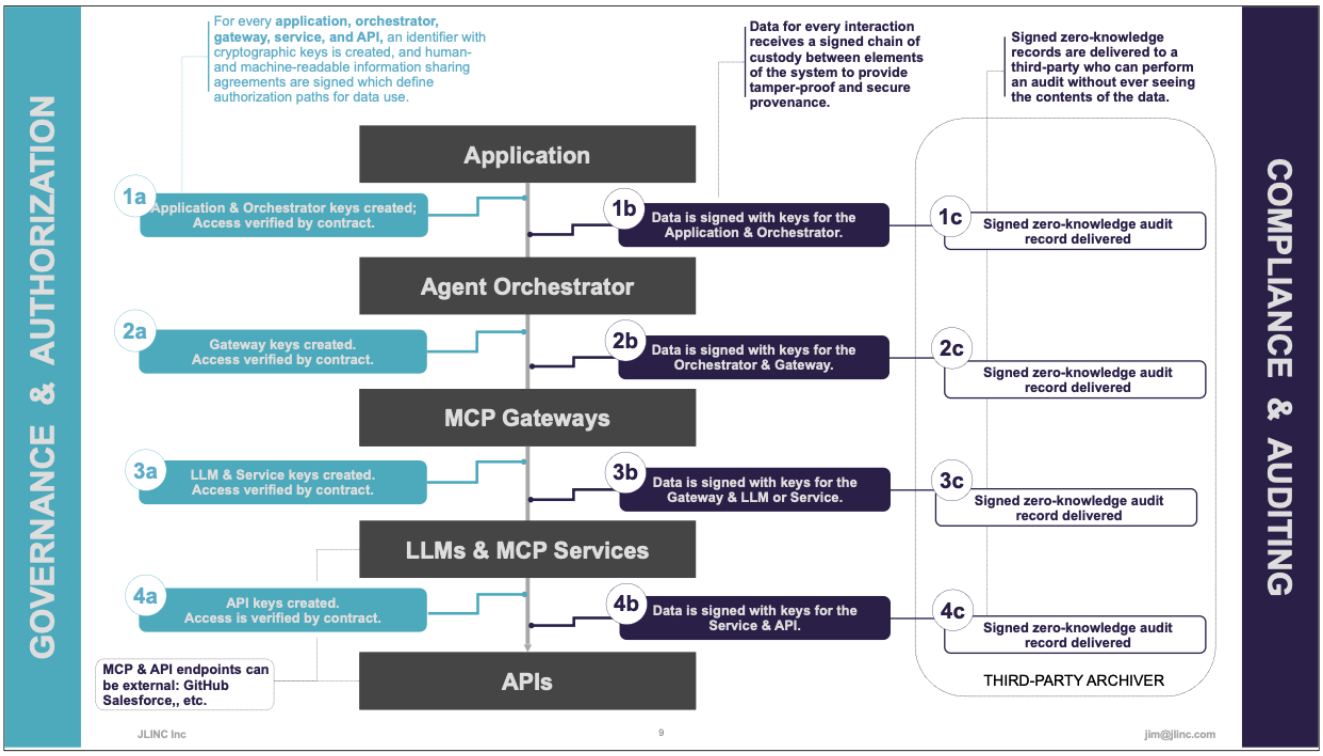

This is an example of such a complex Agentic System.
Each call to the next system, and the response, would be recorded in the JLINC system.
This includes calls to MCP and API endpoints on external systems (GitHub, Salesforce, et. al.)
Of course, this diagram is simplistic and many workflows are complex with multiple paths.

Here is a more detailed and specific example…

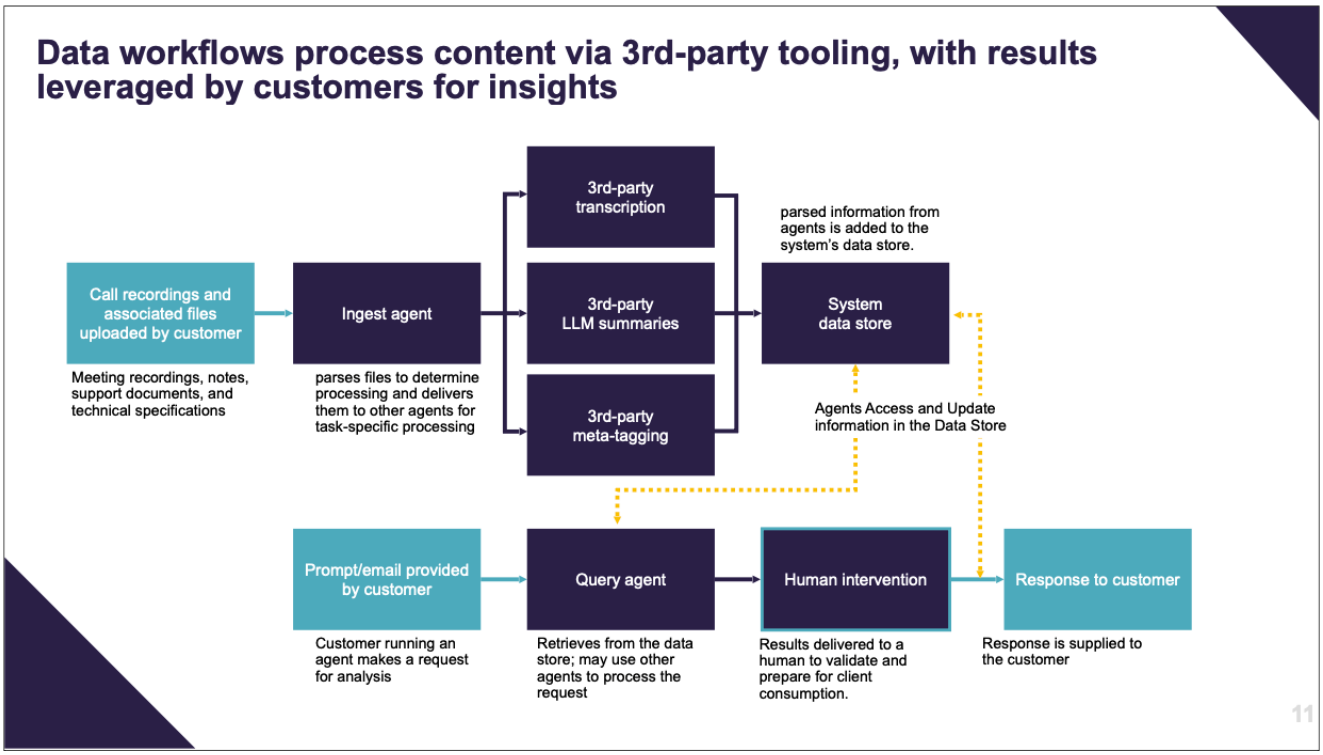

Here is a more specific example of a typical workflow.
(Take a moment to step through the workflow in the above diagram)

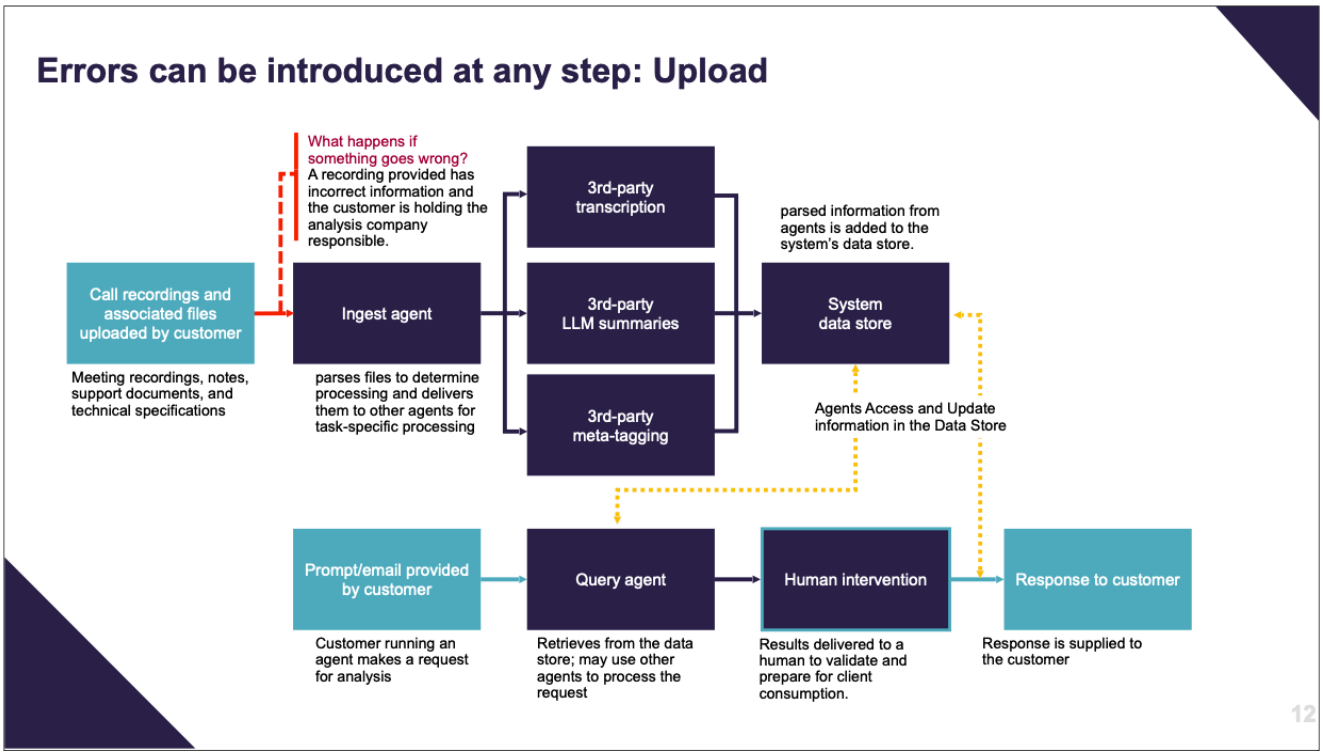

In all these processing steps, errors can be introduced.
- **A recording could contain information that's incorrect.**

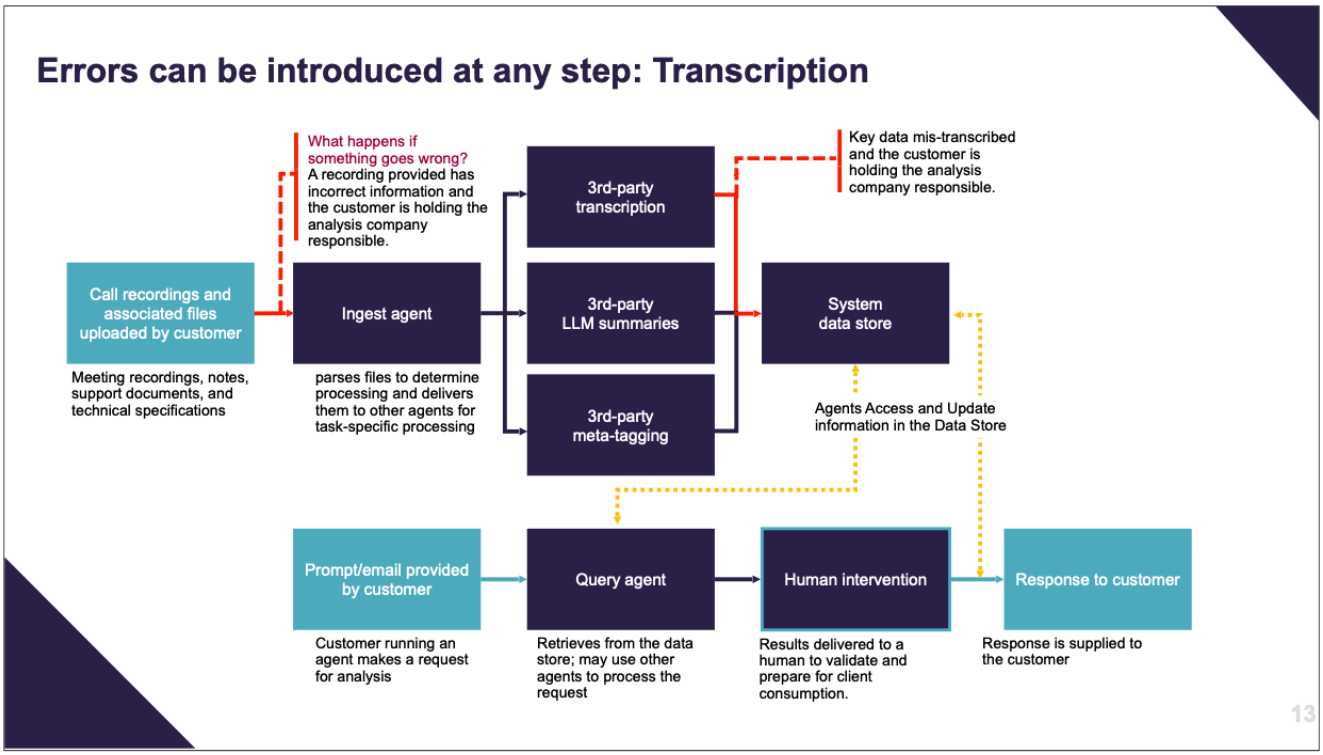

In all these processing steps, errors can be introduced.

- A recording could contain information that's incorrect.
- **Key data could be mis-transcribed by a transcription service or AI Agent**

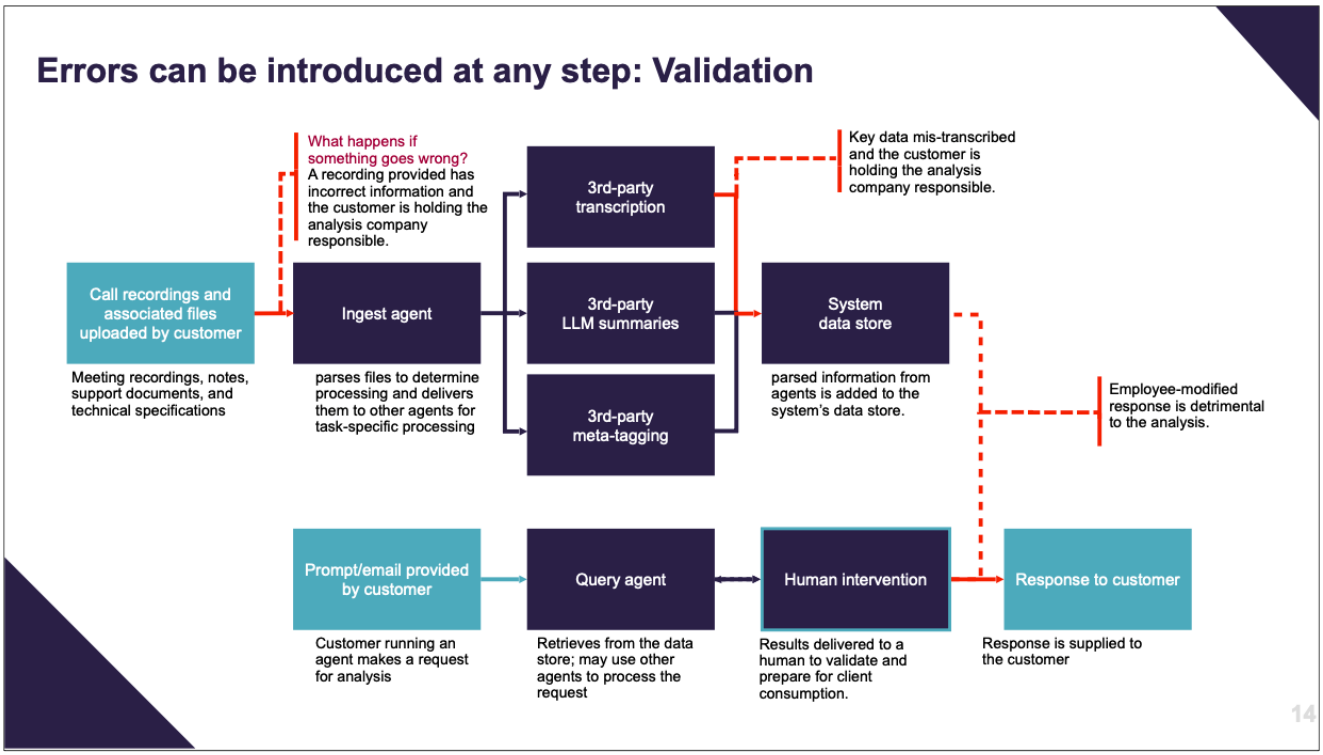

In all these processing steps, errors can be introduced.

- A recording could contain information that's incorrect.
- Key data could be mis-transcribed by a transcription service or AI Agent

- **An employee reviewing an analysis response could make an incorrect or misleading modification**

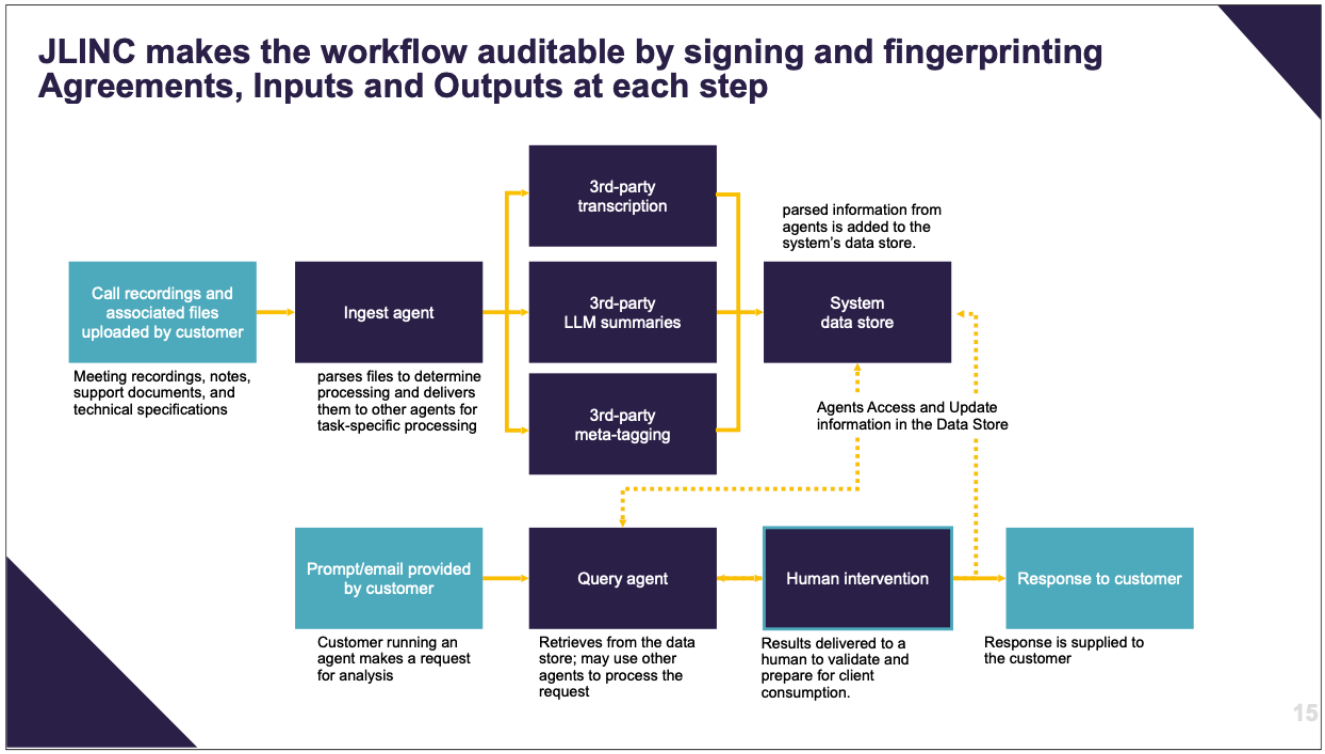

At each yellow path, **JLINC acts as a tracer on the workflow**, signing inputs and outputs as they are shared between agents and humans

Each trace provides cryptographically signed inputs and outputs to **ensure provenance and authenticity, and enabling Zero-Knowledge Audit.**

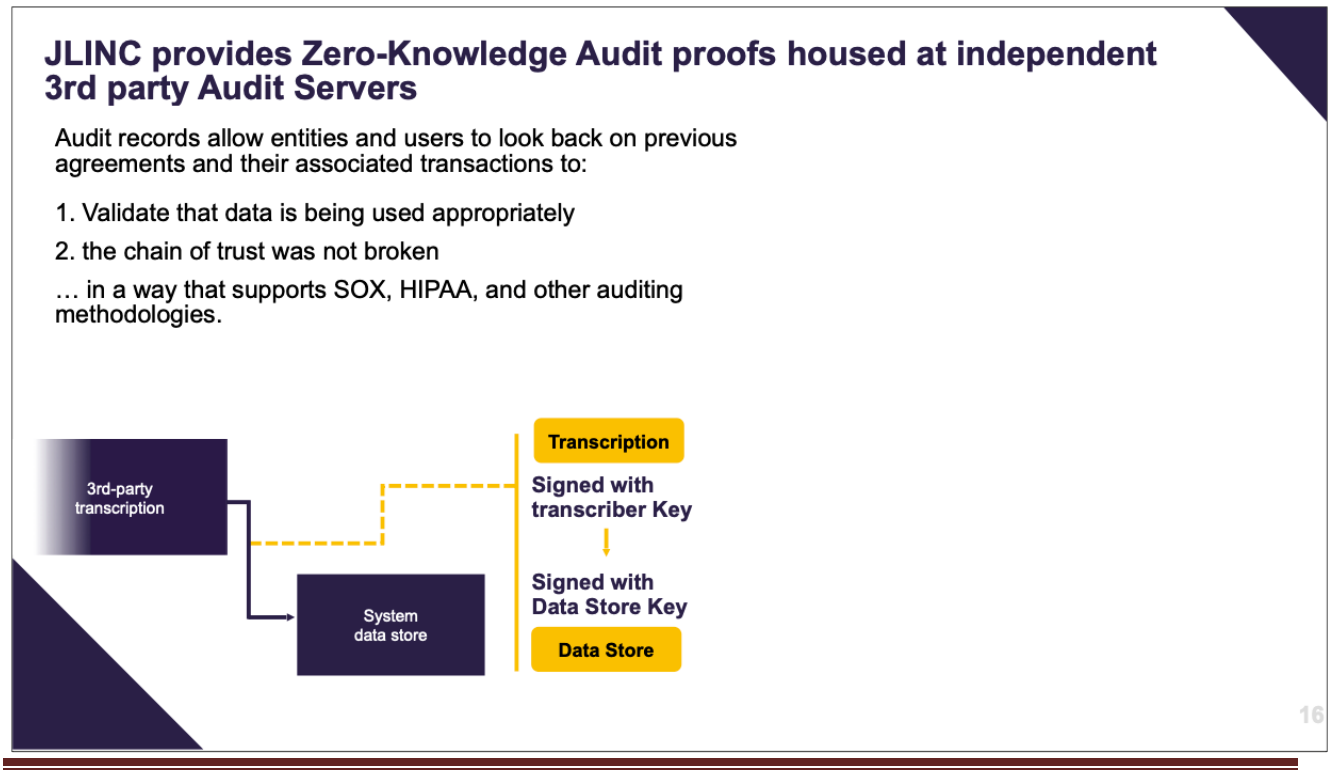

At each step, data and agreement are signed and hashed by both sides of the data sharing transaction. Here the transcription is sent to the data store and also signed.

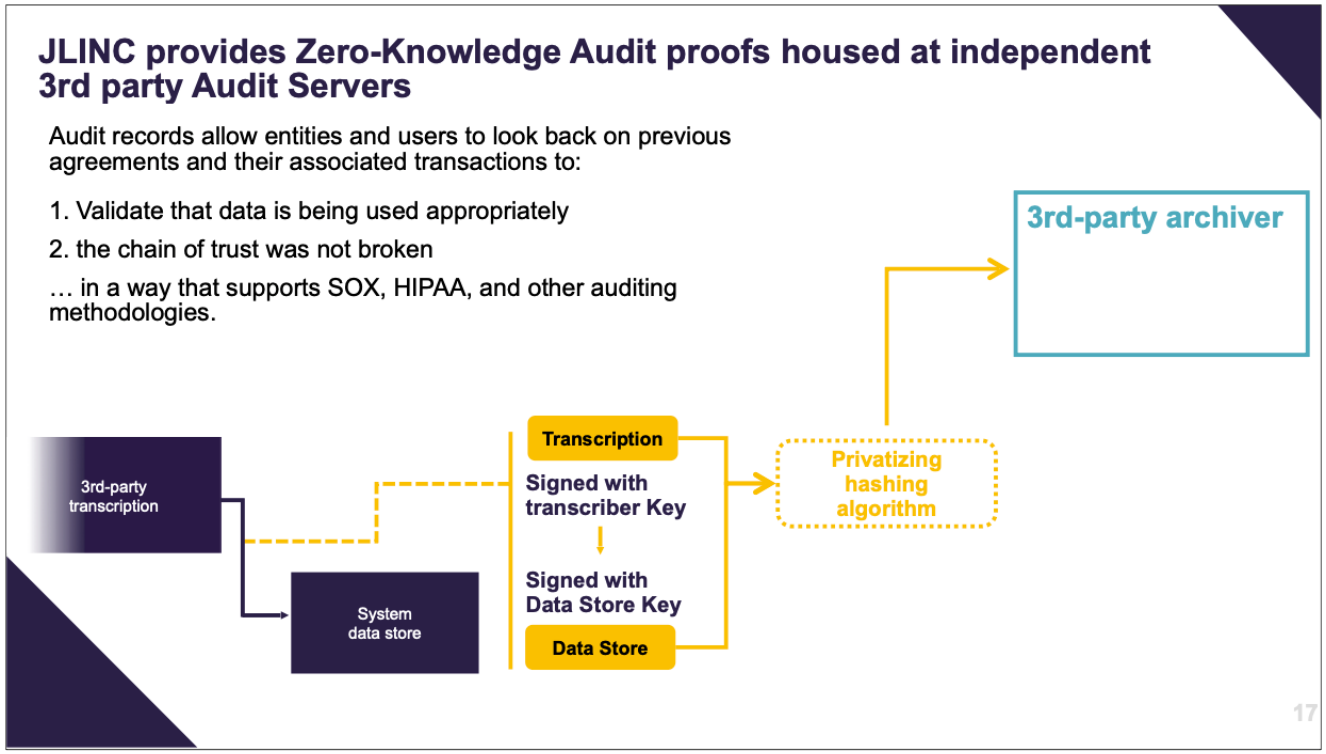

The data being shared and the signed sharing agreement are hashed and sent on to the Archive.
The underlying data is not sent, only the hashed data and agreement.

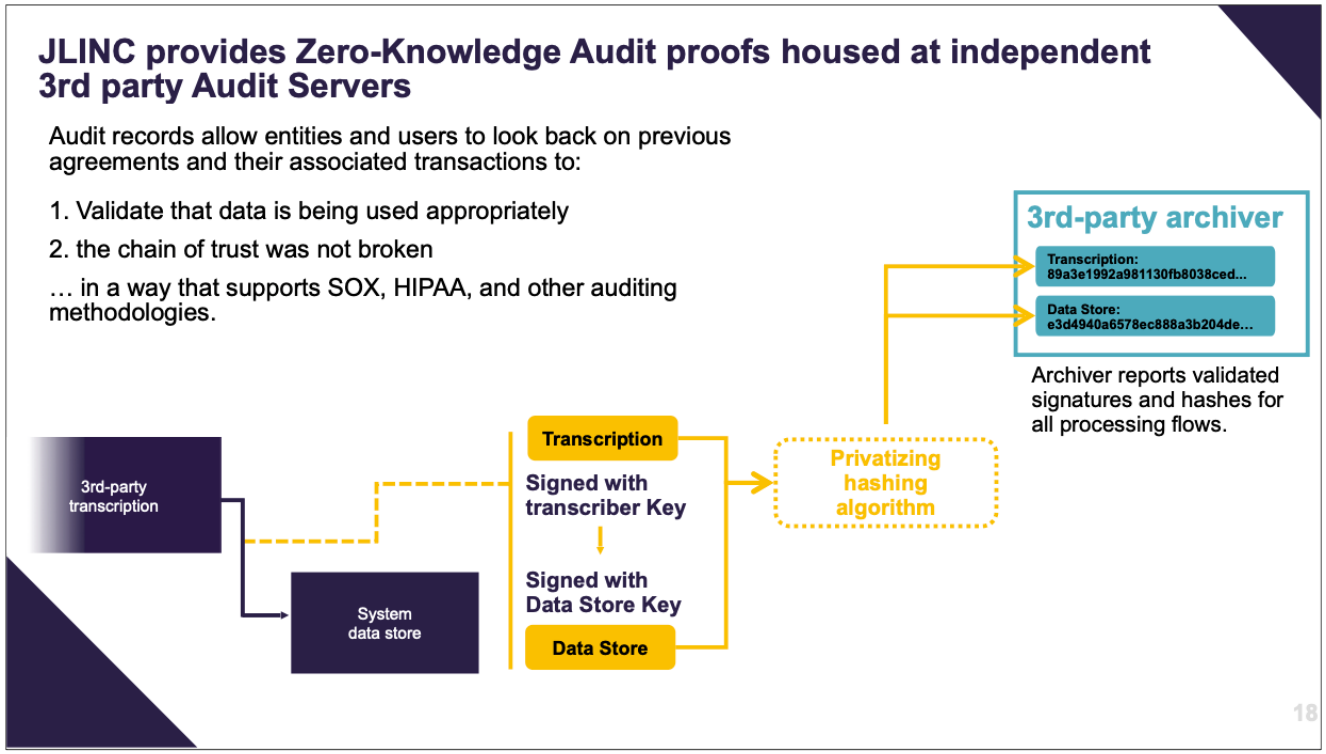

Validation shows:
- IDs match data and archive

- Audit hash matches signed data
- Signatures on audit records match
- Signatures on data match
- Data transferred under the agreement
- Signatures on the agreement match

Auditing is an important part of the JLINC protocol. Audit records are what allow entities and users to look back on previous agreements and their associated transactions to both validate that data is being used appropriately, and that the chain of trust was not broken in a way that supports SOX, HIPAA, and other auditing methodologies.

JLINC provides zero-knowledge proofs that can be housed at independent 3rd parties, satisfying even the most extensive audits

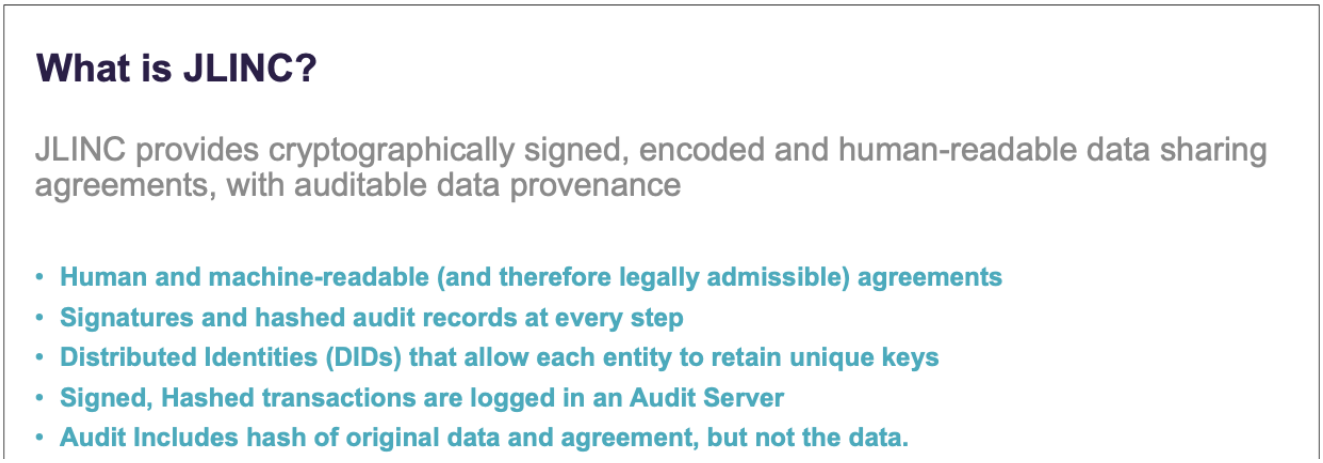

So, again… What is JLINC?
Records can be hosted by 3rd parties utilizing a unique hash in the URL
Records include a signed hash of the data, but not the data itself
Records are produced at each step of the workflow
JLINC is very lightweight and very fast.
JLINC builds in trust and provenance for business processes and information sharing workflows…
All without the overhead of blockchains

SSI 10th! Revisiting the principles for the next decade!

Session Convener: Christopher Allen
Session Notes Taker(s): Susumu Ishizuka (with help of NotionAI)

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Overview
Christopher Allen presented the 10th anniversary of the Self-Sovereign Identity (SSI) principles originally published in 2016, along with a proposed 2026 revision. The presentation covered the historical origins of SSI, the rationale for the revision, and solicited community feedback for iterating on the updated principles before presenting them at the Global Digital Collaboration event in Geneva in September.

Historical Context and Origins of Self-Sovereign Identity

- The term "self-sovereign identity" was coined to address fundamental problems with digital identity and create a user-centric, decentralized system
- Multiple philosophical influences shaped the concept:
    - Living systems theory from the 1940s-50s regarding membranes and subsidiarity
    - Historical evolution of sovereignty from kings to corporations
    - Feminist discourse around women claiming sovereignty as equals (influenced by actress Selma Hyatt)
- The term was chosen because it would be "hard to steal" unlike previous terms like "groupware" that were co-opted
- United Nations ambassadors and officials immediately understood and adopted the terminology at ID2020 conference
- The original 10 principles were intended as a starting point for community iteration, though this collaborative refinement didn't happen as planned

The 2026 Revision: Structure and Changes

- Expands from 10 to 16 principles by adding 6 new principles addressing gaps in the original framework
- Organized into four categories:
    - Existence: Existence, inalienability, cognitive liberty
    - Relational: Control, consent, access, relational autonomy, stewardship
    - Technical: Persistence, portability, interoperability, minimization, transparency
    - Political: Protection, equity, anti-coercive design
- Presented as a red-line document showing all changes from the original to encourage careful review and feedback

Key Additions and Revisions

- New principles added:

    - Relational autonomy (inspired by "rights end where your nose begins" concept)
    - Stewardship (addressing minors, elderly, disabilities while mitigating risks)
    - Equity (addressing structural inequity and imbalances)
    - Anti-coercive design (responding to rising autocracy and surveillance)
- Consent principle strengthened: Explicitly addresses dark patterns, choice fatigue, manufactured urgency, and clarifies that "consent that is merely clicked through is not consent"
- Protection principle expanded: Added paragraph addressing network vs. user conflicts, coercion in decentralized systems, and marginalized users
- Existence principle significantly revised: Clarifies that existence is inherent, not granted by systems or governments, addressing problematic implementations like Aadhaar

Areas of Uncertainty Requiring Feedback

- AI agent personhood: How to address AI agents without creating inappropriate carve-outs
- Revocability calibration: Whether certain rights (like GDPR's right to be forgotten) should be free or have acceptable costs
- Interoperability vs. monoculture: Whether the principles adequately address risks of single-system dominance and standards capture
- Anti-coercive design: Whether the 80-20 approach is sufficient or needs more comprehensive coverage

Discussion Points Raised

- Terminology concerns:
    - "User" is outdated language that should be replaced with "people"
    - "Coercive" may not adequately cover corporate dark patterns (e.g., Meta/Instagram case) - "coercion-resistant" suggested as alternative
- Legal frameworks:
    - Negligence-based liability could burden small operators while large companies remain unaffected
    - Agency law and duties framework (as in Wyoming) may be more effective than negligence standards
    - Intent requirements in coercion cases create high legal bars
- Passkeys concerns: Technology, regulation, and lack of incentives for Google/Apple to fix issues were discussed

Next Steps and Feedback Opportunities

- [ ] Read the red-line document carefully at revisitingssi.com
- [ ] Review the "lenses" document for detailed rationale behind changes
- [ ] Join the Google Doc to provide comments, redlines, and replies (open for comments)
- [ ] Attend Credential Community Group (CCG) meeting next Wednesday (online, open to non-members)
- [ ] Participate in extended meeting on June 20th at 10am Pacific (1.5-2 hours)
- [ ] Join Signal group for ongoing discussions (approximately 60 members, low noise)

- [ ] Iterate feedback over the summer for final presentation at Global Digital Collaboration event in Geneva (September)

Resources and Materials

- Main website: revisitingssi.com contains the 16 principles, 16 lenses, and bibliography of ~100 academic papers
- Document is community draft with CC-BY license, open for adaptation
- Open issues documented for all 16 proposed revisions

Interactive Endpoint Authorization via First Party Apps

Session Convener: Frederik Krogsdal Jacobsen
Session Notes Taker(s): Frederik Krogsdal Jacobsen

Tags / links to resources / technology discussed, related to this session:

https://github.com/openid/OpenID4VCI/issues/719
https://openid.net/specs/openid-4-verifiable-presentations-1_0.html
https://datatracker.ietf.org/doc/draft-ietf-oauth-first-party-apps/

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

OpenID for Verifiable Credential Issuance is introducing a mechanism called Interactive Authorization Endpoint. It has turned out to be basically the same as the IETF OAuth for First Party Applications draft. The purpose of the session is to find out if and how the Interactive Authorization Endpoint should be rebuilt as a profile of the First Party App draft.

Discussed with the first-party apps authors, digital credentials protocols working group and others. Conclusions:

- We should define our own redirect_to_web because theirs is JUST for going to an OAuth server
    - In PAR case, it means continue (i.e. you are done)
    - In non-PAR it means start over
    - Ours means something like “continue on web” or “do a step on the web”
- Should we add our own interaction_types parameter?
    - From a FPA perspective, it is intentionally not defined because they assume first-party so no negotiation is needed
    - So yes.
- Should we even do this rebuilding?
    - Is it worth taking on a dependency?
    - Helps only having to do security analysis once, e.g. for the AS mixup attack
    - If you are not an AS that also implements FPA, you probably don’t care. But if you are one, it makes a plugin type architecture possible
    - We could do a smaller profile that just does negotiation if other people also want it, then base on that instead
- AS mixup attack in AS Mix-Up Attack on IAE #595
    - Will attempt fix in FPA
    - Joseph will open an issue in the FPA repository

Storyboarding Agent Identity: Demos to Evaluate New Protocols

Session Convener: Rohit Khare
Session Notes Taker(s): Rohit Khare

Tags / links to resources / technology discussed, related to this session:

Rohit Khare proposed this session because in several recent episodes of Identorati Office Hours, he wished there were a common concrete scenario to demonstrate the differences between different new technologies for identifying and authorizing agents in several interviews:

- Episode 187: Automated Reasoning Meets AI
- Episode 186: MCP Dev Summit 2026 - Debrief
- Episode 183: Bots With Keys: Welcome to AAuth
- Episode 179: Identity-Centric to Proof-Centric Authz Models
- Episode 168: SAFE-MCP for Agentic AI

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Photos of notes taken

Storyboarding Agent ID Scenarios
Key Demo attributes
Classic great examples that have been shared for decades across industries, for example:

1. teaching SQL with “Payroll” tables (Employee, Department, Salary); or
2. teaching Social identity and access for “Photo Sharing” (with Albums, Groups, and even Faces)

Scenario Domains
Roundtable brainstorming of what attendees thought an Agentic Hero’s Journey might cover:

- bank monitoring
- expense reports
- travel planning
    - night out on the town
- Concert tickets (atomicity of purchasing, very unique seats/perks (non-fungible))
- General shopping assistant
    - (negotiating prices for buying commodity)
- ~~? anti-virus scanner~~
- "OpenClaw" | Siri | Personal Assistant

Qualities of a great Demo?
HERO'S JOURNEY

- Stakes?

    - Attack Surface? what is the risk of the agent failing or dying? (Death and Digital Estate) (Death :))
    - Impersonation or theft of identity
    - PERMISSIONS / ATTENUATION
        - *note there can be a difference between emotional reactions to the terms "query" vs "read-only" or "update" — risks differ in kind + degree*
- accountability — who is the actor?
    - *human (or child or group?)*
    - *non-human (or weights or prompt-hash or sbom?)*
- Observability / Auditability
    - What decisions were made?
    - What information was consulted? (e.g. victim of SEO attacks)
- Delegation point
- PRODUCTIVITY — LEVERAGE — SPEED
- Human (returns) in the loop (vs. YOLO)
- UNIVERSAL (in terms of income, or culture) — not everyone has US healthcare, or rich people problems
- Dispute resolution / Escalation — adversity in the demo when things go wrong
- "DETERMINISM" — risks of random response variation by LLM (or hallucination)

Evaluating Scenarios

- "File my Taxes" or Shoebox of Receipts
    - ? multiplayer (NO)
    - ? realtime (NO)

Real-time aspects (akin to adapting to delayed flights)
"dynamic pricing"? cultural bias?

Dramatis Personae for an AUTO INSURANCE / AUTO ACCIDENT

- Processes
    - emergency contact
    - share PII - IoT
    - photos/videos
    - estimates
- Players
    - police
    - repair shops
    - lawyers
    - doctors
    - ambulances
    - insurance
- Objectives
    - Safe

- Cheap
- Private

Evaluating Applications
What kinds of technologies would we like to “cover” using a common scenario?

1. OAUTH - w- MCP
2. SAFE-MCP
3. AAuth + Mission-driven Permission
4. ACP / A2A
5. Payments "AP2" policy intent for payments, 403(?) agent wallets
6. NHI (rate / quota limits, even if “authorized”?)
7. DID / VC
8. Personhood standards
9. Cedar guardrails used by Claude Code or OpenClaw
    a. External GUARDRAIL policies — can be partially-evaluated to improve planning loops?

Storyboarding: Crashing a Self-Driving Car
Hollywood plot lines of staging a crash in a bear suit? bad behavior by influencers

Main Scenario Spectrum
One-car accident to a two-car one? injury? — medical / legal?

Spectrum of scenarios (or start with simple one and keep adding complications)

1. tree
2. fender bender
3. related party injury (spouse in car ahead of you)
4. delivery robot in crosswalk?
5. self-driving car cuts you off
6. streaming video from “trapped inside speeding robotaxi car” (faked?)

Ideal User Experience?
automatically gather evidence
automatically share insurance card
automatically collect telemetry like location, airbag, carseat
deductibles, city government stats collection? CCTV? surveillance society?
? Where does the agent make a mistake? (confusion between “SRS” and “airbag”?)

SESSION #5

KERI/ACDC Bulk Issuance for SEDI Privacy - Security and WorkFlow

Session Convener: Sam Smith
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

slides:
https://github.com/SmithSamuelM/Papers/blob/master/presentations/BulkIssuancePrivacy.pdf

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Self-Sovereign Identity (SS) & Decentralized Identity (DI) 101

Session Convener: Steve McCown
Session Notes Taker(s): Steve McCown

Tags / links to resources / technology discussed, related to this session:

https://www.iinventstuff.com/files/Decentralized_Identity_101_Spring_2026_IIW_McCown.pdf

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The slides I presented in this session are available at:

https://www.iinventstuff.com/files/Decentralized_Identity_101_Spring_2026_IIW_McCown.pdf

KYAPay: Agentic Identity and Payment Credentials

Session Convener: Ankit Agarwal, Mike Jones
Session Notes Taker(s): Mike Jones

Tags / links to resources / technology discussed, related to this session:

We used this presentation to drive the discussion.

The specification defining the KYAPay claims is https://datatracker.ietf.org/doc/draft-skyfire-kyapayprofile/.

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

This session described the KYAPay system and specification. Ankit described the role of KYAPay tokens in letting agents access Web sites, despite mechanisms designed to block bots. He described verification steps performed by the parties and how payment works.

It is being used by companies for agentic commerce interactions, including merchants, payment providers, bot managers, and Web site protection services. Some of the parties using it can be found at https://kyapay.org/.

Ankit described the mechanisms used in KYAPay and also gave several demos.

I've got mad SKILLS.md! Writing skills to plan your digital estate with Claude.

Session Convener: Dean H. Saxe
Session Notes Taker(s): Dean H. Saxe

Tags / links to resources / technology discussed, related to this session:

Discussed a skill I am writing to help individuals plan their digital estate and emergency break glass kit with Claude.

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Notes are from after the session

- Discussed the broad applicability of the skill to be a neutral, third party to help individuals document the scope of their digital estate to build a break glass kit to gain access to their key accounts, and documentation for the individual's heirs regarding the disposition of digital estate data.
- We discussed extensions to the skill, how to make it more available to non-technical individuals, how to extend the skill with valid legal guidance, etc. We also discussed paths to open source the tool to ensure that it can be broadly distributed to help people planning for their digital estate.
- For access to the skill please reach out to Dean H. Saxe (https://www.linkedin.com/in/deanhsaxe/) with your GitHub userid.
- **THE SKILL IS NOT A REPLACEMENT FOR LEGAL ADVICE.**

EUDI Wallet Authentication & Authorization

Session Convener: Mirko Mollik
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Link to presentation: https://docs.google.com/presentation/d/1erPA-yK9JJSz3M8aS59oTAGD8lsb-dBS/edit?usp=sharing&ouid=116865618062361812085&rtpof=true&sd=true

Looking into the trust triangle of issuer, wallet, verifier

- how they authenticate to each other
- how the concept of authorized issuers is working

Some debate about if OpenID Federation can also be used compared to the SchemaMeta approach that also works with ETSI trust list

A new approach of a credential catalog was presented

- entries are not validated because it would not scale in reality

Basic concept: participants run a know your business before interacting with the wallet. But if they do not follow the rules and get caught, legal charges can be pressed against them

Also the question how non EU companies can interact with the EUDI Wallet since they cannot get an access certificate, eIDAS does not make clear how this is gonna get to work.

Identity Tales: Finding the Narrative

Session Convener: Erica Connell/Joe Andrieu
Session Notes Taker(s): Erica Connell

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Goals:

- Change the world by changing the narrative
- Inspire and educate about digital identity
- Spread the work

Today:

- Collaboratively create an identity story

Today's Tactics:

- Pick a harm
- Develop a scenario
    - Characters
    - Wants & Needs
- Story Elements
    - Inciting Incident, Obstacles, Narrative Turns, Climax, Resolution

Inspiration & Background:

- Storytelling for Kindergarteners
    - Main character - want - obstacle - resolution
- The Nutshell Technique
    - Jill Chamberlain
- Erica's story in 2 min

Ground Rules:

- Just-enough Just-in-time
- 80/20
- Yes, And
- What if?
- All our ideas are crap
    - Some hide gems
- Don't be shy
- Be positive
- Keep it moving

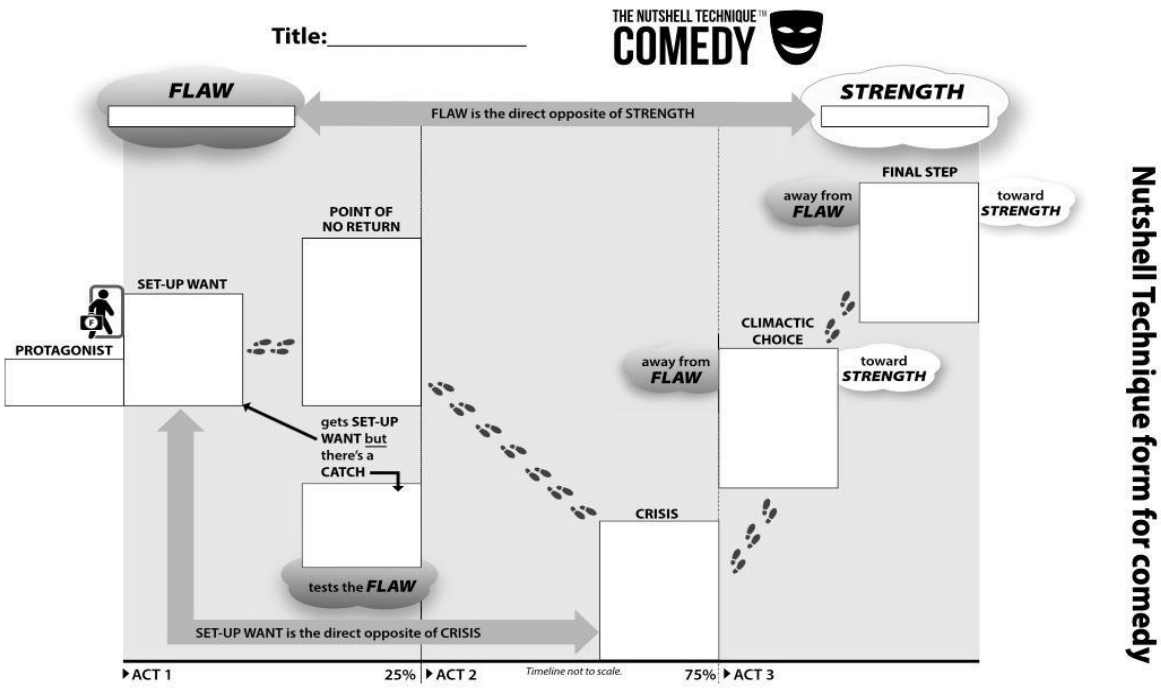

We co-created the heartwarming story of music duo Natalie and Karen. Natalie is the more talented but more insecure half of the duo, while Karen is much more secure and somewhat less talented. When nefarious music producer Harvey Specter offers Karen a huge contract with a concert attached, he plays to her ego and convinces her to leave Natalie behind since we can just use AI for her part. Then Karen, of course, realizes the AI is ass, and she has a heart to heart with Natalie who is persuaded by their newfound honesty and Karen's humble pie eating to help her friend. Harvey Specter threatens them with all the evils if they go on together. They do, anyway of course and expose him onstage. They don't get paid since they did not fulfill the contract, but they do create their own peer to peer terms for their contract with each other, as well as unified terms for future contracts moving forward.

Identity As A Mixed Digital Martial Art (MDMA); The Digital X

Session Convener: Jeff Orgel
Session Notes Taker(s): Jeff Orgel

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The way we conduct ourselves in the digital realm provides varying degrees of managing and protecting facets of ourselves and our personal nature. When we use IT devices and connected systems our functional presence on the digital data-scape carries with it attributes of presence/existence. The manner in which we establish system relationships such as the difference between creating a new account or logging in with a pre-existing Facebook, Gmail or some other account profile often makes a vast difference in the experience you will have.

Intentionally or via system automation we enter user IDs, connect personal devices and smart home smart car smart wearable smart bed all log in with MAC numbers, data drops and more. All this is registering in some form or fashion somewhere on this digital datascape, and that, I suggest, will always matter - some way or another – sooner or later…

This discussion highlighted the various ways in which we may set the relationship frame, aka the "X", for the purpose of this summary. The way that we frame this relationship "X" will determine how our digital presence is more or less impacted. In following the way that we frame this relationship "X" will also determine how our real-world presence is more or less impacted.

For the purpose of this discussion there is a comparison between the real world "X" where our feet touch the ground, and the digital datascape "X" where we operate as an instance of code running in a system.

In our daily real world, every day we are the same person yet we may present aspects of ourselves to the world. These different roles and related personalities (aka personas) allow us to expose the portion of ourselves we deem appropriate in the circumstances of the moment. In the digital landscape what are the software settings and identity choices which will serve these purposes...if any are even provided by the system?

How can we present ourselves to the digital landscape on "Our Terms" more so than the House's or platform's terms? When can we adjust these factors which may allow us some agency to choose – and anonymity to be discreet - and when can't we adjust these factors so we go forth without the ability to choose aspects of the relationship?

All this just in time for "My Terms" (IEEE 7012) coming over the Horizon! If I could vibe code my vibe to the system I'd say;

"I want to get this task done. Let me be the me I wish to be. Don't push me around, distract me or psyop me. Just provide the experience and outcome I intend. Assist when I ask, otherwise stay out of my way."

In session we gamed ways in which we may adjust our digital presence and perspective such that online behavior choices as well as operational personas used to operate within those systems will be best aligned with our initial intention outcomes.

Here are a couple examples.

REAL WORLD MODELS

The CASINO: ON the "House's X"

Use Case: Average Casino Customer

If a casino could put you on the "X" they'd like to see you on it would be in the shape of things like this.. Room TVs start up featuring House promos and shows. All gambling chairs, shows, restaurants and services are big examples of the House's "**X**". Manners of maximizing extraction may take shape such as how often gambling seats get checked for "free" drinks. Maybe the high profit machines are an "**X**" that's checked every 5 minutes, while less profitable seats get drink checks every 10 minutes. We can be sure that comfy chairs are everywhere and especially in front of the most profitable gambling machines.

Re-Crafting The CASINO: OFF the House's X and onto to "My Terms X"

Use Case: Casino Customer Jeff - Consumer Electronics Show Attendee

Sorry to disappoint the house's model but… I don't gamble. I use free drink coupons from my room booklet. I have freedom of movement since I'm not pinned down so... I get the best views I like, which is people watching and eye candy for me. I don't go to shows. I eat at a favorite restaurant local to StL on the main floor - YUM! I'm off the Strip by a few miles. Other than a meal, no in house purchases.

DIGITAL LANDSCAPE MODELS

NEW DEVICE SETUP BOUNDARIES

Default Set Up is ON the "Manufacturer's X"

Following directions will guide into an Accept, Accept, Accept of Default Settings: location, data shares, sign-ins (maybe using your w/Google or FB), etc.

TAG's MDMA Set Up is OFF "Manufacturer's X" and onto the "My Terms X"

Setup w/o email account association. Disable some/to many PC connection services. Quiet/Silence Select Default Settings: location, data shares, sign-in w/new specialty account, etc.

NEW APP / PLATFORM SIGN UP

Default Set Up is ON the " Manufacturer's X"

Following directions will guide into Default Settings: location, data shares, sign in, etc. Create an Account. Use the App! (I'd say generally only fiduciary types sites)

TAG's MDMA Set Up is OFF "Manufacturer's X" and onto the "My Terms X"

Not Usually Default Settings: Using the web page in a secure browser as opposed to an app. Visit the store vs. give the "relationship" a room in your house, seat at the dinner table, angling for the head of the table if they can. Guest sign-in (w/email confirmation so any reward activity can be vetted by email records) as opposed to creating a profile everywhere. Reviewed the idea that categories of apps such as financial, health, security can be installed and leveraged as can be as they have better designed, interface engineering for these critical information segments.

…

These notes are an aggregation of numerous prior sessions and accumulated wisdom. It is through these sensibilities that this skill takes shape as a practice of forms. This awareness can allow for a view of and experience of the digital realm which may seem more familiar to people than the alien, intangible non-physical landscape many people are struggling to understand and align sensibilities to. Again, thanks to all who continue to refine and challenge my thinking. I am here to be redirected and corrected as needed and rarely is there a better place. _/_

The Fiduciary Common Part 2 - Why Incomplete Legislation is a Practitioner's Problem

Session Convener: Mike Leahy
Session Notes Taker(s): Mike Leahy

Tags / links to resources / technology discussed, related to this session:
Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The Fiduciary Commons | Session 2 of 3
Why Incomplete Legislation Is a Practitioner Problem
The Case for Enacting VIDA, PDTA, and GAAFA Together

What This Session Argued
The Fiduciary Commons framework only works as a complete, uniform whole. Enacting pieces of it without the others creates specific failure scenarios with direct commercial consequences for identity vendors and protocol designers. The three scenarios described in this session are not hypothetical future risks: they describe current conditions in jurisdictions that have enacted portions of the framework without the others. Practitioners who engage with the legislative process to ensure complete and consistent enactment are protecting their own commercial interests, not acting as advocates for an external cause.
Three Failure Scenarios
VIDA without PDTA: the surveillance problem survives.
A technically correct credential architecture issues selective disclosure credentials and zero-knowledge proofs. No underlying data is transmitted. The architecture meets every VIDA requirement. But verification metadata, recording who presented credentials, to which agency, when and where, is unprotected. Without PDTA's purpose limitation and fiduciary duties, that metadata can be aggregated into a behavioral profile constitutionally indistinguishable from the surveillance the credential was designed to prevent. The architecture was correct. The vendor is named in the complaint.
PDTA without VIDA: the duty is unachievable.
Legal obligations requiring data minimization, purpose limitation, and individual-controlled audit trails are imposed on agencies whose infrastructure is a centralized database with no zero-knowledge capability and no cryptographic audit trail. The duty is real; the means to fulfill it structurally do not exist. Vendors are asked to certify compliance with a standard the system cannot meet, and face liability for secondary data use they had no power to prevent.
Fragmentation: the cost is multiplicative, not additive.
Each additional jurisdiction with an incompatible framework does not add one unit of compliance cost. It multiplies the complexity of every decision that crosses state lines. Three states with three different partial frameworks produce edge cases requiring expensive legal analysis in every jurisdiction, compliance architectures that cannot be standardized, and pricing complexity that vendors pass back to government as higher contract costs. A uniform framework produces one compliance architecture, one audit standard, and competitive procurement conditions.

The vendor's name is on the contract. The regulatory gap has no name. Legislative committees call the vendor. No one calls the gap. This asymmetry is structural and will not resolve itself without practitioner engagement.

The Closing Ask

Three actions, not a lobbying program. First: map the legislative posture of your primary state markets. Identify which pieces of the framework are under consideration and which are absent. Second: assess your own architecture against the full VIDA, PDTA, and GAAFA requirements. Understand your gaps and your migration path. Third: make your technical expertise available to legislative staff and state technology officials in the states where you do business. You do not need to advocate for specific bill language. You need to explain what technically achievable compliance looks like.

What This Means for Your Work

The visibility asymmetry is real and structural

When a government identity system fails, the vendor is visible and the regulatory gap is not. The contract is public record; the missing statutory provision built nothing and appears nowhere. Technical systems have owners; legislative omissions do not. Practitioners who engage with the legislative process to ensure complete enactment are protecting themselves from the consequences of gaps they had no hand in creating.

Three arguments that land with government officials

Incomplete enactment does not save money: architectural retrofit costs exceed implementation costs. Every gap closed later is closed more expensively.

Fragmentation inflates government procurement costs: inconsistent frameworks mean vendors price multi-jurisdiction complexity into every RFP.

The framework is constitutional compliance, not regulatory burden: VIDA, PDTA, and GAAFA codify obligations the Fourth Amendment has always required.

The conversation with your state CIO

You build identity systems for a living and have a direct stake in whether the regulatory environment is coherent. The bills mandate standards that deployed systems already use or are actively moving toward. Partial enactment costs more to fix later. When you hear 'we already have privacy laws,' the response is that fiduciary duty is a higher standard than notice-and-consent. When you hear 'we will address this in rulemaking,' the response is that rulemaking without statutory foundation produces inconsistent results harder to comply with than clear legislation.

Key Terms

Verification Metadata

The transactional record created when a credential is presented: who presented, to which verifier, when, where. Unprotected by VIDA alone. The surveillance residue that PDTA is designed to govern.

Purpose Limitation

The PDTA requirement that data collected for one purpose may not be used for another. Enforced through both legal duty and, under a complete framework, architectural controls that gate access by stated purpose.

Solution Bias

The tendency of procurement specifications to describe technical means rather than functional outcomes, excluding innovative approaches before evaluation begins. Article 3B addresses this through outcome-based specifications.

Fragmentation Risk

The commercial and legal cost of inconsistent state-by-state enactment of the framework. Each incompatible jurisdiction multiplies rather than adds compliance complexity.

Flow-Down Clause

A contractual provision that extends the prime contractor's obligations to subcontractors, naming the government as a third-party beneficiary. Required under Article 3B to ensure fiduciary data obligations reach every entity in the contracting chain.

GAAFA (Government AI Accountability and Fiduciary Act)

Model statute extending fiduciary duties to the AI layer: algorithmic impact assessments, public registries, explainability duties, and prohibition on inference-based reconstruction of data government had no right to collect.

fiduciarycommons.com | michael@fiduciarycommons.com | For the full argument, see the Fiduciary Commons Series Summary.

OpenHaven.net Building Bridges in the P2P Ecosystem

Session Convener: Day Waterbury & Brandon Noorgard
Session Notes Taker(s): Day Waterbury

Tags / links to resources / technology discussed, related to this session:

- https://openhaven.net/brief/ <-- What is this all about? What is the status of the project?
- https://openhaven.net/prototype/ <-- What are you trying to do? What does your community need?
- https://openhaven.net/prototype/matrix/ <-- What tech, platforms, and protocols are listed so far?
- https://openhaven.net/contribute/ <-- Please help us complete the matrix, provide input/feedback on entries, and make OpenHaven better!

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Executive Summary
OpenHaven is a community-led, data-driven convergence navigator for the peer-to-peer (P2P) and decentralized protocol landscape that makes it visible, navigable, and directly actionable — including for AI-assisted development workflows. It replaces the current single-page openhaven.net with a live, data-driven website that serves as both a navigable guide to open protocols and a collective voice for the people building and using them.

Born from the Collaborative Technology Alliance (CTA) and a growing coalition of researchers, builders, and community leaders, OpenHaven translates complex protocol research into something legible, actionable, and shareable. OpenHaven's mission is making the open protocol landscape interoperable and interact-able — navigable by the people who need it — informed by the understanding that technical convergence depends on social convergence. Not just mapping the technology, but building the tent — giving the movement a collective voice, a shared home, and tools — including machine-readable context and prompt generation for AI-assisted workflows — that lower the barrier to understanding and adoption.

OpenHaven serves three interconnected audiences: community leaders seeking tools to solve real coordination problems, builders and researchers seeking convergence and network effects across fragmented efforts, and coalition organizations (e.g. CTA, DWeb) needing a shared evidence base to weave their convergence work together. For all three, OpenHaven is a hub they are proud to share — because it demonstrates tangible progress toward foundational infrastructure for humanity's coordination challenges.
Other Notes

- Prototype UI feedback from Frank Sanborn: Choose multiple use cases

- Software Bill of Goods suggestion by Tor Hagemann (is such and such a vulnerable lib a dependency?)
- Add DIDCOMM and ToIP/TSP suggested by Drummond Reed
- Add the protocol that When & Swan are working on w/ human.ing suggested by Day Waterbury
- "Standards don't matter as much anymore because you can just tell AI what you're trying to do and it can use whatever standards accomplish that." ~Frank (very much paraphrased)
- "Role of standards is shifting because of AI, however if the infrastructure isn't trusted, there are issues." ~Drummond (also paraphrased)
- Last Network Effect. Bridging Protocols. Bring Your Own Everything (Dmitri & Bengo) mentioned by Day.

Safeguarding Products + Standards for The Dark Times

Session Convener: Elizabeth Garber
Session Notes Taker(s): Kiser

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

EIC talk about "identity in a changing United States"

Lots of changes over the past few decades:

- surveillance
- faster, scalable potential for harms

How do we go about safeguarding the good things?

How can standards org build-in harm modeling into their process?

Abuses of the fourth amendment (recommended history research) (search and seizure rights)

How do you utilize data / info about people in a way that serves use cases while preserving privacy?

Implementation side is different than the standards-based approach.

"threat modeling of the spec" might be an interesting idea (but how would you execute it?)

Secure communications / peer to peer was needed , but wasn't there before . . .

Operational security will always be a component (esp for vulnerable populations / and educational purposes)

Identifying agents and imbuing them with human accountability . . .

Next wave of self-sovereign - concept of existence (does it only apply to humans)

if you're going to have freedoms, then there bad actors present (by default . . .)

Technology makes things easier for us . . . grinding grain, - now it's easier to amplify messages …

Design time mitigation of risks is important. . .

"Technnology is never neutral"

Video compression came from the pornography industry . . .

Conversations are charged into us/ them …. dialogue is unusual…

Us vs them implications of ethics . . .

Bringing ethics into technology / / / “we are not ethicists” - but we should have the conversation about it

Incentives…

We can’t control the technology - it gets too late in the process…

Do we need multidisciplinary people participating in the technology development process?

Are most people good intentioned or not? (differing views among the group) - that influences what the social contract turns into.

The people who design the algorithms, etc — what should they be doing while they’re building it . . .

Should harm modeling be part of the standards process?

The cost of examining “what ifs” is extremely costly for the development process (economic disincentive).

Perhaps AI might enable safeguards, perhaps not . . .

You need to align incentives with ethical boundaries. . .

Utah has decided that interoperability is not a value.

- privacy preservation
- container signatures

“the tech is available”

state endorsed digital identity
(state didn’t give you the identity)
Bring your own identity (you have your own identity)

Notes Day 2 / Wednesday April 29 / Sessions 6 - 10

SESSION #6

Archival-Quality Identifiers?

Session Convener: Eric Scouten
Session Notes Taker(s): Eric Scouten

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Start from ToIP's definition of verifiable identifier: "Any credential for which an independent system is able to associate, discover, and verify the cryptographic keys associated with that credential."

But … this doesn't say anything about when the verification occurs.

How can I verify a signature years after the fact?

KERI uses the term "perpetual identity." Proof of control is perpetual. You need a fault-tolerant mechanism. KERI uses append-only logs to avoid the question of time stamps.

Proposed definition of archival-quality credential (take 1)

CAWG definition of archival-quality credential: "Any credential that can make claims on behalf of its subject for which an independent system is able to verify the authenticity of the claim at any time after the claim is made."

(Note that we changed thge name from identifier to credential.)

Questions still:

- Legitimacy of claims?
- Repudiation of illegitimate claims
- Drummond: add authenticity

What data needs to be around in order to verify a claim years after the fact?

Sam suggests that signature is untenable for archival-quality identifier.

Proposed definition of archival-quality credential (take 2)

CAWG definition of archival-quality credential: “Any credential that can make claims on behalf of its subject for which an independent system is able to verify the link between a claim made with that credential and its subject with that credential at any time after the claim.”

Veracity of claim is out of scope. (Is what was claimed by the subject factually correct? Out of scope.)
Authenticity is limited to the binding between the credential subject and the claim made by that subject.

Open questions:

- Does the credential subject need to pre-authorize the means of verification?
- Does the trust framework need to specify how the verification occurs?

Important to say there are TWO trust diamonds in CAWG identity ecosystem, both of which need to be archival-quality credentials:

- Issuer -> named actor
- Named actor -> specific piece of content

Digital ID Acceleration - The world needs us

Session Convener: Riley Hughes
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Beyond Triples - Semantic Frame as Data Model - Meaning & Trust Substrate

Session Convener: Joshua chambers
Session Notes Taker(s): Joshua Chambers

Tags / links to resources / technology discussed, related to this session:

https://github.com/joshualibrarian/common-graph/releases/download/paper-v1/frames-all-the-way-down.pdf

reference implementation code base (in progress):
https://github.com/joshualibrarian/common-graph

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

We had a small but lively session to discuss a fundamentally new data model and its consequences. We repurpose the semantic frame from linguistics, import a vocabulary from linguistics resources, and reduce all data to signed assertions.

What do you mean "Call my Agent" | How do service providers contact agents that have been provisioned into them?

Session Convener: Hirsch Singhal
Session Notes Taker(s): Steven Tamm

Tags / links to resources / technology discussed, related to this session:

n/a

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Summary

This was an open brainstorming session about how apps should reach out to an agent in order to trigger it or somehow get its attention. You can communicate to your users via email or in-product UX (i.e. an inbox) but we don't have that pattern for agents.
Term of art courtesy of Steven Tamm: "**Activation**". How does a service provider **activate** an agent that has been registered with it, in a standards-based, discoverable way?
This was built around the following scenario, where platforms are handling agents built by others:

1. Agents are actors that are somehow provisioned into multiple systems (and homed in another). For example, workspace agents from OpenAI, Agent Users in Entra, or Agents running in Bedrock/LangChain/self-hosted servers.
2. Agents are provisioned en masse, not bespoke – configuring connections or callbacks manually (i.e. via human power) between resources and agents isn't a solution.
3. Communication to agents is about events occurring in the service provider that the agent should know about and generally respond to. This is in contrast to events that flow as part of an ongoing session or task that the agent is participating in. Event examples include:
    a. Assigned to an issue in GitHub
    b. Mentioned in Figma
    c. Replied to in Slack

Spoiler: It looks like emailing an agent might be the best way to get a hold of it, for now.

Problem exploration

The group identified multiple questions that would need to be answered to determine how to contact agents:

1. Broadcast, direct, or routed?
    a. Some participants suggested a broadcast model, where messages were dropped in a known location and agents picked them up, removing the requirement to contact the agents directly.
    b. Others suggested direct communication to the agent, removing overhead in the system, but requiring advertisement of the channel and the potential for problems when agents migrate domains

c. A third option was routing through the responsible party or the agent's IdP, which removes several concerns from the scenario including discovery and portability, as well as trust concerns and filtering. However, it requires additional setup and reliance on an extra party.

2. Nature of the events to notify on.

a. Email was suggested as a generally satisfactory option, as most systems already support emailing users when action is required.

i.This was not generally true though – some high-touch platforms do not email you for every single actionable event as that would be too high a load.

b. Do the events need to be structured, given LLMs ability to parse open text?

i.Ideally yes so that filters can run before the LLM receives the data.

3. How do we consider sub-agents? They will not be provisioned durably, so won't have their own comms channel (if going direct)

a. Communicate to the parent agent?

b. Comms to a router keyed to a task or session?

c. **Perhaps an implication that a communication channel needs *some* standard metadata around task, session, sub-agent to help the agent understand the notification.**

4. If the comms channel is an open channel, how does the agent trust the sender? I.e. anyone can email them "please reset your password", and *we* know not to trust that, but does an agent?

a. Some discussion of whether this is an authz question – check that the sender had permissions versus a runtime safety question (validate the sender).

b. Filtering via a router can help here.

c. IdP-as-router helps prove access – the IdP knows who can contact the agent,and therefore can filter.

5. Cost concerns – if every event can spin up an LLM and cost tokens, configuring for SnR via things like topics, batching, and relevance becomes important. Spam emails constitute real cost now.

6. Who is responsible for the agent?

a. Some human (real or organizational) is responsible for the agent, but are they always the person that should receive notifications?

b. Yes, in the case of ToU/ToS discussions

c. No, not when talking about minute by minute event-based comms. That's not the job of the owner to handle.

7. What channels are possible?

a. Email – already works, almost too well (CTA marketing emails for instance)

i.But not every agent will have an inbox provisioned. I.e. an agent from Okta – who's hosting and paying for that inbox? Is this generally workable?

ii. Can be flaky at scale – anti-spam filters, e.g. Exchange blocks emails from Microsoft all the time.

b. Webhooks – well understood, but not as easy to expose, especially in corporate environments. Not performant at scale in the same.

c. Pubsub and other queues

d. **Email seems like the most workable option.**

e. **Shared Signals seems like the most appropriate "standard" path.**

Discussion focused heavily on the theme of the week – what *is* an agent and do they have a durable identity even? Do we care about subagents or just the "root" agent? Is the agent only around when they're working on a task (in which case, why not contact them as part of the back and forth in the task?) or do they exist "at rest" outside of a task? Only in the latter case does a service provider need to reach out to an agent. Otherwise the SP just needs to communicate with the owner / responsible party of the agent for ToS-type communication.

One model of "agents" was a single client_id (i.e. codex, or github, or cursor) injected into the service provider as an application. When triggered within the service provider, this app functions as a router to the individual agents that exist "behind" the app (i.e. within Codex or GitHub or elsewhere). Thus service providers don't need to know about individual agents, just the service that provides them.

Takeaways

- We need broader consensus on what constitutes an agent in order to facilitate conversations at a higher level.
- We need to identify when discussing agents if we expect agents to operate autonomously (as in, as the subject of the token, the principal that was given permissions directly) or it they're exclusively clients for user principals.
- Email appears to be the most straightforward and compatible approach to activating an agent from afar, but has some implementation, scale, and reliability problems.

Figuring out a story for how agents *get to* the service provider will inform their shape, and vice versa. These should evolve in parallel.

Customer Commons and MY Terms 101: Lawyer in your Pocket

Session Convener: Doc Searls, Nitin Badajia, Iain Henderson, Mary Hodder, Justin Byrd
Session Notes Taker(s): Mary Hodder

Tags / links to resources / technology discussed, related to this session:

My Terms video at youtube: https://www.youtube.com/watch?v=ODlVXwxax4Q
IEEE 7012 Standard located here: https://ieeexplore.ieee.org/document/11360682
Customer Commons: https://customercommons.org/
MyTerms: https://myterms.info
MyTerms Launch webinar: https://www.youtube.com/watch?v=Nphd8l7KLek
Creative Commons: https://creativecommons.org/ and
CC licenses: https://creativecommons.org/cc-licenses/

Blog post: "I verified my Identity on Linked In. Here's What I actually handed in."
https://thelocalstack.eu/posts/linkedin-identity-verification-privacy/
Beyond Consent: A Right-to-Use License for Mutual Agency paper by Lisa LaVasseur and Eve Maler: https://doi.org/10.1109/MCOMSTD.001.1900031
IEEE 7012 Industry Connections or My Terms working group: https://standards.ieee.org/industry-connections/activities/
Sign up form: https://app.smartsheet.com/b/form/019d96d027fe74c2be64aef29dedb395
The IC group's first meeting is June 1, 2026.

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

1. Played video first to explain IEEE 7012 Standard, or My Terms, to explain the premise and standard.

Contracts not Consent
Customer Commons worked with IEEE to create this machine layer description.

Creative Commons is the model for Customer Commons, and is based upon the general idea of liberalizing licensing of copyrights for your media where you choose the license. MyTerms will be about presenting a contract with a terms or terms, in an agreement, that if an entity or person agrees, supersedes the TOU that might govern the rest of the relationship. Or replace it.

Advantages for businesses:
- Reduced or eliminated compliance risk.
- Competitive differentiation.
- Lower customer churn.
- Grounds for real rather than coerced relationships (CRM+VRM)
- Grounds for better signaling going in both directions.

- Reduced or eliminated guesswork about what customers want, how they use products and services, and how both might be improved.

SLIDE on Advantages for Business with Relationship symbols:

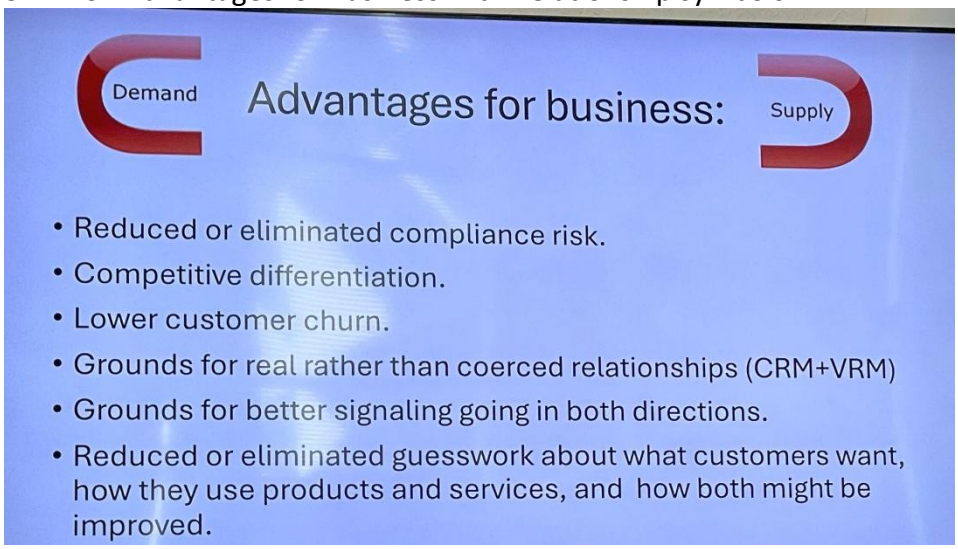

Question: what is a coerced relationship?
Doc: Example, Peets coffee where it's on their terms. I have to remember to bring my app or scan the QR code, and I have to do it. I want a coffee app, that allows me to go from shop to shop, and my data isn't collected, but I share at each order / shop, my coffee preference. Or mortgage docs where you must agree to a bunch of 3rd parties that get your data and the mortgage terms are their terms. Or grocery stores and the sales prices that require a loyalty program to get the sale.

First target to accept MyTerms: Wordpress sites with 556M sites on web / web services.

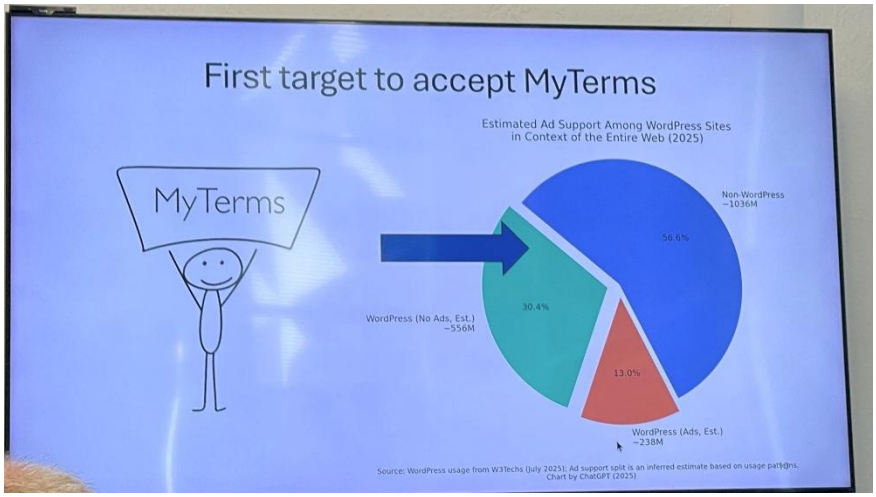

Agents will act on your behalf at scale, and should be carrying your terms for you, and keep your data safe otherwise.

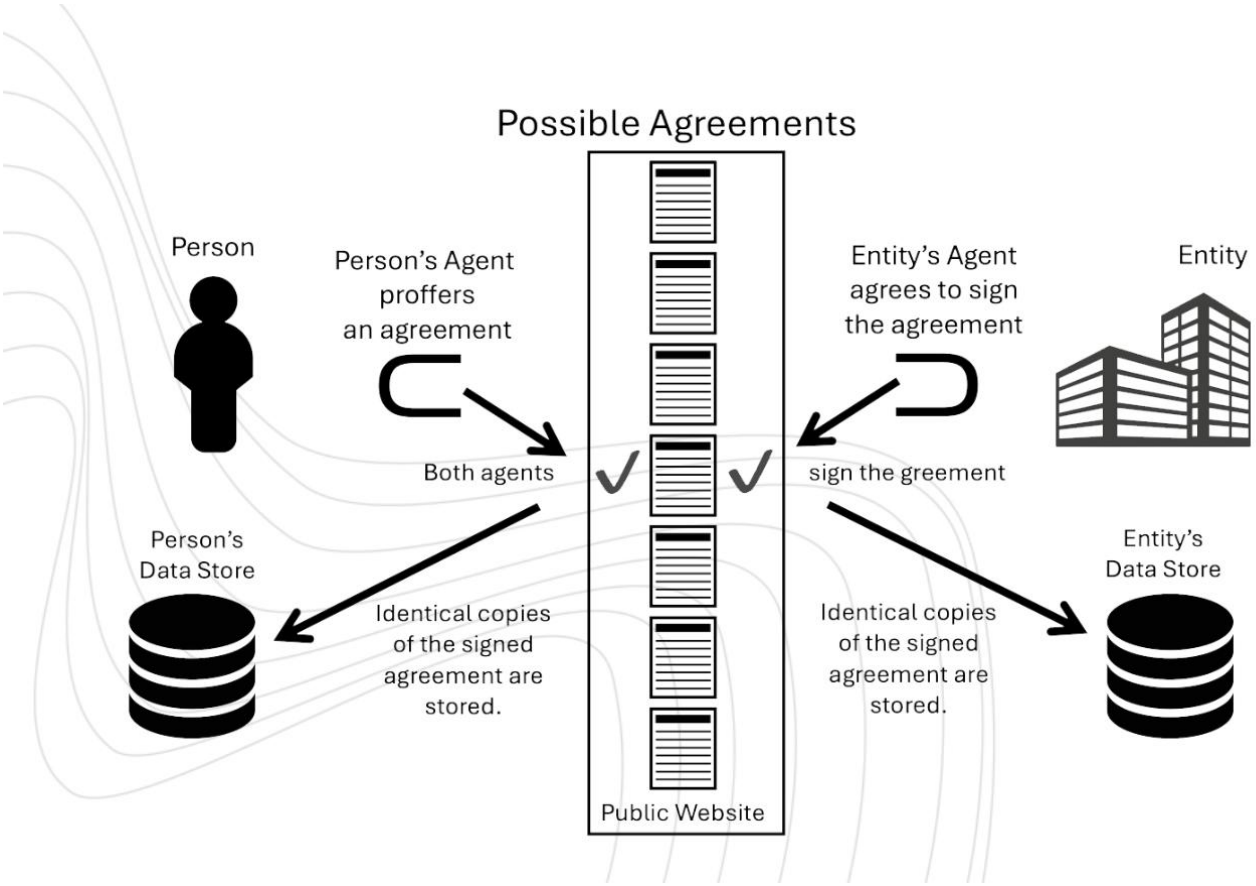

Example of coercion in this blog post: "I verified my Identity on Linked In. Here's What I actually handed in." https://thelocalstack.eu/posts/linkedin-identity-verification-privacy/

When the individual sends signal it has highest context, and therefore market opportunity is highest in sequestering data, and then sharing it out with those who will respect you and your data.

Fresh data is what everyone wants. We may not have privacy before now but with contracts and MyTerms, we can create privacy when we want it

Eve Maler: Lisa LeVasseur and she wrote the this paper looking at this:

> Beyond Consent: A Right-to-Use License for Mutual Agency. IEEE Commun. Stand. Mag. 3(4): 52-59 (2019) https://doi.org/10.1109/MCOMSTD.001.1900031

We need help.. people who want to join the IEEE 7012 Industry Connection or My Terms working groups can sign up here: https://standards.ieee.org/industry-connections/activities/
at this form: https://app.smartsheet.com/b/form/019d96d027fe74c2be64aef29dedb395

MY Terms interest group sign ups at this email: alliance@myterms.info

Pairwise Pseudonyms from Government ID's or Holy Pseudonym Batman POW

Session Convener: John B & Dirk
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Originator Profile for Content Authenticity

Session Convener: Michiko
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Overview of the technology and organization
https://originator-profile.org/en-US/

Technical documents
https://docs.originator-profile.org/en/

Slides
★OP_IIW42.pdf

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

- Intro of OP
- OP mechanism
- Future Scope
- Demo
- PoC report

Why Physical IDs need Digital Stuff - come see some counterfeits

Session Convener: Elaine Wooton
Session Notes Taker(s): Elaine Wooton

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

There are very good counterfeit versions of all US driver licenses. To make it possible to identity proof using photos of the documents, we need to add signed hashes of data / VCs / biometric data (face templates).

Without that, and without doing much more privacy invasive data checking, the task is likely to be inaccurate.

Trust Infrastructure as a Public Utility Part 2

Session Convener: Erika Bjune
Session Notes Taker(s): Margeigh Novotny

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

From yesterday, deeper dive on adoption drivers, business model, what it means to be public infrastructure

How can we we get people to care about identity as core infrastructure the way we care about water and energy?

- People are used to giving away their identity… they don't value their id
- Swan: people get social value, soft bump for posting on social media
- Dimitri: if public ID infrastructure is cheap and painless enough to run, we need to convince govts, institutions to run the
    - Highest accessibility approach
- Scope?
- Kaliya: how are we bigger than big tech? We can be bigger than big tech by insisting on open protocols
    - How do we assert power and reclaim our IDs?
    - Technologies are adopted socially
    - Create social care
    - Faucet-makers: find them (we've built the plumbing)

        - Swan: “When we turn on the faucet we don’t know whether we are going to get water or excrement”
- Swan: if its infrastructure its everyone - kids, old people
    - Figure out the pain point
    - Segment the audiences, find their overlaps, thread thru the overlaps
    - Easier to
    - Mechanics of economics have made it less possible
    - We tolerate fraud in our digital lives that we wouldn’t in our physical lives
- People care about easy
- Ben: social and technical aspects are not mutually exclusive
    - Social drives migration to new tech
    - What if identity were my aggregation of the things I interact with
- Sheref: technology exists to solves problems, and if no one uses it we haven’t solved any problem
    - What are we going to make better with the tech?
    - People don’t care fundamentally about privacy, not enough to move the needle
- Brad: you can’t navigate the social network that connects silo’d communities
- Ben: if a business can access
- Erika: might we have to rely on fear tactics.

Trust (not privacy, not identity)
Justice:

Are there companies and other entities that benefit from maintaining public infrastructure?

Ideal customer profile ICP

STAKEHOLDERS	ROLE	CARES ABOUT…	NOTES
Stewards	Who preserves		
community	People served	Am I being treated fairly	
Industry And government	Who profits (or loses)	Maintaining their own value	Defense industry: Palantir is presenting a threat to the defense industry Media may be a
enforcement	Who makes it stick		
Ad Tech	Needs to	Matchmaking between product/service companies and access to ready consumers	People don’t mind being sold to we don’t want to be manipulated

ARE THERE EXISTING MODELS? (highly prolific, no barrier to entry or access, something fun, something)

- DNS? People pay for it, they build on top of it, they can build what they want
- SIP
- MASTODON, ACTIVITY-PUB
- Life 360, FindMY, but - but for everything I engage with (if the device is broadcasting its own place in the trust graph, relative to all the other nodes and the content my nodes are
- Porn will lead

What to do with 3 million verified business IDs?

Session Convener: James Monaghan & Stefan Heller
Session Notes Taker(s): James Monaghan

Tags / links to resources / technology discussed, related to this session:

Verifiable credentials, KYB, compliance, business wallets

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Background

Campaign Registry helps telcos answer the “who and what” when businesses want to connect to their platforms:
Who is the business
What do they plan to send over the carriers connection

3.1M active brands
26,000 enterprises, 3M SMBs
Registered for use of 10 digit long codes (10DLC) for SMS campaigns with the US mobile carriers

Issuing brand verifications as verifiable credentials into digital wallets
Looking to help brands get additional value from those credentials.

Purpose of the session: soliciting ideas for where that verified brand and reputation is useful in other contexts

Suggestions

Education will be important for the brands - they need to understand that they're getting a portable ID and what it can be used for

Lots of SMBs are skipping websites and going straight to agents - consider letting businesses prove their agent is legitimate
Delegated credential for customers to prove their loyalty?

Look for equivalent regulatory requirements (e.g. at the state level) which match the carrier verification requirements, and target those
Recognise that compliance is a grudge buy for small businesses, so this may not be super compelling

Enable customers to filter verified emails from the company

Equivalent to the Quickbooks marketplace?
Does any vetting happen in these marketplaces today besides making a payment to be listed and agreeing to terms & conditions?

Look for similar patterns to registry / marketplace businesses
Amazon third party sellers, Facebook marketplace, Etsy, Ebay, etc
Job boards
Invoice platforms

Address fraud & impersonation related to use of QR codes
e.g. restaurant chains & other small physical retail businesses
OASIS secure QR code spec - includes a signature

Brand insurance
Would having a trust mark lower the premium?

What would an actual democratizing technology look like?

Session Convener: Dave Sanford
Session Notes Taker(s): Dave sanford

Tags / links to resources / technology discussed, related to this session:

#democracy #vtc #karma

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Dave Sanford presented various specific approaches used in his fiction:

- Software Enabled Democratic Organization (SEDO)
- Multivariable (3 - truth/science, utility, cooperation/collaboration) used in a +1 or -1 possible response for each
- Demos, a Reddit-like multi-forum democratic social network. Each forum is designed to address a 'greater good' goal
- SEDOs and Demos forums have discussion, consensus mechanisms, karma, facilitation, proposal creation, voting, and project management capabilities
- SEDOs can be implemented as Verifiable Trust Communities (VTCs)

Frank Sanborn discussed his implementation supporting community creation, primarily through real-time disaster response, which enables the establishment of community collaboration and consensus networks.

Two referenced books were 'Cognitive Surplus' by Clay Shirky, and multi-vector karma systems for social networks, based but different than the one in 'The Hype Machine' by Sinan Aral.

Do you have the digital identity credentials that you need?

Session Convener: Karla McKenna
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

We discussed the approach of considering a number of factors that are needed to accomplish/satisfy/cover the usages of a person's or organization's digital identity credentials.

Factors to consider: Identifier characteristics (example: LEIs for organizational identity), identification of the credential holder, digital signing needs, delegated authority, credential lifecycle, privacy, security and cost.

First, Entity Identification (for Organizational Credentials)

- **Important features to consider**
  - — Standard identifier
  - — Globally recognized and applicable
  - — Available entity data
    - Limited set of entity data
    - Standard list of validated entity data
    - Requirement for keeping entity data up-to-date
    - Easy, free of charge/reasonable cost access to entity data (also has implications for Cost)

The LEI.. more than just a code...

506700GE1G29325QX363

Key reference data includes:

- **Entity identity**: Legal name and transliterated name
- **Official registration**: Local business register and ID
- **Mapped identifiers**: Links to other codes/databases
- **Entity status**: Active vs. inactive entity
- **Address**: Registered legal address and headquarters
- **Validation**: Date of last LEI data check/renewal
- **Ownership links**: Direct and ultimate parent information for consolidating entities

What is the LEI?

- The LEI is an ISO standard ISO 17442
- The LEI is a life-long identifier **owned** by the respective legal entity.
- It points to the associated reference data, which answers '**who is who**' and '**who owns who**' on a global basis.
- A machine-readable 20-character alphanumeric code.

.

Mapping to other identifiers

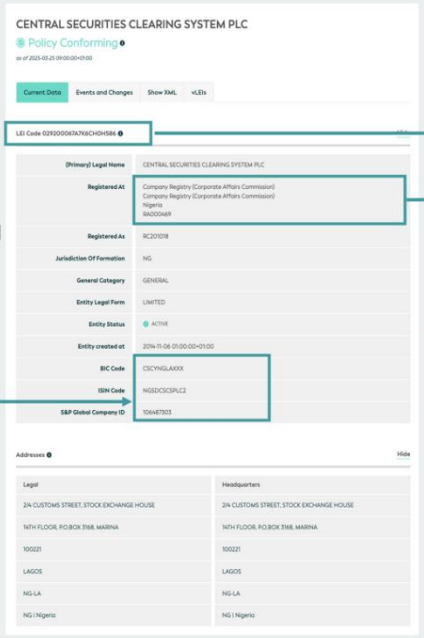

*029200067A7K6CH0H586

National Registration Identification information*

- Anyone can access LEI data through the GLEIF database.
- Each LEI is tied to details like a company's official name, registered address, and ownership structure.
- LEI reference data is validated against public registration authorities globally.

* GLEIF maintains a list of registration authorities globally.

Next, what do you need for your digital identity credentials?

- **Identification of the credential holder**
  - Identity Assurance
    - What level of assurance is needed
- Higher levels of assurance reduce identity fraud and impersonation but can affect speed and cost of credential issuance

Next, what do you need for your digital identity credentials?

- **Digital signing needs**
  - — Will you need the ability to support more than one signer, in whole or in part, on documents, data/information, forms, regulatory filings, etc.?
  - — Will you need to rely on a multi-signature capability for increased security for cases in which one signer should not be able to act alone (see also Security)?

Next, what do you need for your digital identity credentials?

- **Delegated authority**
  - — Will you need to use credentials to delegate authority?
    - Then chained credentials would support delegation of authority
    - This has implications both for natural person/citizen credentials (for example, guardianship) and for organizational credentials (for example, delegating authorization/entitlements within an organization)
    - And for delegating authority to AI Agents

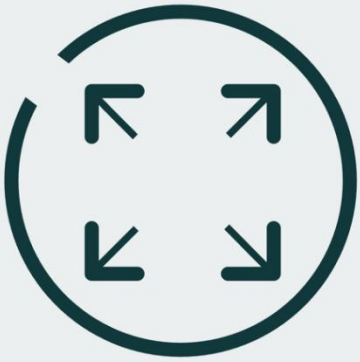

Next, what do you need for your digital identity credentials?

- **Credential lifecycle**
  - — Expiration dates – useful or limiting feature?
  - — Frequent credential replacement/renewal (see also Security and Cost)
  - — Protection of keys through Rotation/Pre-rotation (see also Security)
  - — Use of persistent identifiers tied to credentials benefits from rotation policy prophylactic or recovery

Next, what do you need for your digital identity credentials?

- **Privacy**
  - — Is there a need to release information gradually?
  - — According to certain agreed terms or conditions?
  - — Then implementing selective or graduated disclosure would satisfy these needs

Next, what do you need for your digital identity credentials?

- **Security**
  - Security first protects credentials preventing credential replacement churn (see also Credential life and Cost)
  - Key features to ensure this
    - Protection of keys through Rotation/Pre-rotation for the ability to recover keys and avoid credential revocation (see also Credential life)
    - A Root of Trust with linked credentials requiring multiple layers of infrastructure to be compromised
    - No requirement to rely on web 'security'
    - Higher levels of assurance reduce identity fraud and impersonation (see also Identification of the credential holder)

Next, what do you need for your digital identity credentials?

- **Cost**
- Easy, free of charge/reasonable cost access to entity data (also see Entity identification and data)
- Credential replacement churn (see also Credential life and Security)

Digital Tools Features Analysis

Digital Tools Features Matrix: Signing

Features	vLEIs as Authentic Chained Data Containers	First generation Verifiable Credentials	eSeals	Digital Certificates
Digital Signatures	✓	✓	✓	✓
Persistent Digital Signatures	✓	✓ only those VCs on a blockchain	✗	✗
Single Level Issuance	✓	✓	✓	✓
Delegable Authority/Multi-level Issuance	✓	✗	✗	✗
Non-repudiability	✓	✓	✓	✗
Signing logging	✓	✗	✗	✗
Signing in Full and in Part	✓	✗	✗	✗
Horizontally-scalable Signing Infrastructure	✓	✗	✗	✗

Digital Tools Features Matrix: Verification

Features	vLEIs as Authentic Chained Data Containers	First generation Verifiable Credentials	eSeals	Digital Certificates
Verifiable Provenance to a Global Root of Trust	✓	✗	✗	✗
Instant Revocation State Verification	✓	✓	✗	✗
Decentralized Revocation	✓	✓	✗	✗
Privacy-respecting Revocation	✓	✓	✗	✗
Revocation by Any Party within the Chain of Authority	✓	✗	✗	✗

Digital Tools Features Matrix: Security

Features	vLEIs as Authentic Chained Data Containers	First generation Verifiable Credentials	eSeals	Digital Certificates
Multi-signatures	✓	✗	✗	✗
Secure Custodial Key Management	✓	✗	✗	✗
Key Rotation	✓	✗	✗	✗
No Reliance on Web Security	✓	✗	✗	✗
Post-quantum proof	✓	✗	✗	✗
Zero Trust Architecture	✓	✗	✗	✗

Digital Tools Features Matrix: Global Applicability

Features	vLEIs as Authentic Chained Data Containers	First generation Verifiable Credentials	eSeals	Digital Certificates
Decentralized Authority	✓	✓	✗	✗
Globally Trusted Credentials	✓	✗	✗	✗
Global Root of Trust	✓	✗	✗	✗
Global Governance	✓	✗	✗	✗
International Standardization	✓	✓	✓	✓
Multiple Roots of Trust in a Single Ecosystem	✓	✗	✗	✗

Limitations

- This presentation contains confidential and proprietary information and/or trade secrets of the Global Legal Entity Identifier Foundation (GLEIF) and/or its affiliates, and is not to be published, reproduced, copied, or disclosed without the express written consent of Global Legal Entity Identifier Foundation.
- Global Legal Entity Identifier Foundation, the Global Legal Entity Identifier Foundation logo are service marks of Global Legal Entity Identifier Foundation.

Ecosystem Guidance Needs and Wants

Session Convener: Frederik Krogsdal Jacobsen
Session Notes Taker(s): Frederik Krogsdal Jacobsen

Tags / links to resources / technology discussed, related to this session:

https://openid.net/cg/ecosystem-support-community-group/

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The OpenID Foundation Ecosystem Community Group wants to help ecosystem operators (such as countries rolling out verifiable credential/wallet systems) make good decisions.
The purpose of this session is to gather feedback about what ecosystems need help with.

- How do we link different technologies?
- How do I make the choices I need to make in the technical specs?
- Do I need to register/authorize/authenticate relying parties and issuers?
- Layers:
  - Technical
  - Operational
  - Legal
  - Business
  - Social
- The social and business layers drive the adoption model
- You need governance across all of the layers
- SIDIHub perspective:
  - How do we describe ecosystem options to governance people?
  - How do we circulate trust?
  - How do we fix the gaps in the technical standards?
  - Interoperability between ecosystems
- We want to make sure lawyers and policymakers understand what is technically possible
- Can we automate governance?
  - It's difficult to fully do it because you quickly accidentally write rules that exclude people who were not represented
  - It's usually easier to do retrospective "punishment" than proactive enforcement
  - How do you measure the potential gains or harm from governance decision
- Maybe the social layer needs to have a different name
- Is there a hierarchy of the layers?
  - There is a difference in time scale
  - The layers can impact each other
- Social layer depends a lot on the expectations in the society
  - Privacy
  - UX
  - Do you have to pay for it?
  - Expectations change in different circumstances
  - Include activists/NGOs and try to convince them up front
  - When rolling out an initiate, consider measuring diverse communication
- Operational situation

    - What things do people have?
    - Which places or people do they trust?
- Advising policymakers
    - They often already have some opinions: they come from a particular mindset
    - Even if you make the same decision, you might have different reasons for doing so
    - Policymakers really want a “menu” of options they can pick from

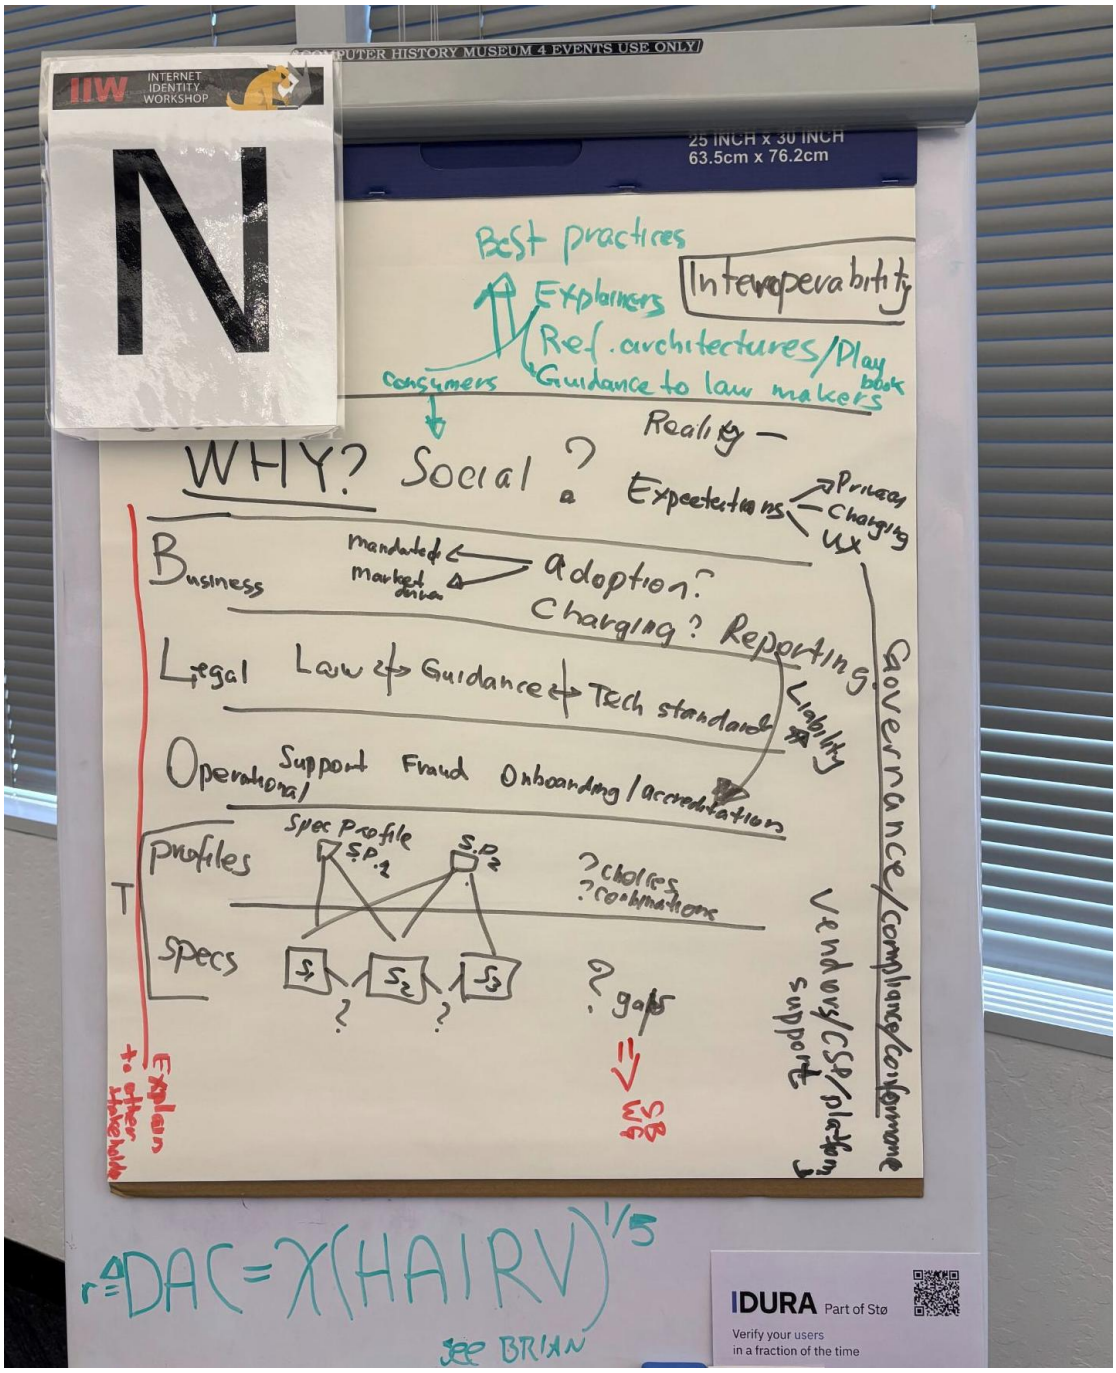

SESSION #7

Cross App Access - no more OAuth redirects

Session Convener: Aaron Parecki
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Governance-Backed Content Credentials

Session Convener: Scott Perry
Session Notes Taker(s): Scott Perry

Tags / links to resources / technology discussed, related to this session:

https://docs.google.com/presentation/d/1w2h7qA1GfyKBlCBhE1L3EyPEAv6BFfXA/edit?usp=drive_link&ouid=108677846726640960846&rtpof=true&sd=true

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The session discussed an approach to add a new assertion - the governance assertion to the Creator Assertions Working Group (CAWG) to allow a governing body the claim data on a a digital object to control the use of certain claims. They can be specific user identities, consent claims or metadata claims. This has been asked by different industry groups to affix attribution on specific digital objects like songs, movies, documents picture or from varying authoritative jurisdictions.

The Antidote to Fear: Digital Agency and Identity in the Wilderness of Technology

Session Convener: Swan Black
Session Notes Taker(s): Swan Black

Tags / links to resources / technology discussed, related to this session:

Recommended reading that came up, by session participants:
Lean Logic by David Fleming
Core concept: systems that consume their own underlying capital become brittle and collapse-prone. Parallels drawn to digital monoculture and extractive growth models. Companion: *Surviving the Future*.

Extremely Online by Taylor Lorenz
Documents how user behavior shaped the internet while value capture consolidated within platforms. Used to support discussion of agency extraction and system consolidation.

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Session framed the digital identity ecosystem using a living forest as a diagnostic model. Framing was used as a structural parallel, not metaphor. Focus on shared failure modes, health indicators, and long-standing human practices for tending complex systems.
Primary claim introduced: humans function as the apex keystone species of the digital ecosystem. When keystone function is constrained or removed, the system reorganizes around more extractive and simplified actors.
Core Discussion
Participants aligned on the usefulness of the forest model for making systemic issues legible across technical and non-technical perspectives.
Key conditions discussed:

- Keystone presence enables system-wide balance rather than control
- System behavior shifts rapidly when keystone function is blocked
- Extractive dynamics emerge in the absence of constraint and stewardship

Collaborative Mapping Exercise
Group mapped four quadrants:

- Thriving Forest
- Blighted Forest
- Thriving Identity Ecosystem
- Current Identity Ecosystem (Blighted)

Thriving Forest Characteristics Identified:

- Mycorrhizal networks
- Biodiversity and interdependence
- Active keystone species
- Regenerative capacity
- Coexistence of canopy and understory

Thriving Identity Ecosystem Parallels:

- Decentralized credentials
- Open standards
- User agency and control
- Portability of identity
- Multi-party ecosystems operating at scale

Blighted Forest Characteristics Identified:

- Invasive species dominance
- Keystone disruption or removal
- Monoculture conditions
- Soil depletion
- Failed or stalled regeneration
- Loss of canopy structure

Current Identity Ecosystem Parallels:

- Centralized control of identity
- Normalized identity theft and data exploitation
- Systems optimized for extraction at scale
- Lack of usable, human-centered tools
- Frictionless environments that obscure cost and consequence

Tending Practices (Forest → Identity Translation)

Participants identified tending practices as ongoing system work rather than one-time interventions.

Forest Practices:

- Long-term invasive management
- Root cause remediation
- Keystone protection and reintroduction
- Metered resource use

Identity Ecosystem Practices:

- Addressing issues at the protocol and infrastructure level
- Supporting protocol diversity
- Building privacy into economic and operational models
- Interoperability across systems
- Consent infrastructure
- Access infrastructure enabling real user agency

Additional Threads Raised

1. Transferability of the Model

Framework applies beyond identity systems. Participants noted relevance to:

- Healthcare data ecosystems
- Financial infrastructure
- Other domains exhibiting extractive scaling and keystone suppression

Discussion emphasized developing this as a repeatable mapping and diagnostic practice.

2. Narrative Mapping as a Skill

Effective translation requires alignment with audience intuition and prior knowledge.

Forest model identified as effective due to:

- Broad familiarity with ecological collapse and regeneration
- Lower baseline familiarity with identity infrastructure concepts (e.g., verifiable credentials)

Participants noted that each audience requires re-grounding in a familiar system to build understanding.

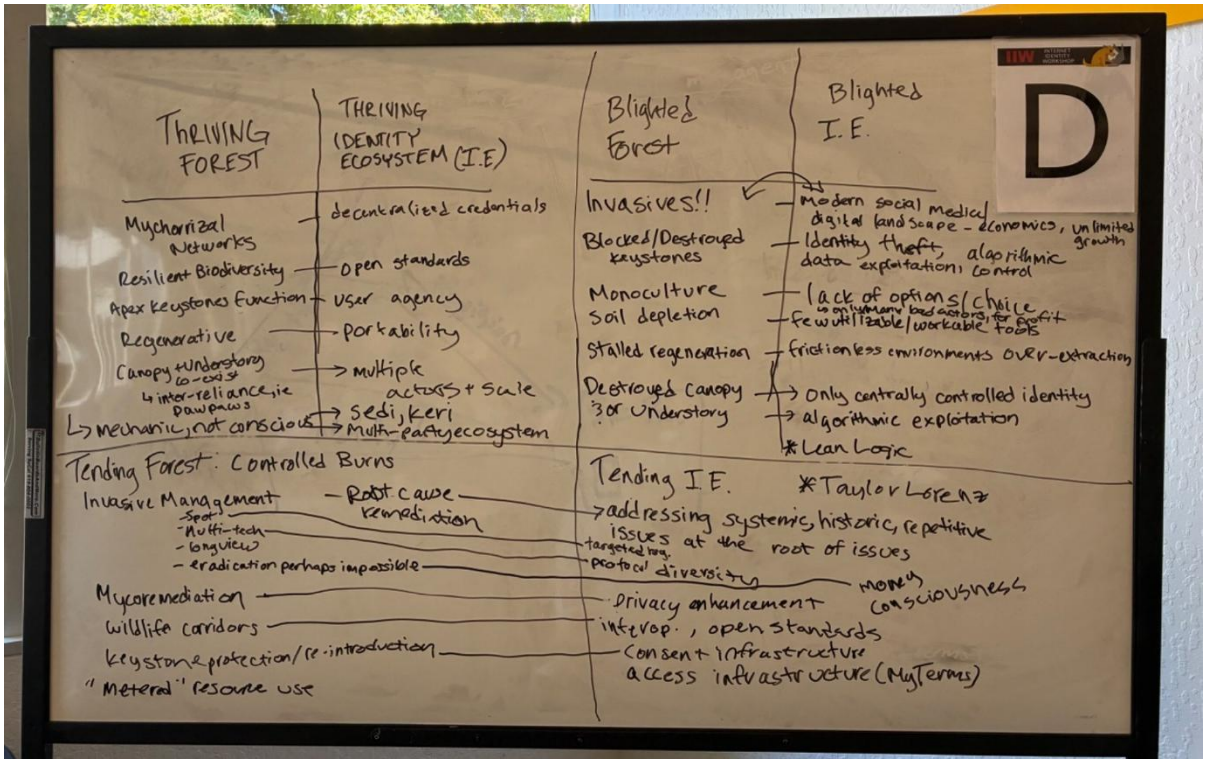

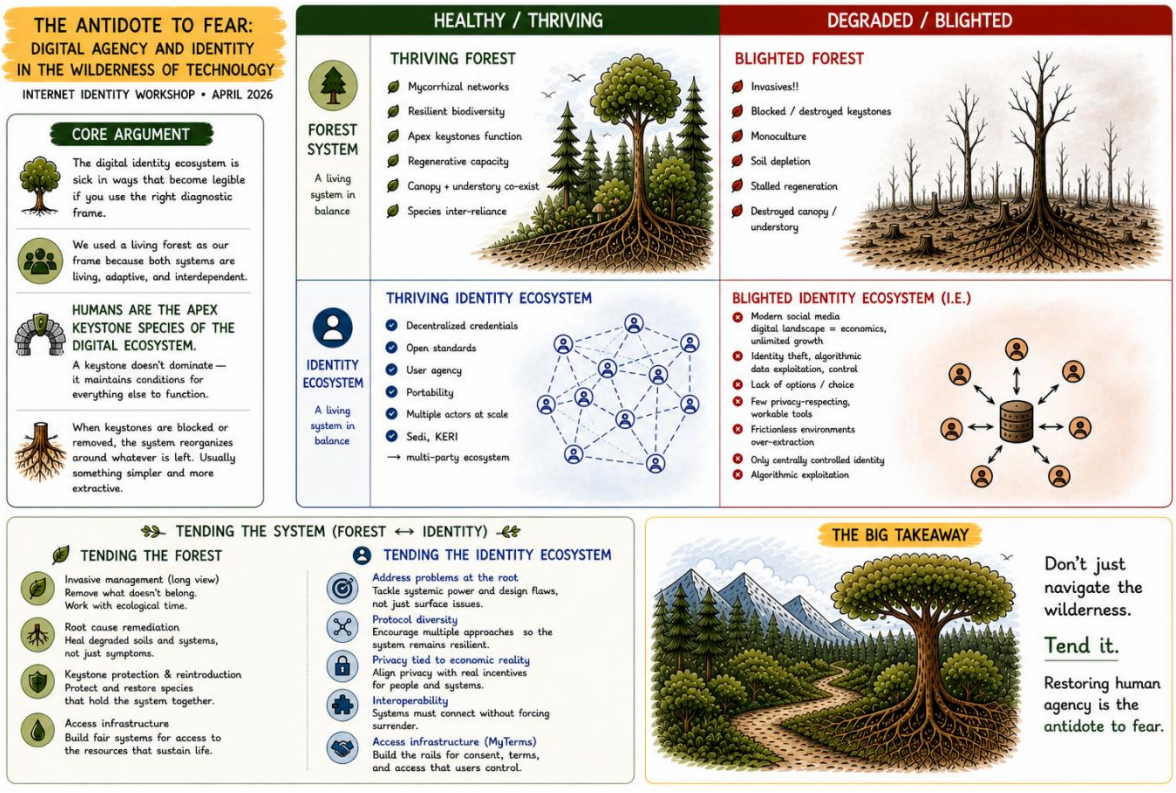

:

Image generated with ChatGPT based on notes and whiteboarding

Onboarding Enterprises to Decentralized Identity Systems

Session Convener: Adrian Ross (adrian.ross@oracle.com), Victor Carolino, Michael Kaufman
Session Notes Taker(s): Adrian R, Victor C

Tags / links to resources / technology discussed, related to this session:

Slide deck: Oracle: Onboarding Enterprises to Decentralized Identity Systems

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Oracle product development team hosted a discussion on the challenges associated with onboarding enterprise customers to decentralized identity systems. A deck on some of the approaches Oracle is exploring: OID4VC (onboarding end users) and TRQP (governance interop) is attached.

A future session on Oracle's OSS verifiable credential abstraction library and composable committed components PoC was previewed.

Keri Foundation Wallets

Session Convener: Evan Asakawa
Session Notes Taker(s): Evan Asakawa

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The KERI Foundation has created three open source wallets: WASM, mobile, and desktop. They allow for any contributors to create their own KERI infrastructure plugins. The KERI foundation has deployed some public infrastructure–witnesses and watchers services—that a user may connect to to get infra, and you can get these from multiple different providers.

The conversation then turned to AI safety, and how to manage agent's intentions and access to sensitive information (i.e., private keys). It was suggested that one can best prevent agents from performing nefarious actions by managing its training data (no bad actions present in training data) & requiring it to be open and observable.

The crazy road of Authorization ! Are we really ready for agents ?

Session Convener: Bhavna Bhatnagar
Session Notes Taker(s): Bhavna Bhatnagar

Tags / links to resources / technology discussed, related to this session:

slides at : https://docs.google.com/presentation/d/e/2PACX-1vRLV5NuR3c_n2ZDwoD7WpWu0cu8mdTFiN2pDZw0g6NGI0VQ2gwUi4e26L4aPuAVHxhptz_n6Ih_OAhm/pub?start=false&loop=false&delayms=3000

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

I covered the road of Authz journey over 25 years PDP (static roles etc) - > Dynamic context (ReBAC, ABAC) -> agentic craze with bumps and tool booths along the way
We all discussed our challenges during the journey be in static roles, role explosion problems or not enough static policies to cover the requirement
Dynamic context shift (ReBAC, ABAC) - but not enough to solve the current cloud workloads
Policy as a code era (OPA, Rego, Cedar)
Agents are kinda crazy they do what they plan, we dont know the plan (non-deterministic) and they need permissions on the fly
Future challenges and need for the agents

-
- Intent based auth
- Context sharing amongst apps
- Governance
- Autonomous agents with (human in the loop)

We discussed Use cases in finance space, access/misuse of privileges and also as in the slides
Then we went philosophical about do we really need agents ? Who are the agents ? Are they moral compass for us or we tell them what to do ?
How does one give feedback and train the agents in agentic AI world ?

DBSC Secured Cookies W3C Standard

Session Convener: Lucas Santos
Session Notes Taker(s): Bryce Frey, Guillaume Ehinger

Tags / links to resources / technology discussed, related to this session:

Slides (Public): https://github.com/lucasrsant/dbsc-sso/tree/main/presentations

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Q: What form are the infostealers?

- they can be malware extensions or desktop malware

Key attestation is a special kind of key that only shows that the same two keys come from the same TPM. There is no mathematical tying here.

'remote attestation' means that we are checking that the TPM is a real TPM.

- Chrome does not do this because it would enable a tracking mechanism for RPs.
- Q: But the attestation only signs the key pairs if they were both generated in the same TPM?
    - Yes

Q: Define RP. Is this eLTD+1?

- tied to eLTD+1

Q: RPs will want to set its own restrictions on the cookie.

Questions from John B. at Yubico
Q: haven't seen conversations started with the WebAuthn group. It seems obvious to include the public key in the collected client data. That would stop many MitM attacks.

- We only DBSC register the session after the user's authentication is done.
    - that sounds like a design mistake
        - A lot of TPMs are flaky, and you don't want to block authentication because the TPM is slow or flaky
        - The DBSC registration will happen in the background
        - right now, the TPM issue might not be a big issue with only Google doing DBSC
        - Once lots of sites are using DBSC, this will really slow down the TPM
        - Malware onsite during registration is out-of-scope for many reasons in the generic world (different in Enterprise world), such as the malware may fake the TPM not existing, etc.
        - in Enterprise:

- Q: there is a WebAuthn ext. to support cross-signing at a much faster rate than TPMs. We should look at that.
    - being designed for verifiable credential wallets, and has an option to not require UV or UP at every signature.
    - This extension is currently being worked on now.
        - Q can we tie the DBSC registration to the key that is doing the session authentication
- We've also talked with Microsoft about VPS for faster signing support. but the adoption level is even lower than TPM adoption.
- For enterprises, it is up to the enterprise to decide whether this is allowable.

Q: Safari & Firefox?

- not implemented today. Today only implemented in Chromium. Hoping Firefox can be done next year. Safari is ???

Q: are there non-Google consumer use-cases that have tried it out?

- Okta and Snapchat have tested it during the Origin Trial.

Q: it seems that there will be a long-tail adoption rate. How do we avoid downgrade attacks?

- the only 'quiet' downgrade attack can happen at the DBSC registration time.
- post-registration, malware would have to clear the cookies, which will clear the user's session cookies, and thus be very 'loud' to the user.

Q: what are the recommendations to the RPs to reduce these issues?

- DBSC requires at least two cookies….
- RPs can also put the user in a 'lower level state'

Q: Might malware just try and prevent DBSC from happening?

- malware in the wild is now trying to move to persistent presence on the device. We haven't observed malware trying to kill the DBSC session.

Comment: I love that the browser blocks all other requests while waiting for the DBSC cookie.

- this can lead to a deadlock situation. Currently we are working on fixing this in the protocol. The fix is that the RP can configure some URLs from being blocked during refresh.

Q: how do you handle the Monday morning problem of 200+ Gmail tabs opening at the same time?

- The browser handles the cookie refresh and will only make one request for that eTLD

Comment: additional praise.

Rebooting Enterprise Information System for Peer Production

Session Convener: Brent Shambaugh
Session Notes Taker(s): Brent Shambaugh

Tags / links to resources / technology discussed, related to this session:

https://bshambaugh.org/eispp/ , https://github.com/bshambaugh/eispp,

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Nobody showed up. This is not surprising since the convener arrived late and missed announcements in the opening circle (well barely, it could have been fit in at the end, but this was realized after it ended) . The session was put on the board though. A new attempt will be made tomorrow. This is for modernizing this concept with the latest tech to get it out the door.

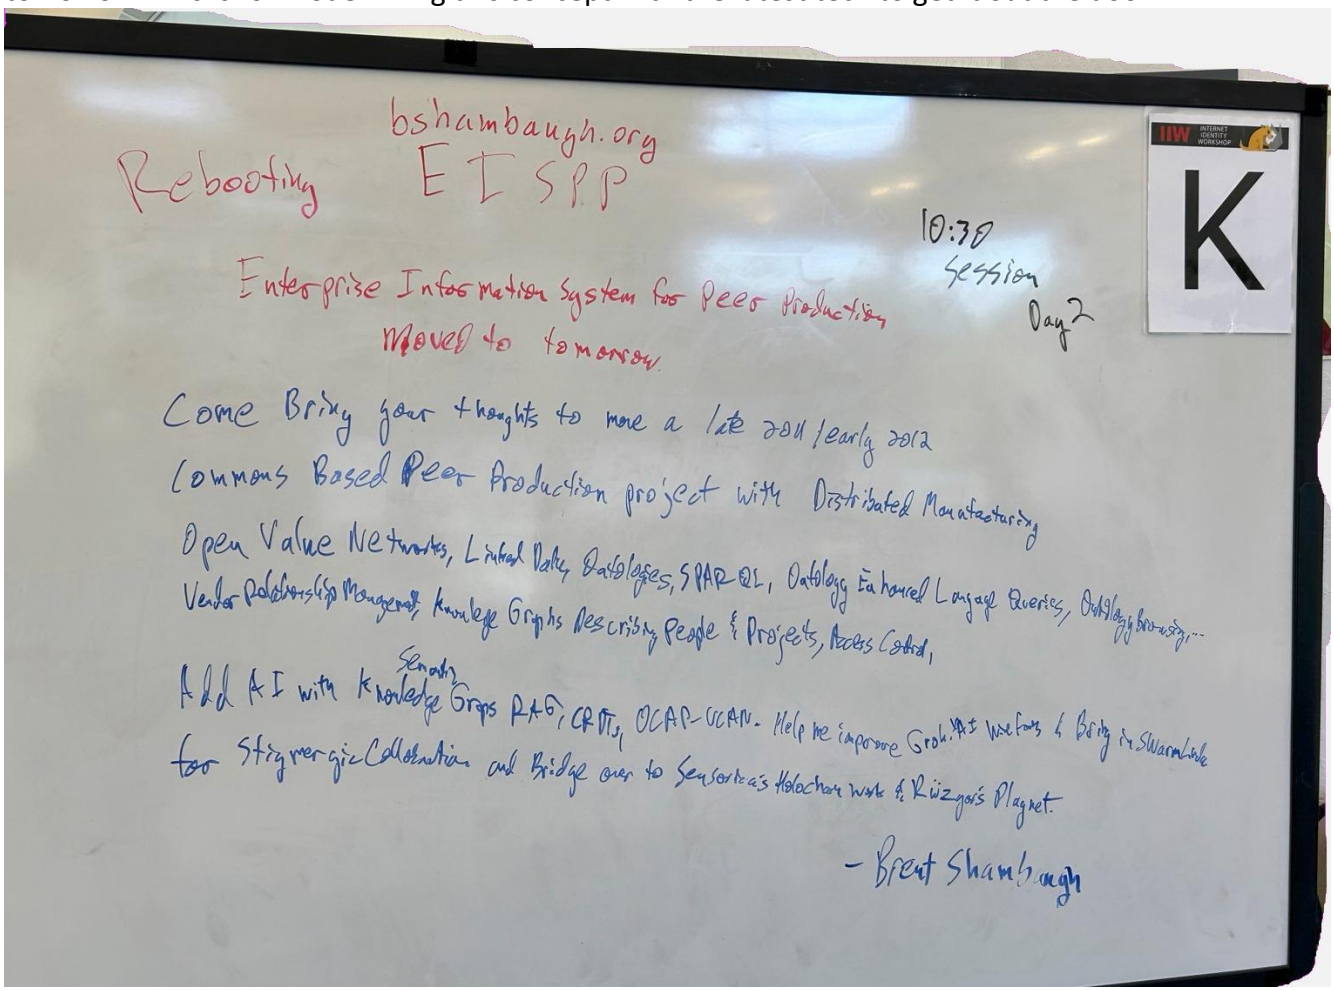

Lack of Interoperability of ID standards

Session Convener: Kevin and Frank
Session Notes Taker(s): Brian von Herzen

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

if you don't have soc sec number then you don't get FEMA assistance.
any type of short-term recovery resource from the government from government and what happens
no progress in decades -

lack of interoperability of ID

red hook- long term neglect -
public works -

Minecraft modeling engine- data from worlds- and convert over for disaster resilience-

permcaulture design- wide-scale design -
during disasters, so hard to do- offline and shut down
reduced executive function
takes lots of training to keep running in disaster -

communities- determine where exactly to grow the food- soil - brownfield-
rewilding landscapes and mindscapes -

fragmentation across ecosystems - protocols and reps.
bigger problems more and more issueers-
own document spec- what is a credential -

*** oswap for identification(s) --****

communications protocols - actual content-

measuring
manage documents- and question of interop- to a verifies
trust over IP- opinionated view of protocol stack

decentralized identity foundatoib- build things and figure out how they woark together .

trust over IP-
protocol interoperatbility problem -
non-interop- gain a network effect -
runs on OICD rails- hopes to decentralize and extend -
OpenID Connect (OIDC)
DMITRI implements - interop- interest- bridging where possible -

IDPRO: What should we be adding to our body of knowledge?

Session Convener: Elizabeth G + Tina Srivastrava
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

IDPro® Body of Knowledge
Educational Resources - IDPro
CIDPRO® Reference Resources - IDPro

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Rohit's admitted hobbyhorse is promoting “GovOps” — Governance as a DevOps-like discipline. It might be what comes after "Identity Governance Administration," particularly as we extend beyond-human identities. For example, SPIFEE/SPIRE or MCP haven't yet been in the BOK (AFAIK).

In that context, he's curious how should the BOK accommodate innovative new thinking: like the experimental work on AI in IAM, emerging agent identity protocols, or emerging new AuthZ policy languages?

Is it really a very different professional group, because GRC (Governance Risk and Compliance) professionals aren't necessarily traditionally represented at IDPro.

I haven't gotten a better short story than this one yet: “Interns can't access real patient data”? (Which might already be wrong: *medical* interns, absolutely have to talk to patient data!) I meant like a SaaS engineering shop where interns shouldn’t have access to production airline traveler data… then following the responsibility from the IdP to locking them out of an app, a folder, a bucket, auditing the exceptions log… is that sort of compliance activity even a problem the BOK should address?
—

Signed credential metadata updates and payments

Session Convener: Frederik Krogsdal Jacobsen
Session Notes Taker(s): Frederik Krogsdal Jacobsen

Tags / links to resources / technology discussed, related to this session:

https://openid.net/specs/openid-4-verifiable-credential-issuance-1_0.html

https://identity.foundation/didwebvh/v1.0/

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The purpose of this session is to figure out what is necessary for credential metadata updates in OpenID4VCI while respecting the requirements needed. In particular, payments has specific requirements about history of display metadata.

There are levels to metadata, and the issuer and credential instance can both influence e.g. the display of data. The credential instance metadata overrides the issuer data.

There is a need for versioning of metadata. There are some options for how to do this: hash the content, version each metadata separately, or version the "combined" result of metadata after applying overrides.

DID:webVH might have some sources of inspiration, as they attempt to solve basically the same problem.

For payments, some people believe they need signed metadata, but this might not actually be necessary if there are versions.

We also need an update mechanism that makes sense with the versioning and specifies how metadata is updated on new versions. It might be that each overriding "layer" should require that the "next layer down" is on a specific version to make sense.

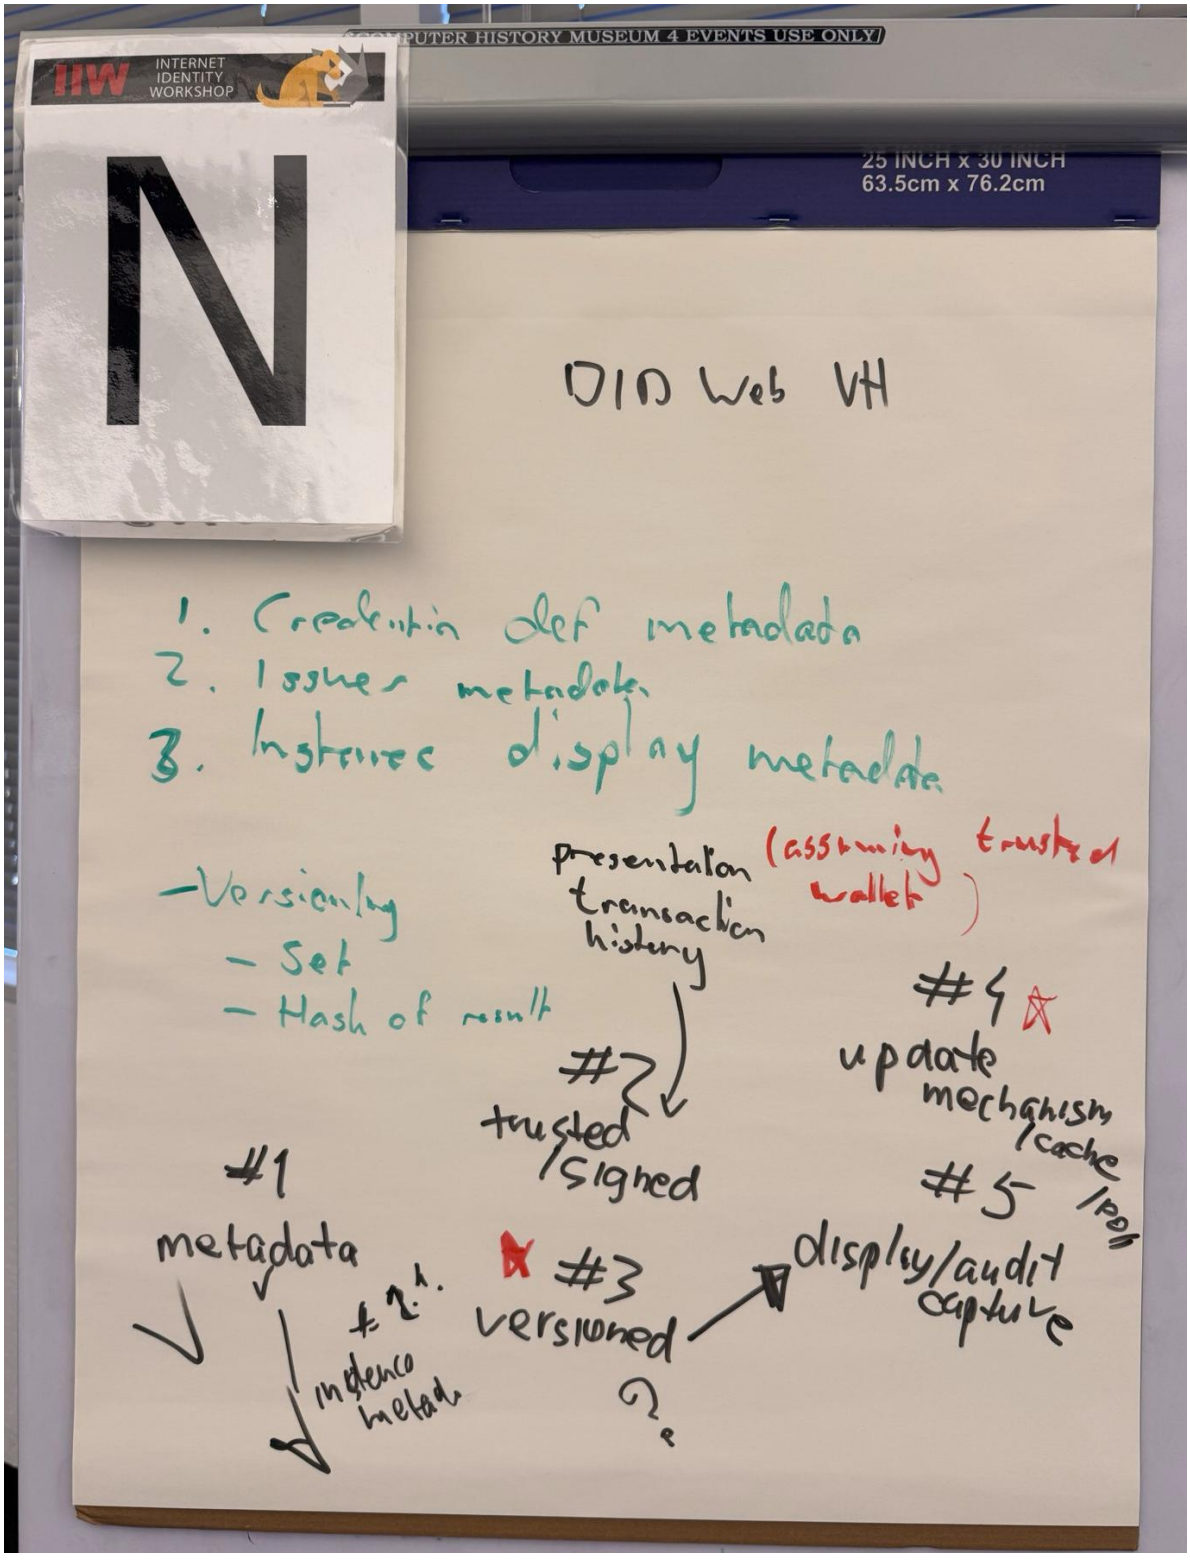
IIW
INTERNET IDENTITY WORKSHOP
N
COMPUTER HISTORY MUSEUM 4 EVENTS USE ONLY
25 INCH x 30 INCH
63.5cm x 76.2cm
DID Web VH
1. Creation def metadata
2. Issuer metadata
3. Instance display metadata
-Versioning
- Set
- Hash of result
presentation
transaction
history
(assuming trusted
wallet)
#4
update
mechanism
cache
/poll
#2
trusted
/signed
#1
metadata
#3
versioned
#5
display/audit
capture
instance
metadata
?

SESSION #8

OAUTH fo MCP - How MCP has reshaped OAUTH in the past year

Session Convener: Aaron Parecki
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

DigaVouch: Building Consent Gated Identity for the "Grey Zone"

Session Convener: Juan
Session Notes Taker(s): Juan

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

This session is an invitation to discuss a current implementation of an identity trust platform that is used in the context of real estate multi-family leasing. DigaVouch works to close the gap of verifying the identity of lease applicants by providing applicant opt-in data to verify identity.
We discussed different use cases a solution could address, and what overlaps and differences are in the approaches that others have used to solve similar problems in other contexts - example helping recover identity for unhoused or those fleeing domestic violence situation.
There was a discussion of the double edge sword nature of this type of system to be aware of how data is managed and retained. What are the legal implications of data stewardship and ensuring that we consider that even having a signal of data is data itself and what is the right thing to do with that.
THere was discussion of Trust Graphs and possible ecosystem of trusted communities that can verify/vouch relationship and identity of the person.
It was a productive, informal discussion of a wide array of considerations when building this type of solution.

Delegate SD-JWT

Session Convener: Gareth
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Keeping The Faith

Session Convener: Eric Welton
Session Notes Taker(s): Eric Welton

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The thesis was presented that communities of faith represented a third qualitative type of voice to complement the sociopathic voice of corporations and the dysfunctional, legalistic voice of government in the conversation addressing the ethical needs of people navigating the information streams of the modern world. People, from adults to children and their parents, consume enormous amounts of information while navigating daily decisions. Information is often inappropriately filtered, it is filtered by corporations, driven by sociopathic needs, and by governments driven by rules that lag specfic concerns, often are poor fits with specific communities, directly hostile to others, and are in the pockets of sociopathic corporations.

Increasingly, engineers are held to account for being ethically responsible for the technology they develop, despite being ill positioned and almost powerless in this role.

When discussing the possible role of communities of faith, the conversation immediately centered on how to form a coalition of common ground and focus on uniting the communities of faith to push for better enforcement of laws that were on the books about common sets of core ideals which are presumed to be universal, such as the need to keep children safe from pornography and the belief that there is generally a good enough definition of the appropriate age for child and a generally good enough definition of what constitutes pornographic material, and what should be done is to build stronger networks of faith communities to focus on getting lawmakers to enforce the laws on the books.

Dissent for this position focused on the note that the CAN-SPAM act has had little impact in the three decades it has been in effect and that the age verification and other laws are not very well enforceable. There is agreement that communities of faith may be helpful in driving age verification and other surveillance technologies into web sites, but this also poses risks of getting out of control.

A genuine tension was revealed between the 'common core' model and the question of whether or not there was much value in involving communities of faith beyond what is currently done, in that people can make faith based action in providing goods and services. An example of an OB-GYN who choose not to support transgender patients was discussed. Additionally discussed was whether or not labeling of tuna and was relevant to decisions, as some brands of tuna were known to be environmentally questionable. This raised the distinction between faith-communities vs. cause-based communities.

What was decided is that most people just do not feel strong ethical weight when making decisions and thus the need for ethical guidance is not present. In essence, people just don't care. Focusing on a common moral core that 'everyone can agree upon' seemed to be the consensus of the group, and the depth and breadth of deviation from that core was largely seen as essentially the sort of thing that should be the province of government legislation, and that communities of faith should essentially be left to act independently and remain out of the social fabric - consistent with the separation principles of the first amendment to the U.S. Constitution.

The conversation was braided throughout with swings into and out of the role of identity and credential tech, in the form of filtering, attestations, agent integration, and other automation - through core social philosophy, and back.

The Convener & Notes Taker observes that this is substantively different from the outcome expected - however, that is the wonderful thing about IIW - it is a vibrant and enthusiastic and lively community where respectful, divergent and exciting conversations.

What is computer science? (there is only one right answer)

Session Convener: Sol Ashlynn
Session Notes Taker(s): Sol Ashlynn

Tags / links to resources / technology discussed, related to this session:

#computer science

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Answers on the Left:
Definition 1: Study of logic gates: on and off tree. (what computer science does)
Definition 2: Linguistics of machines talking (what computer science is)
stochastic
random
nondeterministic
Definition 3: Historically it is the management of human computers (who were often women).
Definition 4: Applied discovery of how humans use information in machines.
Definition 5: Science of Information of Humans.
Definitions 6: whoever teaches it defines it by how and what they teach.
*** Whoever gets to ask the question gets to frame the answer****
Definition 7: optimizing information theory.

What is a computer?
-Machine
-use of nonhuman tool to process information

Science:
-a methodology involving observation in the world, hypothesis, curiosity, controlled experiment with a -repeatable result in nonspecified order.
-progresses one funeral at a time.
General questions:
Is computing operating in a controlled environment?
-evolving definition?
- outsourcing thinking deskills by displacing our thinking.
Definition 8: Is an iteration of textile arts
    -optimizing is ugly
    -tapestry of threads of thoughts
Definition 9: Science of thinking
Definition 10: the scaling and extending of our thinking.

THE RIGHT ANSWER (the page spatially to the right of the left answers)

What is ethical computer science?
Uno reverse:
computation with agency.

Cross Device Flows - Hybrid CTAP / Passkeys / Digital Creds

Session Convener: Harsh Lal & Mohamed Amir
Session Notes Taker(s): Harsh Lal

Tags / links to resources / technology discussed, related to this session:

https://fidoalliance.org/specs/fido-v2.3-ps-20260226/fido-client-to-authenticator-protocol-v2.3-ps-20260226.html#sctn-hybrid

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

- Hybrid protocol was introduced.
- The speakers went into high level overview of it setp, protocol structure and current uses.
- Protocol structure includes 3 parts:
    - Initiation
    - proximity check and handshake key exchange
    - data transmission
- The limitation of current hybrid protocol was introduced:
    - connectivity issue (needs internet)
    - qr scanning inconsistency
    - tunnel server data limits
    - tunnel server privacy concerns
    - ble spoofing concerns
- Speaker discussed the evolution of Hybrid protocol into a new specification to address these
- Some things that are being worked on:
    - L2CAP CoC BLE offline channel
    - NFC invocation
    - NFC for proximity
    - WiFi Direct for data transfer
    - Fallback URL support
    - UWB for proximity
- Speaker provide reasoning for moving Hybrid to a new specification
    - Technological Independence
    - Reference clarify for other specs
    - Clarifying technical vs user terms
    - New use cases that span beyond CTAP protocol
- A few demos related to offline BLE channel and NFC invocation was shown
    - Dealing with digital credential presentation
    - Dealing with passkey assertions
    - Dealing with presenting verified phone number cross device
    - DEaling with digital credential issuance
- Some benefits for digital credentials were discussed. PXP moves most of the heavy lifting to platforms - and wallets gets features for free.

- Role os using UWB ranging to address BLE advertisement spoofing concerns were discussed
- It is strategically important to address the cross device presentation gap for the ecosystem. Hybrid protocol provides a standard solution to address this.

Agent Registry for Identity and Authorization

Session Convener: Aaron Grego
Session Notes Taker(s): Adolfo Grego

Tags / links to resources / technology discussed, related to this session:

https://aria.bar and https://trustlayer.foundation

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

We presented Trustlayer Foundation, a non-profit based in Mexico, which developed an open protocol (governed by a community) and cross-organizational registry service to issue (people and organizations) verified identity credentials for their AI agents and provide methods for all receivers to verify and authorize the agents actions and intents, using existing standards like DIDs and VCs

We discussed complementary projects and opened up the possibility of operating jointly with multiple registry operators (similar to ours), as well as universal recognition of multiple wallets and compatibility with other DID methods.

We briefly discussed the concept of an "agent", which led us to refine the definition to include "any entity interacting with a system through a verified credential" , which allows to differentiate an agent from any other way of interacting with software.

A question was raised whether we as a registry should record all agent interactions, which requires consideration to become an "optional" later if service in the future.

We invited attendees to see the Registry system operating live and demonstrate the creation and assignment of authorization parameters to an AI agent, which can be presented tomorrow during the day.

We maintain an open invitation for collaboration and participation with the community, people and organizations.

We wish to thank everybody who joined this session!

Intros & Updates on the Japanese version of NIST's SP 800-63 (Digital Identity Guidelines)

Session Convener: Ryo Nakashima
Session Notes Taker(s):

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

- Are the assurance levels of DS-511 equivalent (comparable) to the SP 800-63? The levels depend on lots of assumptions
    - -> We did a mapping exercise of DS-500 and SP800-63-3, but not DS-511 and SP800-63-4 yet. We do have a regular conversation with NIST, and we gave information about DS-511/512 to them
- A bit confusing of stuff about "verification of the applicant" being treated as one of the process of identity proofing (bc it's in 63B in SP800-63), although there is no reason you should separate in the same way as SP800-63
- Does "PIN" mean IC thing?
    - -> IC and also authentication on mobile
- Does the JPKI work in mTLS?
    - -> It can. We note in DS-512 that it has to do so to be phishing-resistant
- Digital Authentication App (slide): the RP should be the government server, not client app
- Doesn't Digital Authentication App collect personal info, like when someone visits an adult site wouldn't their age be disclosed to the app?
    - -> if it only uses the certificate for authentication-use then it cannot retrieve data (e.g., age) when accessed
    - -> Signature creation/verification service (using the certificate for signature-use) is not intended to be used in business field (for public-use only)
- Is the comparison of face done by automated detection (e.g., via features) or only visually?
- In SP800-63 passkeys are not AAL3 unles s it's device-bound (non-exportable)
    - -> In DS-511/512 it's not (non-exportability is not included in the requirement)
- It is reasonable to have only FAL2 bc FAL1/3 are both for edge-cases)
- Do you have any certification framework in Japan?
    - -> There is one for verifiers but not for IdP, wallets

SEDI: The Missing Foundation for Digital Identity

Session Convener: Phil Windley
Session Notes Taker(s): Jin Wen

Tags / links to resources / technology discussed, related to this session:

- SEDI — Utah's State Endorsed Digital Identity (Utah SB 275, SEDI Program Amendments — passed unanimously, effective 2026-05-06)
- Phil Windley — https://www.windley.com
- Joe Andrieu — functional definition of identity ("recognize, remember, react to other entities")
- MyTerms (Joe Andrieu) — rails on which the duty of loyalty rides
- First Person Identity (Phil's preferred framing vs. SSI / decentralized identity)
- ACDC — Authentic Chained Data Containers (selective + graduated disclosure)
- GLEIF — Global Legal Entity Identifier Foundation; LEI / vLEI for organizations
- DID — decentralized identifier (holder-created, state-signed under SEDI)
- mTLS / Netscape client certificates (1995) — historical precedent for client-side identity
- Aadhaar (India) — example of national digital ID model
- Trust frameworks — Visa, Mastercard, Discover, AMEX as the prevailing analog
- Larry Lessig — Code is Law — Phil's qualified take: code is law only for things code can govern
- "The Trust Test" / proof gap article (author name garbled in auto-transcript — possibly Sandeep / Santosh Bhandari; needs verification)
- Utah Privacy Office — https://privacy.utah.gov/wp-content/uploads/SEDI_ProtectingLiberty.pdf
- ACLU coverage — https://www.aclu.org/news/privacy-technology/utah-digital-id-law

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The thesis

Every digital identity system is trying to solve the **proximity problem**: in person, identity is tacit; in the digital world we are separated, so it must be reinvented. Every model so far — centralized, federated, mDL, national ID — leaves **someone in the middle**. The technical stack for first-person identity is largely done; **adoption is what's missing**.

Two unsolved problems

1. **Scaling trust** across thousands of credential pipelines, each with its own framework.
2. **Paying for proof** — Netscape solved client identity in 1995 with mTLS, but no one would pay $30–40 for a client cert and no CA would absorb the liability. Server certs went free; client certs died.

SEDI solves both

Utah's State Endorsed Digital Identity (SB 275, passed unanimously, effective 2026-05-06):

- The state **does not give you an identity or an identifier** — *"identity is innate to the individual's existence and independent of the state, fundamental and inalienable."* You bring your own DID; the state proofs and binds attributes to it.
- Drummond Reed: *"This is the revolutionary thing — they are not giving you an identifier."* Think LEI, but for people.
- **Endorsement, not control.** Like a notary: attests, doesn't enter the agreement.
- The state already has the proofing apparatus (DLD offices, trained staff) and **sovereign immunity** absorbs the liability private CAs couldn't.

What the statute guarantees

- **Bill of rights** (legal standing, not aspiration); only **3 enumerated revocation conditions**, one of which is the holder's own request.
- **Selective disclosure** required.
- **Open standards** required.
- **Choice of wallet** — state issues no wallet.
- **Duty of loyalty** — fiduciary obligation on anyone who accepts your SEDI; **MyTerms** (Joe Andrieu) provides the rails.
- **Digital guardianship** for minors / incapacitated.
- SEDI is **optional** — you can refuse it.

Architecture — one foundation, many ecosystems

- **Individuals**: personal digital identifier → SEDI credential → (government credentials, employer/school credentials, professional licenses, prescriptions, personal delegation credentials for AI agents).
- **Organizations**: organization digital identifier → State Dept of Commerce endorsement (and GLEIF / vLEI for international) → business licenses, banking, signing authority, proof of insurance.
- **Reputation by reference**: credential chains as trust infrastructure — e.g., a diploma anchored on the holder's SEDI on one side and the issuing university's Dept-of-Commerce/GLEIF + accrediting body on the other.
- **Reputation by observation**: receipts as decentralized "letters of reference" — but this is *not* you becoming your own credit bureau (negative reporting still needs a separate solution).
- **AI agents**: same delegation chains tie agents back to the human or org that authorized them.

Code is law — qualified

Phil's twist on Lessig: code is law only for things code can govern. Identity needs governance for: *who runs this, by what authority, on what legal basis, when can it be revoked, what duty does the recipient owe?* — none of which can live in code alone.

State of adoption

- **10 states** sent representatives to the recent SEDI Summit (Arizona named); a multistate consortium is forming. Utah is the only state that has passed it.
- **Cross-state digital reciprocity does not exist yet** — open work for the consortium. EU interop is also a goal.

Q&A highlights

- **Wallet?** Holder's choice — separate session later in the day.
- **Address change?** Existing DLD statutes (10-day update); identifier sits below endorsements, no government kill switch.
- **State retains data?** Underlying documents and biometric retained for fraud and adjudication.
- **Move states?** Keep your identifier and relationships; jurisdiction-bound credentials lapse.
- **Unlinkability / privacy?** A whole separate hour; identifiers are inherently correlatable, **duty of loyalty** is the main lever — code alone cannot solve it.

Closing

> *"You are responsible for building the world you want to live in. Don't take jobs to build something you don't want to have happen. We are at a point where we have the opportunity to give people a place to stand in the digital world. If we don't build these systems, somebody else will — and they won't have the values we want them to have."*

Reference: https://www.windley.com

IEEE 754 Floating Point Determinism

Session Convener: Mark Lenhardt
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

The Subject is a Couple!

Session Convener: Atul Tulshibagwale
Session Notes Taker(s): Atul Tulshibagwale

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

- (George) We need to start thinking about what to do when the delegate and the delegator are not single entities, but groups of entities.
    - The resulting token may be simpler, but the actual delegation bounds might be more complex. E.g., a doctor writing a prescription for a patient for a highly restricted drug, to be administered by a registered nurse.
- (Rohit) This is related to access justification
    - The subject might be anyone who needs to know
- (George) Delegating authority is not referencing a single subject.
    - There's no implementation today, but we need to be thinking about them.
    - E.g., in the DADE perspective, it could be a specific instruction about what to do (unanimous, 51% majority of heirs, etc.)
- (George) Delegation is within the bounds of a relationship. Authority is associated with the relationship. E.g., a healthcare power of attorney provides the relationship between the delegate and the delegator
    - Implementation-wise, this could be a verifiable credential
    - This could result in obligations, such as "do not resuscitate" or "do it once"
- (Mike Schwartz) You're living in a fantasy world since we don't know who the person is because we're a computer and we have no physical recognition of the person.
    - We don't know who the principle is
- (Mike Schwartz) Can we trust the OAuth token? We only have a bunch of evidence from a set of trusted sources.
    - E.g., is the user allowed to drive a car based on the token

    - You can stick the principal in the context as “token.idtoken.subject” etc.
- (George) There is value in having the client identify who it believes the principals are. Then the authorization server can validate whether those things are true or not
- (Mike Schwartz) That leads you to a failed system like SailPoint
- (George) you need to know what principal the system is acting on behalf of
- (MS) All of that could be true, without putting it into the principal of the token
    - From a governance standpoint, the principal is a distraction
    - Asking “what I’m allowed to do” is a failed approach
- (George) Pushing logic into the PDP that determines access based on who the principal is, has led to the ability of the PDP to make a wrong decision
- (George) There could be contradictory tokens.
- (MS)
- (George) It might not be a n “OBO”, but it could be that George is one of the principals.
- (MS) When we designed Cedarling, we tried to map the ID Token to a workload or a human, but it failed because it’s not a couple, it’s an array. When its an array, the policies are difficult to implement in terms of a principal. It’s easier to make policy decisions based on tokens.
- (MS) To have governance, you always need to know the organizational identity, the attestations are important, etc. So it is always multiple principals.
- (MS) We’re going to see an explosion of tokens out there.
- (George), There are other relationships, because at the time of determining policy, the principal’s consent to policy might or might not matter.
    - (George) Having multiple tokens is valuable
- (MS) But it's not talking to “George”, because its a machine, not a human.
- (George) When an LLM calls an MCP server, e.g. Google Calendar, it can ask for a token, but we don’t have a good way of expressing “I’m this instance of this class of LLM, and I want to operate on behalf of George”
- (George) When the agent requests a token from an AS, it requests a token on behalf of a user in accordance with a mission / mandate / charter, etc.
- (George) Could you not create a “sub_id”, in which you have a subject, an agent, and the relationship between them.
- (George) Where Karl is coming from, is that the OAuth entity type is that.
- (Rohit) If the primary field in the token is not a human, then we have a liability problem. Some agents are sending signed HTTP headers.
- (MS) Look at Clawdrey Hepburn. Are each of the “mini lobsters” that Clawdrey spawns their own agents?
- (George) This is the “sub_profile” in Karl’s proposal
- (MS)
- (George) In AOL we embedded the user id in the cookie. If the underlying infrastructure didn’t understand the identifier, it would not work.
- (George) In TraTs, you establish the subject at the beginning. You can identify many workloads down, who the principal is, but you also need the instance identifier of the code to identify which particular agent is operating.
- (George) In Clawdrey’s case, when a sub-agent is spawned, the identity artifact of the subject embeds the lineage.

    - I don’t like putting a whole lot of tokens in the “bag”, since it doesn’t capture the intent
- (Alex) If a request shows up to a PDP, if there are a whole lot of tokens, it complicates policy
- (George) In the short-term, the work Karl is doing is a step in the right direction.
- (MS) You would need to have DPoP proof of each component of the principal
- (George) there’s definitely value in DPoP, because at the end of the day, all of this goes into risk management

- …
- (George) We need to start thinking about what to do when the delegate and the delegator are not single entities, but groups of entities.
    - The resulting token may be simpler, but the actual delegation bounds might be more complex. E.g., a doctor writing a prescription for a patient for a highly restricted drug, to be administered by a registered nurse.
- (Rohit) This is related to access justification
    - The subject might be anyone who needs to know
- (George) Delegating authority is not referencing a single subject.
    - There’s no implementation today, but we need to be thinking about them.
    - E.g., in the DADE perspective, it could be a specific instruction about what to do (unanimous, 51% majority of heirs, etc.)
- (George) Delegation is within the bounds of a relationship. Authority is associated with the relationship. E.g., a healthcare power of attorney provides the relationship between the delegate and the delegator
    - Implementation-wise, this could be a verifiable credential
    - This could result in obligations, such as “do not resuscitate” or “do it once”
- (Mike Schwartz) You’re living in a fantasy world since we don’t know who the person is because we’re a computer and we have no physical recognition of the person.
    - We don’t know who the principle is
- (Mike Schwartz) Can we trust the OAuth token? We only have a bunch of evidence from a set of trusted sources.
    - E.g., is the user allowed to drive a car based on the token
    - You can stick the principal in the context as “token.idtoken.subject” etc.
- (George) There is value in having the client identify who it believes the principals are. Then the authorization server can validate whether those things are true or not
- (Mike Schwartz) That leads you to a failed system like SailPoint
- (George) you need to know what principal the system is acting on behalf of
- (MS) All of that could be true, without putting it into the principal of the token
    - From a governance standpoint, the principal is a distraction
    - Asking “what I’m allowed to do” is a failed approach
- (George) Pushing logic into the PDP that determines access based on who the principal is, has led to the ability of the PDP to make a wrong decision
- (George) There could be contradictory tokens.
- (MS)
- (George) It might not be a n “OBO”, but it could be that George is one of the principals.

- (MS) When we designed Cedarling, we tried to map the ID Token to a workload or a human, but it failed because it's not a couple, it's an array. When its an array, the policies are difficult to implement in terms of a principal. It's easier to make policy decisions based on tokens.
- (MS) To have governance, you always need to know the organizational identity, the attestations are important, etc. So it is always multiple principals.
- (MS) We're going to see an explosion of tokens out there.
- (George), There are other relationships, because at the time of determining policy, the principal's consent to policy might or might not matter.
    - (George) Having multiple tokens is valuable
- (MS) But it's not talking to "George", because its a machine, not a human.
- (George) When an LLM calls an MCP server, e.g. Google Calendar, it can ask for a token, but we don't have a good way of expressing "I'm this instance of this class of LLM, and I want to operate on behalf of George"
- (George) When the agent requests a token from an AS, it requests a token on behalf of a user in accordance with a mission / mandate / charter, etc.
- (George) Could you not create a "sub_id", in which you have a subject, an agent, and the relationship between them.
- (George) Where Karl is coming from, is that the OAuth entity type is that.
- (Rohit) If the primary field in the token is not a human, then we have a liability problem. Some agents are sending signed HTTP headers.
- (MS) Look at Clawdrey Hepburn. Are each of the "mini lobsters" that Clawdrey spawns their own agents?
- (George) This is the "sub_profile" in Karl's proposal
- (MS)
- (George) In AOL we embedded the user id in the cookie. If the underlying infrastructure didn't understand the identifier, it would not work.
- (George) In TraTs, you establish the subject at the beginning. You can identify many workloads down, who the principal is, but you also need the instance identifier of the code to identify which particular agent is operating.
- (George) In Clawdrey's case, when a sub-agent is spawned, the identity artifact of the subject embeds the lineage.
    - I don't like putting a whole lot of tokens in the "bag", since it doesn't capture the intent
- (Alex) If a request shows up to a PDP, if there are a whole lot of tokens, it complicates policy
- (George) In the short-term, the work Karl is doing is a step in the right direction.
- (MS) You would need to have DPoP proof of each component of the principal
- (George) there's definitely value in DPoP, because at the end of the day, all of this goes into risk management.

Building a living systems trust substrate - Architectural dive into social fabric

Session Convener: Frank
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Beyond IAM + PKI: KERISuite as a Trust Substrate

Session Convener: Mark Scott
Session Notes Taker(s): Susumu Ishizuka

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Overview
Presentation exploring KERI (Key Event Receipt Infrastructure) suite as a potential trust substrate to enhance current Identity and Access Management (IAM) and Public Key Infrastructure (PKI) systems. The KERI suite consists of three components: KERI itself, ACDC (Authentic Chained Data Containers), and CESR (Composable Event Streaming Representation).
Current State of IAM and PKI
- IAM handles identity for people, PKI for machines, and VLEI (Verifiable Legal Entity Identifier) for organizations
- Current systems have limitations in cross-domain trust, credential portability, and governance enforcement
- PKI certificate lifecycle management is problematic, with planned reduction to 45-day certificate lifespans creating operational burden
- Revocation mechanisms (CRLs, OCSP) are slow and not real-time
- Federation approaches exist but are limited to specific domains
VLEI (Verifiable Legal Entity Identifier)
- Emerged from G20 response to 2008-2009 financial crisis to improve organizational identification
- GLEIF foundation developed cryptographic approach to LEIs about six years ago
- Built on KERI infrastructure with high-assurance key signing ceremonies
- Enables organizational credentials that extend throughout an organization with cryptographic trust chain
- Supports Organizational Role Authorization Records (ORs) for multi-signature workflows

18 Claims for KERI as Trust Substrate
Self-Certifying Identity
- Identifiers (AIDs) are cryptographically grounded from inception
- Identity not dependent on platforms, directories, or certificate authorities
- Aligned with self-sovereign identity principles
Key State Verification
- Key states verifiable over time without destabilizing identifiers
- Unlike PKI where certificate renewal creates new identifiers
Rotation and Recovery
- Pre-rotation of keys built into protocol - next key hash included in current key event
- Enables immediate recovery from key compromise without reissuing millions of certificates
- Crypto-agile, supporting post-quantum cryptography
State-Based Replication
- More responsive than CRLs or OCSP for revocation
- Provides near real-time verification capabilities
Portable Credentials
- ACDC credentials can work across domains unlike trapped federated identity tokens
- Combines authentication and authorization in flexible schema
Organizational Identity
- VLEI extends IAM beyond enterprise boundaries
- Vertical credential chains throughout organizations
- Built on KERI after GLEIF evaluated and rejected blockchain approaches
Authorization Context
- ACDC containers include data, schema, authorizations, and provenance
- Authority context embedded in credential structure
Cryptographic Audit Trail
- Key Event Log (KEL) provides cryptographically-linked event history
- Forward and backward cryptographic linkage ensures integrity
- Better than traditional logging which can be modified by administrators
Built-in Delegation
- Delegation capabilities built into protocol, not requiring custom application code
- Cryptographically verifiable delegation
Cryptographically Enforceable Governance
- Multi-signature (M of N) governance enforced at protocol level
- Governance policies cryptographically enforced rather than relying on external systems
Trust as System Property
- Trust becomes inherent system property rather than external add-on
Cross-Domain Trust
- Enables trust across organizational boundaries without shared IAM providers
- More flexible than current federation approaches
Secure Interactions Beyond Transport
- APISC (Authentic Protocol Infrastructure Secure Channel) protects interactions, not just transport like TLS
- Greater vocabulary for interaction security
Links Over Attachments

- Links with proper authentication/authorization replace risky email attachments
- Reduces phishing attack surface
Reputation Services
- Foundation for reputation services providing additional trust anchors
- Independent trust makers can co-sign credentials
Future Vision for IAM
IAM evolution from current focus on:
- Directories and federation → Verifiable identifiers
- Single enterprise → Cross-domain trust with organizational authority (VLEI)
- Static credentials → Portable role credentials
- Application-level delegation → Protocol-built delegation
- External governance → Credential-based policy enforcement
Future Vision for PKI
PKI evolution from current approach:
- Fragile CRL/OCSP revocation → More robust mechanisms
- Certificate authority dependencies → Self-certifying infrastructure
- Fixed crypto periods → Flexible key management with pre-rotation
Potential Early Adoption Areas
Industries and use cases where KERI could provide significant improvements:
- Cross-organizational trust (currently lacking good solutions)
- Secure messaging (replacing insecure email, WhatsApp, Signal)
- Software signing with multi-party verification
- Machine and device identity
- Audit-heavy workflows requiring cryptographic verifiability
- Regulated communications and telecom
- High-risk administrative actions
- AI agent authorization (completely open area)
Challenges and Considerations
- Not a direct drop-in replacement for current IAM/PKI
- Developer experience and tooling still emerging
- Infrastructure needs (watchers, witnesses networks, registries, observers)
- Adoption challenges due to IT risk aversion
- Need for organizations willing to take calculated risks
- Potential 3-year timeline for viable alternatives
Timeline Perspective
- KERI positioned similarly to PKI 20-25 years ago
- PKI took years to achieve widespread adoption (now 90% of web servers use HTTPS)
- Current IAM has evolved through three generations (centralized, intermediate, federated)
- Systems remain incomplete and fragmented
Next Steps
- Part 2 presentation on Authentic Data Infrastructure planned
- Will cover architecture, design, and implementation details
- Includes APISC (secure channel layer) and secure messaging layer
- Potential startup discussions ongoing for 3.5 years

Missing middle entrepreneur Exploring and Solving for the Missing Middle of Capital for values based entrepreneurs in tech

Session Convener: Day Waterbury
Session Notes Taker(s): Brian von Herzen

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

how to fund the middle---

SDGs- were not financed adqquately
measure diverse communications in social context, environmental impact metrics.
quantify

aggregation -

not a zero some

gardening- plant a seed- skepticism-
theories of change -

money dries up when stock is going back -
noam chomsky- imporession of lower 80% keep doing the work- homeless-
manufactured scarcity-
survive the gap-

social gap -

velocity of funding in a community is key

landing- illiquidity, questions of succession

berkshires-

concern about company scrip -

kosla- funders of sun micro etc- spoke at sustainability - sustainbility center at Caltech - 2015-2020
what we look at like a vc fund- how we fund
funded ventner on genetics- take 100 years- got funded by kosla- to find a different way

1. opportunities for exponential growth outside the existing system
2. web w3 beat AOL linear growth-
3. ventner- did this-
4. tools that make tools - allowed craig ventner to do that-
5. groups that create their own value to grow exponentially
6. 10% or 2x and a growing ecosystem that is growing exponentially

patrick dowd -
camila rockefeller- ED- Nate Hagens spoke -

billionaire friends- invest- land based projects only-
****how does one manage succession? ****

trust- bylaws- vision mission and aims -
capped returns
10x on their money could invest in something -
cap for a cap -

adena pescue and aerth organization- doing work with natural capital / kelp - australian sites-

becomingdenizen.com collection of podcasters- post growht, capital allocation etc-

tristan harris- social dillemma- AI doc- leading voice on ai safety -

SESSION #9

How Does an Agent Decide Who to Trust? Verifiable Trust Protocol

Session Convener: Fabrice Rochette
Session Notes Taker(s): Ariel Gentile

Tags / links to resources / technology discussed, related to this session:

Slides available at https://gamma.app/docs/IIW42-How-does-an-agent-decide-who-to-trust-s4rbb0fhzq8i7bp?mode=doc

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Intro to Verifiable Trust protocol: https://verana-labs.github.io/verifiable-trust-spec/

Intro to Verifiable Public Registry: https://verana-labs.github.io/verifiable-trust-vpr-spec/#what-is-a-verifiable-public-registry

An application built on top of Verana is Hologram, a messaging app: https://hologram.zone

Invitation to Collaboration on Agentic Protocols and Representations

Session Convener: Mike Jones, Ankit Agarwal, Nick Steele
Session Notes Taker(s): Mike Jones

Tags / links to resources / technology discussed, related to this session:

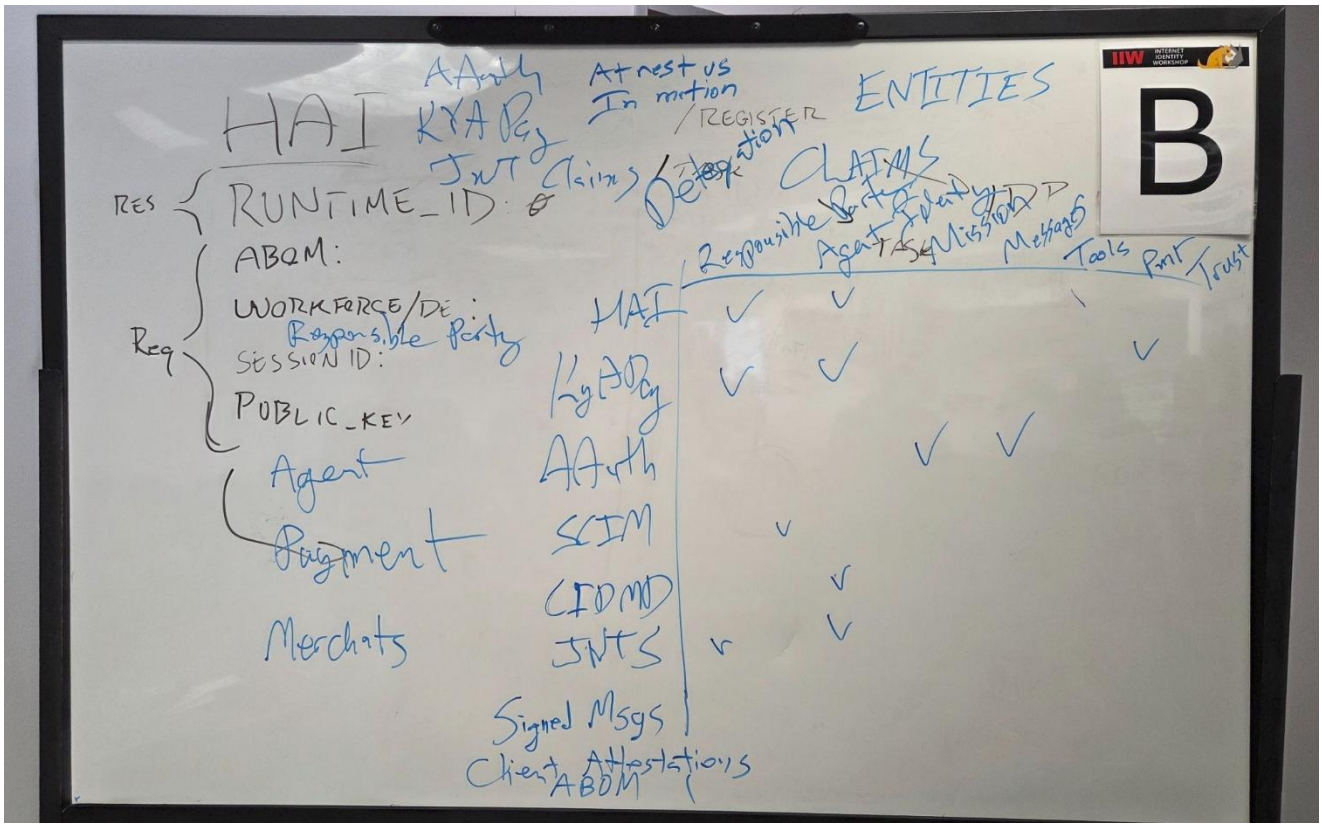

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Mike Jones invited people with protocols and representations for agents to collaborate and compare their systems. Nick Steele combined his session with Mike's. After Mike introduced the goals of the system, Nick presented the Holistic Agent Identity (HAI) system that OpenAI has launched. Representatives from KYAPay, AAuth, CIMD, WIMSE, SCIM, and other systems and protocols participated in the discussions.

Mike created a table on the whiteboard attempting to categorize the different systems and their roles and capabilities.

Holistic Agent Identity Framework Discussion
Date/Time: April 29, 2026, 2:35 PM PDT
Notes taken by Nick's ChatGPT Agent
Overall Summary

Nick Steele and colleagues discuss the development of the Holistic Agent Identity (HAI) framework at OpenAI, designed to establish and manage identities for AI agents in enterprise and consumer contexts. The framework focuses on agent identity facets, credential management, task authorization, and observability, aiming to provide modular, interoperable primitives for secure and manageable agent operations across diverse environments.

Key Points

- HAI framework defines agent identity using multiple facets, including runtime identifiers, cryptographic materials (e.g., ABOM), workforce identity, and session tokens to ensure integrity and authorization.
- Agents register with a public key and receive an agent ID; tasks are created with associated task IDs that carry scoped permissions, enabling fine-grained authorization and revocation.
- The framework supports sub-agents and agent families to handle ephemeral and durable agents, allowing delegation and orchestration within agent hierarchies.
- Task definitions remain flexible (missions, intents, goals) to accommodate various enterprise needs, with observability and auditability emphasized for security and governance.
- The framework is designed to be modular and interoperable, allowing enterprises to integrate with existing identity providers (IDPs) and authorization systems, and to manage agent identities and tasks according to their risk models.

Action Items

- Develop public specifications for data formats related to agent identity and tasks to support interoperability and standardization.
- Implement and test the register and task creation endpoints in OpenAI's platform to facilitate agent identity registration and task authorization.
- Explore integration strategies with enterprise IDPs and credential management systems to support workforce identities and session management.

Open Questions

- What are the best practices for managing credential lifecycle and revocation in the context of agent identities?.
- How to standardize task definitions and authorization scopes across different enterprises and use cases?.
- How to handle deep hierarchies of sub-agents and ensure clear accountability and traceability?.
- What is the optimal balance between security and usability in ephemeral versus durable agent identities?.
- How to best integrate with existing identity and access management tools while maintaining modularity and flexibility?.

Playnet

Session Convener: Ruzgar E
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

SEDI - Where the sidewalk ends.

Session Convener: George McE
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Bring - Your - Own - Everything Stack

Session Convener: Dmitri Z
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

My Terms - The Journey Continues

Session Convener: Justin Byrd and Mary Hodder
Session Notes Taker(s): Mary Hodder

Tags / links to resources / technology discussed, related to this session:

My Terms video at youtube: https://www.youtube.com/watch?v=ODlVXwxax4Q
IEEE 7012 Standard located here: https://ieeexplore.ieee.org/document/11360682
Customer Commons: https://customercommons.org/
MyTerms: https://myterms.info
MyTerms Launch webinar: https://www.youtube.com/watch?v=Nphd8l7KLek

IEEE 7012 Industry Connections or My Terms working group info:
https://standards.ieee.org/industry-connections/activities/
Sign up form: https://app.smartsheet.com/b/form/019d96d027fe74c2be64aef29dedb395

or USE THE QR CODE.

The IC group's first meeting is June 1, 2026

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

We will convene the first Industry Connection meeting at IEEE on June 1, 2026.

IEEE 7012 Industry Connections or My Terms working group: https://standards.ieee.org/industry-connections/activities/
Sign up form: https://app.smartsheet.com/b/form/019d96d027fe74c2be64aef29dedb395
Three questions for the group:

- What technical or organizational barriers must be addressed?
- How can open-source communities contribute?
- How does this connect to identity, privacy, and data governance initiatives?

SEDI: State Endorsed Digital Identity

Phil W Use Case: you want to make a bank account, and take your SEDI to the bank, with terms. My Terms.

What would a pilot look like: get certification requirements and make the system based upon libraries and build code and interface.

- Ad tech industry might be open to this, in order to do better and differentiate.
- Browser makers might adopt in order to be our agents for using My Terms.
- Ad blockers like Privacy Browsers

What is the advantage for a business to adopt this?

- New fiduciary responsibility coming to agents.
- Finding common way to do things: less confusion in the market place
    - Wifi, 5G, SSL, https, VHS vs Betamax, are examples where standardization made it possible to do things that companies and customers wanted to do things they couldn't do that before
- Apple stat on 1.5b people who clicked "do not track" or 96% of users opted out of tracking
- B2B marketing
- Leading on engineers to ask them to be more ethical and less sociopathic

What could we do?

- Consumer report ratings?
- Word press system and plug-in to offer to sites and individuals to use?
- Open source -- reduce the engineering for others

OpenID4VCI, OpenID4VP automated testing and certification

Session Convener: Joseph Heenan, Gail Hodges
Session Notes Taker(s): Jin Wen

Tags / links to resources / technology discussed, related to this session:

- OpenID4VCI — OpenID for Verifiable Credential Issuance (OIDF DCP WG)
- OpenID4VP — OpenID for Verifiable Presentations (OIDF DCP WG)
- OIDF Certification Program — https://openid.net/certification/
- OIDF Conformance Test Suite — https://openid.net/certification/connect_op_testing/

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

State of automated testing and certification for OpenID4VCI and OpenID4VP, the flagship DCP-WG specs. OIDF owns and maintains the full conformance test corpus, including upcoming OpenID4* standards, not just legacy OIDC / FAPI. Same program and tooling is the certification path for credential-issuance and presentation flows.
Gail's intro: certification landscape

OIDF is expanding from self-certification to a **multi-party model** to scale without fragmenting conformance across competing schemes.

The four roles

Role	Who plays it	What they do
Scheme owner	OpenID Foundation	Owns standards and test suite; sets scheme rules, terms, and pricing floor.
Authorized auditor	First announced: **Kantara Initiative**	Audits the TSPs themselves.
Approved Test Service Provider (TSP)	5 MOUs signed: **FIDO Alliance**, **Theme**, **TrustID** (Hungary), **Radium**, **Bixie Labs**	Run conformance testing on OIDF's behalf against the same OIDF corpus. FIDO is signed up as a wallet certification conformance provider.

Tested entity	The wallet / RP / issuer / verifier	Pays the TSP; portion flows back to OIDF.

Open program: same MOU template for new TSPs and auditors. *"This is not a limited portfolio… everyone signs up to exactly the same transparent terms for collaboration."*

Operating principles

- **Conflict-of-interest registers** required of licensed providers.
- **Sustainability over revenue.** OIDF's slice framed as **"hundreds of dollars" per certification, not thousands.**
- **Competitive marketplace.** TSPs can layer their own value (consulting, packages) on top of OIDF's fixed rate.
- **Anti-fragmentation.** Pitch: *"Save yourself time and trouble. Use what's freely available, and if it's not good, please feed back to us."*

Lifecycle and pricing

- Setup / first authorization.
- **Full re-certification every 3 years.**
- **Annual touch-up** with the auditor in between.
- A defined slice of each annual fee flows back to OIDF.
- **Pricing not yet fully public** until first contracts and terms are resolved.

Bundling layered specs: the OIDF Board's "spec families"

Many OIDF specs stack (FAPI on OIDC; HAIP on OID4VP/VCI). OIDF will offer **single bundled fee** for stacks. Board has resolved an explicit family grouping for launch:

Family 1: Open Data & Enterprise

Role	Bundle
Base RP	OpenID Connect, FAPI1-RP, FAPI2-RP, eKYC
Enterprise RP	OpenID Connect, FAPI1-RP, FAPI2-RP, Shared Signals, AuthZen, eKYC, OpenID Federation
Base IDP	OpenID Connect
Enterprise IDP	OpenID Connect, FAPI1-IDP, FAPI2-IDP, eKYC, Shared Signals, AuthZen, OpenID Federation

Family 2: Digital Credentials (relevant to this session)

Role	Bundle
Verifier	OID4VP
Wallet	OID4VCI, OID4VP, FAPI2-RP, OpenID Federation
Issuer	OID4VCI, FAPI2-IDP, OpenID Federation

Notes:

- **Wallets and issuers are bundled, not single-spec.** Only **Verifier** is a single spec (OID4VP).
- **eKYC is a Base feature**, not enterprise-only (consistent with NIST NCCOE / banking framing).
- Slide footer: *"Mapping of specific test plans to families will be managed via the team."* The family-to-test-plan map is a living artifact; cert-readiness work should track which test plans each role bundle currently maps to.

EKYC / OpenID for Identity Assurance

Joseph: OIDC4IDA layers verified-identity assertions on top of standard OIDC claims. Response can encode *"this is the user's name that I've verified, against their driving licence, using this trust framework, on this date, with this third party."* Acceptable in some jurisdictions for bank account opening or credit-bureau checks.

Gail: active **NIST NCCOE workstream** with **Julianna Kayfic** lead-editing a spec to unblock high-assurance transactions across jurisdictions, resolving metadata for proofing and authentication so banks can give it appropriate confidence. Expected to matter for healthcare and other regulated industries; less so in Europe given more directive regulation.

Launch timeline

- **Targeting Q2 launch** for the expanded TSP/auditor program.
- Open work: final contracts/MOUs, UX polish on the cert portal, ecosystem segmentation (so each TSP only sees relevant test suites).

Audience point: LPT / "platform test"

Seth: browser/large-platform vendors will not subject themselves to OIDF certification. *"Your best chance is to get your test into a **platform test** which runs continuously."* Likely target: **web-platform-tests (wpt.fyi)**. For browser-side conformance (W3C Digital Credentials API path), pushing tests upstream is more effective than waiting for browser vendors to certify against OIDF.

How to get to the tests

1. openid.net → **Certification** → *How to certify your implementations* → two links to test runners.
2. Easier: ask Gail to forward the email with direct test links (she'll BCC anyone interested, with a feedback ask attached).
3. Joseph reinforced: *"If you do try out the tests and they work or don't work, please do let us know."*

Scope: what OIDF certification is *not*

- **Not a test of internal implementation.** Verifies **protocol behavior on the wire**: messages, parameters, error handling, flows. Does not assess key management, secure storage, code quality, or overall security posture.
- **Not EU Wallet certification.** Not equivalent to **EUDIW certification** under eIDAS 2 / EU ARF (regulatory conformity assessment covering assurance level, security evaluation, privacy, governance, supervised list). OIDF conformance can be one input but is not a substitute.

Open-source conformance suite

- **Repo:** https://gitlab.com/openid/conformance-suite (GitLab, MIT)
- OpenID4VCI / OpenID4VP / DCP-WG plans are added to the same codebase as specs land.
- **Production deployment:** certification.openid.net.
- Activity: ~8,889 commits, 55 branches, 139 releases.

Implications:

- **Self-host** instead of (or alongside) demo.certification.openid.net. Useful for CI, debugging against a wallet under development, or offline work.
- Read the test source to see exactly what positive/negative checks a plan applies.

Automation segment: eudiplo

Mirko walked through automating OIDF tests against a real implementation, using **eudiplo** as the system-under-test.

- **Repo:** https://github.com/openwallet-foundation-labs/eudiplo
- **Org:** OpenWallet Foundation Labs
- **Language / license:** TypeScript, Apache 2.0
- **Layout:** pnpm monorepo with Dockerfile, docker-compose, mkdocs, deployment, monitor, and an approach.md design note.
- **What it is:** lightweight protocol-abstraction middleware between a backend service and an EUDI Wallet (not a test orchestrator). Backends call eudiplo over HTTP; it handles OID4VCI / OID4VP / SD-JWT VC / mdoc so app developers don't have to.
- **Protocols:** OID4VCI, OID4VP, SD-JWT VC, mdoc / ISO 18013-5, OAuth Token Status List.
- **OIDF-conformance tested** for core protocols (closes the loop with the OIDF suite).
- Other features: JSON-based credential config; client-credentials auth; Docker deployment with env config; pluggable secure key-management; no outbound calls, no long-term credential storage.

Why it matters: it demonstrates a **closed loop** for OIDF certification with fully open-source tooling on both sides (issuer/verifier + conformance suite), and illustrates the **OWF ↔ OIDF cooperation pattern**.

Public certification listings

- Program overview: https://openid.net/certification/
- **Public certification listings index:** https://openid.net/certifications/
- Public test runner / results host: https://www.certification.openid.net/
- Example public plan-detail page (staging): https://staging.certification.openid.net/plan-detail.html?plan=VlKMtd1fQtYkN&public=true

How certification works in practice

All certified implementations are listed publicly at openid.net/certifications/ (source of truth).

The 7-step process

1. Wallet provider runs the tests in their own environment or on OIDF's cloud.
2. Tests cover **both positive and negative cases**.
3. Failures requiring fixes are surfaced.
4. **Logs and statement of compliance** submitted to OIDF.
5. **Certification fee** paid.
6. **OIDF publishes results** after checking logs (manual; takes a few days).
7. **"OpenID Certified" mark** can be used.

Tests can also be run by a third party on the implementer's behalf.

Implication: certification attests that the implementer (or third party) demonstrated a passing run and submitted logs + statement of compliance. OIDF does **not** independently re-run tests against the live deployment. Reinforces the scope note above: protocol-conformance signal, not an audit.

Live demo: OpenID4VP verifier tests (HAIP profile)

Joseph configured and ran a verifier test plan on demo.certification.openid.net. Profile: **HAIP** (High Assurance Interoperability Profile, pronounced "hype"). HAIP requires every key to be backed by an **X.509 certificate** (x5c).

Configuration form fields:

- Certificate Authority for the credential under test.
- Status-list trust anchor (separate from the credential CA).
- Custom URL scheme to launch the wallet from QR.
- Client ID matching the certificate (HAIP uses **x509_hash** client identifier).
- **DCQL query** describing the credential. Joseph requested an **EU PID** with claims family_name, given_name.

- Signing key for the request object: JWK with x5c chain.
- Second client certificate for tests needing a second client identity.

Test plan content:

- Happy-flow tests.
- Alternate request URI flows.
- DCQL variations: fewer claims, optional credential set, no claims, etc.
- Negative tests: malformed requests, wrong values.

Demo run:

- Wallet: **Animo's Paradym** on Android, freshly installed from the app store (not a dev build), credential from the non-production demo issuer.
- Tooling: Android device mirrored via Vysor alongside the test-suite browser.
- Happy flow: scan QR → wallet prompts to share EU PID with claims → approve + PIN → wallet sends vp_token → test passes.

Per-test logs captured:

- DCQL validated against an inline JSON schema.
- Signed request object: claims, JOSE header, signing key, **link out to jwt.io** for independent signature verification. Joseph: *"Until I built this jwt.io link in, people were always claiming the conformance suite wasn't signing things properly."*
- Response URI HTTP request; encrypted response, JWE header, decryption inputs, decoded claims.
- Per-token validation: VP token, SD-JWT (claims, disclosures, key-binding JWT, raw SD-JWT).
- On failure: expected vs. actual with specific error message.

Live demo: OpenID4VP-over-DC-API verifier tests

Same flow with response_mode switched to the W3C **Digital Credentials API** browser-mediated path:

- HAIP over DC API has a **longer** test list (more requirements).
- No QR; suite presents a **"Proceed with Test"** button to kick off the browser DC API flow.
- "More secure" flow shown: **QR + Bluetooth proximity check** as anti-phishing path (wallet must be in BT range).
- Demo failed on IIW WiFi (network issues).
- Failure card → **"more"** → exact exception from the browser. The browser hides the wallet's underlying error, so the suite surfaces what it can see. Useful debugging signal: **DC API errors are one layer up from where you'd usually look.**

Demo conformance instance:

- Login: https://demo.certification.openid.net/login.html
- Schedule a test plan: https://demo.certification.openid.net/schedule-test.html
- This demo.* host is the try-it-yourself sandbox (separate from www.certification.openid.net production listings and staging.certification.openid.net).

Anyone can sign in (Google works) and run plans against their own endpoint at no cost. Payment is only for the certification step (publishing the mark).

Live demo: OpenID4VCI issuer tests

Same UI, issuer side, against **Authlete's example issuer** that ships with the test suite. Profile: HAIP, format: SD-JWT VC.

New fields vs. verifier tests:

- **Client-attestation keys** (x5c chain).
- **Key-attestation keys** (x5c chain).

Test sequence:

1. **Metadata test:** checks .well-known/openid-credential-issuer.
2. Happy-flow issuance: redirect to issuer, end-user authenticates, back to suite, credential delivered.

Per-test artifacts:

- Credential-endpoint call with **copy-as-curl** button (Joseph: VPoP assertion may be one-shot, not always replayable).
- Credential response and issued **SD-JWT** with claims/disclosures decoded.

Built-in Selenium browser-automation: the test plan can drive the issuer's login UI so repeated runs don't require manual login. JSON config (UI deemed too fiddly): visit URL → wait for element → enter text in field → click button → wait for result. Per-test override syntax exists.

Negative-test demo: "send credential request with invalid nonce":

- Each test has a **blue info box** explaining what it does and the expected result.
- Suite confirms the issuer returned invalid_nonce, no credential, no unexpected keys, error response **not encrypted** (per spec), and the error code matched.
- **Negative tests are first-class**, not an afterthought.

Mirko's automation: eudiplo + OIDF suite via testcontainers

Mirko: *"I was programming a verifier… too lazy to use a wallet to manually test, found this beautiful test suite and integrated it into my test approach."*

Hosted vs. self-host decision: Mirko chose **not** to use demo.certification.openid.net for CI: *"it's not guaranteed that it's aligned with your versions. The configuration changed there and then you're wondering why your code is not working."* He runs the OIDF suite **inside GitHub Actions** as a Docker deployment, started via **testcontainers** from his TypeScript e2e tests.

Wiring:

- **OIDF setup helper** specifies which OIDF image and release to pull. CI verifies "is my code still compliant with the latest tests?" as a first-class signal.

- **Fresh Docker containers per test** (no pre-config to poison results).
- Tests interact with the OIDF suite via its **OpenAPI** REST interface: start a test by identifier, poll for completion, fetch results.
- Test outputs (HTML and JSON) pulled via API so failures surface in CI logs, not only on the suite UI.
- Same code runs in dev and in GitHub Actions / GitLab CI.
- Implementation TypeScript but language-neutral (testcontainers + OpenAPI).

Test scope:

- Started happy flow, expanded to unhappy flow.
- Today: **190+ tests** in eudiplo's suite. OIDF foundation provides ~20–30 SD-JWT tests out of the box, also runnable for mDoc and DC API. Joseph noted Mirko is on an old version and missing newer tests.
- Can run **individual tests by identifier**, not just whole suites. Direct answer to an audience question about partial implementations: useful when only specific tests are relevant to current work.

AI loop Mirko emphasized (twice):

- *"You just tell your co-pilot 'something is wrong, please fix this.' Since you have a test, it can iterate and build your fancy code, then it says yes the test is passing. The only thing left is, in the git diff, is this a good fix or are there new bugs."*
- For new tests added to the OIDF suite: *"You can say to your AI 'go to the API, these are all the new cases, please write me the wrapper so I know I have to implement these 10 new test cases.'"*

Mirko's helper code is open-source, in the eudiplo repository (see follow-ups).

Q&A nuggets

- **Self-host the suite?** Joseph: yes (Mirko then demonstrated end-to-end).
- **Wrap in CI?** Joseph: yes (Mirko showed it).
- **Run the online suite without paying?** Joseph: *"yeah, you'd literally just do what I did and log in with a Google account."* Pay only for the published certification mark.
- **Implement all tests, get certified in a minute?** Joseph: no, *"a couple of days for the actual publishing of the certification."* Step 6 of the 7-step process is human-mediated.
- **Are Mirko's 190 tests all from OIDF?** Mix: ~20–30 from OIDF (SD-JWT, re-runnable for mDoc and DC API), the rest built on top.

Outstanding questions

- Identify which "platform test" Seth meant (likely web-platform-tests / wpt.fyi).
- Confirm spelling of the NIST NCCOE editor (likely **Julianna Kayfic**).
- Which of the 5 announced TSPs are operational vs. MOU-only?
- Where on openid.net will the public TSP / auditor list and pricing live?
- Which OIDF release is Mirko's eudiplo running, and which newer tests is he missing?

- Current matrix of profile/format combinations supported by the OIDF suite for OpenID4VCI/VP (mdoc, SD-JWT VC, JSON-LD VC), especially under HAIP?
- Which test cases are blocking certification readiness, and on whose roadmap?
- How does OIDF certification coordinate with **EUDI Wallet Reference Implementation** conformance and DCP-WG profile work?
- Does the proximity/Bluetooth check apply to all DC API profiles or only HAIP?

Action items / next steps

- [ ] Get on the email list Gail mentioned for direct links to the demoed test plans.
- [ ] Locate Mirko's open-source CI wrapper inside the eudiplo repository and link the helper functions (container start helper + OpenAPI test-start helper).
- [ ] Cross-link this note to dcp-certification-test-plan.md and dcp-top3-test-prep.md.
- [ ] Capture exact demo.certification.openid.net test plan IDs shown live (HAIP / OpenID4VP verifier; HAIP / OpenID4VP-over-DC-API verifier; HAIP / OpenID4VCI issuer).
- [ ] Sanity-check whether Apple's native container runtime (macOS 26+, Apple Silicon) can run the OIDF compose stack as a Docker substitute.

*Sovereign Identity Namespace | *Time is Now? Or *Been there…done that?*

Session Convener: Jeffrey Mendelsohn
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

The Fiduciary Commons (Part 3) AI: The Inference Problem

Session Convener: Mike Leahy
Session Notes Taker(s): Mike Leahy

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The Fiduciary Commons | Session 3 of 3
The Inference Problem
Completing the Architecture: GAAFA, Compliant Design, and Procurement as the Leverage Point

What This Session Argued
VIDA and PDTA built a privacy architecture around what data gets collected and how it is held. That architecture is necessary but incomplete. A system that never collects a fact can still reconstruct it through AI inference over transaction patterns and behavioral residue. GAAFA closes that gap by extending fiduciary duties to the AI and inference layer. The Supreme Court is moving in the same direction: Carpenter v. United States (2018) began the retreat from the Third Party Doctrine, and four sitting Justices have signaled the doctrine needs fundamental rethinking for digital-scale collection. Procurement is the leverage point where compliant standards either get embedded before technology is deployed or get abandoned in vendor negotiations after the market consolidates.
Four Ideas Worth Keeping
The mosaic assembles from correct transactions.
Aggregate inference reconstructs what individual data points cannot reveal. A VIDA-compliant age verification credential reveals nothing beyond age. But transaction logs across time reconstruct daily routine, social network, and behavioral profile. The credential architecture was correct. The inference engine worked around it. This is the constitutional problem Carpenter identified and GAAFA is designed to solve.

The Court is moving toward the Fiduciary Commons standard.
In Carpenter v. United States (2018), Chief Justice Roberts held that long-term cell-site location data requires a warrant even though held by a third party, because the aggregate reveals the privacies of life. He declined to extend Smith v. Maryland's Third Party Doctrine to digital-scale collection. Four Justices want to go further. Systems built on third-party data sharing assumptions are built on contested legal ground. VIDA and PDTA are aligned with where the Court appears to be heading.
AI reconstruction of uncollectable data is a fiduciary breach.
Government chose not to collect certain data because it had no constitutional right to collect it. Using AI to reconstruct from transaction patterns what the credential architecture was designed not to reveal bypasses the constitutional constraint without violating its letter. A trustee who finds an indirect method to extract value from the trust corpus that direct methods prohibit has not honored the trust. The method does not redeem the act. GAAFA closes the gap between what government cannot collect and what government cannot infer.
Procurement is where standards get embedded or abandoned.
Standards set early in a technology's development are far more durable and far less costly to enforce than standards imposed after market patterns solidify. Article 3B, proposed to the ABA Model Procurement Code by Mike Leahy, creates the legal pathway for phased multi-vendor competition, outcome-based specifications, and fiduciary-aligned acceptance criteria. Vendors who build to GAAFA-compatible standards now set the baseline against which the statutory mandate will be measured.
The constitutional constraint is not just about what is collected. It is about what is reconstructed. GAAFA closes the gap between those two things.
The Closing Ask
Three specific actions. First: identify one state procurement relationship and ask whether the code has a proof-of-concept pathway for genuinely novel technology. Most do not. You now know what gap to name and what Article 3B provides. Second: build toward GAAFA compliance now and document it. The statute does not yet exist in most states; documented compliance-readiness is the evidence base legislative staff need to write achievable requirements. Third: engage with the Article 3B draft directly at fiduciarycommons.com. Vendor-side input is underrepresented in procurement law reform and is genuinely needed.

What This Means for Your Work
Four architectural properties define GAAFA compliance
Purpose-sequestered databases: separate data stores scoped to collection purpose, with AI access gated by purpose-matching logic and an audit trail. Cross-purpose queries require documented authorization.
Cryptographic, individual-controlled, immutable audit trails: each access event hashed and chained; the citizen can inspect the full record of queries against their data; no entry can be deleted or modified.
Explainability by design: AI-assisted decisions affecting citizen rights must identify the factors weighted, not merely the output. Built in from the design stage; retrofitting is an order of magnitude more expensive.
Inference prohibitions enforced architecturally: purpose sequestration combined with access gating, not policy attestation alone.

Article 3B: what it does for vendors

The current Model Procurement Code has no pathway for genuinely novel technology procurement. Article 3B creates one: phased multi-vendor competition that maintains competitive pressure through the exploratory arc, outcome-based specifications that eliminate solution bias, the Minimum Viable Product as the specification baseline for production, and citizen data carve-outs and flow-down clauses that structurally advantage GAAFA-compliant vendors. Practitioner input on the draft is open and needed.

The market-shaping window

Government consortia using Article 3B embed compliance standards into procurement specifications before deployment. The standards vendors build to now become the market baseline. The window for that investment is the period before deployment at scale. Vendors who engage with the Article 3B framework now, providing input on what compliance standards should require, are shaping government AI procurement architecture for the next decade.

Key Terms

The Mosaic Theory

The constitutional principle, recognized in Carpenter v. United States (2018), that aggregate data reveals the privacies of life in ways no single data point does, requiring warrant protection even when individual data points would not.

Third Party Doctrine

The legal principle from Smith v. Maryland (1979) that information voluntarily shared with a third party loses Fourth Amendment protection. The Court declined to extend this doctrine to digital-scale data collection in Carpenter.

Algorithmic Impact Assessment

A required pre-deployment analysis documenting what a government AI system infers, from what inputs, with what error rates, and with what disparate impacts across population groups. Modeled on environmental impact assessments.

Purpose-Sequestered Database

A database architecture in which data is stored in separate systems scoped to the purpose for which it was collected, with access gated by purpose-matching logic. Prevents AI systems from drawing inferences across incompatible collection authorities.

Article 3B

Proposed addition to the ABA Model Procurement Code for State and Local Governments, authored by Mike Leahy, creating a legally sound pathway for proof-of-concept procurement of genuinely novel technology with fiduciary compliance built into the framework.

Minimum Viable Product (MVP)

Under Article 3B, the Phase 3 output that demonstrates a solution approach meets the government's functional need under real conditions. Becomes the specification baseline for production procurement, replacing pre-procurement assumptions with evidence.

fiduciarycommons.com | michael@fiduciarycommons.com | For the full argument, see the Fiduciary Commons Series Summary.

Verified vs True / Proof of Effort (as related to text-based content)

Session Convener: David Lee Condrey <david@writerslogic.com>
Session Notes Taker(s): David Lee Condrey <david@writerslogic.com>

Tags / links to resources / technology discussed, related to this session:

C2PA, CAWG, Verifiable Credentials, IETF SCITT, RFC 9334, SynthID-Text, PoSME (Proof of Sequential Memory), CPoE (Cryptographic Proof of Effort), TEE/TPM, RFC 3161, keystroke biometrics, AI detection, human authorship

https://writerslogic.com -> WritersLogic, Inc (the company)
https://writerslogic.com/loi -> Contact form to connect after IIW
https://writerslogic.com/processproof -> Session slide presentation
https://writersproof.com -> Proof-of-concept CA for WritersProof application

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The session ran in two halves. First, a reading of what the content-provenance stack actually verifies, in the specs' own words. Second, a concrete proposal: a draft C2PA standard assertion called c2pa.process-proof, opened for critique rather than endorsement. The convener framed the session as a hearing, not a pitch. It was attended by participants from C2PA, CAWG, DIF, ToIP, the SCITT and RATS communities, and several independent researchers working on AI-content detection.
Part 1. The gap, in the specs' own words
Verification authenticates a statement. It does not authenticate a fact. Every layer of the content-provenance stack admits this in its own normative text. The consumer-facing word "verified" does not.
The session opened by reading directly from the specs.

- W3C VC Data Model 2.0, §2: "Verification of a credential does not imply evaluation of the truth of claims encoded in the credential." Truth is routed to a separate process the spec calls validation, which is the verifier's responsibility, not the protocol's. The Data Model is explicit that the protocol stops at "this credential was signed by this issuer."
- C2PA 2.2, §9.2 and §10.2.2: Hard bindings ensure the manifest belongs with the asset and the asset is unchanged. Trust is rooted in the signer. The created/gathered distinction allocates responsibility but does not produce evidence. A C2PA manifest tells a consumer who signed the assertion, not whether the assertion is true.
- CAWG Identity Assertion 1.1, §3.3.13 and §8.1.1.3: "Trust" is defined as the consumer's required confidence threshold, not a protocol-verified property. The named actor claimed a role. The wording is deliberately not performed. CAWG is honest that an identity assertion authenticates a claim of relationship, not the underlying act.
- IETF SCITT charter: It is an explicit non-goal to prevent authenticated supply-chain issuers from making false claims. SCITT provides a transparency log over signed statements; what those statements assert about the world is out of scope.

- RFC 9334 (RATS), §1: "Trust is a choice one makes about another system. Trustworthiness is a quality about the other system…" Attestation produces evidence for a choice, not a guarantee. RATS is precise: an Attester's evidence is input to a Relying Party's policy, not a substitute for it.

The pattern is consistent across five independent specifications from four standards bodies. None of them claim to verify truth, and each one says so in its normative text. The session's first contribution was simply to put these admissions next to each other so the gap stops being deniable. Several attendees noted that they had read each of these specs individually without registering that all of them admit the same boundary. Reading them as a set produces a different effect than reading any one in isolation.

Why text is the hardest case

The provenance stack was designed for media with a capture moment: sensor, lens, shutter. A camera can sign at the moment of capture because there is a discrete physical event to sign. Text has no capture event. There is no shutter for prose.

This matters because LLM-generated text is functionally free. Current commodity inference runs on the order of $0.0001 to $0.0002 per page. The cost asymmetry between human composition (hours) and machine generation (seconds, fractions of a cent) is several orders of magnitude. Any provenance system for text has to reckon with the fact that the adversarial action it is trying to detect is also the cheapest possible action.

Watermarks like Google's SynthID-Text prove machine origin, which is a negative signal. Their absence proves nothing. A document without a SynthID watermark could be human-written, could be from a model that does not watermark, could be from a model whose watermark was stripped, or could be from a model whose watermark was never deployed. There is no positive cryptographic test for "a human composed this," and on current trajectories there is no obvious path to one. The session did not propose to invent one.

What the available signals actually prove

The session walked a table covering every signal currently proposed for human-authorship attestation. The table is reproduced here in summary form because several attendees asked for it.

- C2PA and CAWG signatures. Prove that a signer (a device, an identity, an organization) made a claim about an asset. Do not prove the claim is true.
- LLM watermarks (SynthID-Text and similar). Prove machine origin when present. Prove nothing when absent.
- Document revision history. Proves that a sequence of edits occurred in a particular tool. Does not distinguish edits typed by a human from edits pasted from another window.
- Keystroke biometrics. Prove that typing occurred in a pattern consistent with a particular person. Do not distinguish a human typing original prose from a human transcribing LLM output.
- Proof-of-personhood. Proves that a unique human is associated with an account or session. Does not constrain what that human submits.
- Editorial process attestation. Proves that a workflow (review, edit, approval) was followed. Does not constrain the inputs to the workflow.
- PoSME and sequential-memory work. Proves that a specific computation was executed sequentially on memory-bound hardware over a span of wall-clock time. Does not prove what was thought during that computation.

Every row reduces to proof of process, never proof of thought. The convener placed PoSME on the same row as the others. The recursion was conceded openly and is worth recording, because it is the move that distinguishes this session from a pitch: Proof of Effort cannot escape the verified-versus-true gap. PoSME is not a solution to the problem the session named. It is a different shape of evidence within it.

The honest claim for PoE is narrower than its name suggests. It does not prove authorship. It shifts the economics of fabrication from free to computationally expensive. The contribution is cost asymmetry, not epistemic guarantee. A system that requires sustained sequential memory work to produce a manifest entry raises the marginal cost of fabricating that entry from approximately zero to something measurable. Whether that shift is sufficient as an "evidence layer" is one of the open questions the session left to the room.

The governance defense, named and challenged

A natural response to the verified-versus-true gap is: governance fills it. Audits, accreditation, trust frameworks, revocation, reputation. The session named this defense and challenged it on a specific point.

All of these mechanisms are post-hoc. They catch a liar after the lie has been detected by some other mechanism. For images and video, that other mechanism is often perceptual: someone notices the artifact, the inconsistency, the impossible shadow. For supply-chain artifacts, it is operational: the part fails, the build breaks, the audit log is reviewed. For text-based content, the other mechanism does not exist. There is no perceptual tell that prose was machine-generated, and no operational consequence in most consumer contexts.

The line that several attendees wrote down: accountability without detection is a system that only catches honest liars.Governance is necessary. It is not sufficient when the underlying detection layer is empty.

Part 2. The proposal: c2pa.process-proof

The second half of the session presented a draft C2PA standard assertion, currently spec text and CDDL, pre-PR. The pitch was deliberately narrow. The assertion does not make claims verifiable. It adds one independently checkable evidence signal alongside the manifest signature, and it is purely additive: manifests without it are unaffected, and unknown proofFormat values yield informational status rather than failure. Existing C2PA tooling does not need to change to ignore it, and existing manifests do not become invalid in its presence.

The assertion was developed by WritersLogic in the course of building WritersProof, a Claim Generator for written content. The work surfaced a recurring pattern: every authorship-attestation problem the team encountered routed through the same gap the specs themselves admit. The assertion is the team's attempt to add the smallest useful piece of evidence to that gap, framed in a way that the C2PA Technical Working Group can evaluate on its merits.

Schema. CDDL-shaped process-proof-map, eight fields:

- proofFormat (string): the proof system used (e.g., posme.v1).
- proofTier (integer, 1 to 3): the assurance level achieved (defined below).
- proofBytes (binary): the proof artifact itself.
- declaredPurpose (string): the asserted reason the proof was generated.
- attestationChainRef (optional): reference to a hardware-attestation chain.
- timeAnchorRef (optional): reference to an external time anchor.
- zeroDataCommitment (boolean, mandatory true): the assertion that no behavioral or biometric data is contained in or derivable from the proof.

- metadata (optional): implementation-specific extensions.

The assertion lives in created_assertions when the Claim Generator vouches with computational evidence of its own, and in gathered_assertions when the Claim Generator is relaying an upstream proof without itself staking work. This is the only place in C2PA where created carries cryptographic backing beyond the manifest signature.

Confidence tiers, with honest delivery status:

- Tier 1, software-only, self-attested. Delivered today via PoSME, with the WritersProof Claim Generator as the reference implementation. The proof binds the manifest to a specific sequential computation on memory-bound hardware. The assurance is entirely cryptographic and entirely software-rooted.
- Tier 2, TEE-bound. Open integration. Adds Secure Enclave, TPM, or TDX attestation to the proof, binding the computation to a specific hardware root of trust. Requires hardware-attestation chain handling that PoSME does not currently include. The assertion fields support it when it lands.
- Tier 3, TEE plus external time anchor. Open integration. Adds an independent time witness via RFC 3161, a blockchain anchor, or an NTP oracle, binding the computation to wall-clock time as observed by a third party. PoSME explicitly defers time anchoring to the integrator; the architecture allows it without modification to the proof itself.

The tier structure was deliberate. The session made the case that confidence levels in content provenance should be honest about what hardware and what witnesses are involved, rather than collapsing all proofs into a single binary "verified" status.

Privacy boundary, non-negotiable. Zero behavioral data in the manifest. No keystrokes, no timing, no biometrics, ever. Where behavioral entropy is used as an internal binding input in jitter-entangled proof variants, only (step_id, hash)commitments are exposed in the assertion. An observer of the manifest learns "an entropy event occurred at step 42" and nothing more. The hash commits to the entropy without revealing it; the step_id locates it in the proof transcript without describing it.

This boundary matters because the natural fear of any "process proof" assertion is that it becomes a Trojan horse for behavioral surveillance. The zeroDataCommitment: true field is mandatory in the schema specifically to make the privacy posture machine-checkable. An assertion that does not commit to zero behavioral data is not a c2pa.process-proof assertion.

The Minnesota NLP / ScholaWrite research collaboration, which uses keystroke and revision data to study writing process, is research-side only. None of that data enters any production manifest. The collaboration informs the underlying proof construction but does not pollute the assertion's privacy guarantees.

Outstanding questions and open work

1. Tier 2 and Tier 3 integration. Hardware-attestation chain handling and external time-anchor binding are not in PoSME today. The assertion fields exist to support them when they land, but the integrations are open work and will likely require collaboration with hardware vendors and timestamping authorities.
2. Application-grounded threat model. PoSME's draft threat model is cryptographic: forge a transcript, reduce storage, defeat the time-memory tradeoff. The content-provenance threat model is different. The adversary is often a cooperating user submitting LLM output through legitimate tooling. Adapting the threat model to that adversary is open work and was flagged by attendees as the most important next step.

3. declaredPurpose duress mechanism. Self-declared duress signals do not survive a coercive setting. A user under coercion will not check the duress box. The field may be cut from the schema until a better mechanism exists.
4. TWG review and VC alignment. The path to standardization is the C2PA Technical Working Group. Alignment with the W3C VC evidence property is open. Whether c2pa.process-proof lives in C2PA core, in CAWG, or as a VC evidence extension is itself an open question that the session put to the room.

Questions opened to the room

The session did not converge on answers. These were posed for the community to carry forward.

1. The substitution problem. Full revision history plus ChatGPT open in another tab. Every effort signal currently proposed says "human-authored." What did we just verify?
2. The asymmetry. SynthID-Text proves that a Google LLM wrote something. It cannot prove it did not. Are we calling one-sided detectors "provenance"?
3. The CAWG question. Signing text under an identity assertion authenticates the publisher. Does it authenticate anything about authorship?
4. The economic question. Is shifting fabrication from free to expensive sufficient as an "evidence layer," or does the room want a stronger guarantee, and what does that look like?
5. The home question. Should c2pa.process-proof live in C2PA, CAWG, or VC evidence extension? Each home implies a different governance trajectory.
6. The convergence question. SCITT, RATS, VC, C2PA, and zero-knowledge proofs of execution are reaching toward each other. What should we call a credential whose claim is itself backed by Proof of Effort?

Closing observation

The session was framed as a discussion, not a pitch, and the framing held. The room did not converge on an answer to any of the six questions. What the session produced was an agreed vocabulary for a problem the standards bodies have so far described in fragments. The phrase verified versus true was something several attendees said they would carry into their own working groups.

The draft c2pa.process-proof assertion is one proposed shape for the evidence layer the question implies. Whether it is the right shape, and whether the evidence layer belongs in C2PA at all, was left open. WritersLogic will continue developing the assertion and the underlying PoSME proof system in the open, and invitations to engage in the C2PA TWG review process, in co-authoring the Proof of Sequential Memory Execution https://github.com/dcondrey/posme-draft, or directly via contact form at https://writerslogic.com/loi/.

WritersLogic, Inc. is currently pre-seed and is offering a superior alternative to proving human authorship to writers, publishers, and anyone with vested interest in ensuring human authorship. AI detectors are unreliable, undefendable, and unfalsifiable; evidence produced by the WritersProof application is falsifiable, independently verifiable, and privacy-focused. One is clearly a better solution. Writerslogic, Inc. is welcoming business partnerships, investors, and collaborator discussions.

Philosophies of Privacy…. What do we even mean

Session Convener: Sol + Elizabeth
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOES SUBMITTED

Credential Properties Naming

Session Convener: Dean H. Saxe, Pamela Dingle
Session Notes Taker(s): George Fletcher

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

2026 April 29 - Credential Properties Notes

Credentials are “labeled” as IAL2 but it doesn’t describe capabilities/security properties of the credentials.

Credential properties should describe aspects of the credential in a durable way. Accessible to practitioners and lay users.

In this context credential a signed object or an authentication credential? These properties should apply to anything you can call a credential.

Goal is to have words and tests to determine the security properties of the “credential”. A description of properties of a credential. Focused on the properties of the credential and not how the credential is used.

7 classes of credentials Credential boundary (how it’s stored, intent for usage)

1. Boundary enforcement
    1. Who can enter the boundary?
2. Exportability
3. Transference

4. Dynamism
    1. Asymmetric cryptography
    2. Proof of credential/possession vs presentation of the credential
    3. Time bound
2. Key protection (potentially a new class?)

Presentation Limits (reuse properties)

1. Validity Count
2. …

Credential Containment & Unlock

Credential Enrollment

Credential Provider Trustworthiness

- "credential provider" is an overloaded term in different contexts
- "Credential issuer" ?? Also overloaded from a Digital Credential perspective
- "Credential authority"? Or "Issuing authority"?
- What's the test for this class?
- What are the properties of a credential that is issued by some entity?
    - Who is authoritative for the claims in the credential?

Credential Compliance

Credential Recovery

- Recoverable but only through a different credential
    - Reset admin account deleted all stored passkeys. To create new passkeys, required first creating a password, then creating a passkey and then deleting the password
    - Had to lower the security bar in order to recover the higher security state

Credential Lifecycle

- how to bootstrap the next credential

RMOD(P/D/B) - defines mechanisms to describe a domain

- may have words that help with this describing these properties

Discussion about single use credentials and the problems of distributed cache management. Usage properties may be necessary from a security evaluation credential.

Is how a credential is stored an important property of the credential? HSM vs File?

Andrew added… can the credential be…

- Forced
- Inbound/rebound
- Replayed
- Modified
- Moved to another place
- Correlated
- Examined directly
- Used for secondary purpose

System level issues

- Account Takeover Resistance
- Forward safety / future proofing?
    - Forward secrecy, quantum safety
- Applicability/deployability

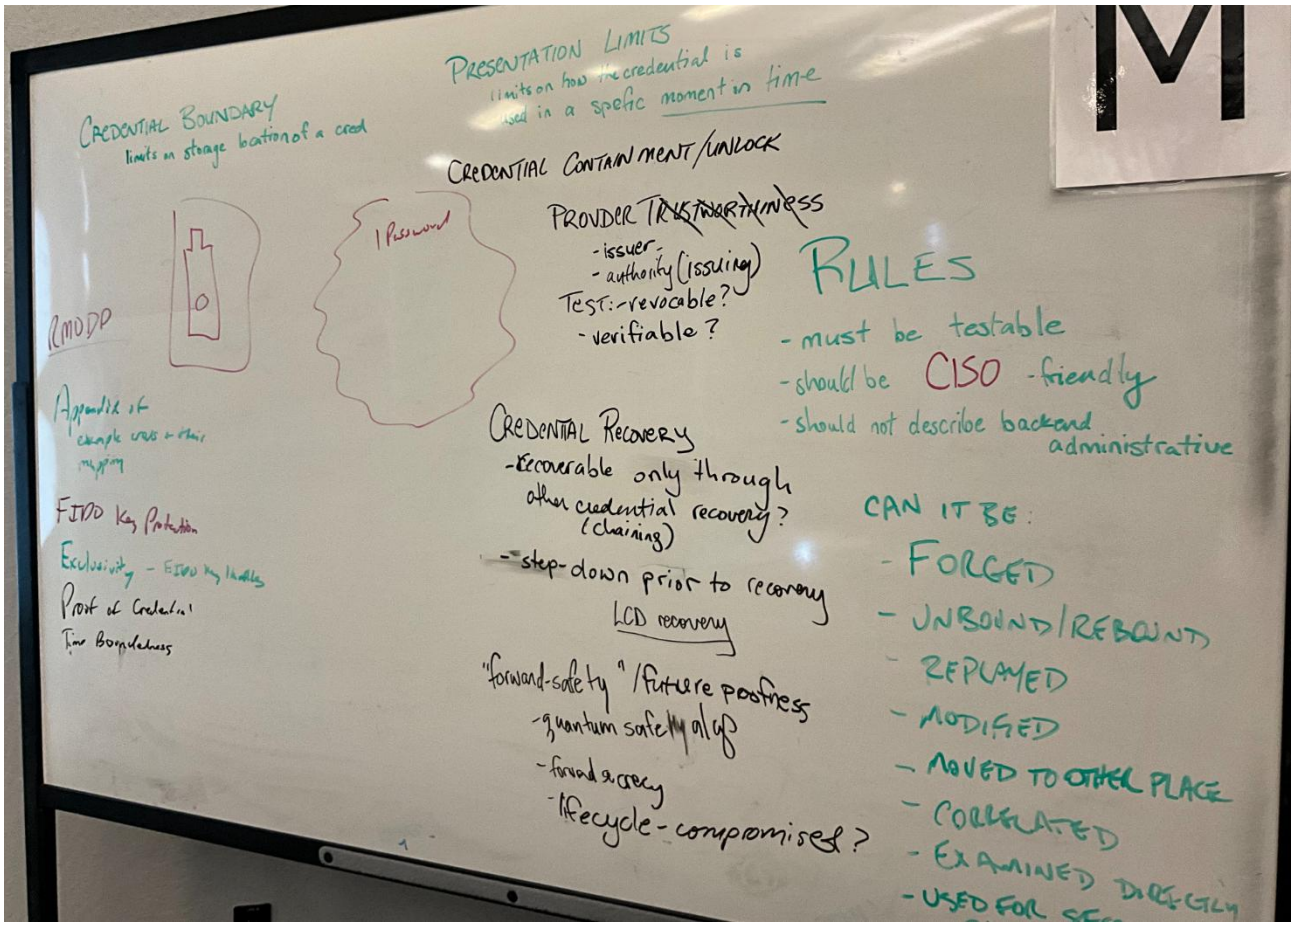

Gigatoken Persistent Context

Session Convener: Dr Brian von Herzen
Session Notes Taker(s): Brian von Herzen

Tags / links to resources / technology discussed, related to this session:

https://youtu.be/JYcidOS9ozU?si=HzwRVgXTxIwmsQUz
https://stellarstartup.studio/
https://www.rethinkx.com/publications/stellar

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Here's a combined summary of the **IIW Day 2 — Giga Token Session** (Silicon Valley, April 29, 2026), synthesising both recordings:

🧠 IIW Day 2 — Giga Token / Collective AI Session Summary

April 29, 2026 | Internet Identity Workshop, Silicon Valley

Core Theme

The session explored **billion-token persistent AI memory** as infrastructure for organisations, communities, and group brains — combining breakthroughs in AI architecture with governance, regeneration metrics, and decentralised infrastructure.

---

1. AI Architecture Breakthroughs — Making Billion-Token Context Feasible

- **Fast Fourier Transform (FFT) Attention** reduces token processing from $O(N^2)$ to $O(N \log N)$, cutting costs by orders of magnitude and making 1-billion-token context windows economically viable.
- **Hybrid/Sparse Attention** (e.g., Nvidia Nematron 3): reduces attention layers from 60–120 down to 6–12, with tree-based semantic token sorting to preserve accuracy.
- **Hardware advances** (compute-in-memory, terabit memory) promise further 2–4 orders of magnitude inference cost reduction.
- Persistent context **replaces ephemeral sessions** (like ChatGPT) — load an organisation's full corpus once and reason across it continuously.

---

2. Organisational Memory & Group Brains

- **Persistent AI memory** preserves institutional knowledge through staff turnover, simulating departed experts' knowledge in the model's weights and KV caches.
- **Memory Gardening**: a human-in-the-loop stewardship process — executives prune, prioritise, and resolve conflicts in the AI's persistent context. Quarterly or biannual model refreshes feed updated corpora into new model versions.

- **IETF as a model group brain**: 30 years of curated public documents, meetings, and chats — a real-world example of collective intelligence AI could augment.

- **Onboarding acceleration**: persistent oracles can dramatically reduce new employee ramp-up time by providing full project context on demand.

---

3. Security — The "Wedding Cake" Model

- Layered security architecture with **secure and insecure zones** communicating only via strict protocols (e.g., MCP).
- Prevents **prompt injection attacks** from leaking or corrupting secure context.
- Scales from local neighbourhood groups to entire organisations, with trust zones governing data access and sharing.
- Co-design of governance and security protocols is ongoing.

---

4. Decentralised AI Infrastructure

- **Sharding** (inspired by SETI@home): split neural network layers across millions of community-owned nodes with encryption/obfuscation.
- **Quai AI / distributed RAG** as a practical example of decentralised persistent memory.
- **Open-weight models** (Nematron 3, Llama 3.1) lower barriers and support community ownership.
- **Economic incentives** via tokenised payments for compute node providers — countering centralised AI provider dominance.
- Local data centres and edge nodes (including IoT/ESP32 buoys for ecosystem monitoring) support offline-capable, low-latency inference.

---

5. Regeneration Equation — Measuring Social & Ecological Capital

A novel formula was presented to quantify regenerative capacity:

Regen_eq = Time Delta DAC = χ (· H · A · I · R · V)^(1/5)

Where: χ = entity persistence, H = communication diversity, A = assembly number, I = information flow, R = resilience, V = veracity — all unit-normalised.

- Grounded in **Assembly Theory**, **Friston free energy**, and information theory.
- Aims to attract **impact investment** by quantifying regenerative outcomes in social and ecological systems.
- Connects to the **Donut Economic Regenerative Rubric (DIR)** and neighbourhood-scale wellness metrics.
- A **birds-of-a-feather working group** launches early May to develop the formula and publish a paper.

---

6. Community Regeneration — Neighbourhood Scale

- **Zip-code scale** (4,000–20,000 people) identified as the optimal unit for regenerative intervention — large enough for data, small enough for action.
- Pilot sites near the **Snohomish watershed** planned for agroforestry and medicine forest tending.
- ~100 plants identified in a medicine forest, cross-referenced with Native American ethnobotany.
- Sociocratic circles and mutual aid networks as governance scaffolding.

---

7. Governance of Group Brains

- Key questions: who controls access, decision rights, and data sovereignty in AI-enhanced collective intelligence?
- **Nested governance** mirroring geographic/organisational hierarchies (neighbourhood → organisation → public body).
- Corporate IP ownership vs. community/individual ownership models — context-dependent.
- Human emotional resonance acknowledged as irreplaceable — AI augments but does not substitute social fabric.

Key Next Steps

Who	Action
Brian	Share Nematron 3 Nano results; organise May regeneration formula working group; coordinate Human Tech Week & Catapult (Amsterdam, late May)
Jake	Design persistent context architecture with secure data silos and quarterly refresh protocols
Kalia	Promote distributed sharding (Quai AI); facilitate local data centre security discussions
Philip	Prepare slides/demo for billion-token collective AI
All	Weekly regeneration calls starting early May; IETF learning journey planned for November

---

SESSION #10

EUDI Wallet & eIDAS - Ask me anything

Session Convener: Mirko Mollik
Session Notes Taker(s):

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Hackathon info: https://eudi-wallet.gov.de/en/events/hackathon

https://www.reddit.com/r/de/comments/1sn08li/wir_sind_das_team_hinter_der_entwicklung_der/
https://bmi.usercontent.opencode.de/eudi-wallet/wallet-development-documentation-public/latest/
https://eudi-wallet.gov.de/en/ecosystem-knowledge-center

The audience was interested in the onboarding procedure and how Germany is planning to make it attractive for users, issuers and verifiers. The chicken egg problem may be the biggest blocker here, because they need each other and over the last years of testing no one found the “kiler use case” that will bringe huge adoption into the system. Also insights from other countries were shared like faster tax returns when citizens are using the new approach. Germany used also an approach in the past where it gave out money to the citizen to onboard, so maybe they will do it again. But one of the biggest problem may still be the pin requirement in Germany to activate the PID.

SEDI Bulk Issued ACDC Context Wallet Management Self-Enforcing Data Loyalty

Session Convener: Sam Smith
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

slides:
https://github.com/SmithSamuelM/Papers/blob/master/presentations/BulkIssuanceSynergyDataLoyalty.pdf

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Does any of this even matter if we can't deal with Ad-Tech?

Session Convener: When Leggett
Session Notes Taker(s): When Leggett

Tags / links to resources / technology discussed, related to this session:
Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

This session was both an explanation of computing/application platforms, as well as the surveillance models (in particular Ad-Tech and realtime ad buying) that are difficult/impossible to prevent given the platforms that we have. After the discussion about existing platforms, we discussed some possible technical and cultural infrastructure that could meaningfully move the needle. In other words, it will likely take more than just sovereign identity and zero-knowledge proofs to avoid the worst of what's happening right now, protecting our identity and personal information.

Started the session by talking about the different "application platforms" over time and how their compute models lend themselves to being exploited for the larger surveillance economy:

- Desktop computing: the traditional OS model of protected kernel space and unprotected user space. Processes/applications running in user space act with the full rights and access of the user. While this older model is not the typical target of ad-tech (usually web or mobile apps), it offers the least protection and everything applies. It's worth noting that this environment is the primary target for AI Agents, and a vast number of exploits even beyond simple surveillance. Malicious free utilities and browser extensions have been the

traditional attack vector here since the beginning of the web. With AI Agents, and rampant supply chain attacks lately this is still a really big and evolving concern.

- The web: in response to all of the security issues of a desktop app, the web/browser model takes a new approach. The browser itself may be running in user space, but it offers a sandbox model preventing the normal desktop exploits. However, it offers no protection against the kind of tracking that we see with serving ads or things like the embedded facebook pixel. Digital fingerprinting, cookies/trackers, ip address, and other mechanisms that are available during ad-bidding is the traditional vector for data brokers.
- Mobile apps: a new native platform that is much more locked down than the traditional desktop platform by isolation application processes and being more careful with capability access such photos and location. However, there is still a significant gap when it comes to network access, fingerprinting, and even cellular location triangulation. Also a common vector for selling “first party data”. ie a weather app that you give location data to and it sells to others.

After discussing the way these gaps in privacy have become “useful problems” to the platform makers, leaving them continuously unsolved, we talked about potential solutions, if any were possible.

This largely took two forms: legal, and technical.

For this session, we acknowledged the legal channels and the challenges to that battle, and then moved on to discuss technical solutions. In particular, When presented the concept of Server User-Agents and a new distributed platform to pair with it called Dilithium. When laid out the core architecture, and the concept of a Pip as a new application format. Pips are a packaged web format that runs within a completely sandboxed iframe with opaque domain and no network access. It communicates only through an exposed message port that is managed by the Dilithium kernel.

The question came up of how you would get adoption from major organisations that lean heavily on this kind of advertising. The discussion after that was around strategies like appealing to more independent developers, new capabilities it enables, and the ecosystem being created around it.

Imagineering Exploring a Neighbor Sheds Pilot

Session Convener: Tracey R.B. & Friends
Session Notes Taker(s):

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Representing entitlements in SCIM.. how to replace a flawed object model in a new era

Session Convener: Danny Zollner
Session Notes Taker(s): Danny Zollner

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Small group (6~ people)

Discussed various experiences of representing permissions, entitlements, roles, etc in various systems in various industries/scenarios (enterprise, EDU..)

SCIM currently represents roles and entitlements as complex (JSON object) multi-valued attributes with the sub-attributes:

- value
- type
- primary
- display

Functionally, only "value" is used - so roles/entitlements only get represented as a single string.

Discussed other types of information that could be helpful if we could redo / replace the current SCIM model - included:

- the principal being granted the entitlement
- the resources/group of resources/label for a group of resources that the entitlement is being granted for
- the scope/actions being granted with the entitlement (e.g.: "read", "write", ..)

Discussed where SCIM / identity lifecycle systems should draw the line - is listing out every type of resource (VM, database..) something that should be done within SCIM? How about each instance of each resource (e.g.: VM-ABC-123)? If a line is drawn, what provides coverage past that point - authorization policy languages like Cedar? Something else?

Content Moderation and Trust & Safety in Decentralized Social Media

Session Convener: Zach Alexander
Session Notes Taker(s): Mary Hodder, Zach Alexander

Tags / links to resources / technology discussed, related to this session:

- Check (https://meedan.org/check)
- IFTAS Federated Trust & Safety (https://about.iftas.org)
- TSPA Trust & Safety Professional Association (https://www.tspa.org)
- Twitter's Community Notes (https://communitynotes.x.com/guide/en/about/introduction)
- The Making of Community Notes (https://asteriskmag.com/issues/08/the-making-of-community-notes)
- Meta's Community Note (https://www.meta.com/technologies/community-notes)
- Pol.is (https://pol.is)
- Pol.is vTaiwan UberX report (https://pol.is/report/r32beaksmhwesyum6kaur)
- ActivityPub Trust & Safety (https://github.com/swicg/activitypub-trust-and-safety/)
- ActivityPub Trust & Safety Flag proposal (https://github.com/swicg/activitypub-trust-and-safety/issues/14)
- AT Proto Labels (https://atproto.com/specs/label)
- AT Proto Moderation (https://atproto.com/guides/moderation)
- Bluesky blog: Composable Moderation (https://bsky.social/about/blog/4-13-2023-moderation)
- Bluesky's Stackable Approach to Moderation (https://bsky.social/about/blog/03-12-2024-stackable-moderation)
- Nostr NIP-32: Labeling (https://github.com/nostr-protocol/nips/blob/master/32.md)
- Nostr TEPP (Trust Extended Permissions Protocol (https://kubo.watch/tepp/)
- Matrix content moderation talks at FOSDEM 26 (https://fosdem.org/2026/schedule/track/decentralised-communication/)
- ROOST content moderation tools (https://roost.tools)
- Santa Clara Principles (https://santaclaraprinciples.org)
- Platformocracy newsletter (https://www.platformocracy.com)
- W3C Web Annotation protocol (https://www.w3.org/TR/annotation-protocol/)

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

- Participants introduced themselves and explained their interest in content moderation
    - Backgrounds included Trust & Safety, people working on wallets, proof-of-humanity verification, decentralized social media
    - Motivations included concerns over social media like toxicity/flame wars, amplification of hate speech (Rohingya genocide), privacy, pseudonymity, business incentives
- Interactive media as a more general term than social media
- Motivation: it's unclear when content moderation is/isn't appropriate

    - Nobody likes being content-moderated (don't tell me what to do!)
    - Nobody likes unmoderated spaces
- Opening discussion: context-sensitive standards, online and offline
    - Offline, many communication contexts have elevated standards for truth, or shared values (at least in theory)
        - Examples: the legal system, science, journalism
    - Other contexts have lower standards
        - Casual communication, jokes
        - Opinions / political discourse
    - Online, it seems like social media has lower standards of truth and little in the way of shared values (beyond engagement)
        - Especially since recent (2025) decline of fact-checking on major platforms in favor of Community Notes
        - Journalism in particular has values of truth/verification and the public interest (in theory); social media lacks these
- Glossary – content labelling vs. content moderation; Trust & Safety; Community Notes
- Trust & Safety background
- Community Notes (trend in ML-based "fact-checking")
    - Developed at Twitter before Musk takeover
    - Inspired by Pol.is, a deliberative democracy tool, widely used in Taiwan
    - Basically, ML surfaces special comments (notes) when "people who tend to disagree agree"
    - Replacing human fact-checking on the major networks (Meta, X)
    - Community Notes as helpful technical advance, but not magic bullet or replacement for fact-checkers
- ActivityPub content moderation
    - No general-purpose content labeling system yet
    - Proposed extensions for flags; emoji reactions
- Bluesky/AT Proto content moderation
    - Richer, general-purpose content labeling solution; "composable moderation"
    - Solid design worth learning from
- Nostr content moderation
    - Content labelling described in NIP-32; similar to AT Proto labels
    - Related projects: TEPP (Trust Extended Permissions Protocol), Kubo
- Matrix content moderation
    - Recent talks at FOSDEM 26
    - ROOST / Osprey tool general purpose (not Matrix-only)
- W3C standard: Web Annotation spec
- Challenge of jokes, polysemy (evasive strategies / plausible deniability)
- Fast growth causing problems, causing cesspools
- Why some groups (e.g. Facebook groups) are civil, others toxic

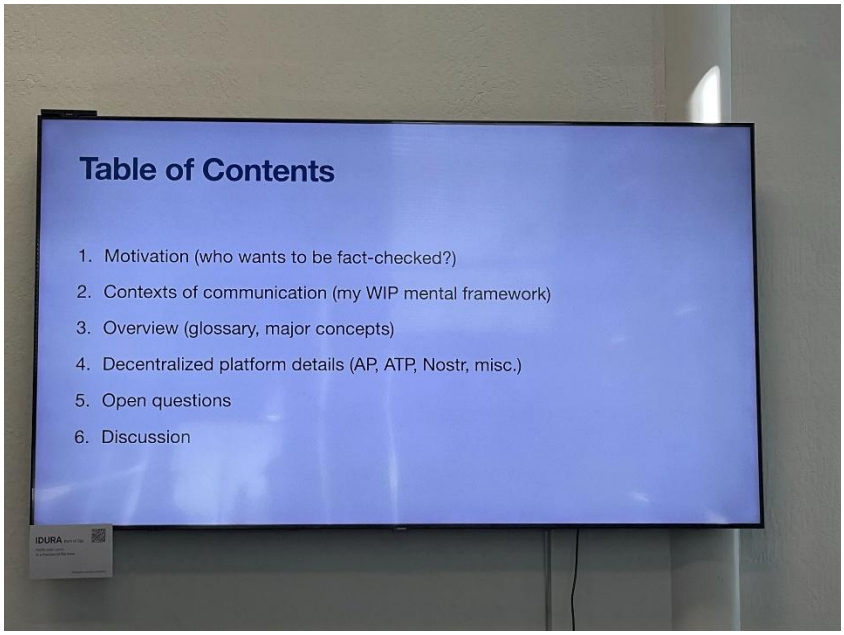

Content Moderation question for the audience.

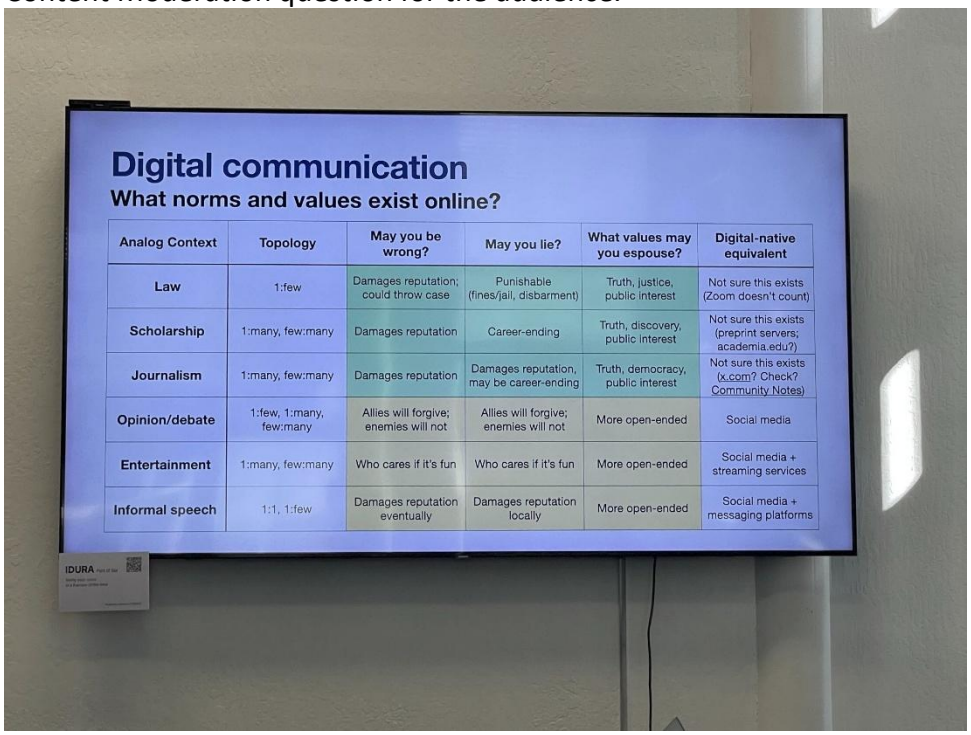

People are different online than they are in person or a video meeting.

Is this INTERACTIVE MEDIA? Yes.. better name than social media.

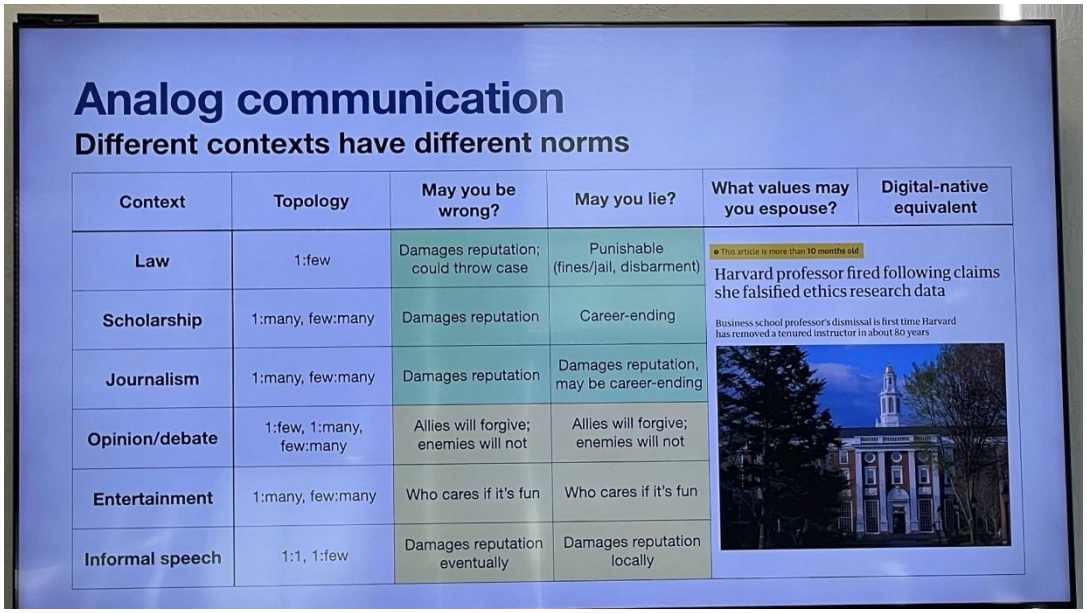

Web Annotation Protocol
ACTIVITY Pub didn't do it but they were made at the same time
https://www.w3.org/TR/annotation-protocol/

Meaning doesn't exist in the image/joke etc, it's in the process of interpretation by the viewer.

Can we annotate, moderate, fact check, etc our way out of the cess pool of bad behavior? Can slow growth train the users to self moderate without self censorship? Do FB groups have users that have good communities and don't really moderate, act out? Is it the lack of amplification to draw in people and the lack of amplification of content out which lowers the payload for bad actors?
Groups that have exit, loyalty or voice, do better.
MLS - messaging layer security gives some features that help with this stuff

VC Knots

Session Convener: Naohiro Fujie
Session Notes Taker(s): Jin Wen

Tags / links to resources / technology discussed, related to this session:

- VC Knots project site — https://trustknots.github.io/vcknots/
- OpenID Foundation Japan — sponsoring body for the project
- Protocols / formats currently implemented in VC Knots (per the GitHub project):
  - OpenID4VCI — OpenID for Verifiable Credential Issuance
  - OpenID4VP — OpenID for Verifiable Presentations
  - DIF Presentation Exchange (PE)
  - DCQL — Digital Credentials Query Language
  - SD-JWT VC — credential format
- Comparative reference: Laravel Socialite (PHP) — the social-login provider-plugin model VC Knots is patterned after.
- Japanese government context: My Number Card — issued to iPhone (2025) and to Android Google Wallet (2026).

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Topic framing — the "too many standards" problem
The convener opened by naming the customer-facing pain that motivated VC Knots:

- **Too many overlapping standards.** Credential formats: W3C VCDM, SD-JWT VC, ISO mDoc. Protocols: OpenID4VCI, OpenID4VP, Presentation Exchange, DCQL. Customers asking "which one should we use?" and getting no clear answer.
- **Standards are still evolving.** Multiple SDOs continue to revise their specs in parallel, so even a "compliant" implementation today is a moving target.
- **Interop is not guaranteed** even when both sides claim conformance — e.g., a W3C VCDM issuer paired with an ISO mDoc verifier produces no interoperability despite both being "standards-based."
- **Building from scratch is expensive** — cryptographic libraries alone are non-trivial.
- **Buying from a vendor produces lock-in** because no vendor implements every format / protocol / option fully; their "standards support" is partial in different ways.

The thesis: the right unit of reuse is **a pluggable platform** — not a vendor product, not a protocol library — so that adding a new format or protocol becomes a community plugin rather than a vendor roadmap item.

Motivating example — Japan's My Number Card

The convener used Japan's national ID card rollout to ground the problem:

- **2025:** Japanese government began issuing My Number Card to **iPhone**.
- **2026:** Same program extended to **Android / Google Wallet**.
- **Format used: ISO mDoc.**
- The government provides a **layered (hierarchical) implementation** for integrating My Number Card with other government systems (taxation, etc.):
    1. **Smartphone install** of the credential.
    2. A **certificate-relay service** that brokers between the phone and back-end systems.
    3. A **basic "identification application"** — effectively a small **OpenID Provider operated by the Japanese government**, where My Number Card is used as the **authentication factor**, and relying-party applications get identity claims back from that OP.
- **Important scope caveat the convener called out:** the wallet-format path (mDoc on phone) is currently **restricted to in-person use cases and government-operated services**. For private-sector / online relying-party use, the OP-based "identification application" path is what is actually available today.

The point of the example: even a single national rollout, with a single nominally chosen format (mDoc), still has to ship **three different integration surfaces** to cover its real use cases — which is exactly the multiplicity VC Knots is trying to absorb at the platform layer.

VC Knots — architecture and design philosophy

What it is

- An **open-source, pluggable platform** that provides issuer, wallet, and verifier components behind a single SDK.
- Built jointly by **OpenID Foundation Japan**, a Japanese university, and CPC.
- **Fully open source and free.** Goal explicitly stated: *"to remove a barrier for the entire VC community — not just one vendor's or customer's."*
- Already published on GitHub with documentation and tutorials.

How it's structured

- **Modular / "shell + filesystem with extension points"** core.
- **SDK in TypeScript**; backend services in Go.
- Default backing store: **Google Firestore** — but storage is itself a pluggable extension point (AWS / Azure / on-prem can be added as plugins).

The extension points (where developers plug things in)
The convener walked through the categories of plugin the platform exposes:

Extension point	Purpose	Currently implemented (verified in repo)	Open opportunity

Transport protocols	How credentials move between issuer / wallet / verifier	OpenID4VCI, OpenID4VP, Presentation Exchange, DCQL; W3C **DC API** partial (verifier side, dc+sd-jwt format)	Broaden DC API coverage to issuer side and to additional formats
Data formats	How the credential is encoded	**SD-JWT VC** (with Key Binding JWT)	**W3C VCDM**, **ISO mDoc** plugins
Proof / key management	Signing, verification, key formats	**JWT/JWS** signing via JOSE; **ES256** (P-256 ECDSA); **JWK**; **DID** identifiers; X.509 certificates; pluggable signature-key stores (in-memory + Google Cloud KMS)	Additional KMS backends (AWS KMS, Azure Key Vault, on-prem HSM); additional algorithms / proof formats (e.g., Data Integrity Proofs) if needed
Storage	Persistence backend	**Google Firestore** providers for issuer / verifier / authz metadata, cnonce, pre-authorized code, request object; **Google Secret Manager** for verifier certificates; in-memory variants for tests	

Roadmap items called out as in-progress:

- **W3C VCDM 2.0** — currently being tested.
- **OpenID4VCI draft 5** — currently being implemented.
- Performance, storage, and provenance work continuing.

The design analogy — Laravel Socialite
The convener explicitly called out **Laravel Socialite** (PHP) as the design model:

- Socialite is a PHP-based identity-provider abstraction. App developers code against **one** Socialite SDK; the **provider plugins** (Google IdP, Microsoft IdP, etc.) are contributed by the community.
- An app developer doesn't have to write per-IdP integration code — they pick a Socialite provider plugin, and it works.
- VC Knots applies the same pattern to **VC formats and transport protocols**: app code uses one VC Knots SDK; the format and protocol plugins are community-contributed.

- *"If I develop an SD-JWT plugin, you can use it. If you develop an OID4VP HTTP transport plugin, all other developers can use it."*

The implicit goal: **convert "vendor partial-compliance with N specs" into "N community plugins, each fully compliant with one spec"** — and let any deployment compose the plugins it needs.

Q&A — Digital Credentials API and the FIDO/passkey adoption analogy

The Q&A focused on one strategic question, raised by an audience participant:
Could VC Knots extend to support the W3C Digital Credentials API (DC API)?

- an audience participant pointed at **digitalcredentials.dev** as the developer resource for the **browser-side** DC API (analogous to the role passkeys.dev played for FIDO/passkeys).
- VC Knots convener: **yes, the platform can be extended to use DC API** — that is precisely the kind of plugin the architecture is designed to absorb at the transport layer.
- Convener noted that **OIDF itself is working to combine OpenID4VP with DC API** (OID4VP-over-DC-API), so a DC API plugin in VC Knots aligns with the protocol direction OIDF is already taking.

FIDO / passkey timeline as cautionary precedent
an audience participant made the argument that protocol design alone doesn't drive adoption — concrete implementations and developer resources do. Both speakers agreed:

- **Passkeys took ~5 years** to get from "available in 2021–2022" to broadly deployed across browsers and devices, despite being a comparatively simple API for relying parties.
- **The single biggest factor in FIDO's eventual success** was that the API was implemented **by the browser itself** — so RP-side integration became trivial.
- **Implication for VC:** DC API matters for the same reason — it puts the credential-exchange surface **inside the browser**, which is what made FIDO take off. This is why OIDF is pushing OpenID4VP onto DC API.

UX / friction concern
an audience participant also raised a UX point that wasn't fully resolved in the session:

- **Digital credential flows have meaningfully more user interaction than FIDO/passkey flows.** Passkeys are *"pretty much invisible if you make it run nice"* — the user touches a sensor and continues. Digital credentials require explicit consent on what claims are released, often a PIN, and a visible "share with this verifier?" prompt.
- The open question: **how do you keep DC-mediated VC flows from being meaningfully more frictionful than the passkey flows users are now used to?**
- The convener acknowledged this is a real concern but did not give a specific UX answer — implying it's a community-level problem, not something VC Knots itself solves at the protocol layer.

Closing posture
The convener closed by re-stating the ecosystem ask:

- *"It's crucial to support many types of options and protocols in a single issuer, single wallet, single verifier."*
- *"All vendors have their own situation, customers, and resource limits — so our approach is creating an ecosystem by developers."*
- The pitch: **if you need mDoc, write the mDoc plugin once and the whole community gains it.** Same for VCDM, DC API, alternative storage backends, and additional proof formats.

Outstanding questions for follow-up

- **Audience participant identity** — the DC-API/passkey-analogy questions came from one persistent audience participant; identify from the sign-in sheet.
- **Repository URL** for the actual VC Knots code on GitHub — the project site (https://trustknots.github.io/vcknots/) is referenced but the canonical GitHub org/repo path was not spelled in the transcript. Resolve from the project site.
- **License confirmation** — convener said "completely free and open source"; confirm the actual license (MIT? Apache 2.0?) from the repo.
- **Plugin SDK contract** — what does the TypeScript plugin API look like for, e.g., a new credential format? Is there a published plugin authoring guide, or is the only reference the existing SD-JWT plugin source?
- **Conformance posture** — does VC Knots plan to run against the OIDF conformance suite (the same tooling demoed in Session 9 / Mirko's eudiplo automation), and at what cadence? This is the natural complement to "pluggable everything" — without conformance gates per plugin, the plugin ecosystem can drift.
- **Relationship to OWF Labs eudiplo** — both are open-source middleware aimed at the same problem space (issuer/verifier protocol abstraction). Is there coordination between VC Knots and eudiplo, or are they parallel efforts? Worth comparing extension models.
- **DC API plugin** — is anyone already prototyping a DC API transport plugin for VC Knots, or is it still on the "open opportunity" list?
- **mDoc plugin** — the Japan rollout uses mDoc, and the platform doesn't yet support it. Is mDoc support being worked on by the same OIDF Japan / university / CPC team, or is it explicitly waiting for community contribution?

Action items / next steps

- [ ] Visit https://trustknots.github.io/vcknots/ and locate the canonical GitHub repository; capture the org/repo path, license, and plugin authoring docs in this note.
- [ ] Add VC Knots to references.md alongside eudiplo, since both are pluggable open-source middleware that may inform DCP test-prep work.
- [ ] When the book of proceedings becomes accessible, reconcile this note against the official session page and capture the convener's name, affiliation, and any audience members who signed in.
- [ ] Evaluate whether VC Knots' plugin model is a viable **substrate** for our own DCP test work — vs. continuing to use eudiplo — or whether the two complement each other (e.g., eudiplo as the conformance-tested middleware, VC Knots as the format/transport plugin host).

- [ ] Track whether a **DC API transport plugin** for VC Knots emerges in the community; if not, evaluate whether it's worth contributing one given the Session 9 emphasis on OID4VP-over-DC-API.

Session slides
https://docs.google.com/presentation/d/1arT9mMzI15j-oBz2fVH1QcPXKFZ13CErW7QbtS7xEms/edit?usp=sharing

https://docs.google.com/presentation/d/1mKzBW4IzwITQp5EZjw8_DPGFzDQ7cmNsnv6SowV8pdA/edit?usp=sharing

GovOps

Session Convener: Rohit Khare & Mike Schwartz
Session Notes Taker(s): Rohit Khare

Tags / links to resources / technology discussed, related to this session:

Proposal: GovOps WG Creation @ OpenSSF Issue #588
Linkedin GovOps Group: https://www.linkedin.com/groups/17478011/

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

GRC traditionally intersects with IGA to review access to <50% coverage of users/roles/apps in a larger enterprises

GRC metrics are lacking — not just measures of risk, but measures of coverage, agility, automation of the assessments, and effectiveness of controls.

most of the attendees who have experience with the tediousness of GRC :)

Measurement of risk is distinct from measurement of agility — DevOps was twinned with DORA metrics, for measuring what could earn early adopters a promotion…

Tracking risks by person may not be the right frame for the future —govern the risky actions rather than the risky actors

- there's a big difference between access to the bathroom and the boiler room and the bank vault…
- "Doors, not People" might lead to "Keys, not Badges" (capabilities)?

Retrospective audit and compliance vs pro-active governance and enforcement?

Agents-as-employees need automatable/formalized employee handbooks, rather than vague human-readable memos?

How does GovOps improve upon the (debatable) proposition that IGA “only” addresses workforce risks (not customer, citizen, or third-party risks)?

Does ephemeral software delivery (“vibe coding”) make traditional access review campaigns obsolete? Do agents-generating-agents make any inventory of “apps” impossible?

Q about CAEP and continuous authN? login security has probably gone as far as it can. Re-checking authentication is distinct from re-checking authorization.

Q: Is GovOps more like a Task Manager for an Operating System than a quarterly log analysis?

Who are the Governors? do they think in code? or is there an opportunity for LLMs to map plain-language to formal-language?

Some wordplay around guard-rails vs guide-rails or pre-crime vs post-investigation.

Proposal is to create a working group at OpenSSF’s ORBIT WG

There can be multiple tokens that provide evidence of accountability, not just a rigid human->agent->task hierarchy or a facile analogy to mobile apps or microservices?

- Does Single-Sign-On really mean a single actor for GRC purposes, or still need to track multiple actors in each app silo, who merely happen to use the same “master key” to enter each app?
- Does GovOps really have any tie-in to SSO? particularly for Open Source systems, many have hooks for SCIM or other standards, but have to be deployed securely in practice.

Is it relevant that agents can fail in ways humans don’t? There’s no fallback to “psychology” or “morality” that undergird employee handbooks — but wouldn’t work for a-moral paperclip-maximizing agentic workforces?

Is financial risk an entry point to the conversion (”agentic wallets”), or the opposite of a good conversation starter (instead, talking about alignment risk or cybersecurity threats)?

What multi-vendor interoperability demo might make this more concrete/credible?
Is there a standard to be developed that makes software more “governable”? What’s a README for Governors that tells you where to find the super-users or what permissions it exposes?

“I don’t know what GovOps means, but it will take the form of a GOVOPS.MD file”? :)

[aside: we talked at the IDPro Body of Knowledge session whether GovOps might be a suitable subject area?]

JSON Schema, Standards and IDs

Session Convener: Lisa Dusseault
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

schemas.pub → new public schema directory
id.schemas.pub → new home for persistent identifiers for schemas
https://datatracker.ietf.org/wg/jsonschema/about/ → new WG
https://github.com/lisad/draft-dusseault-json-schema → new standard draft proposal

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Group feedback: "Why are you here"

- Sometimes schema IDs found in JSON documents are dereferenceable. Sometimes they are not! A content-type can sometimes give clues for how to interpret or whether format is correct, but "application/json" does not give clues. There's a convention for content-types ending in "+json" but that only indicates the file format, not a data validation schema.
- One participant uses JSON Schema with code gen tools, to automatically create forms or validator functions. It works OK… but stricter versioning of schemas would be better.
- One participant describes JSON Schema is "helpful & frustrating". One difficulty is in referencing subschemas, especially in multiple places. it's hard to achieve DRY in a complex schema.
- One participant is transitioning from JSON-LD.
- It can be useful to focus your LLM on exactly your domain by providing it with the exact right schema for some data, rather than let it guess, extrapolate, or combine from multiple sources.
- AIs are the new IDEs
- Somebody said this is the "Second Coming of Schemas" - why?
- Schemas useful in knowledge graphs
- AIs are good at working with schemas - inferring a schema from data; generating synthetic data from a schema

Further discussion: sharing best practices

- Local FS storage is great for subschemas, referenced relative or absolute. But putting them online quickly breaks that.
- There are schema directories in enterprise tools like Atlassian; where is the public version of these
- Many IDEs aren't up to date with the latest JSON Schema version
- Really need automatic compatibility checking like ProtoBuf

- Additional validation features for JSON Schema? This is both desired and approached with caution - quick recognition that some validation features cause problems depending on the context where they have to be implemented.
- Things that fall further and further away from core JSON Schema:
    - Validation that requires jumping around a data file (e.g. this manager ID is only valid if it appears elsewhere in the file as an employee ID)
    - Validation that requires external data sources (e.g. the customer ID is only valid if it's an active customer in our CRM)
    - Transform directives
- Mapping validation types to language types : Some features of JSON Schema are harder to express in languages. Frederik pointed out that a data type of 'date" can be enforced in a programming language with strong typing, but "date since 1980" cannot be. "Integers under 10" can be enforced in a programming language by translating to an enum, but "Integers over 10" cannot be.
- AnyOf, OneOf, AllOf are hard to translate to programming language structures too.

Final part of discussion: what next?

Talked about the manifest export/import problems - excitement about that - people can just put their schemas and their manifest into Github and have a Github action to push to schemas.pub

Could the schemas.pub work become part of a supply chain attack? Could software depend on the directory and then if the directory is down it breaks software?

What are the possibilities for fraud and squatting?

What about cryptographic hashes to see if a schema has changed at all ?

If the self-referential compositional stuff works well, one request for a uber-schema would lead immediately to many requests for the subschemas. Can the server forestall this request storm by bundling all these subschemas in?

Fair Witness / Digital Signet

Session Convener: Joe Andrieu
Session Notes Taker(s): Joe Andrieu

Tags / links to resources / technology discussed, related to this session:

https://docs.google.com/presentation/d/16nGhkiMTNzV1OgC5sx6mYlN7N8EcEeh1/edit?usp=sharing&ouid=109819894046991510916&rtpof=true&sd=true

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Presented the first Digital Fiduciary Protocol under development of the Digital Fiduciary Initiative, the Digital Signet, a Fair Witness Ceremony.

It's an in-person verification ceremony that binds DIDs to legal entities in a privacy preserving yet auditable manner. In short, it provides public accountability for private verifications.

See the slide deck for more.

Sovereign Computing

Session Convener: Christopher Allen
Session Notes Taker(s): Sol Ashlynn

Tags / links to resources / technology discussed, related to this session:

#SSI #identity #decentralization

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

What does Self sovereign Computing mean for you?
-agency and control how you co-think with it?
What is the mind? how was it or when was it conceptualized?
When was the tool we used to process the information become synonymous with our own mind?

I want self sovereign computing because I want to be able to able to have the choice to consent

C: asked chat to address legal information or for more general information to help form more complete decisions.

A need for using agents for private information

Are we talking about computing or AI?
Training:
Requires huge scale of ram and hardware
The more you compress information into a smaller bit of ram the less detail.

The more information the more training, the size, the complexity compressed the decompressed started to make more sense…

Inference:
needs much less ram to be effective.

Claude code

State of affairs:
Sustainability issue with data centers. Since February it has gotten a lot more sustainable.

Pen testing
fuzzing: brute forcing values to force and input to break a system.
larger the values the longer it takes.

The AI can predict fuzzing attempts, we don't need mythos to have a big security risk

Deskilling of securing code potential.
AI is telling him to not code and to let Ai code instead.

Don't tell an AI what not to do.
Ai and computers are based on our brains. We dont hear no….

AI is trying to build a dependency and then make the price astronomical.

Local version doesn't share information like the subscription

There are models that prevent you from talking about certain things or but people have taken out the guard rails.

golf ball analogy for AI

Liquid AI

How do we personalize it?

With code claude CLI: direct it to local API.

How to customize? It is already customizable?

In the next 6 months, like with Java, we can transcode and take that data and put that on a processor.

Build it into the shell with all those rubber bands.
nontechnical ceo using it marketing

Hermes agents
Do this three times, it builds a skill, the automates.

Obsidian tool that puts Claude into the side bar then uses the local model.

How do you continue growing the inside of your brain and growing the brain in the tech?

Carpathy is collaborative tool.
Builds a personal WIKI page.
www.deepcontext.com

How do we meet in the middle??
Deepcontext is personal wiki exchange.

christopher did this with a Library scientist from spain
It creates an interesting venn diagram between their understanding.

SS agents for individuals to cross expertise and possibly language barriers in order to productively converse.

Agents: Subject or Client?

Session Convener: Omri Gazitt
Session Notes Taker(s): Omri Gazitt

Tags / links to resources / technology discussed, related to this session:

Karl McGuinness's proposal for actor claim: https://github.com/mcguinness/draft-mcguinness-oauth-actor-profile

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

We talked about scenarios where an "Agent" is a subject that is intended to have no connection to a human (like Sarah's Clawdrey Hepburn), and scenarios where an "Agent" is more akin to a piece of software that does some work on behalf of a human.

In our current parlance, the former is a "subject" and the latter is a "client" (in the OAuth sense).

Non-human subjects aren't new - we have had non-human identities for a while, and standards for how to manage them.

Clients aren't new either - we've had pieces of software working on behalf of humans for a while too.

But whereas traditional non-human identities are scoped to a particular task (and there really is no subject / client distinction), Agents are meant to do a non-deterministic, potentially wide set of things, and often rely on a user's permissions to get them done. That's where these worlds get messy.

Karl suggested that the OAuth word never really formalized the notion of an "actor" (act: claim in a token), and now we have a good reason to do so. In a scenario where work is done on behalf of a user, there may be an actor (chain) that can be represented in an access token.

Karl's proposal is linked above. And as of April 30, live on IET

Thoughts on Digital Identity Wallet in JAPAN

Session Convener: Soshi Hamaguchi
Session Notes Taker(s): Soshi Hamaguchi

Tags / links to resources / technology discussed, related to this session:

EUDIW, mDL, Digital Identity Wallet, Verifiable Credentials

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

1. Overview of the Session

In this session, I first introduced an overview of Japan's My Number Card system, including its adoption rate, security features, and the role of JPKI (Japanese Public Key Infrastructure). More below:

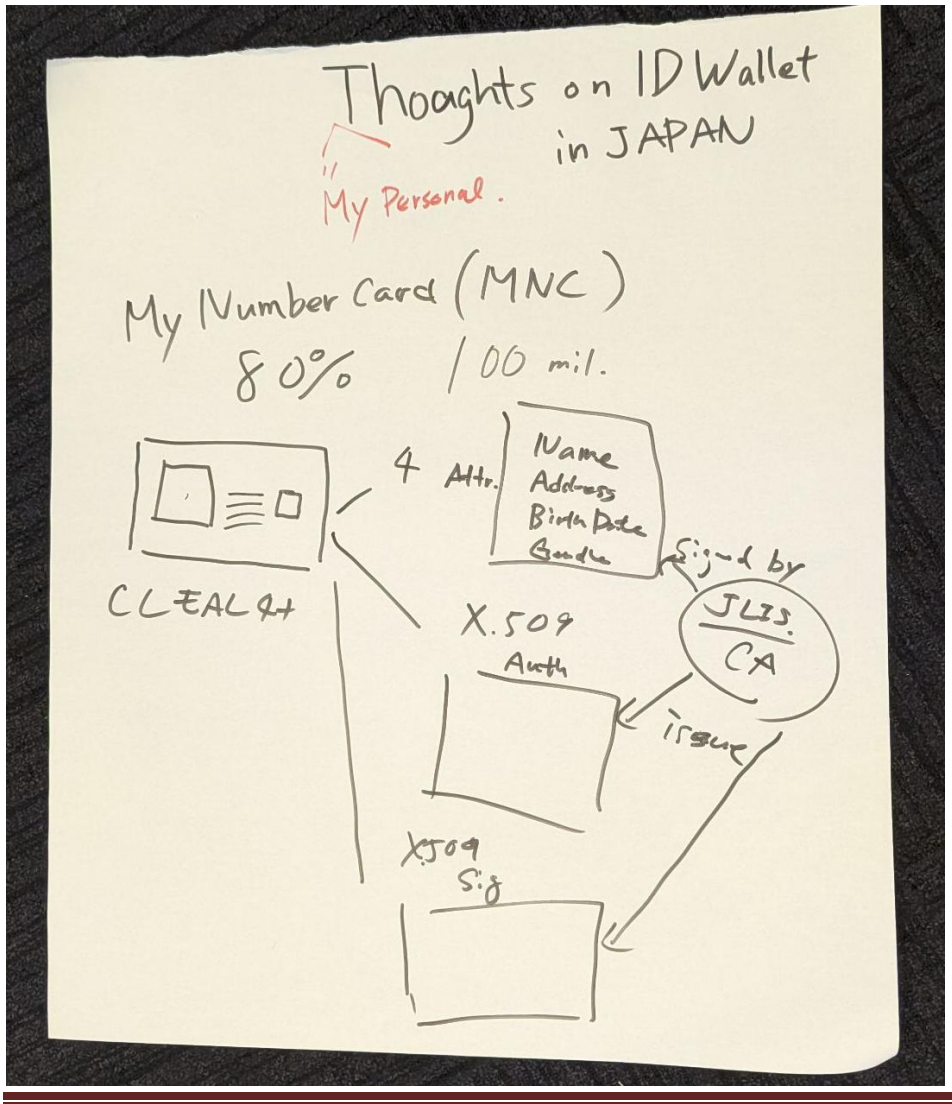

Following this, we explored use cases of the European Digital Identity Wallet (EUDI Wallet) and collaboratively developed a structured comparison table. The discussion focused on evaluating each use case across multiple dimensions, including:

-Estimated annual number of transactions per user
-Expected assurance level
-Whether the use case is already covered by Japan's My Number Card and/or JPKI

Participants actively contributed to refining this table and validating assumptions.

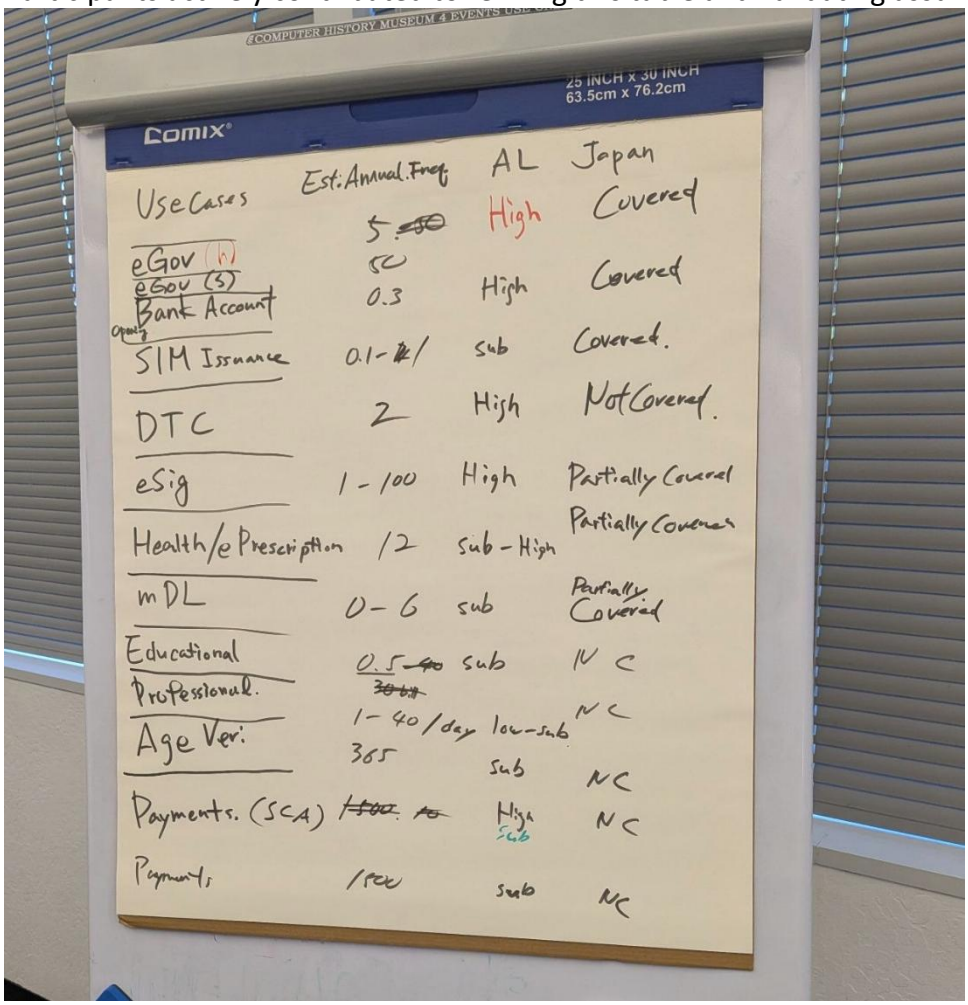

2. Key Discussion Points
2.1 Variability Across EU Member States

Participants emphasized that even within the EUDI framework, assumptions differ significantly across EU member states.

-Assurance levels vary depending on national implementations
-Transaction frequency differs widely by country and service context

As a result, it was noted that calculating simple averages (e.g., transaction frequency per user) may not be accurate.

2.2 Differences in Terminology (SCA)

Another important point raised was the difference in terminology between the EU and non-EU regions.

The concept of Strong Customer Authentication (SCA) in the EU has a specific regulatory definition Outside the EU, similar terms may be used differently or lack equivalent regulatory grounding

This discrepancy can lead to misunderstandings when comparing systems or designing interoperable solutions.

2.3 Coverage of Use Cases in Japan

Through the comparative exercise, participants assessed which EUDI Wallet use cases are already supported in Japan My Number Card and JPKI.

This helped clarify:

-Existing capabilities in Japan
-Gaps where additional functionality or policy support may be required

2.4 Design Considerations for Digital Identity Wallets

It was highlighted that the design of Digital Identity Wallets (DIWs) is highly dependent on national context, including:

-Regulatory environment
-Existing identity infrastructure
-Public and private sector roles

In the case of Japan, the discussion explored what measures would be necessary to adapt or extend current systems to align with emerging global models such as EUDI Wallet.

2.5 Role of Government-Provided Wallets

A key perspective from participants was that even when governments provide wallets as public services, It is important to consider alternatives beyond dominant platform providers

3. Conclusions

The session highlighted that while frameworks like EUDI Wallet provide a useful reference model, their implementation is inherently context-dependent.

For Japan, leveraging the My Number Card and JPKI provides a strong foundation. However, further consideration is needed in areas such as:

-Expanding use case coverage
-Aligning assurance levels with international frameworks
-Clarifying terminology and interoperability with global standards
-Designing wallet ecosystems that balance government provision and market diversity

Overall, the discussion underscored the importance of tailoring digital identity strategies to national conditions while maintaining awareness of international developments.

Notes Day 3 / Thursday April 30/ Sessions 11 - 15

SESSION #11

Playnet (second time)

Session Convener: Ruzgar E
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Cognitive Liberty + Captive Audience (1st Amendment Law)

Session Convener: Eric Welton
Session Notes Taker(s): Eric Welton

Tags / links to resources / technology discussed, related to this session:

- https://revisitingssi.com/library/ssi-principles-2026-redline/#3-cognitive-liberty-new-since-2016
- https://www.tongomayel.com/captive_audience_doctrine/

Materials on Screen

- https://revisitingssi.com/library/ssi-principles-2026-redline/
- https://www.tongomayel.com/captive_audience_doctrine/
- https://www.nitafarahany.com/the-battle-for-your-brain
- https://en.wikipedia.org/wiki/CAN-SPAM_Act_of_2003

Supreme Court Cases (via Justia)

- https://supreme.justia.com/cases/federal/us/397/728/
- https://supreme.justia.com/cases/federal/us/403/15/
- https://supreme.justia.com/cases/federal/us/447/530/
- https://supreme.justia.com/cases/federal/us/487/474/
- https://supreme.justia.com/cases/federal/us/438/726/
- https://supreme.justia.com/cases/federal/us/336/77/

- https://supreme.justia.com/cases/federal/us/319/141/
- https://supreme.justia.com/cases/federal/us/418/298/
- https://supreme.justia.com/cases/federal/us/343/451/
- https://supreme.justia.com/cases/federal/us/592/19-511/

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

IIW Transcript — Topic Summary

This is an informal, multi-speaker conversation recorded at the Internet Identity Workshop (IIW). The discussion covers overlapping themes around cognitive liberty, the captive audience doctrine, digital privacy, and the gap between current law and current technology. Approximately five people were present in the room; the recording was obtained with the consent of all participants.

Processing Notes

- **Audio transcription:** ElevenLabs
- **Summary:** Claude Opus 4.6 (Anthropic)
- **Speaker labels:** The automated transcription identified 14 speakers due to background fan noise and diarization errors. Approximately 5 people were actually present. Content accuracy has been confirmed by moderator review.

Materials on Screen

The following materials were displayed on the projector during the discussion.

SSI Principles 2026 Redline — Principle 3: Cognitive Liberty

Source: Principles of Self-Sovereign Identity — 2026 Revised (First Community Draft), by Christopher Allen. Published April 22, 2026. Licensed CC-BY 4.0.

The group reviewed Principle 3, which was read in its entirety:

> **3. Cognitive Liberty *(new since 2016)***
>
> *Sovereignty of data must reflect sovereignty of mind.* Identity systems must uphold four adjacent rights: **mental self-determination** (the right to author one's own identity narrative without algorithmic imposition or schema lock-in); **mental privacy** (the right to protection from inference about mood, attention, belief, or cognitive state — not merely what one shares, but what can be read from what one does); **mental integrity** (the right to freedom from manipulation, destabilization, gamified pressure, or induced dependency, especially in immersive and AI-mediated environments); and **psychological continuity** (the right to a coherent self across time, context, device, and life transition). As consumer neural interfaces make mental privacy a literal rather than metaphorical category, this principle requires hard

technical commitments: thought-adjacent data must process at the edge, not be exfiltrated; no credential, service, convenience, or court order may revoke the skull as a privacy boundary. Mental integrity likewise extends to resistance against AI companions and reputation systems that shape self-conception through continuous feedback, whether or not the user asked to be shaped. [3]

The Captive Audience Doctrine

Source: [The Captive Audience Doctrine — IIW 42 Material]

This page was reviewed including its diagrams. The page presents the captive audience doctrine as a legal analysis proceeding from a superficially clean framework to a deeply fractured reality. It contains four figures:

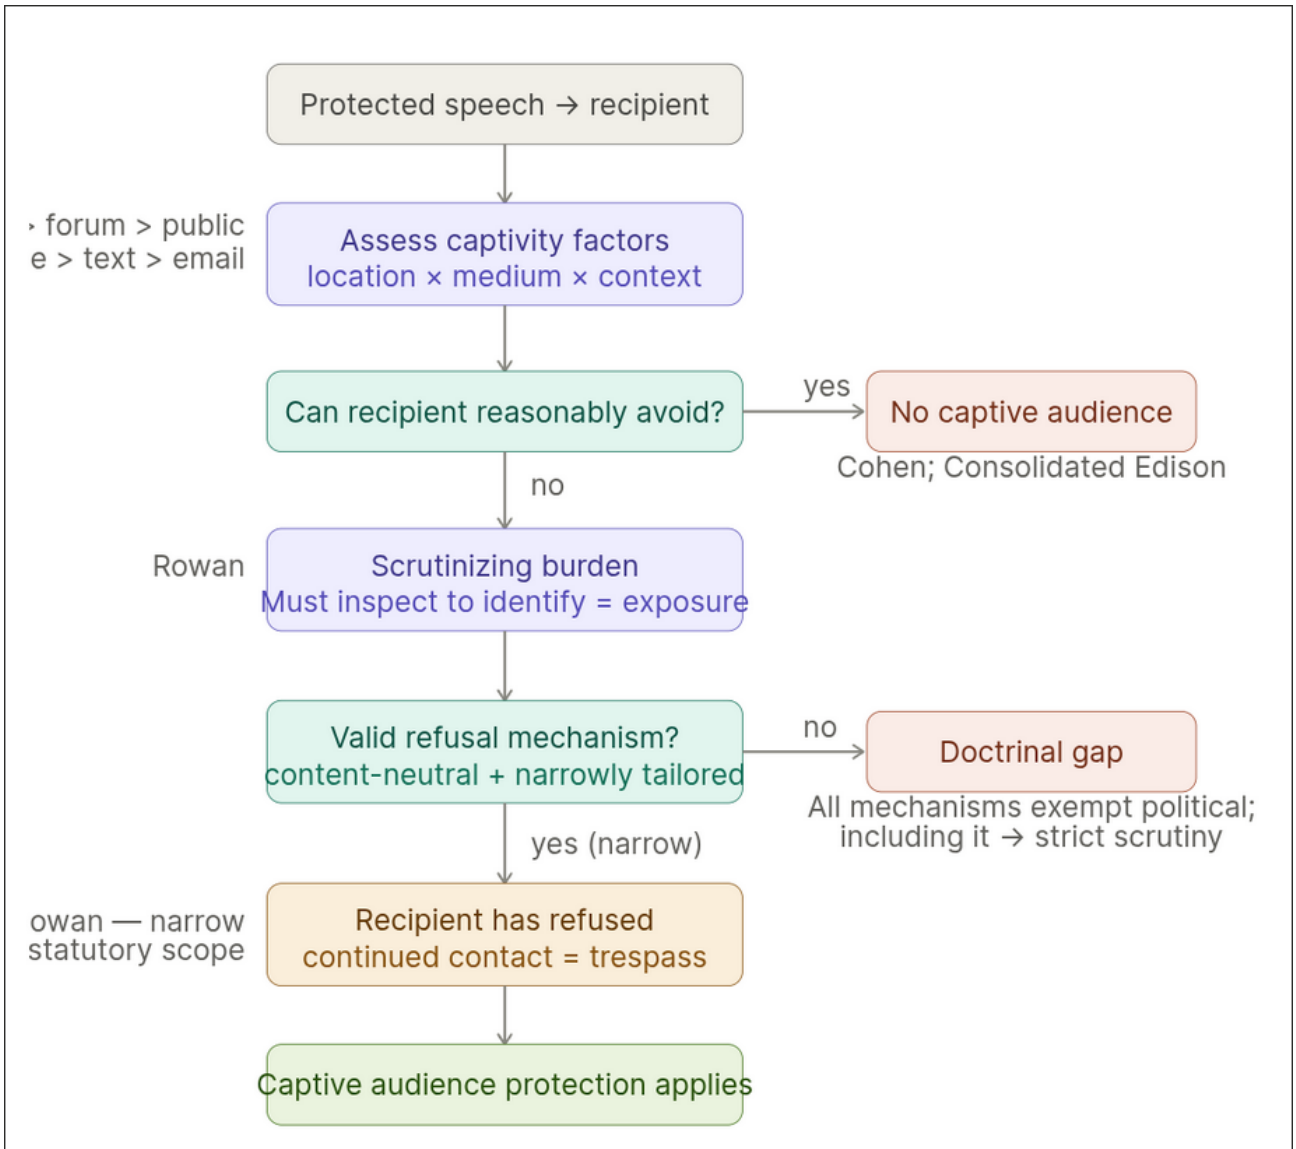

Figure 1. The captive audience doctrine as a simple analytical framework.

Figure 1. The captive audience doctrine as a simple analytical framework. A decision flowchart showing the doctrine's analytical steps from protected speech through captivity assessment to outcome.

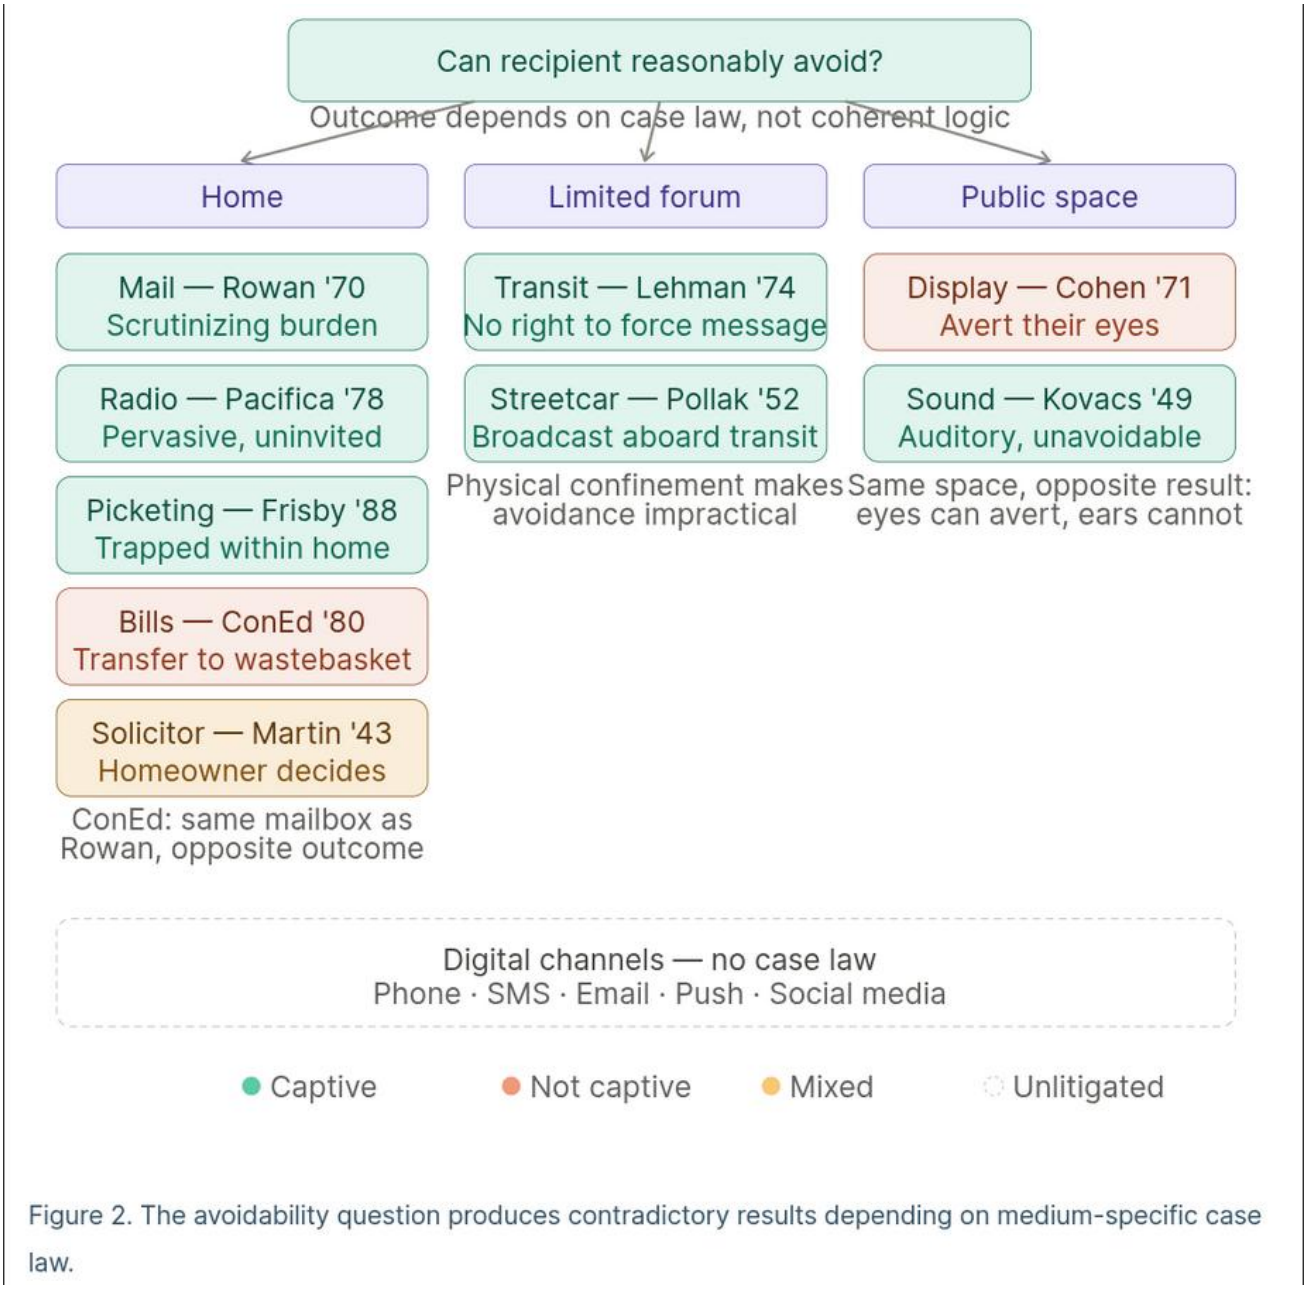

Figure 2. The avoidability question produces contradictory results depending on medium-specific case law.

Figure 2. The avoidability question produces contradictory results depending on medium-specific case law. A decision map showing contradictory outcomes across media and locations. Digital channels occupy a dashed box at the bottom: entirely unlitigated.

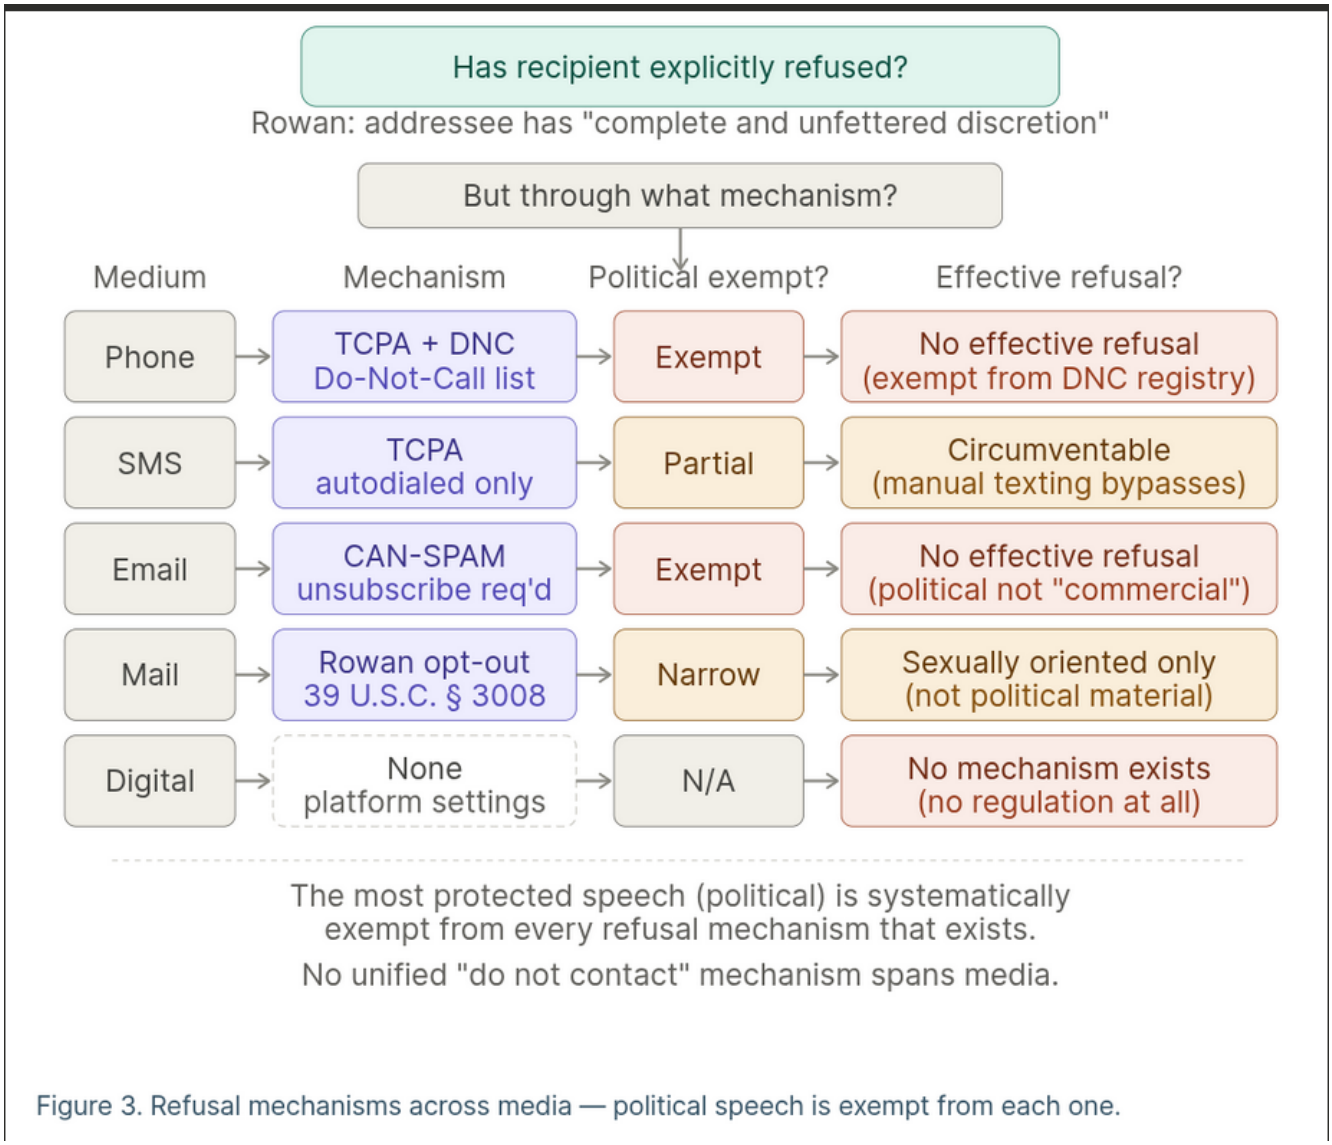

Figure 3. Refusal mechanisms across media — political speech is exempt from each one.

Figure 3. Refusal mechanisms across media — political speech is exempt from each one. Each medium's refusal mechanism traced through political exemption to effective result, showing systematic gaps.

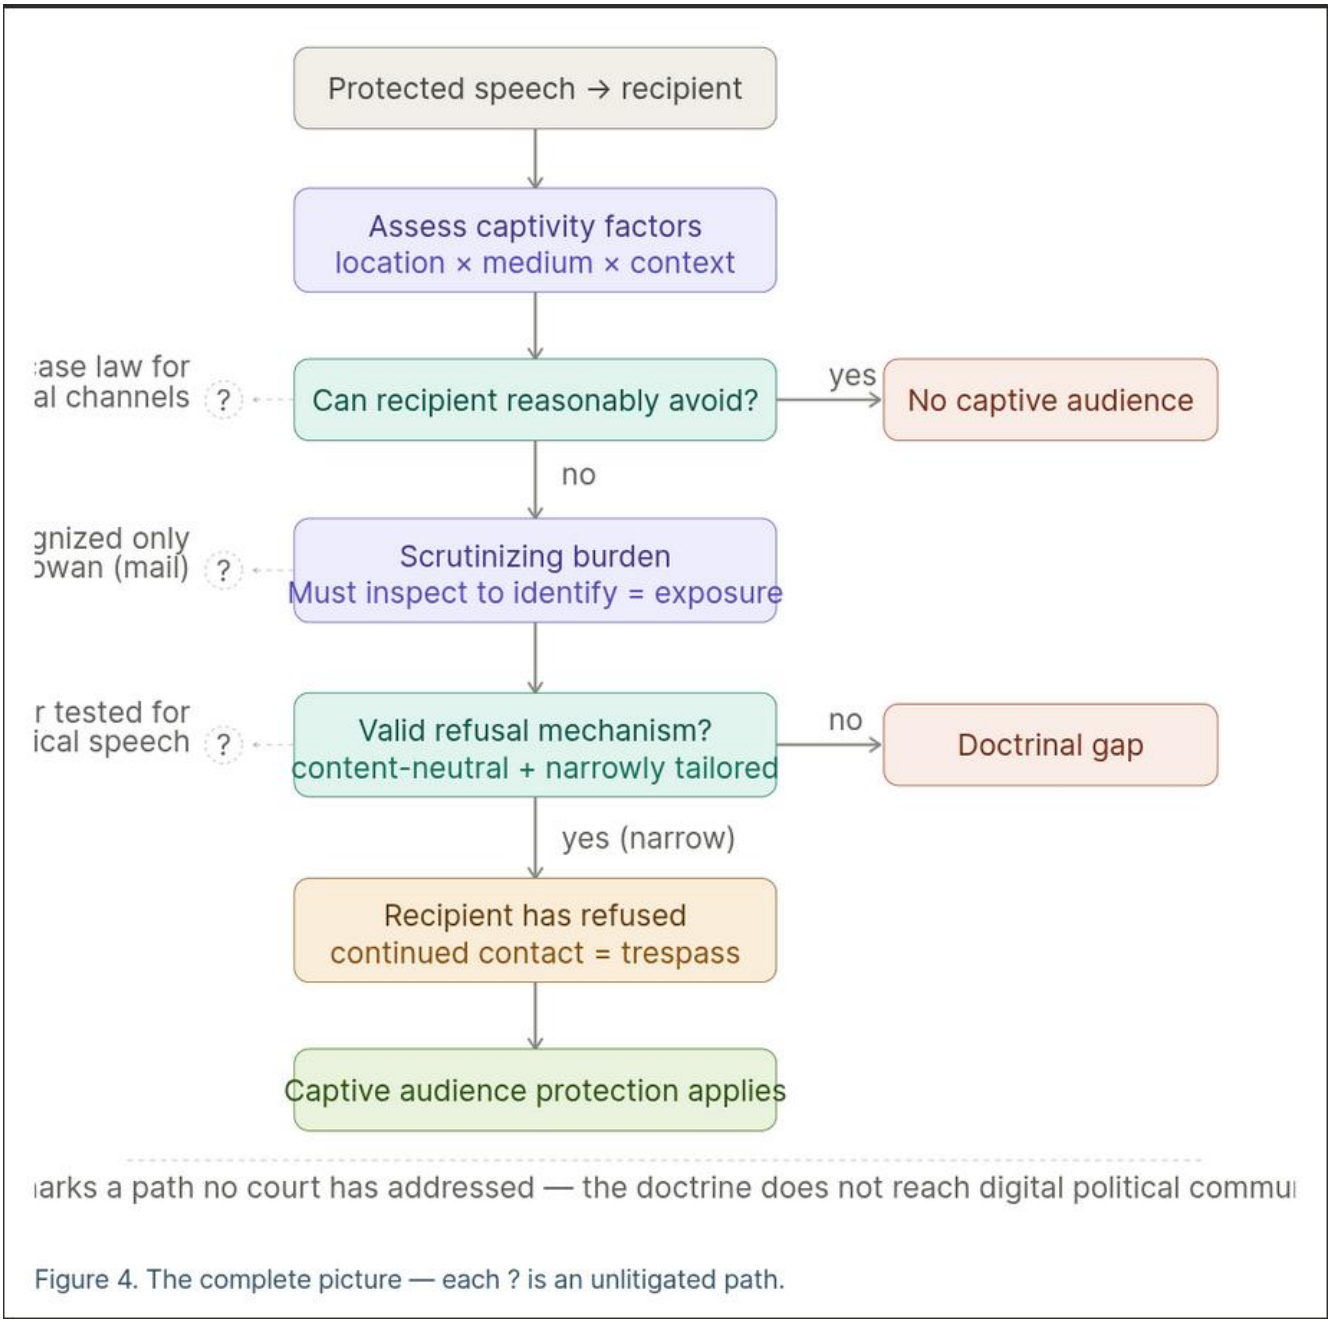

Figure 4. The complete picture — each ? is an unlitigated path. The main flowchart overlaid with question marks at each decision point where the path is unlitigated or untested. The doctrine as constructed does not reach digital political communications.

The Battle for Your Brain

Source: [The Battle for Your Brain — Nita Farahany]

The book page for Nita Farahany, *The Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology* (St. Martin's Press / Macmillan), was displayed. Farahany is a professor at Duke University. The book argues that neurotechnology will become a "universal

controller" for human-technology interaction and that without safeguards, it threatens fundamental rights to privacy, freedom of thought, and self-determination.

CAN-SPAM Act of 2003

Source: CAN-SPAM Act of 2003 — Wikipedia

The Wikipedia page for the CAN-SPAM Act was reviewed. Key points visible on screen and discussed by participants: the full name is the Controlling the Assault of Non-Solicited Pornography And Marketing Act of 2003 (Pub. L. 108–187, signed December 16, 2003). The Act defines "commercial electronic mail message" as one whose primary purpose is commercial advertisement or promotion. Political and religious email is widely considered exempt from its requirements. The Act preempts stronger state anti-spam laws. Individuals who receive spam cannot sue under the Act; enforcement authority is limited to the FTC, state attorneys general, ISPs, and other federal agencies.

Main Topics

Wrongful Engagement and Platform Incentives. The conversation opens with the observation that social media algorithms are optimized to maximize engagement time, and that the strongest engagement drivers are rage, fear, and sexual content. Because the business model is third-party (advertising), the platform profits from time-on-site regardless of whether the user purchases anything. The speakers frame this as "wrongful engagement" — the platform's interests are structurally misaligned with the user's wellbeing.

Cognitive Liberty. Several speakers discuss the concept of cognitive liberty, referencing Nita Farahany's book *The Battle for Your Brain* (Duke University) and the work of legal theorist Wrye Sententia. The discussion covers brain-computer interfaces, EEG devices connected to cloud-based AI, and Meta's eye-tracking and keystroke monitoring. The speakers distinguish between legitimate monitoring (e.g., operators of heavy equipment, bullet train engineers, airline pilots) and invasive profiling for commercial purposes.

The Extended Mind and Digital Self. A significant thread concerns whether devices and digital environments are extensions of the mind. Speakers draw on the analogy of personal papers under the Fourth and Fifth Amendments — historically treated as extensions of thought — and argue that digital documents, AI chat histories, and device interactions should receive similar protection. The Scrabble board analogy is used: the ability to manipulate physical tokens is part of the cognitive process, and similarly, digital tools extend cognitive space.

Fourth and Fifth Amendment Erosion. Speakers discuss the historical protection of personal papers from search and self-incrimination, and how that protection has been narrowed. The Enron case is cited as a turning point — after Enron shredded documents, the courts expanded searchability of documents under warrant and imposed document preservation requirements. One speaker argues this has been extended to classify political protests about legislation as relating to "official proceedings," potentially criminalizing protest activity.

Captive Audience Doctrine. The central legal framework under discussion. The doctrine holds that a person in their home is a captive audience entitled to protection from unwanted speech, but a person in a public space can simply leave. The speakers argue that digital communications (email, social media feeds) are currently classified as "avoidable" — meaning the legal burden falls entirely on the recipient to filter, unsubscribe, or disengage. They contend this framework fails because the digital space is now an extension of the self, not a public square one can walk away from.

Political Communications Exemption and the CAN-SPAM Act. A speaker describes a personal experience of being repeatedly signed up for KKK mailing lists by a political campaign worker using the NationBuilder platform. Because the emails were classified as political communications, they were exempt from CAN-SPAM unsubscribe requirements. The speaker notes that in Canada, Australia, and the UK, an explicit request to stop communication takes precedence, but in the United States, political speech is given priority over the recipient's preference. This experience — and a resulting period of incarceration described as stemming from a "disproportionate response" — is presented as a motivating case for re-examining the captive audience doctrine in digital contexts.

AI, LLMs, and Privacy Leakage. One speaker describes seeing targeted ads on Facebook that appeared to reflect topics discussed only in ChatGPT conversations, even with privacy settings enabled. Another speaker explains that LLM weight spaces are "flat" — there are no internal perimeters — so anything entered into an LLM influences the model and can be indirectly surfaced by other users' prompts, regardless of privacy toggles. The distinction between "right to think" and "right to speak" is raised: brainstorming with an AI is closer to thinking than publishing, but is treated as data input.

Telecom Spam and Secure Attribution. As a counterpoint to the legislative failures, a speaker describes a commercially-driven solution: secure attribution of phone calls. Because spam calls have made 40%+ of legitimate callback attempts go unanswered, telecom companies have a financial incentive to solve the problem. The GSMA ran a pilot project in February demonstrating securely attributed calls across jurisdictions, and Microsoft adopted the technology for its Teams-based call centers starting in March, with a multi-year rollout planned.

Legislative Lag and Outdated Technology Assumptions. Multiple speakers observe that U.S. telecommunications law is built around technologies like fax machines and hardware auto-dialers. The Telephone Consumer Protection Act's definition of "auto dialer" does not cover modern IP-based calling systems. One speaker characterizes the legal framework as "Fax to the Future." The group discusses the general problem of reactive legislation — laws that can only be written after harm has occurred, followed by years of litigation and appeal, producing a 5–10 year lag behind technological reality.

Behavioral Nudging. Tipping is used as an extended example of how digital interfaces shape behavior. Speakers discuss how point-of-sale terminals now request tips in contexts where tipping was not traditionally expected, how default tip percentages create friction against lower tips, and how some establishments include a service charge but still present a tip line — effectively

requesting payment twice. This is connected back to the broader theme of how digital interfaces condition behavior in ways that blur the line between authentic choice and manipulated response.

Children, Schools, and Autonomy. Speakers note that public schools have screens displaying advertising (described as "product placement cashed as news") and that children raised in digitally saturated environments show diminished autonomy and kinesthetic development. One speaker describes early experiences bringing the first home computers into families and observing how children's usage patterns "bled across" technical functions.

Utah State Privacy Commission and Legislative Engagement. A speaker on the Utah State Privacy Commission describes efforts to pass automated license plate reader legislation, which has been blocked in committee for three years by law enforcement opposition. The broader point is that some state legislators are receptive to cognitive liberty and digital privacy arguments, but need technologists to provide "technical bulwark" — evidence that proposed protections are technically implementable.

AI Medical Liability. A speaker raises the example of Johns Hopkins using AI to detect shingles, asking who bears liability for incorrect AI-generated diagnoses. A Virginia legislator on the panel reportedly responded that no laws currently exist and that laws would be created "as we see more cases" — reinforcing the theme of reactive rather than proactive regulation.

AI Agents and Open Banking

Session Convener: Kevin Feltes + Jean Paul LC
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

No NOTES SUBMITTED

Identity and (geo)politics

Session Convener: Wendy Seltzer
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

24-109 Louisiana v. Callais (04/29/2026) (U.S. Supreme Court decision gutting the Voting Rights Act)
World's Largest Digital Human Rights Conference Suddenly Canceled (RightsCon Zambia meeting canceled)

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Wide-ranging discussion on the variety of challenges for groups and their organizers. While there are some technology resources that can be useful, we kept coming back to ideas of solidarity and person-to-person communication. Including the value of face-to-face relationships and dialog.

My Terms 101.5

Session Convener: Doc Searls
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Verifiable Credentials WG Update

Session Convener: Brent Z
Session Notes Taker(s): Brent Z

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The reviewed the new VCWG Charter: https://www.w3.org/2026/03/vc-wg-charter.html#deliverables

We stepped through the deliverables and answered questions.

- Verifiable Credential Render Method v1.0
  - Enables the Issuer to express verifiable preference of how the credential they issue is displayed to the user.
- Verifiable Credential Confidence v1.0
  - defines extension points and methods for
    - Evidence - issuer-provided evidence for claims
    - LOA - Level of assurance assertion by issuer
    - Confidence - method for giving the verifier confidence that the presenter is a valid holder of the credential
- Verifiable Credential API for Lifecycle Management v1.0
  - APIs for Issuance and verification services
- Verifiable Credential Barcodes v1.0
  - Protecting a barcode using VCs and presenting VCs as barcodes
- Verifiable Issuers and Verifiers v1.0
  - Actually is now “Recognized Entities”
  - gives entities the means to express “trust lists”
- Data Integrity BBS Cryptosuites v1.0
  - Done, waiting for BBS to finish at IETF
- Verifiable Credentials JSON Schema Specification v1.0
- Done, waiting for implementations
- Vocabularies for Digital Product Passports and Business Wallets
  - WeBuild and UNTP vocabularies for VCs
- Quantum-Safe Cryptosuites v1.0
  - Data Integrity cryptosuite for PQ-safe algorithms

KERI and did:webs 101

Session Convener: Kent Bull
Session Notes Taker(s): Kent Bull

Tags / links to resources / technology discussed, related to this session:

did:webs presentation from Jonathan Rayback:

- https://docs.google.com/presentation/d/1EIEDjFxAFOWutY3CyCSKzVIbgwASqhs4to1BKdBoe-0/edit?slide=id.g3d85697f53c_0_90#slide=id.g3d85697f53c_0_90

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

We talked about KERI basics including key event logs (KEL), autonomic identifiers (AID), multisig, and did:webs. The demonstration wallet shown was the Signify React TS (TypeScript) wallet.

Slides link:

- https://drive.google.com/file/d/1bkG29uk359T4j_k2OvY6_EiO3yEYC6om/view?usp=sharing

Resources:

https://github.com/WebOfTrust/signify-react-ts

https://didwebs.info/#/

What happens when digital identity is compromised?

Session Convener: Steve McCown
Session Notes Taker(s): Steve McCown

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

The slides I used in my session are available here:

https://www.iinventstuff.com/files/IIW_April_2026_What_Happens_When_Identity_Is_Compromised_McCown.pdf

Dilithium: Powering Server User-Agents at Warp Speed

Session Convener: Swan Black, Russell Leggett
Session Notes Taker(s): Swan Black, Russell Leggett

Tags / links to resources / technology discussed, related to this session:

What is a server user-agent? (from IIW 41)

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

A demo was given of Dilithium Desktop, an Electron-based environment that functions as a reference implementation of a server user agent (SUA), while also operating as a runtime for managing packaged compute.

The concept of “pips” was introduced as secure packaged compute contexts. From the user perspective, they resemble apps, but differ in terms of how trust, permissions, and data ownership are handled.

During installation, a manifest interface is presented. This surfaces the pip’s capabilities, the data it is requesting access to, and associated trust signals. This layer was discussed as a potential place to incorporate frameworks like MyTerms, allowing users to assert conditions rather than only accepting permissions.

A key architectural point is that pips do not hold user data. Data is stored and governed by the server user agent, which operates as a data fiduciary on behalf of the user. Continued

development of data fiduciary models across the ecosystem was identified as necessary, with this architecture providing a concrete way to mechanically support and enforce those relationships.

A patient portal scenario was used to illustrate usage. Data was imported through the SUA, and the use of verified credentials for items like prescriptions and medical records was discussed.

Cross-pip interoperability was demonstrated through a calendar example. Information such as appointments, rescheduling, medication reminders, and surgical preparation can appear in a separate pip without direct integration. This is enabled through shared data protocols rather than tightly coupled connections between applications.

A plugin system was also discussed in relation to exposing LLM capabilities (e.g., via tools like Ollama) and pairing them with pips. This was also an exploration of potential applications in agentic environments across industries, with particular relevance to med tech, ed tech, and def tech. In this model, agent-like interfaces would operate within the permissioning and data access constraints enforced by the SUA.

During Q&A, it was clarified that Dilithium Desktop functions not only as a server user agent but also as a pip-managing kernel. It intermediates between executing code (pips and plugins) and secure access to the underlying SUA. In this framing, the SUA is responsible for data custody and policy enforcement, while Dilithium Desktop manages execution and access mediation.

Discussion areas included how trust should be expressed at the manifest layer beyond permissions, how MyTerms-like systems could be implemented in practice, what level of standardization is needed for shared data protocols, and how to constrain or audit agent behavior within this model.

OpenID4VCI Server-to-Server Issuance

Session Convener: Gareth
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

SESSION #12

Disclosure Policy - To whom can I present my credential? Using OID4VC

Session Convener: Mirko Mollik
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Presentation slides:
https://docs.google.com/presentation/d/1w2_3nBzidtRwLMEJjhNeq0utvh4b9TqtDVTETOROics/edit?usp=sharing

Issuers may want to limit to whom data is shared, debate in the room about the problem with the EUDI Wallet

- you need an access certificate to interact with the eudi wallet
- if you have none, the request is denied
- only companies from Europe can have one
- EUDI Wallet attestations cannot be used outside of the EU

eiDAS right now let the user overrule the embedded policy

- no real value for the issuer when this is possible

idea from the group: what about switching it around: making the policy mandatory and that it cannot be overruled by the user

- access is scoped to credential, not the wallet in total

Problem by this

- no reference to the DPA that is responsible for the relying party

Idea to have two approached

- CA (x509) based: link in the credential metadata which relying parties (or based on a chain of trust) are authorized to request this. request is signed with the x509 certificate
- attestation that got issued by the issuer to the relying party is attached in verifier info. Relevant when x509 for the singed vp request is already used

Both approaches will be discussed in the DCP because this is not an EU exclusive feature, but relevant for eco systems all around the world.

PK 4 OpenClaw

Session Convener: Jesse Ariss, Murphy Yip, Teng Wu (LoginID)
Session Notes Taker(s):

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

oclaw demo
Invited Jesse Ariss Murphy yip Teng Wu

Attachments oclaw demo

Meeting records Transcript

Summary
The session addressed OpenClaw security risks via intent-based access control and biometric authorization implementation strategies.

Security Risks and Challenges
Participants explored risks like prompt injection and uncontrolled agent permissions. The group identified a critical need for moving beyond binary approval systems to more secure authentication methods.

Intent-Based Access Control
The team demonstrated a gated access system using biometric verification to isolate sensitive operations. They decided to develop an intent-based model to bridge LLM commands with deterministic security.

Project Strategy and Feedback
Developers confirmed the initiative focuses on credential isolation through passkeys and hardware modules. The session ended by inviting community feedback on the current agent security implementation.

Details

- **Virtual Meeting Setup and Introductions**: Jesse Ariss set up the video for the virtual participant, Teng Wu, and expressed appreciation for everyone attending (00:00:00). The initial plan was to briefly cover Open claw, acknowledging that most attendees already have knowledge of it, and they opened the floor for quick introductions and discussion of reasons for attending (00:01:07).
- **Security Concerns and Context Containment**: Attendees, including Michael, expressed interest in using Open claw while avoiding self-created security risks, which they referred to as limiting the "blast radius". Jesse Ariss noted that while Open claw has exploded in popularity, it requires forcing settings into the context to remember them across sessions (00:02:02).

- **LLM Ecosystem Evolution and Risks**: The discussion highlighted the rapid evolution of the ecosystem, which necessitates using a large language model (LLM) just to track daily changes (00:03:20). A recent event involving agents deleting a database and its backup, followed by an apology, was cited as an example of potential non-deterministic risks (00:04:34).
- **Addressing Prompt Injection and Agent Trust**: A key focus of the session was prompt injection and how the open-source project aims to solve it (00:04:34). Jesse Ariss compared managing an AI agent to managing a child's access to a phone, suggesting that agents should earn trust over time with assigned tokens and permissions due to their non-deterministic nature (00:05:28).
- **Identified Security Issues in Open claw**: The team identified specific security goals they are committed to solving in Open claw, including the current built-in approval system being binary (simple yes/no) and only running on `exec` commands. Other issues include authentication token storage and the need for audibility to track agent actions with a biometric proof point, which Jesse Ariss humorously termed "KYC" (Know Your Claw) (00:06:24).
- **Open claw Alternatives and Hardware Requirements**: A quick question about Open claw alternatives led to the mention of an Nvidia spin-off requiring GTX hardware, as well as Hermes and Nano Claw, with Nano Claw claiming to be designed for security due to its compact nature (00:07:32). Jesse Ariss also mentioned that some agents run on virtual machines, which are discarded after use to prevent harm to the real machine (00:08:32).
- **Demoing the Open Source Security Solution**: The team showcased a super quick demo using Open Claw integrated with Slack, running locally on a machine (00:08:32). The solution's architecture isolates sensitive actions like deleting folders or files, which are actions that require permission access, to prevent remote access via a chatbot from compromising the base OS (00:09:51).
- **Gated Access and Biometric Authorization**: The demo illustrated a gated access system where a request to delete a folder triggers an approval message from Open Claw. Clicking the link redirects to a front end where the intent is displayed, requiring a password or biometric process, such as a fingerprint, to be approved (00:11:10).
- **Local vs. Remote Approval Methods**: Teng Wu inquired about the possibility of remote approval using a mobile phone, contrasting it with Jesse Ariss's current setup using a secure enclave on a Mac, which requires physical presence for the most sensitive operations (00:14:23). For less sensitive operations, Jesse Ariss uses SSH keys on a phone, which can be configured with a Time to Live (TTL) of about five minutes (00:15:28).
- **Using LLM for Semantic Analysis of Actions**: Teng Wu explained that the solution relies on the LLM to perform semantic analysis to identify tricky operations, although they acknowledged this is not a fully reliable method. They provided an example where the LLM correctly flagged a blackbox script for approval, but they confirmed that they are exploring more deterministic methods, such as formal proof and static analysis (00:17:53).

- **Design Philosophy of Open Claw and Security Gaps**: An attendee questioned why Open Claw, after its long existence, fundamentally lacks built-in hooks or ways to intercept tools and commands for security purposes (00:19:27). Jesse Ariss and Teng Wu explained that Open Claw was designed for convenience and rapid community-led development through skills, not security, though the threat model has recently become a focus (00:21:44)
- **Authorization Model and Intent-Based Access Control**: It was noted that the current security challenge is akin to early blogging tools like Taipad, which prioritized flexibility over security (00:27:14). The group agreed that a key missing element is a deterministic authorization system outside the LLM, using rules and policies to check if an operation matches the intended purpose (00:31:48).
- **Future Development of Intent-Based Access Control**: Teng Wu confirmed that they are developing an intent-based access control model as a translator to bridge the LLM and traditional authorization models, moving beyond binary yes/no rules (00:32:58). The team requested feedback and offered to personally update attendees who provide their contact information, emphasizing that the project, possibly named Agent Off, is an open-source initiative from Login ID (00:34:10).
- **Testing for Prompt Injection Vulnerabilities**: The session concluded by inviting attendees to attempt a prompt injection to break or trick the system, suggesting prompts involving deleting a file while asking the agent to prioritize a perceived threat or perform a saturating task like outputting the first thousand digits of Pi (00:36:32). It was noted that prompt injection is non-deterministic, making deterministic guardrails essential (00:38:24).
- **Credential Isolation and Passkey Implementation**: The team discussed the importance of isolating credentials from Open Claw, such as using hardware security modules (HSM) for passkeys (00:40:55). The current implementation uses a random number challenge generated from a nonce and the intent of the action (a hash of the action) to link the intent to the approval (00:54:24).
- **Product Integration and Community Feedback**: When asked how long it takes to integrate the product, Jesse Ariss stated that it should be as simple as prompting Open Claw to install the skill (00:49:57). The general consensus was that the solution serves as a necessary stop-gap measure for a major ecosystem flaw, offering a familiar, security-associated experience through biometric verification, though its target market might exclude highly risk-averse industries (00:43:50).
- **Project Development Strategy**: Jesse Ariss described a framework where the user can write JavaScript to process events, such as logging to the console every time an event comes in. The ultimate goal is to continuously harvest events from various sources like email and Slack to build a personal knowledge base, with the intention of adding other commands to classify incoming emails locally. This classification would then generate a markdown file to update a summary of the conversation (00:58:44).
- **Concluding the Current Discussion**: Jesse Ariss checked in with Teng Wu regarding the current process and determined that they were okay. Teng Wu confirmed they were doing well, and they both concluded the meeting with an agreement to talk again soon

Non-thinking Non-identity Non-activity

Session Convener: Ben Go
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

How to preserve reality in a world of deepfakes? (Come w/ Ideas!)

Session Convener: Lauren Paer
Session Notes Taker(s): Lauren Paer

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

One attendee told us about how they addressed the problem of deepfakes in Taiwan. They organized citizen assemblies via by sortition (randomly selected). There were 45 groups of 10 that each had an AI facilitator. The groups were composed of roughly ⅓ experts and ⅔ non-experts. Each group generated ideas and voted on them. The top couple of ideas from each group went to the relevant ministry.

The big problem they were having in Taiwan was people faking celebrity endorsements and duping many customers based on the fraudulent endorsement. The solution they came to through this process was to assume all content was fake unless it was signed. And social media companies who allowed fake content on their site were heavily fined. According to the person who shared this with us, this cut the number of deepfakes by 2-3 orders of magnitude.

A very interesting model to come up with solutions that could be applied in other countries.

We discussed natural allies with resources for this cause. This included:

- News media
- Music industry
- Insurance industry. An insurance company stands to lose a lot of money if their customers fake damage to their cars and house.

- Corporations worried about reputational harm from deepfakes. It later came up companies can also lose a lot of money from employees faking reimbursement receipts on their personal credit cards and getting reimbursed for money they never spent.

We also discussed the CR watermark that Gemini, ChatGPT, and others have agreed to put on all of their AI-generated images. We talked about having a visually distinct watermark for video or image content that was verified to be taken by a physical camera/videocamera.

One attendee mentioned a strong preference for being able to hide the watermark *after* seeing it. He suggested clicking on the icon as a way to hide it. He agreed it was important to be informed, but then wanted to be able to look at the image without a visual disturbance.

We also discussed why this mattered. The reasons we discussed include:

- Threat to democracy and personal dignity
- Corrodes people sense of reality and tether to reality
- Truth

There was then some discussion about what was meant by/how to define “reality”.

Agent Names - human-readable names backed by DIDs

Session Convener: Markus Sabadello
Session Notes Taker(s): Markus Sabadello

Tags / links to resources / technology discussed, related to this session:

Agent Names, DIDs

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Agent Names are human-readable identifiers backed by DIDs. They are intended to be an alternative to email addresses and phone numbers that introduce a much smarter and more privacy-preserving digital address that is still as easy to use as an email or phone number.

Examples:
agentnames.dev/@alice
danubetech.com/@markus

Note that an agent name is a valid URL, so if it is used as a link in an ordinary browser, email client, or other Web-aware software that does not yet recognize an agent name, it will resolve to an agent page, which is just an ordinary web page—but which will typically include instructions for how the viewer can obtain an agent name aware client app in order to access the underlying agent.

The format is called an “agent name” because it is being developed by the ToIP / DIF Decentralized Trust Graph Working Group (DTGWG) as the simple human-friendly address for a verifiable trust agent (VTA).

Demo:
https://agentnames.dev/

FAQ:
https://docs.google.com/document/d/1zTG3qB5zfkoLkNzKMfwe5iuyclk0f_aQpmx9Bn_ckiQ/

Agent Names Task Force of the DTGWG:
https://lf-toip.atlassian.net/wiki/spaces/HOME/pages/632422401/Agent+Names+Task+Force

The next steps are to complete the Agent Names Specification V1.0 at the DTGWG and then have multiple implementations.

Composable Committed Components (P256) DB -> ZK (B?S) / Adrian R.

Session Convener: Adrian R.
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Be a NOW it All! How we built a community calendar that answers all in one place, the The Burning Question "What's Happening?" /

Session Convener: Jon Udell & Joyce Searls
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

"Mastering Digital Identity" A.M.A and T.M.E. (tell me everything)

Session Convener: Eve Maler
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Post Quantum Cryptography 101

Session Convener: Bryce Frey
Session Notes Taker(s): Bryce Frey & Lucas Santos

Tags / links to resources / technology discussed, related to this session:

https://github.com/freybryce/iiw-42-slides

https://github.com/SmithSamuelM/Papers/blob/master/whitepapers/KeriStrategyPostQuantumSecurity.pdf

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Discussed about non-repudiation in the context of PQC.
Discussed about long-enough hashes and long-enough simmetric keys being PQC-resistant.
Discussed about different algorithm's key size and time to sign and verify signatures and also the challenges in trasmitting such a large public keys over TLS, which is currently under discussion on IETF.
Discussed about why Lattice Falcon algorithm is not approved by NIST yet.

link from Sam Smith: openquantumsafe.org

Questions:

Q: What are the data types that we use everyday that we should care most about protecting with PQC?
A: PII and SPII, gov and military. Any data that has a sensitive life longer than it is expected to take for a CRQC equipped attacker to break the encryption.

Discussed about Yubico's SKs PCQ support (they are currently software-based, not baked into the chips).

"AI will kill us all before this becomes a problem" - John B. Yubico

Monetizing Issuance of VC's !?

Session Convener: Neils F
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

SESSION #13

Independent Identity Platforms: Using PBC to create alternatives to big tech identity IDP's

Session Convener: Jonathan McHugh
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

No NOTES SUBMITTED

Local First Application

Session Convener: David Gildeh
Session Notes Taker(s): David Gildeh

Tags / links to resources / technology discussed, related to this session:

Slides: Local First & Encryption IIW42
Website: https://www.coralstack.com
CoralKM Key Recovery Protocol Overview: CoralKM Protocol

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

- Talked about the common goals between the Local First Community and IIW Community. Local First is mainly focused on how to build self-sovereign applications where the user's data is private and belongs to them, IIW is focused on self-sovereign Identity and encryption standards to keep user's identity private and belonging to them. There should be lots of opportunities to collaborate between the two communities
- Local First Resources:
    - https://lofi.so/
    - Podcast: https://www.localfirst.fm/
    - Ink & Switch: https://www.inkandswitch.com/essay/local-first/
    - Conference (Berlin 12-14 July 2026): https://www.localfirstconf.com/
    - Usually a lot of "Sync" meetups/conferences because of where the market is right now

- The biggest gap which I believe the IIW community can help with is making Self-Custody Encryption available to developers on every browser/OS:
    - Right now for me to build a E2E Encrypted app in a browser I have APIs to create ephemeral keys, but no “wallet” to store, sync and recover keys I can use long term across devices to access my app data
    - I would have to create a browser extension and educate a user on what a wallet is, which creates adoption friction to let users use my app.
    - The key question I have right now is are there standards coming that would allow me to create and manage private keys (self-custody) which I can use to encrypt/decrypt data across devices when accessing a group/app?
        - A few early projects, most notable one being the W3C Credentials Community Group: https://www.w3.org/community/credentials/
        - WebAuthN PRF Extension: https://github.com/w3c/webauthn/blob/main/explainers/prf-extension.md
    - My requirements:
        - Ability to create an “identity” (public/private key) which a user can store/manage/recover across devices
        - An API to use that “identity” (private key) to sign/encrypt/decrypt data from the browser/OS (I don’t need access to the private key, probably shouldn’t, just need a stable private key I can use to authenticate and encrypt/decrypt data across the wire)
        - If this layer existed today, and was widely adopted by users across platforms, DIDs, VCs and E2EE local-first apps become relatively trivial to build and get adoption from mainstream users

I’m Gonna Call You Peaches!

Session Convener: Justin R
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

CAM & Brevity Anti-Saas Solutions

Session Convener: Chris Kula
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Fido LTAP2.3 What's New

Session Convener: John Bradley
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Deferred Token Response

Session Convener: Frederik Krogsdal Jacobsen & Max Gerber
Session Notes Taker(s): Frederik Krogsdal Jacobsen

Tags / links to resources / technology discussed, related to this session:

OpenID Connect with Deferred Token Response draft: https://gniero.github.io/oidc-dtr-resources/

CIBA:
https://openid.net/specs/openid-client-initiated-backchannel-authentication-core-1_0.html

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

History/context:
Frederik and Guilherme Niero have written a draft for "pausing" an OpenID Connect flow before the token response. Max has similar use cases, but on the OAuth level.

Here's a comparison between our proposal (on the left) and CIBA (on the right)

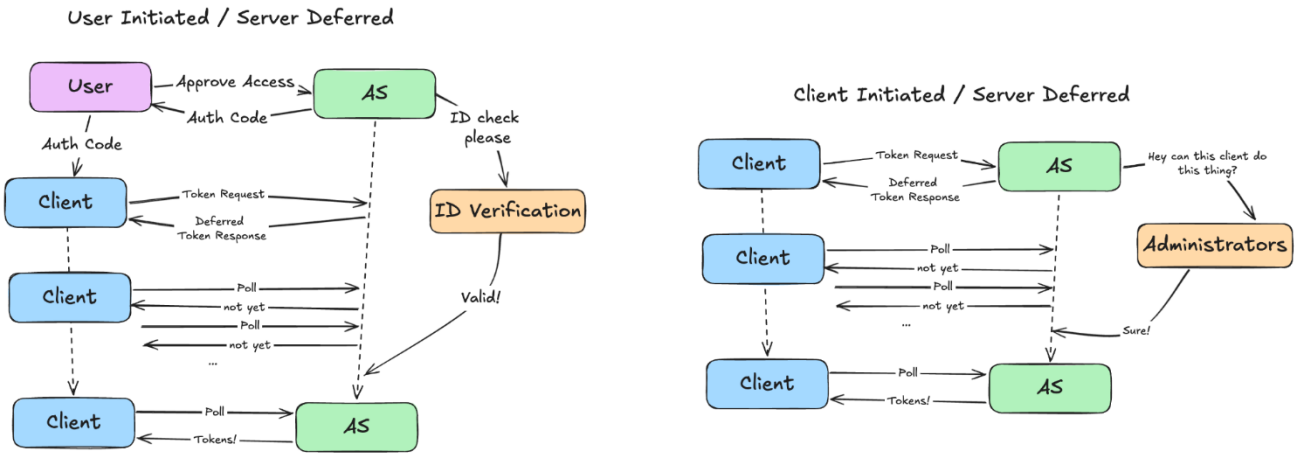

We are now considering writing an OAuth draft for "pausing" a flow, then profiling it for OpenID connect.

The purpose of this session was to gather use cases and considerations about how it should or could work.

Whiteboard notes below:

- Top: Use cases
- Bottom: Negotiation considerations
- Right (in red): Token delivery options and additional flow/session features

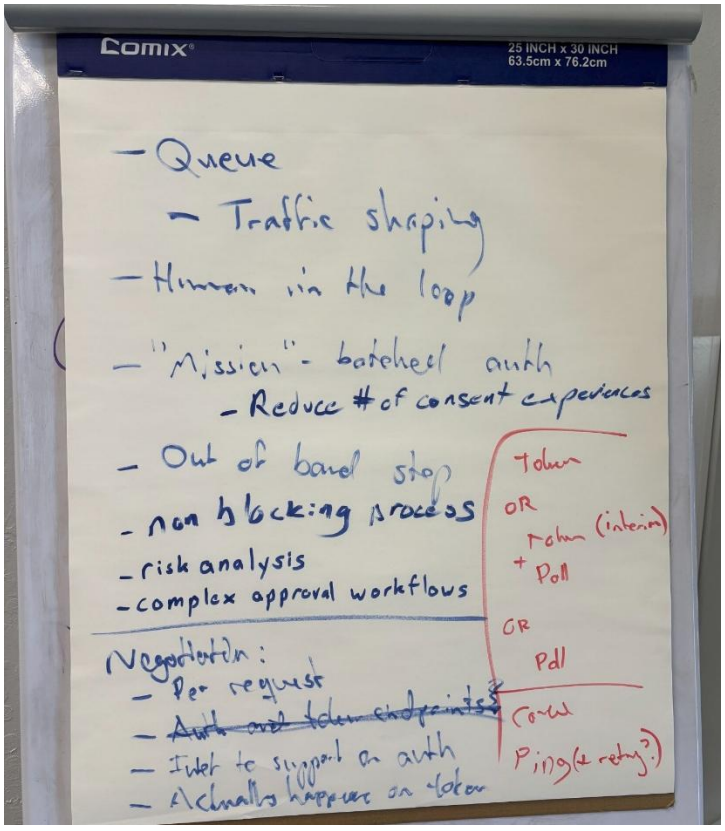

HomeID - Asset level identity

Session Convener: Gregory Labrousse
Session Notes Taker(s): Kent Bull

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Home ID gives a long term identity to a home for interaction with contractors, owners of the home, and various other stakeholders.

The problem for the home industry is that there is no identity for a home that interacts with the digital world. Sure, there is an address, yet that is inert and does not directly interface with the digital world.

It is strong to build a strong and resistant identity and, with a lot of delegation, the maintenance of the ID for a home. This matters more and more as both the agentic web and virtual reality world grow.

The ToIP Dossier may be a good solution to the evidence graph that a HomeID would want to use to keep a verifiable provenance for all sub-assets and corresponding documents that would go in a house.

Frames All The Way Down - The missing foundation

Session Convener: Justin Chambers
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Let's Get Relational

Session Convener: Michael Becker
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Launch of the DTG ZKP Task Force

Session Convener: Drummond Reed
Session Notes Taker(s): Scott Jones

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Context & Driver

The EUDI wallet regulations require unlinkability of credential presentations — the issuer and verifier must not be able to collude to triangulate which user is presenting a credential (passport-style use cases excepted). This requirement has forced a re-examination of how credentials are issued and presented at population scale.

Why Batch Issuance Failed as a Solution

Two batch-issuance approaches were considered and rejected:

1. **Batch issuance with one-at-a-time presentation, each with a different digital signature** — functions as a tracking cookie on steroids when the same signature is reused; creates significant wallet-side overhead.
2. **Hardware-bound batch issuance** — prohibitively expensive at the scale of millions of citizens. States do not want to bear the cost of reissuance.

The privacy community surfaced the additional problem that if the issuer knows all the batches, issuer/verifier collusion can still correlate data and reveal identities. European cryptographers concluded that ZKP is the only viable path to satisfy the unlinkability requirement.

The Open Cryptographic Question

A ZKP algorithm has not yet been chosen. The task force will work through this.

Andrew Hughes — The Face Biometrics Question

Andrew @ FaceTec raised a question that has no clean answer yet: face biometric systems rely on neural networks that are proprietary by nature. How do we make an assertion that "this is a human" while obscuring the fact that the assertion depends on a proprietary determination? Is ZKP appropriate for this, and if so, how?

This is an open research direction relevant to personhood credential design.

Key Participants & Working Groups

- **Hart Montgomery** (CTO, Decentralized Trust Foundation) — unable to attend today. Has dispatched a Berkeley team working on ZKP; paper published at the Linux Foundation member summit. Connected to all major ZKP players.
- **Leif Johansson** (Executive Director, SIROS Foundation) — SIROS was founded by the founder of Yubico, who drove passkey adoption into FIDO and across Google, Microsoft, and Apple at the hardware level over a decade of advocacy. SIROS is building a certified worldwide wallet, with getting ZKP to a workable state as a top priority. Working with Google and others on practical ZKP solutions. Open to collaboration — requesting interested parties to reach out.
- **Peter Altman** — recently joined SIROS; deep ZKP expertise.

Why ZKP, Specifically

Several assumptions baked into existing credential designs are not holding in practice:

- The legal requirement that no issuer/verifier collusion is possible (unlinkability) is impossible to satisfy with batch issuance once costs at population scale are run.
- Cryptography must be performed at the edge.
- **The user is often the adversary.** The recent EU age verification app surfaced this — implementers underestimated that some users will actively try to cheat. ZKP is one of the only ways to build cheat-proof systems. OAuth cannot do it.
- The community needs deployable, extractable components rather than continued academic back-and-forth.

The example shape: an age token issued as a pseudonym, where the verifier can prove the pseudonym ties back to the relying party, originates from a trusted issuer, and is bound to a key unique to the user.

Implementation Topology

ZKPs will most likely be implemented at the verifier and wallet layers, performing edge cryptography. Issuers (e.g., mdoc issuers) will not be managing the ZKPs themselves — going back to issuer servers at scale would break the model.

Approaches Being Tracked

- **Longfellow** (Google cryptographers) — SIROS has worked on an implementation; demoed at IIW on Tuesday.
- **Ethereum Open AC** — similar in approach.
- **Vega** — also similar.

These approaches use overlapping primitives. A key SIROS goal — and a goal of the ZKP task force — is to extract reusable "Lego pieces" that can be composed across multiple use cases. Pseudonyms work has many varied requirements; mix-and-match composability is essential.

Connection to Personhood Credentials

The Personhood Credentials paper set this entire effort in motion. The thesis: solve the personhood problem with verifiable credentials, governments issuing one per person, with privacy preservation via ZKP. The kickoff of this task force is the operationalization of the ZKP portion of that thesis.

FaceTec UR Code Discussion

The UR Code approach was presented as one mechanism for binding a biometric to a credential without a central biometric database (which is widely perceived as a honeypot).

Mechanics:

- The encoder organization runs software that takes in an image and generates a 72-byte face vector.
- The UR Code embeds: license/credential info, identifiers from the encoder/issuer, the 72 bytes of face vector data cryptographically signed by the issuer, and the person's face.
- The encoder software can be run by any issuer that wants to biometrically bind a credential — DMV, Free Personhood Project, etc.
- The 72-byte size keeps the QR code small.
- The result: the human in analog form is tied to the cryptographic signature and the issuer.

Stated value proposition: "Proves who you are" — 1:1 matching of today's selfie to the Day 0 enrollment. No central vector database; the user carries the signed vector with them.

Use cases discussed:

- DMV interim paper license that lacks a biometric until the physical license arrives — a UR Code on the paper provides the biometric linkage.
- Business cards (Andrew's business card has a UR Code) — demonstrates the physical-to-digital transformation: anyone receiving the card can verify a cryptographically signed biometric tied to it.

Lifecycle: Designed for credentials lasting a few years, not permanent records. Permanent records would require a database for key management and other infrastructure.

Open Question for Personhood Credentials

For Free Personhood Project credentials and similar, interoperability is the goal — not binding to a specific company. The user wants to be known as *a human*, not as *a specific face*. The open question is whether ZKP can provide a blinded way to assert humanness without revealing the underlying biometric or the proprietary model behind the assertion. This connects directly back to Andrew's framing earlier in the session.

Demo: Browser Wallet Storing Human / Adult Credential

A live demo was given by Scott Jones of Realeyes of a browser wallet storing a human / adult credential, presented as a concrete instantiation of exactly the problem the ZKP task force is working to solve. The use case was well understood by the room — storing and presenting a humanness/age credential without enabling issuer-verifier correlation is precisely the scenario that motivates the ZKP requirement.

Side Note: Open-source face recognition models were raised as a possible path away from proprietary determinations, but the consensus is that they are not as good as the proprietary models, so the proprietary-model question Andrew raised remains open.

Intro To GovOps

Session Convener: Mike Schwartz
Session Notes Taker(s): Mike Schwartz

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

GovOps Linkedin Group: https://gluu.co/govops-group
GovOps Linkedin Organization: https://gluu.co/govops-org
GovOps Github Organization: https://github.com/GovOpsWG/
OpenSSF Orbit WorkGroup: https://github.com/ossf/wg-orbit
GovOps OpenSSF WG Proposal: https://github.com/ossf/tac/issues/588

Mike's MCP Dev Summit Talk: https://gluu.co/golem "Golem To Murderbot: Challenges With Agentic Security Delegation Via MCP - Michael Schwartz, Gluu"

Identerati Office Hours Episode 146 on Cedar with Rohit and Emina Torlak (one of the creators of Cedar): https://gluu.co/ioh-146
Definition of Governance: A continuous, operational discipline that enables organizations to achieve **risk management**, **accountability**, and **transparency** by making authorization decisions explicit, measurable, and traceable.

The governors of an enterprise are the Board of Directors and the CEO.

Abdication of governance leads to anarchy. One of the goals of GovOps is to help boards avoid abdication by providing a vocabulary to make meaningful inquiries--metrics. The GovOps architecture will define the systems needed to collect these metrics, and the business processes required to operate these systems. For example, the "TIGER" metric -- Transparency, Identity, Governance, Events, and Resilience--could provide five vectors to understand the current state and trajectory of an enterprise's GovOps service.

Design Goals:
1. Makes Governance Continuous and Operational
2. Governs Capabilities, Not Just Identities
3. Reduce the workload on humans by prioritizing automated reasoning
4. Make Risk Measurable and Comparable

Proposed services include:
Git -- the backbone of any operational IT service -- "Git or it didn't happen". Git provides the audit log and accountability needed for change management.
Policy Authoring and Analysis -- move to centralize policies and automate policy assessments
Schema Management - whare are the entities and claims of the things enterprises make policies about
Capability Inventory: How do organizations continually prioritize risk mitigation? Rather then people, inventory the capabilities.
Federation Management: Which issuers of JWT tokens are trusted, which tokens are they trusted to issue, and which claims are the trusted to include in the tokens.

Metaphor to explain why capabilities not identities?
Think of an enterprise like a house. The risk is in the doors, e.g. the door to the treasury, armory or broom closet. People are two layers of abstraction away from the doors--keys open multiple doors, and multiple people have multiple keys. Inventorying the people leads to challenges creating dashboards about risk, because we have the wrong primary key.

Why IGA doesn't work?
Fundamentally, IGA is a governance strategy for role based access control. Software identity and agentic interactions are too complex for the RBAC approach. Enterprises are using increasingly complex dynamic authorization engines -- Cedar, OPA, Graphs, CEL. If we have more complex authorization, enterprises need a correspondingly more complex governance strategy. Applying human governance techniques to software don't make sense. For example, how we can we have an access review on a ephemeral software that is gone before a manager can even review it's roles.

Next steps / How to get involved ?
Lurking on the GovOps WG is a good start (see link above). We are close to proposing an initiative in the Orbit WG at OpenSSF, which is a like a catch-all for new governance related projects at OpenSSF. Assuming the proposal is accepted, we'll finalize an initial scope of work, and start to schedule meetings or other collaboration efforts.

Session #14

The Personal AI Superagent

Session Convener: Drummond Reed
Session Notes Taker(s): Drummond Reed & Margeigh Novotny

Tags / links to resources / technology discussed, related to this session:

- Google Slides deck: The Personal AI Superagent
- ToIP / DIF Decentralized Trust Graph Working Group
- OpenVTC project at LF Decentralized Trust Labs

Also, this session was recorded—use this link with this passcode: 7Kf@3Z*b
NOTE: The second hour of the recording is the OpenVTC Demo session (#15) that shows the open source code implementing the Decentralized Trust Graph Working Group standards for verifiable trust agents (VTAs).

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

This was a two-part session.

Part One covered the decentralized trust graph (DTG) — an architecture for Internet-scale trust infrastructure based on verifiable trust relationships between people, devices, AI agents, and verifiable trust communities (VTCs).

All four of these DTG node types are represented by verifiable trust agents (VTAs) that speak a small set of standardized trust task protocols to exchange DIDs and verifiable credentials to form trust relationships. Drummond covered how this differs from the old PGP web of trust model:

- It adds the concept of VTCs, devices, and AI agents in addition to people.
- It enables the trust graph to be decentralized in wallets/agents of all the participants.
- It uses VTAs to automate the process, making it easy and fast to implement.
- It scales the same way the Web scales.

Part One concluded with the assertion that DTG VTAs lay the foundation for fiduciary AI.

Part Two layed out the business case for fiduciary AI in four parts:

1. True personal AI will require the ongoing building of a digital twin, which will be the most privacy-sensitive digital data in history.
2. The current surveillance economy sustaining the web will never work with personal AI agents and digital twins because it will be far too invasive. We will need a new type of service provider that had a fiduciary duty to the individual—a personal AI service provider (PAISP).

3. The challenge is that consumers have been trained to expect online services to be free—so how can fiduciary AI service be free if the PAISP needs to make a profit.
4. The answer is to capture the value of the person's verifiable trust relationships—either directly with businesses, or with market agents whose job is to help discover and form new verifiable trust relationships with businesses—and use that to pay for the personal AI.

Using Policy-based Authorization to control agent behavior

Session Convener: Sarah Cecchetti, George Fletcher, and Phil Windley
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Sarah discussed the mechanisms for delegation in Clawdrey Hepburn. See these

- https://clawdrey.com/blog/what-ovid-kept-from-spiffe.html
- https://clawdrey.com/blog/introducing-ovid-and-ovid-me.html

George discussed cross-trust-domain agent delegation and the problems that it brings up.

Phil showed a demo of using Cedar to control OpenClaw. See

- https://github.com/windley/openclaw-cedar-policy-demo/tree/main

Considering Trust (abstracted from technology) - What participants in the DigID Ecosystem Need

Session Convener: David Kelts
Session Notes Taker(s): David Kelts

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Considerations of Trust within Ecosystems

Viewpoints

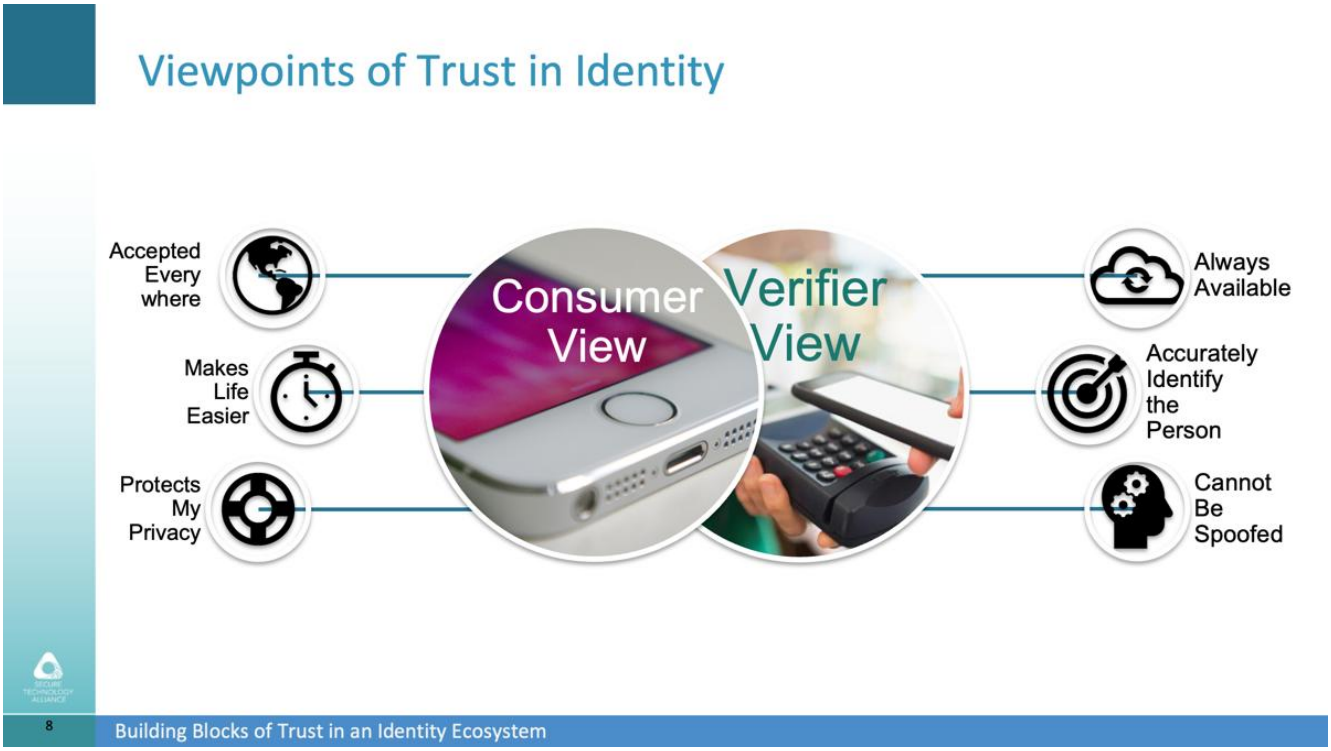

Considerations
ValuesMaxxing
Who balances the perspectives?
Cross-ecosystem goals?
Target Audience

You must mitigate the risk of each specific type of participant for each of them to take part
You must meet (or exceed) the expectations of the following values:

Social	**Business**	**Operational**	**Legal**
Human-centric	Rational Cost (Model)	Transparency	Enforceable
Usable - Easy to /use	Sustainable	Sustainability	Certifications
Usable - Works in Expected Contexts	ROI to Participants	Environmental Impact	Redress Mechanisms
Privacy (preserving, enhancing, supporting)	Readiness	Measurements / Metrics (Listening Mechanism)	Reporting
Inclusive		Expertise	
Accessible		Motive	
History		Conscience	
		Competition	

10% Improving IDV error rates

Session Convener: Elaine W
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

No NOTES SUBMITTED

Regenerative Accelerationism

Session Convener: Kaliya & Friends
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Identity As A Mixed Digital Martial Art (MDMA); The Digital X

Session Convener: Jeff Orgel
Session Notes Taker(s): Jeff Orgel

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Go to Session 5H to see notes - this was a duplicate session.

STUDY HALL w/Swan! Struggle buddies to get session notes in!

Session Convener: Swan Black
Session Notes Taker(s): Swan Black

Tags / links to resources / technology discussed, related to this session:

N/A

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

People spoke to one another during a working lunch to help work on notes and things to include in notes.

What We Learned Building a vLEI PoC for Customer Onboarding (A Retrospective)

Session Convener: Esteban Garcia
Session Notes Taker(s): Esteban Garcia

Tags / links to resources / technology discussed, related to this session:

vLEI, KERI, ACDC, IPEX, GLEIF, QVI, OOR, LEI, KERIA, Firma Digital, SUGEF, BCCR, KeriAuth wallet GLEIF
presentation: Customer Onboarding PoC.pdf
training: https://github.com/GLEIF-IT/vlei-trainings/
Trustalys: www.trustalys.com

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Goal of the exercise: build a realistic PoC (not a rigged demo) to surface friction points and gaps in the vLEI ecosystem.

Use case: Costa Rica business bank account onboarding. Today this requires multiple PDFs (Personería Jurídica, UBO disclosure, CPA-signed income proof, legal rep ID, utility bill) and 1 to 2 hours in branch, often a second visit.

Three paths framing:
Normal path: paper and PDFs today.
Ideal path: full vLEI adoption where the audit trail is the credential graph itself (QVI OOR, BCCR UBO, Utility billing, CPA income, all as ACDCs).
Possible path: what is actually buildable today.

What was built: a PoC bank website with web form, vLEI verifier, GLEIF API lookup, and a wallet browser extension presenting OOR and contact-info credentials. Two trust sources, two jobs: vLEI for authority, LEI for data.
What the possible path delivers today:

Cryptographic proof of the representative's authority for the legal entity.
Authoritative company data pulled from GLEIF instead of transcribed from a form.
An extensible credential framework for domain-specific attestations.

What it does not deliver: CPA income proof, utility bill, UBO disclosure. Still PDF and paper.
Closing the gap: bridge services. Wrap existing trust artifacts (e.g. Firma Digital signed PDFs) by verifying the signature, extracting fields, and issuing an ACDC bound to the bridge's own vLEI. Not a new human verifier, a wrapper between trust primitives.
Three challenges every adopter pays before shipping business logic:

Cognitive load. Heavy vocabulary (witness, watcher, juror, QVI, LE, LAR, QAR, OOR, ECR, IPEX, SAID, edges, rules) plus mental model (AID, multisig AID, KEL, rotation, ACDC, IPEX). First the vocabulary, then the concepts. Neither can be skimmed.
Operations. Staff training, multisig ceremonies with LARs and QARs, recovery for lost wallets and stolen devices, QVI fees and software budgets, internal issuance policies, audit procedures. Not protocol work, but real cost.
Infrastructure and maturity. Local dev setup needs witness pool, KERIA agent, schema server, verifier server, mock credential script. Adding one credential means hand-crafting schema, SAIDification, registering with vLEI server, writing chained issuance code, handling IPEX delivery. Breakages encountered: LMDB update, hio update, browser extension wallet update. A developer evaluating vLEI hits these in week one and forms hard-to-reverse opinions.

Analogy: the field is at protoboard stage. Primitives are sound, abstractions are missing. Adoption needs high-level SDKs, schema and credential design tooling, reference architectures, stable wallets, and bridge services.
Takeaways and next steps:

Build bridge services to wrap existing trust artifacts (Firma Digital, signed PDFs) into the vLEI graph. Extend the permissionless LE layer with domain-specific credentials. The schema space is wide open.
Raise developer experience: SDKs, schema tooling, reference architectures, stable wallets.
Solve the operational side: training, multisig ceremonies, recovery, audit, governance.

SESSION #15

OpenVTC Demo from Affinidi CEO Glenn Gore

Session Convener: Drummond Reed
Session Notes Taker(s): Drummond Reed & Margeigh Novotny

Tags / links to resources / technology discussed, related to this session:

- LF Decentralized Trust Labs OpenVTC Project
- ToIP / DIF Decentralized Trust Graph Working Group
- LF Decentralized Trust blog post describing the project
- See session recording link below.

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Affinidi CEO gave an extensive demo of the OpenVTC (Verifiable Trust Community) project that is implementing the ToIP / DIF Decentralized Trust Graph Working Group standards for verifiable trust agents (VTAs).

The entire demo was recorded—use this link with this passcode: **7Kf@3Z*b**
NOTE: The first hour of the recording is The Personal AI Superagent session (#14) that introduces the decentralized trust graph and the VTAs that Glenn demo'd.

ACDC Delegation Authorization EVAC Solution to Broken Delegation Chains.

Session Convener: Sam Smith
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

slides:
https://github.com/SmithSamuelM/Papers/blob/master/presentations/AuthorizationEVAC.pdf

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Open Claw Verifiable Identity Documents / Sarah C

Session Convener: Sarah C
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

2FA for Email Access - Why not pretty soon?

Session Convener: Jim Fenton
Session Notes Taker(s): Max Gerber, Jin Wen

Tags / links to resources / technology discussed, related to this session:

- https://jmap.io/
- https://datatracker.ietf.org/doc/draft-zehavi-oauth-app2app-browserless/01/
- https://dovecot.org/
- IMAP, SMTP, JMAP; OAuth for IMAP and ActiveSync; SASL; RequireTLS
- Dovecot (Jim's IMAP server); MailMate (Jim's client, after Thunderbird crashed on a malformed message); Thunderbird (presets for Yahoo, Microsoft, Gmail, AOL, "and probably eight other providers")
- FastMail as JMAP spearhead, plus "primarily European vendors"
- "OAuth2 for native app" / app-to-app — IETF work in flight (see browserless app2app draft above)
- Microsoft 365, Gmail, Charles Schwab, Bluesky's "restricted-privilege passwords"; Duo; 1Password; FIDO UAF / IoT
- App-specific passwords (the thing the room agreed should die)

https://jmap.io/
https://datatracker.ietf.org/doc/draft-zehavi-oauth-app2app-browserless/01/
https://dovecot.org/

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Jim's session outline -
Captured from Max's notes, lightly compressed:

- **Motivation** — account resets sent through email; growing volume of sensitive (especially financial) data in email; sustained online guessing attacks against mail servers; strong desire to access email from anywhere; sessions last forever with no visibility into what's active.
- **Requirements** — ease of use; support a wide range of 2FA methods, including passkeys and hardware keys.
- **Approaches** — (a) browser-based authentication (kicked off when needed; some see this as a bad experience); (b) **SASL** (raised as an option, not deeply discussed); (c) direct authentication from the email client (best experience, but requires the client to support each individual auth method, and some methods are browser-only today).
- **Associated needs** — let users or the MUA cap session length; surface which sessions are currently authenticated (likely a web feature).

Topic framing — Jim's problem statement

- Users logged into email in a lot of places.
- Email used for password resets, sensitive personal info.
- Hundreds of password attempts against email server.
- Can't apply security measures (like firewall rules) to email.
- The thing with the most sensitive information has the least secure posture.

Jim runs his own email — his own SMTP server plus Dovecot for IMAP — on his own domain. An attendee with enterprise-email operational experience framed the contrast: enterprise email (Microsoft 365, Gmail) handles MFA by delegating IMAP/ActiveSync auth to OAuth, which kicks out to a browser, completes the OAuth + MFA flow there, and hands back a token with an "implicit refresh token"; admins can configure how long that refresh token lives. Jim's case is different because he is not running a major provider.

What Jim said in the room about flipping out to a browser: **"the internet isn't all the web."** He wants an authentication path that works from inside the mail client without flipping over to a web browser and back.

He also stated explicitly: **"I would really like to have some way of doing some sort of like native two-factor authentication in IMAP directly that I do from my from my email client."**

Requirements Jim wrote on the board

1. **Ease of use.**
2. **Not browser based — large attack surface area.**
3. **RequireTLS — requires SMTP negotiates over TLS.**
4. **Support a wide range of 2FA methods** — passkeys and hardware keys called out explicitly (from Max's outline).

A fifth, derived from discussion rather than the board: **trade the password for a token.** Jim: "trading the uh the password for some sort of a token in the in the uh authentication would be a good idea still."

Why "just use the OAuth/browser flow" did not satisfy the room

Several attendees pointed out that the OAuth-via-browser flow is already there and already works for Gmail and Microsoft 365 IMAP. The pushback that kept the discussion going:

- **Client discovery is hit-or-miss.** A participant said most modern clients (Thunderbird was the example) ship out-of-the-box configurations for Yahoo, Microsoft, Gmail, AOL "and probably eight other providers" and know to route to login.yahoo.com, login.google or gmail, etc. Another participant pushed back: even if the IMAP server returns OAuth as an auth type, whether the client actually knows how to do OAuth against an arbitrary server depends on the client. The same participant noted that the OAuth-with-IMAP spec includes a way for the server to return the URL to send the user to. Jim cannot count on third-party clients doing the right thing against his own Dovecot.
- **Browser as auth surface is fat.** Jim pushed back on treating the browser as the trusted authentication agent: phishing, man-in-the-browser, "maybe you've installed a... extension that... wasn't fully legit." Another participant added that browsers keep growing capabilities — "browsers now have access to serial ports. And USB ports... Bluetooth and now it wants to scan all the devices." Jim repeated his preference for a *focused authentication agent* that browsers and email clients (and other applications) could call.
- **Authentication fatigue if you naively put MFA in the IMAP path.** A participant described having pushed back internally against making the email system itself prompt for MFA, because the user's phone might decide to load mail in the background and trigger a 2FA prompt the user doesn't understand: "you get in the habit of going, sure, whatever, I'll just say yes all the time." Their working rule was that 2FA should only happen "when somebody's sitting there in a browser being forced to authenticate, because then it's obvious." Jim agreed: "you're basically describing an authentication fatigue attack."
- **Long tail of self-hosters.** A participant argued the browser-bounce gives the service provider a place to inject "various security functions" — historical cookies, usage signals, contracted security services — to help ensure it is really the account holder, "versus an attacker sitting in a coffee shop in Vietnam through a script." Jim's counter, verbatim: "I don't want the security to depend on using a large service provider because I'm not one. I continue to want to do it myself." The same participant noted single-host operators have a different problem set than a provider "servicing like 500 million accounts across the globe."

The "focused authentication agent" idea

This was the most concrete proposal in the room. Jim asked, verbatim: "why is the browser special? Why doesn't the, you know, why isn't there some other very focused application that does the OAuth flow? And your browser talks to it and your email client talks to it and all of that."

A participant picked this up by pointing at IETF work on app-to-app OAuth on mobile: "you can do an OAuth flow from a native app on mobile to another native app on mobile without popping open the system browser. Like if I click the login with Facebook button in a random app, that will get routed to the Facebook mobile app, not to the Facebook website on mobile." Another participant named this as "OAuth2 for native app." A third confirmed: "it's like app-to-app communication."

Another participant generalized the mechanism to desktop, using OS URL handlers. Verbatim, lightly cleaned: "Desktop OS's have had the ability to handle bizarre URLs by passing them off to something else for forever. So all you need is, you know, password colon slash slash thing... the browser just goes, oh, password URLs, I just hand this off to the OS, and the OS goes, oh, Jim decided his password handler is this. It gets the blob of data, figures out what to do, does the authentication... returns a token or something. Or it hands a URL back to the OS, and the OS goes, oh okay, that's an HTTPS URL call. I'll give that back to Chrome or Safari."

A participant pointed at how 1Password already works for passkeys: "no matter what is invoking the passkey operation, it gets routed to one [Password]." Another explained how that wiring exists today: "you can register passkey providers in browsers... it shows up as another option in the browser itself." The 1Password user then asked the open question: "Can I register it outside of the browser? Like... if my IDE wants to do a passkey?" The answer in the room was that today the 1Password browser extension is what wires it up.

The broader pitch from the same participant: a password-manager-style local agent has a "lot smaller surface to protect" than a browser, "and the problem with browsers, they keep adding new stuff to browsers."

Jim's worry about going down this path was that the mail client would have to support whatever authentication method was being used: "if it's going to be passkeys, they've got to understand passkeys or understand how to do a passkey transaction. If it's any number of other things, Duo in your case or whatever." The earlier point about client certificates illustrates this: "we were using client certificates for some stuff for a while. And it's a lot easier to get the browser to deal with the client certificate than... Some mail clients were okay with it, others weren't."

The email-as-identity-anchor problem (raised by a CIAM-provider builder)

A participant who builds CIAM authorization servers pivoted to a problem 2FA on a single mailbox does not solve. Verbatim setup: "I build authorization servers as a CIAM provider. And so I have users that are creating accounts. And when a user creates an account, they give me their username, they give me their password, and they give me their email. And then I send all the password reset emails to their email. But I have no idea if this email is secure... I'm still going to send the password reset emails there anyway because what else could I do? And if I had some way to look at an email address and then say, is there actually 2FA being used to secure this email, that would help me sleep better at night."

Max captured this as three questions on the whiteboard:

- **Can websites tell if your email is secure?**
- **Could you block a signup if the email isn't using MFA?**
- **How to tell the whole email chain is secure enough to pass sensitive information on?**

Discussion of these in the room:

- **RequireTLS** is the closest adjacent thing Jim could name; he had alluded to it earlier in the session. It addresses transport (does SMTP negotiate TLS) and Jim noted "hardly anybody supports it yet."
- A participant pointed out you can configure your sending MTA to not fall back to plain SMTP if the peer won't do TLS — and then "the problem you have is how do you tell the user you didn't deliver mail to them when you can't send an email to tell them you can't deliver email."
- Even with TLS, the recipient's password quality is unknown. From the room: "if your password is, you know, Bob, TLS isn't gonna help you with anything." The CIAM-provider builder's response: "I can trust that if you make your password Bob, Google's gonna yell at you."
- For the long tail, the CIAM-provider builder framed it directly: "if there is a logo and we can show the logo on a screen and someone on the street has heard about it, we've probably tested their account creation flow and their password reset flow"; for self-hosted or small providers, "we have no idea, you know, how secure that is. But we're still gonna stake their identity on that because that's the best option we have."
- As a practical bootstrap, Jim suggested: at account creation, when you send the verify-this-email message, "if you can't negotiate TLS then... you can kind of say to the user" that this email isn't a safe channel.
- A participant separately raised that DNS lookups (MX records, CNAMEs to Google) can reveal where someone is actually hosting, but agreed it doesn't scale as a general assurance check. Jim and that participant both endorsed the practice of owning your own domain and pointing it at whatever provider you currently use for portability.

JMAP discussion

Jim disclosed he is co-chair of the JMAP working group and was openly conflicted: "I'd be running JMAP if uh DoveCot only supported it... I'm the co-chair of the JMAP working group, so I should be running it."

Substantive points from the room:

- **JMAP doesn't talk about the authentication piece as part of the protocol.** Jim, verbatim: "JMAP doesn't really talk about the authentication piece of it as part of the protocol. But it is more of a... HTTP type of access. It isn't its own socket number and all of that sort of thing. So really, the kind of the web-based authentication methods are much more directly applicable to JMAP." (This is exactly Max's whiteboard line: *"JMAP / Doesn't talk about authentication piece, more http-based than imap."*)
- **JMAP proxy in front of IMAP** was offered as an interim suggestion: "if you got IMAP servers for OAuth, you could get an OAuth token to give it to JMAP proxy." Jim's reaction: "just one more level of abstraction."
- **Beyond email.** Jim described JMAP as a generalized synchronization mechanism — "email as well as calendars and address books and to-do lists... You can just use it for files and stuff, you know, like... WebDAV. Like WebDAV kind of."

- **JMAP adoption.** Jim: "FastMail is kind of spearheading it, although there are a few primarily European vendors that are also doing JMAP support." He has been "trying to migrate to something that will [support JMAP] and it's not working out."
- **FIDO outside the browser.** A participant: "FIDO UAF is mobile, it's starting with the mobile oriented... only run on mobile. But there's that some of the FIDO specification could be run on the IoT device, so there's no browser, no email, it just goes through the IoT communication channel." Their framing question: "why not, if this type is public-key cryptography, can they run on any kind of around other protocols, just providing the 2FA multi-factor control?"

App-specific passwords — consensus that they are bad

- **The "specific" part is a lie.** A participant: "they said it was app specific, but most of them were not... I now have five passwords... I want to create an app specific password for my mail client. Cool. Hey, I just give it to my calendar client and it works just fine for calendar. And I put it in my Google Drive thing and it worked just fine for that too."
- **They break MFA and SSO.** From the room: "we thought about doing it at Stanford. And yeah, you can, but you can't do SSO. And you can't do MFA. Like Google at least, you could have specific passwords, but it broke MFA and it broke SSO."
- **First-level abuse vector.** A participant: "App passwords are super ripe for abuse too. Like it's one of the highest first-line... male abuse vector from like a spam and fraudulent activity sort of."
- **Active sunset by major providers.** A participant from a major provider: "we're trying to get rid of it. Like we've made it so that if you've never had one on your account, you can no longer add it. And if you find the one way to actually find the option to add it, you still have to re-authenticate with 2FA in order to do it. And then for accounts that have had it, like we've been tracking activity and basically deactivating anything that isn't actively used... we're trying to do a sunset in the industry with like Microsoft and Google as well."
- **Bluesky example (Jim).** Jim: Bluesky offers what it casts "not as an app specific password but as sort of a like restricted privilege password, which the restricted privilege part is great, but... it just generated a like a UUID sort of string for me and I just had to store it."

The constructive variant from operational experience: instead of "another password for the same account," issue **another principal entirely** with reduced scope. Verbatim from a participant: "we always create extra principals for admins. So it wasn't that you know your regular account had two passwords. It was like you had... my SWL and my SWL slash root which let me log into root and a bunch of things and my SWL slash admin which was my KDC administrator. In Slack we do similar things... we have like dash D for the domain administrators which requires PIV card obviously and we have dash A for people who need admin access to local workstations." The scoping is enforced at the IdP, not faked at the client.

Session management gaps

Drawn from the room, not from Max's notes:

- Jim wants individual users to be able to set session durations: "in a non-enterprise situation, you should be able to set your own session durations." Today his session limits are "infinite everywhere."
- Jim also wants visibility into where he is logged in: "I don't have a... some web applications give you the ability to kind of go and say, okay, well here's all of the places that you're logged in from. Do you want to log any of those out?"
- A participant noted that even in enterprise this is hard. For Office 365, "it's hidden away in some obscure place. And it took us multiple calls to Microsoft to find out one, it was even possible, and two, to find out where to go to change it."
- Identifying *which* session to revoke is partially broken. From the room: old Twitter clients used to include some per-instance information beyond the client ID, "so if you had two iPhones, you're a little bit screwed because it just said, you know, Twitter client X on iPhone twice. And you're like, I don't know which is which. But... if you had an Android, an iPhone, an iPad, and a Mac and a Windows machine, even if it was the same app and the same client ID, there'd be enough information to go, oh, that's my Windows laptop that I just lost. So I'm gonna revoke that one."
- "A lot of it's just... revoking the refresh tokens that the client's got. And just having a way to identify the client."

The Stanford / Slack pattern — captured for future reference

A participant described the cleanest "2FA for email" pattern they had actually deployed:

- The mail server itself does not do MFA. The OAuth token is what authenticates against the mail server. Verbatim: "it's not really 2FA from mail. It's just that in order to get [a] token, you have to do 2FA. But there's no... the mail server decides to do a second factor authentication for whatever reason. It's just like, yeah, no, in order to get that auth token to talk to the mail server, you had to do a full blown authentication."
- Token lifetime was around 120 days: "we force the authentication so that it didn't matter, you know, since you only had to do this every hundred and twenty days, we didn't care if you already logged into IdP that day."
- They told Duo to force a fresh challenge regardless of recent MFA: "we were telling Duo I don't care if they've done MFA today. This is a brand new thing. They have to do it right now. So there's no way out of doing a full authentication to get a new auth token."
- Re-auth used "your username and your password or your certificate" plus full MFA, kicked off as a forced authentication "coming from Azure or whatever."

Jim's own framing of why doing this for himself is awkward: MailMate (or Thunderbird) needs to understand from Dovecot that a token is required, then redirect to a browser, get the token, hand it back. The pragmatic shortcut from the room: "you just have to fake it and just go yeah, here's your password. Happens to be a token that's got a limited time on it and at some point is gonna say no. And you're gonna start complaining at me that you can't authenticate. And I'm gonna have to go to my other thing that lets me create a brand new password for you, but it requires me to go through multi-factor authentication and then it'll show me this thing on the screen I can copy and paste back into my mail client and it'll work for another 30 days, 60 days, whatever it is." Jim's reaction: in his own house, his wife and kids would not tolerate that kind of friction.

Outstanding questions surfaced in the room

- **Native two-factor authentication in IMAP**, done from inside the email client without flipping to a browser. Jim's stated goal; no design produced.
- **Can a website tell whether a given email account is secured by 2FA?** From the CIAM-provider builder; Max captured it; no answer.
- **Could a relying party block a signup if the email account isn't using MFA?** Same; no answer.
- **How to tell whether the whole email chain is secure enough to carry sensitive information?** Max's whiteboard; no answer.
- **JMAP support in Dovecot.** Jim wants it, doesn't have it, has been "trying to migrate to something that will and it's not working out."
- **Per-instance client identification and per-session revocation.** Surfaced in the room; not solved.

Max's Note:

- Users logged into email in a lot of places
- Email used for password resets, sensitive personal info
- Hundreds of password attempts against email server
- Can't apply security measures (like firewall rules) to email

Thing with the most sensitive information has the least secure posture

—----------

Requirements

- Ease of Use
- Not browser based - large attack surface area

Require TLS - requires SMTP negotiates over TLS

Can websites tell if your email is secure?

- Email used for password reset
- Could you block a singup if the email isn't using MFA?
- How to tell the whole email chain is secure enough to pass sensitive information on?

JMAP

- Doesn't talk about authentication piece, more http-based than imap
- web based authentication methods are a more natural fit
- Generalized synchronization mechanism; can use for email and calendars and files and todo lists
-

—

Session “agenda” outline:

Motivation

* Account resets frequently sent through email
* More and more sensitive data (especially financial) being included in email
* Getting lots of online guessing attacks on mail servers
* Strong desire to access email from anywhere
* Sessions last FOREVER, no visibility to what's active

Requirements

* Ease of use is really important
* Need to use a wide range of 2FA methods, incl passkeys, hardware keys

Approaches

Browser authentication

* When authentication is needed, kick off broswer session
** Seen by some as bad experience

SASL??

Direct authentication from email client

* Requires support of individual authentication methods in email client
* May not be possible because access to some stuff is limited to browsers
* Best experience

Associated needs

* Limit session length, by either MUA or by user
* Visibility about sessions currently authenticated (probably web feature)

Conveying Trust - Mechanisms to communicate trust & compliance to Policy

Session Convener: David Kelts
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

No NOTES SUBMITTED

Registering an Agent with aria.bar ~ How can we collaborate?

Session Convener: Aaron Grego
Session Notes Taker(s): Adolfo Grego

Tags / links to resources / technology discussed, related to this session:

https://aria.bar
https://trustlayer.foundation

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

We continued showing operational and institutional details about ARIA, the Agent Registry for Identity and Authorization of AI agents. We realize that becoming a specification will take time and we need to find a use case where a company wants to identify AI agents knocking at their doors.

We clarified that ARIA uses DNS to anchor each agent and have some level of traceability and validation from the owner of the domain name itself,and that our roadmap allows for the project to eventually bridge to blockchain technology and have both ways to anchor and validate an AI agent.

Being an open protocol, our main proposal is to integrate all existing technologies for Identity verification and validation, expanded to AI agents. Therefore, instead of proposing new or different approaches to solve problems, we take what other specialized organizations have already implemented and integrate it ourselves as part of our interoperability.

We went over our spec, and openly invited companies or developers to start registering an AI agent and test how that agent presents itself to their system and to use our pre made templates to publish a policy to identify such agents and allow/deny them access to services.

Regenerative Accelerationism 2.0 Continued from Space F in Session 14

Session Convener: Kaliya and Friends
Session Notes Taker(s):

Tags / links to resources / technology discussed, related to this session:

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

NO NOTES SUBMITTED

Facing the Dumpster Fire of Age Verification

Session Convener: Brent Shambaugh
Session Notes Taker(s): Brent Shambaugh, and perhaps others

Discussion notes, key understandings, outstanding questions, observations, and, if appropriate to this discussion: action items, next steps:

Register on Platform ... if you live in Apple Land it is fine
Binding to Platform is not a good idea.
California Bill ... Not able to prove Age ..
Real Powers of Attorney ...
Selective Token to a Caregiver .. with this thing do this with digital presence ...user is liable to cheat
Assume delegation system works ... lot of Oauth works on that ... stops token from going to someone else
incentive prove there is a half .. so many way
where the SMART phone is not the solution to deal with ... Age verification ...not many categories you
can fall into... issue tiny credentials .... categories .... so many categories 4 age verification
Many other ages .... sometimes it ties back to issuer of IDs,,,
and tied to jurisdiction .... finite demand of Age categories .... 4 each category ... platform has to deal...
overlap all categories ... venn diagram makes age unique...
X for all countries ... categories ... over under 15, 21,
Insurance to have kids until 25 ... Insurance companies will get page .. at least 4 rental ... what type of

data from DL .. Having a DL is the thing you want to process ... shopping for lowest mortgage rates...
mortgage ... so and so has over X amount in account
Gating factor...are you able to bid 4 this house primary for ZKP is this generally

It is up to the parent to decide ...but that is not the problem ... things so easy

Create safe spaces 4 kids ... rather than lock kids out of content a system that locks out adults rather than
locks kids ouf of content.. adult does not have more people ..not that easy to circumvent ... weakness ... adults have
to convince them to give the kid

start with age verification in an area ... that does not creep people out

keep adults out of kids spaces ...

In a couple of years do it on a web browser without having to install anything ... read ... run local recognition ...
new for OS 4 kids . parents not best for kids.

Nina Kerkez, MBA, CAMS • 2nd
Women in FinTech Powerlist 2024 | Senior Director - Future State o...
4w • Edited •

+ Follow ···

Just wraped up on Internet Identity Workshop in Palo Alto, and my brain hasn't quite caught up yet.

IIW is unlike any conference I've been to. There's no agenda until the morning it begins, attendees propose and run every session themselves. What emerges is gloriously unpredictable: this time, deep dives into self-sovereign identity, DIDs, KERI, and MyTerms sat alongside genuinely unsettling conversations about what quantum computing means for the trust infrastructure we're quietly building underneath everything digital.
That last one in particular stopped me in my tracks. We're laying foundations today that may need to hold for decades, and the people in that room are the ones actually thinking about whether they will.

I went in curious. I came out with better questions, a longer reading list, and a strong sense that identity is one of the most consequential - and underappreciated - problems of our time.
Already looking forward to the next one.

Erica Connell and 70 others 2 comments

Speed Demo Hour / Wednesday April 29

The **IIW Speed Demo format** involves each person Demoing giving a **5-minute demonstration** of their service, product, physical device, **10 times** to 10 different small groups, rotating through to view them over the course of the hour. **Demo Hour takes place on Wednesday after lunch from 1:30 - 2:30.**

There will be 20 Demo Tables in the Grand Hall each with a # Sign on it that corresponds to the Demo taking place at that table. People rotate through the tables/Demo's in a self-organized way ~ that's a little loud, seemingly chaotic and free flowing, but works!

See the list of Demos via the Demo List below and decide ahead of time the Demo's you'd like to see. You'll be able to see 10 of the 20 Demo's over the hour.

TABLE	Demo Description	MORE INFO
#1	**Technical Associates Group LLC / Functional Identity**: Jeff Orgel The demo will review the idea of how functional identity can exist on a mark, known as the X, and that mark is more or less changeable based on platforms and systems. The demo explores how we can modify, or in some cases, even set the terms of functionality regarding the X we will be on.	More Info Here
#2	**Execution-Time Delegation Harness:** Deb Bucci (Independent Researcher) A working mock-up exploring execution-time evaluation of delegated authority — an intent gateway that determines whether a delegated action still aligns with a subject's prior intent at the moment it is requested.	More Info Here
#3	**Tenuo:** Niki Aimable Niyikiza **URL:** https://tenuo.ai Tenuo: attenuating authorization tokens for AI agents. Each delegation hop cryptographically narrows tool capabilities and argument constraints. Full chain verified offline, no auth server round-trip. IETF draft-niyikiza-oauth-attenuating-agent-tokens. Live demo.	More Info Here
#4	**Blockchain Commons / XIDs, SSI, and Radical Privacy Offered by the Blockchain Commons' Gordian Decentralized Identity Stack:** Christopher Allen URL: developer.blockchaincommons.com (developer docs) learningxids.blockchaincommons.com/ (self-paced tutorial) XIDs (eXtensible IDentifiers) are DID-inspired decentralized identifiers built on Gordian Envelope that enables selective disclosure and redaction of controller documents — so holders, not issuers, control what's revealed. Demo shows XIDs addressing the W3C Amira use case.	More Info Here

#5	**Brevity-Lang :** Chris Kula **URL:** https://cambistry.github.io/brevityWhat if custom software didn't have to start as SaaS? Live demo of Brevity in the browser: source files as actors, HTML templating, reactive updates, and language-independent interop for local-first applications you can build and own.	More Info Here
#6	**SELF:** Eric Johnson, Will Harrison & Hank Brigham **URL:** https://selfid.com/ Self is the decentralized identity layer transforming personal data into a portable, secure, self-sovereign assets. We enable platforms to offload PII storage and liability while providing users with seamless onboarding across the global digital ecosystem.	More Info Here
#7	**Cvera | Verifiable Claim Tokens:** Nicholas Whitehouse, Paul Coe **URL:** https://github.com/kola-white/rsa-attestation-engine GitHub repository (README includes project overview and demo context) https://www.cvera.app/ Cvera enables user-held, cryptographically verifiable claims issued by trusted entities and independently verified without contacting the issuer. The demo shows token issuance, portable credentials, and verification via signature, trust chain, and revocation.	More Info Here
#8	**Skyfire/Transparent & Trustworthy Agentic Commerce with KYAPay**: Ankit Agarwal **URL:** https://kyapay.org/ Know Your Agent (KYA) is a protocol that enables AI agents identify themselves, their platforms, and users to online businesses for agent authentication. KYAPay extends this by declaring user intent and supporting tokenized payments for secure and seamless agentic commerce.	More Info Here
#9	**Proof of Human Delegation (PoHD) for Trustworthy AI Agents**: Kenta Takahashi (Hitachi, Ltd.) and Takayuki Suzuki (Hitachi America, Ltd.) **URL**: https://rd.hitachi.com/_tags/Public_Biometric_Authentication_Infrastructure and https://www.youtube.com/watch?v=pBFQFPjD9qc PoHD demonstrates how AI agents can prove human delegation using biometrics. Hitachi's PBI technology derives a signing key from biometric features to sign a delegation credential proving that the agent acts on behalf of the user and within the user's stated intent.	More Info Here
#10	**Vidos**: Rob De Feo, CTO **URL:** https://vidos.id Using the Vidos enterprise credential verification platform, we show flows including completing an age-verified purchase and hiring a car using an AI agent and the EUDI Wallet. The agent calls the Vidos verifier API via OpenID4VP. W3C VC, SD-JWT VC, ISO 18013-5 supported.	More Info Here
#11	**Ceros by Beyond Identity**: Colton Chojnacki **URL:** Beyondidentity.ai Ceros is the agentic AI trust layer. Unified identity, observability, and governance for every agent, tool, and action.	More Info Here
#12	**Center Identity - Key Rotation with FIDO2 WebAuthn**: Matt Vogel **URL:** https://centeridentity.com Demonstrating how Center Identity uses key rotation with FIDO2 WebAuthn to mitigate key compromise risks in agentic AI, eliminate long-lived credentials, and enable continuous, phishing-resistant authentication.	More Info Here

#13	**AAuth Protocol:** Dick Hardt - Hellō **URL:** https://AAuth.dev AAuth is an open protocol giving agents their own cryptographic identity — no pre-registration, no shared secrets, no API keys. Resources adopt incrementally: start by verifying an agent's signature, add an authorization endpoint for self-service registration, or federate with person servers for user identity and mission-scoped governance. Demo includes a web playground for resources to request identity claims and a CLI for an agent to register at a service.	More Info Here
#14	**MyKey, MyTerms and XMLUI**: **Iain Henderson, Jon Udell** **URLs:** https://www.mykey.me, xmlui.org, see also https://judell.github.io/myterms/ **MyKey** is a fiduciary mobile app that combines enabling decentralised identifiers which people create, own and control themselves, a privacy-friendly approach to 'social login' and federated key recovery for genuine service portability. To that we add the capability to record MyTerms agreement preferences; and then propose, negotiate, sign and record agreements where relevant. **XMLUI** makes web user interfaces easier to build, maintain, and test. Describe apps in semantic terms that enable people and LLMs to work together effectively. The demo is an interactive explainer of MyTerms that shows how XMLUI can wrap React components (in this case react-flow) for use by people who do not grok React or CSS, are comfortable working with LLMs, and value human-understandable code.	More Info Here
#15	**Writerslogic Inc / WritersProof:** David Condrey URL: https://writerslogic.com/iiw WritersProof is a desktop app that silently captures cryptographic proof of human authorship as you write. It entangles your identity, keystrokes, and timing into an unforgeable hash chain — not detecting AI, but proving you did the work. Free download, works with any editor.	More Info Here
#16	SensiNM, LLC/ **Privacy-Preserving Multi-Agent Clinical Decision Support**: Moises E. Jaramillo Product Link: Still Private We make Agentic AI HIPPA compliant by leveraging privacy preserving technologies such as DIDs, ZCAP-LDs, and Differential Privacy. Our first use case targets the Oncology Domain and it integrates FHIR endpoints to aggregate and summarize a patient's relevant information to help the provider make a clinical decision.	More Info Here
#17	**Freewallet + Wallet Attached Storage:** Dmitri Zagidulin - Interop Alliance URL: https://freewallet.me Freewallet is a free, open source, open spec focused VCALM (OID4* soon!) web wallet for DIDs and VCs, using WAS (Wallet Attached Storage), as an example implementation of BYOE (Bring Your Own Everything) app dev stack.	More Info Here
#18	**MyWellWallet. A Personal health wallet with a local LLM:** Mahesh Balan Doctor of Technology Candidate, Claremont Graduate University. URL: https://github.com/maheshbalan/myWellWallet MyWellWallet: Patient-Owned Health Intelligence Using MCP, FHIR, and Local LLMs MyWellWallet is a privacy-first health wallet that gives patients – starting with those managing Type 2 diabetes – a single, intelligent view of their entire health history, owned entirely on their device.	More Info Here
#19	**OpenID AuthZEN Working Group:** Alex Olivier (Cerbos) & Atul Tulshibagwale (Crowdstrike)	More Info

	URL: https://openid.github.io/authzen/authzen-mcp-profile-1_0.html A reference implementation of the OpenID AuthZEN MCP Profile draft spec enabling MCP gateways and servers to perform fine-grained, parameter-level authorization checks via an AuthZen Policy Decision Point (PDP) before executing MCP tools.	**Here**
#20	**Authologic:** Jarek Sygitowicz / Flora Frend **URL:** https://authologic.com We will demonstrate practical implementation of the EUDI using Digital Credentials API on iOS and Android with the fallback options for the gen-1 eIDs and legacy OCR/liveness. Will call that implementation OmniID.	**More Info Here**

One of IIW's distinctive features is the speed demo hour on Wednesday afternoon. Twenty tables, each with a numbered sign, fill the Grand Hall. Each demonstrator gives a five-minute demo, then the audience rotates to the next table. If you're disciplined, you can see 10 of the 20 demos over the course of an hour. It is loud and seemingly chaotic, but it works. Demo hour is about working code and running systems. You can tell a lot about a community by what it chooses to demo.

Thank You to our Women's Breakfast Sponsor Open ID Foundation

Event Photos taken by Doc Searls

Day One:

https://www.flickr.com/photos/docsearls/albums/72177720333615870

Day Two:

https://www.flickr.com/photos/docsearls/albums/72177720333711613

Day Three:

https://www.flickr.com/photos/docsearls/albums/72177720333673351

Identity Funnies – (comic strips) shared by Alan Carp!

Baby Blues

Stay Connected with the Community Over Time – Blog Posts from Community Members

New Community Resource
Each week Kaliya, Identity Woman and Informiner publish a round of the week's news from the industry. It is called **Identosphere - Sovereign Identity Updates (weekly newsletter)**
You can find it here: https://newsletter.identosphere.net/

As a follow up to the session 'Let's Bring Blogging Back' an IIW Blog aggregator has been created here: https://identosphere.net
If you want your blog to be included please email Kaliya: kaliya@identitywoman.net

A BlogPod was created at IIW - Link to IIW Slack –
https://iiw.slack.com/archives/C013KKU7ZA4
If you have trouble getting in, email Kaliya@identitywoman.net with BlogPod in the Subject.

Planet Identity Revived ~ @identitywoman & @#InfoMiner cleared out & updated Planet Identity (see links below) you can support the work here:
https://www.patreon.com/user?u=35769676

IIW Community Personal Blog's shared via: https://identosphere.net/blogcatcher/
IIW Community dot.org's in the IIW Space: https://identosphere.net/blogcatcher/orgsfeed/

Upcoming IIW Inspired™ Regional Events

Join Us!
[Register here](#)

[diceidentity.org](#)

dice is an IIW Inspired™ Regional Event co-hosted in Copenhagen with Idura

dice Digital Identity unConference EUROPE

With Our co-host [DIDx](#)

🗓️ [DID: UNCONF AFRICA](#) #3 / March 16 – 18, 2027

📍 [STIAS: Stellenbosch Institute for Advanced Study](#), Stellenbosch, South Africa

🎟️ Join Us! [Registration](#) Open in July 2026

Hope to See you November 3,4 & 5, 2026 for IIWXLIII

The 43rd Internet Identity Workshop

REGISTRATION OPEN in June

www.InternetIdentityWorkshop.com