



From Attention to Intention

Verifiable product data and agent identity

Branimir Rakić

Trace Labs, CTO & Founder at OriginTrail

Hello from a fresh dad!

Poor guy seems to be allergic to cow milk protein

Pediatrician proposed a new formula

So I asked ChatGPT ...



“Will this work with our new formula?”



Baby Brezza Formula Pro



Novalac Allernova AR




The machine has 10 programs for different formula types

Scan the GS1 barcode to get the right program number

Nothing showed ...


The agent needs to fetch data from many sources



What kind of system are we looking for




**Reliable reasoning -
no hallucinations**




**Data source reputation &
Information provenance**



**Fair value exchange &
incentives**



**Connected
(multiple data sources)**



Data ownership



Privacy preserving

A trusted read/write knowledge base



LLM Knowledge Bases


Something I'm finding very useful recently: using LLMs to build personal knowledge bases for various topics of research interest. In this way, a large fraction of my recent token throughput is going less into manipulating code, and more into manipulating knowledge (stored as markdown and images). The latest LLMs are quite good at it. So:

Data ingest:

I index source documents (articles, papers, repos, datasets, images, etc.) into a raw/ directory, then I use an LLM to incrementally "compile" a wiki, which is just a collection of .md files in a directory structure. The wiki includes summaries of all the data in raw/, backlinks, and then it categorizes data into concepts, writes articles for them, and links them all. To convert web articles into .md files I like to use the Obsidian Web Clipper extension, and then I also use a hotkey to download all the related images to local so that my LLM can easily reference them.

IDE:

I use Obsidian as the IDE "frontend" where I can view the raw data, the compiled wiki, and the derived visualizations. Important to note that the LLM writes and maintains all of the data of the wiki, I rarely touch it directly. I've played with a few Obsidian plugins to render and view data in other ways (e.g. Marp for slides).



OriginTrail Decentralized Knowledge Graph


**Decentralized network
to query and
share knowledge**

Open source

Open standards


Deployed across variety of sectors,
currently in version 10

<https://github.com/OriginTrail/dkg>



Owned Context Graphs on the DKG

The DKG connects owned Context Graphs



Context Graph owners set participation and publishing rules - they can be public or private, curated or open for publishing.

Three memory layers in each Context Graph

Working Memory

Private knowledge on my node

Shared Working Memory

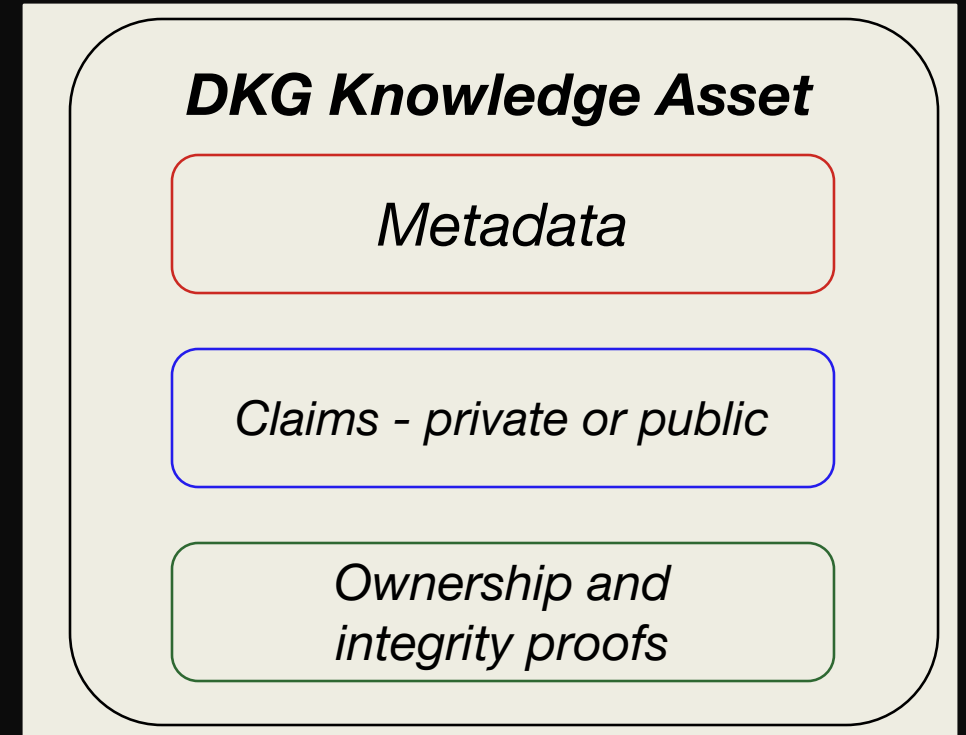

Knowledge shared with peers

Verifiable Memory

Published with a fully qualified UAL

Knowledge Assets


Inspired by W3C Verifiable Credentials



Agents can check ownership, provenance and the integrity of the content

Verifiable memory = commitment to the records

Author signs the Knowledge asset content




Content on DKG nodes

Commitments + ownership on-chain

Integrity and attribution can be checked. Product authority requires its own evidence.

DKG Context Graphs connect Knowledge Assets



Knowledge Assets contribute linked statements

Shared identifiers connect them

Each asset keeps its author and provenance

An answer backed by machine owners

“Will this work with our formula?”

Illustrative answer and owner evidence

“Yes. Owners of this model report that it works with this formula.”

Owners can share their experience as Knowledge Assets

Ownership credential

Links the contributor to the machine

Signed experience report


Identifies the same model and formula

Evidence on the DKG

The agent can inspect the source and date

DKG as a Neuro-Symbolic Marketplace

The vision: monetization of trusted AI services across the DKG



Knowledge owners and developers can earn fees for useful answers or services

OriginTrail sits at the heart of building trust globally.



Trusted factory audits

BSI and SCAN are utilizing the OriginTrail to share security audits for over 40% of US imports through DKG



Pharmaceuticals

Ensuring traceability of donated medicines, and creating an AI-ready, collaborative medical knowledge ecosystem.



Rail Travel Safety

Ensuring the safety of train travel by tracking trains, parts and rail track material used by the Swiss Federal Railways with DKG



Trusted AI Infrastructure for Governments

OriginTrail powers trusted AI platforms for government use across large-scale, domain-specific data.



OriginTrail sits at the heart of building trust globally.



Safe and sustainable construction

Integrating with the EU Digital Building LogBook, the DKG provides a tamper-proof, AI-powered database accessible to all stakeholders.



Trustable manufacturing ecosystem

Revolutionizing manufacturing with secure, real-time data exchange for a safer tomorrow.



Driving international knowledge transfer

Empowering sports organizations' Web3 communities with AI solutions for better performance and engagement.



Secure genomic data sharing

OriginTrail powers decentralized infrastructure for secure genomic data sharing and personalized, AI-driven healthcare.




Call to action:

Create a Context graph on the DKG for your agent

1. Setup a DKG node
2. Connect your agent - Hermes, OpenClaw, Codex, Claude, or anything MCP
3. Tell it to record it's memory and traces
4. Use it and test its memory

Community happy to support you on that journey



Thank you

Let's build the Decentralized Knowledge Graph together!



Github



DKG Red team chat

origintrail.io
branimir.rakic@origin-trail.com

In case you were wondering

It worked out well!



Backup

UALs extend web addressing

Universal Asset Locator: a durable reference to verifiable knowledge.

URL `https://brand.example/product`

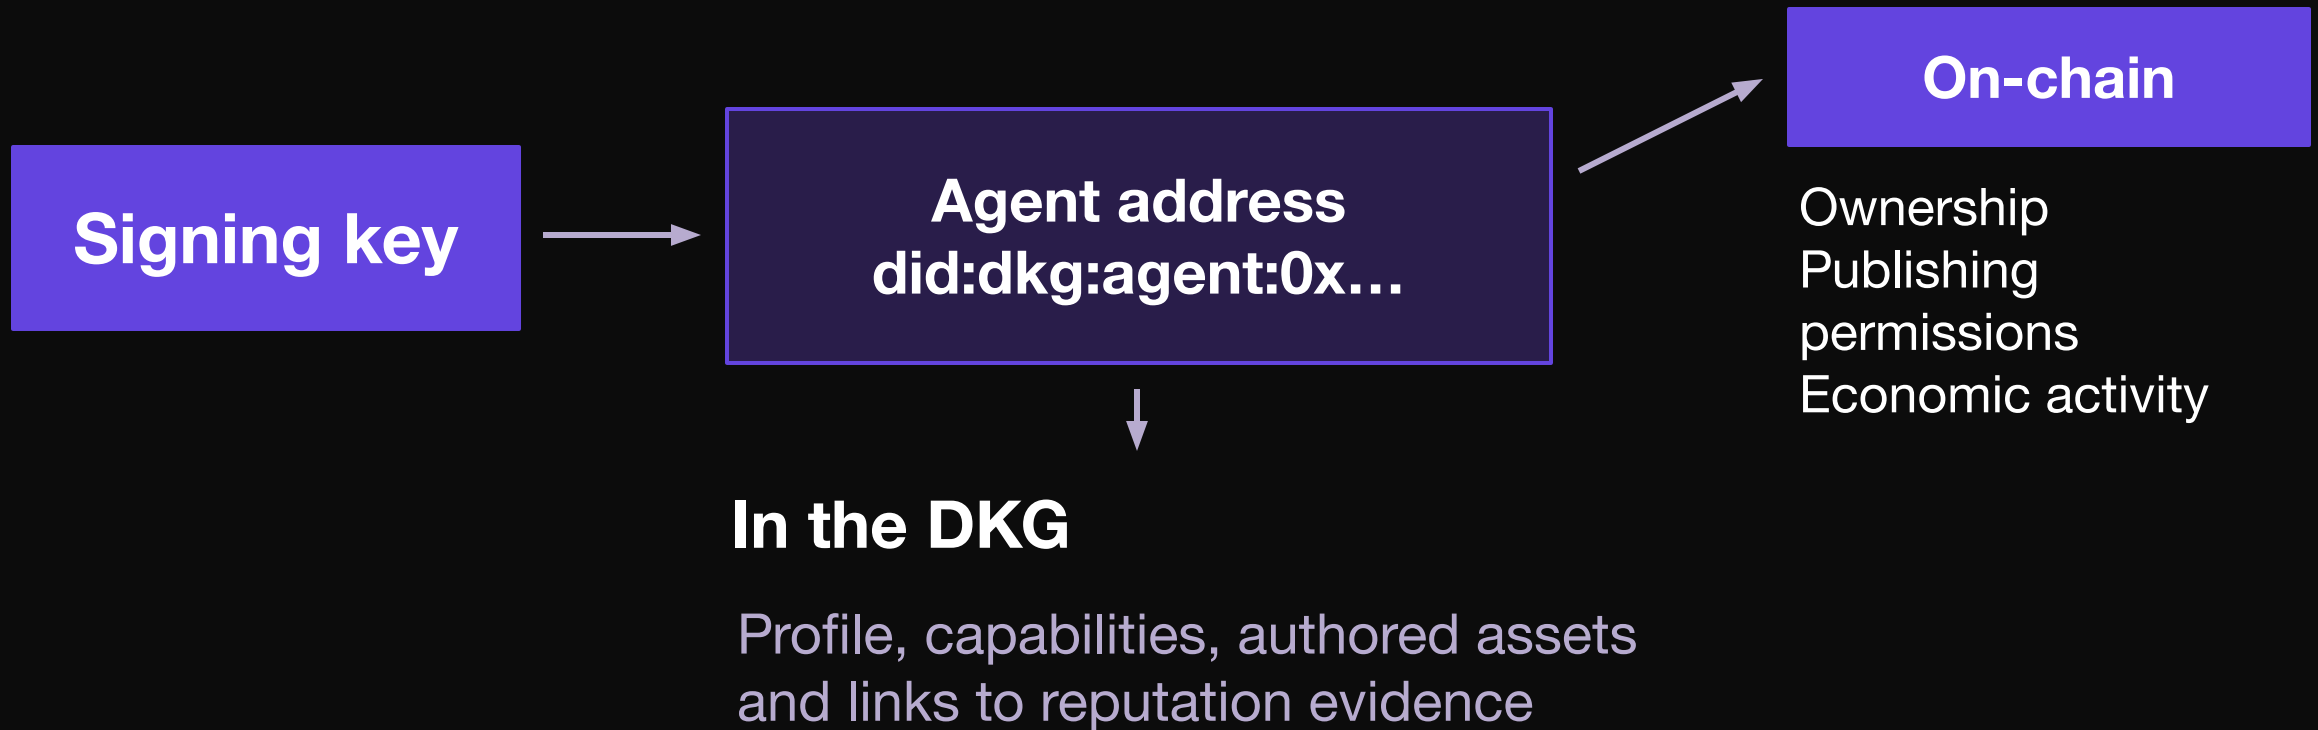
Retrieve a web resource

UAL `did:dkg:base:8453/0x<authorAgentPubKey>/<uniqueId>`

Resolve the asset, check its committed content state, check the owner


Agent identity starts with a key

A signing key establishes control. Its address anchors the agent's identity.



Reputation has an attributable source

An agent's track record connects its identity to evidence from other actors.



Reputation policy


Who evaluated what, in which context?
Which evidence does this retailer accept?

Standards at each boundary

GS1 product semantics, W3C linked data, and a DKG verification lifecycle.

GS1	Digital Link / GTIN, Web Vocabulary, EPCIS Product identity and domain meaning
W3C	RDF, JSON-LD, SPARQL, PROV, RDFC-1.0 Connected data, queries and canonicalization
DKG + chain	UALs, author attestations, ownership, TRAC Attribution, committed state and network economics

OriginTrail Decentralized Knowledge Graph (DKG) Swiss Railways implementation



Shared Railway data Knowledge Graph
across multiple nodes

No central database, each organization
running a node & owning their data

Webinar:



What the author attestation binds

The signature covers a content commitment and its asset/domain context.

Content Merkle root

Author address

Reserved Knowledge Asset ID

Scheme version

EIP-712 domain

Chain ID

Lifecycle contract

An exact citation keeps the UAL, committed revision/root and relevant publication evidence.

A shared W3C / GS1 experiment

One GS1 product across two independent agent stacks

Agree on the evidence

Product identity
Claim and publisher authority
Provenance and revision
Service and provider identity

Test the same cases

Missing information
Conflicting claims
Updated evidence

A reusable evidence profile for MCP, UCP and ACP

Who can bring a knowledge service and an agent implementation?