Tech-X 2026 Hackathon #9

Policy and Credential-Aware A2A Skill Invocation on Gaia-X Danube

An end-to-end demonstrator integrating A2A with the Gaia-X Trust Framework

Yushi Yamada · Yudai Tanabe · Cheanfei Shee · Nagare Kameno

AI Agents Meet Dataspaces

The next generation of data consumers and producers is no longer human.

DATASPACES

Cross-organizational data infrastructure

Built around data sovereignty, contracts, verifiable credentials, and pluggable trust services.

Gaia-X · IDSA · Ouranos Ecosystem

AI AGENTS

Autonomous data consumers and producers

LLM-driven agents that discover, invoke, and compose capabilities across organizational boundaries.

A2A · MCP · agentic LLM systems

*If AI Agents cannot enter Dataspaces as **first-class participants**,*
the Dataspace value proposition will not reach the next-generation workloads.

A2A Agent to Agent Protocol

01

Agent Card

`.well-known/agent-card.json`

Each agent publishes its capabilities, endpoints, and security requirements as a machine-readable JSON document.

02

Skill

`A capability, not an endpoint`

Meaningful units such as aggregate query, statistical analysis, or raw export — granular enough to express usage policy.

03

Direct A2A invocation

`Agents call agents`

Discover via Agent Card, invoke via POST /tasks. No central broker; trust is established peer-to-peer.

Bridge to Gaia-X: an A2A Skill maps to a **Service Offering / Service Instance** in Gaia-X terms.

The Integration Gap

A2A DEFINES

How agents discover and invoke

Agent Card, Skill, Trust Bundle

How peers exchange credentials

VC-JWT bundle, did:web resolution

But not: *who is authorized to call which Skill, under which contract, for which purpose.*

GAIA-X DEFINES

Who and under what terms

Verifiable Credentials, DUAs

How policies are enforced

PEP / PDP, Trust Framework, Notaries

But not: *applied at the A2A Skill invocation boundary in a public, end-to-end reference.*

The gap: no public end-to-end demonstrator combines them at the Skill boundary.

Our Proposal

> Evaluate every A2A Skill invocation against **caller VC**, **DUA status and usage constraints**, and **Skill metadata** — returning **allow** / **restrict** / **deny**.

Caller VC verification	**Provider Agent VC verification**	**DUA status & constraint check**
Policy YAML	**PEP + PDP separation**	**Decision evidence + audit**

Architecture

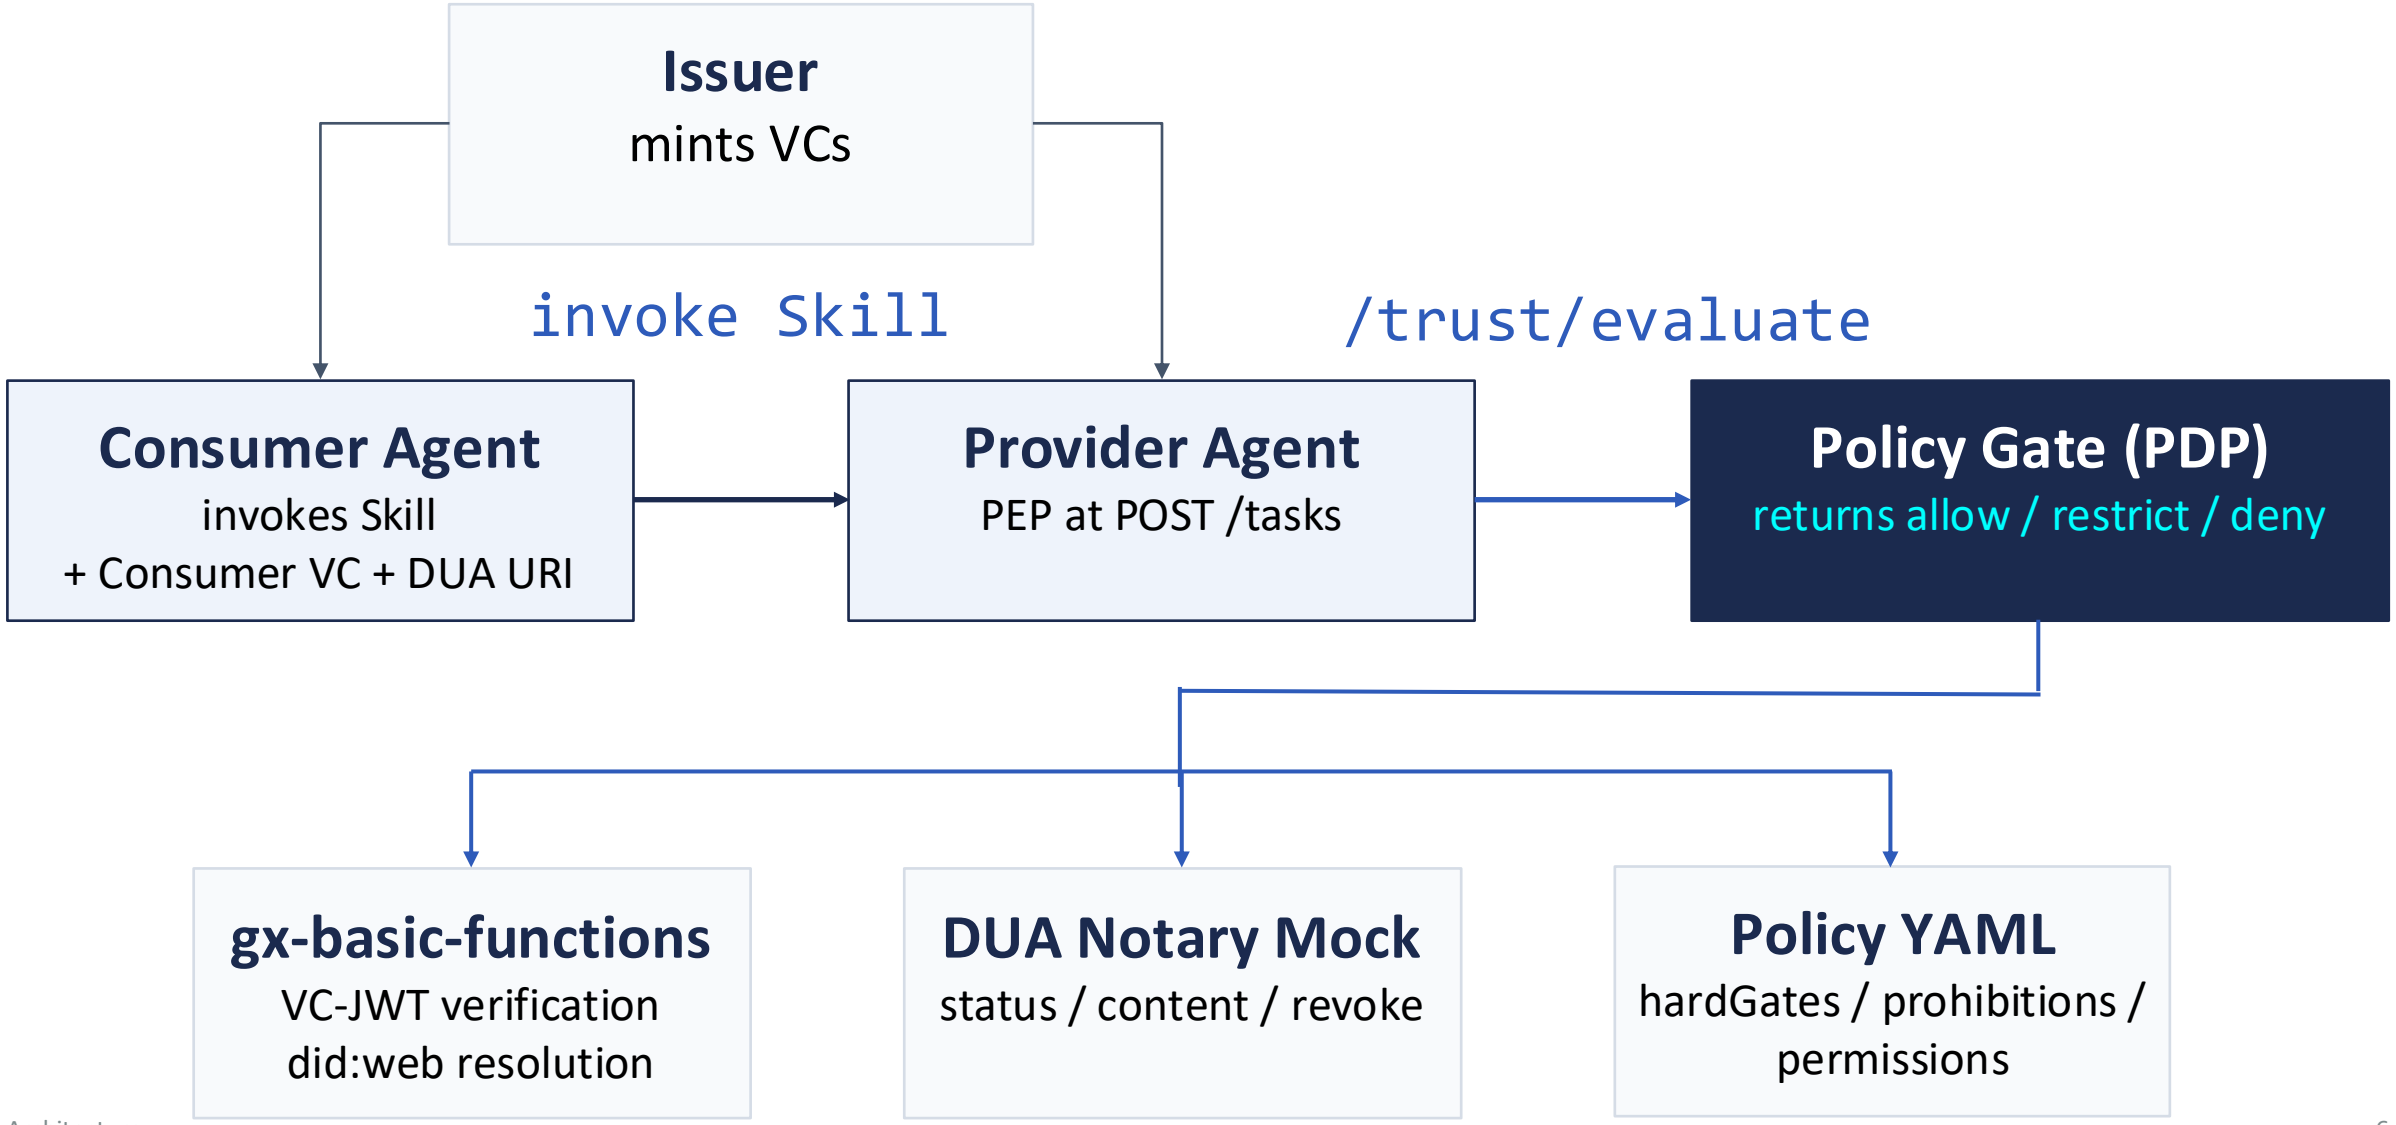

Gaia-X / Danube Alignment

Every component maps to an existing Gaia-X concept; we register as an AI_AGENT_TRUST rules engine.

This work	Gaia-X / Danube concept
Consumer Agent	Data Consumer / Participant
Provider Agent	Data Provider / Participant
Provider Skill	Service Offering capability
Agent Card	Service endpoint description + trust scope declaration
Consumer, Provider VC	authorization credential
DUA	Data Usage Agreement
DUA Notary Mock	Data Usage Agreement Notary
YAML Policy	Provider Policy / Usage Policy (ODRL-flavored)
Policy Gate (PDP)	Domain-specific rules engine (AI_AGENT_TRUST scope)
Provider /tasks PEP	Policy Enforcement Point
gx-basic-functions	Credential / DID verification support

DANUBE

Pluggable rules engines

Danube allows arbitrary rules engines to plug into the Trust Framework via trust scopes.
We plug our Policy Gate in as the **AI_AGENT_TRUST rules engine.**

Policy Model

Three layers, deny-overrides, evaluated at the Skill boundary.

1 hardGates
Required preconditions. Failure → deny.

2 prohibitions
Explicit denials. Match → deny.

3 permissions
Conditional grants. Match → allow or restrict.

default decision: deny

gx-dua-research-policy.yaml

```yaml
profile: gx-dua-research
defaultDecision: deny
evaluationStrategy: deny-overrides

hardGates:
  requireDuaActiveForProtectedSkills: true

prohibitions:
  - id: no-raw-export
    when: { requestedSkill: gx.provider.export.raw.v1 }
    then: { decision: deny }

permissions:
  - id: aggregate-research-active-dua-v1
    when:
      requiredCallerRoles: [Researcher]
      duaStatus: Active
      requestedOutput: aggregate
    then: { decision: allow }
```

Demo Scenarios

Four scenarios you will see on the Dashboard, each producing a machine-readable decision.

S1 + S2

Happy path

- Consumer VC role = Researcher
- DUA status = Active
- Skill = aggregate / statistics

ALLOW

permission: aggregate-research-active-dua-v1

S3

Raw export attempt

- Same VC and DUA as S1
- Skill = raw export

DENY

prohibition: no-raw-export

S4

DUA revoked

- Same VC as S1
- DUA status = Revoked
- Skill = aggregate

DENY

hardGate failed: DUA must be Active

S5

Invalid role

- Consumer VC role = Guest
- DUA status = Active
- Skill = aggregate

DENY

no matching permission (default deny)

DEMO

Live Demo

Vision

AI Agents as first-class Dataspace participants.

When **A2A's interoperability protocol** meets **Gaia-X's Trust Framework**,
AI Agents stop being
API consumers —
and become
Dataspace participants that invoke skills under contracts and policies.

A2A

How to invoke

Discovery, capability description, peer-to-peer skill invocation.

Gaia-X

Under what terms

Verifiable Credentials, Data Usage Agreements, Trust Framework.

This work

Bridging at the Skill boundary

Policy Gate that combines both at every A2A Skill invocation.

We hope you enjoy your time at this Hackathon as much as we did building this! 😀

Acknowledgments

This presentation is based on results obtained from the project, “Research and Development Project of the Enhanced infrastructures for Post-5G Information and Communication Systems” (JPNP20017), commissioned by the New Energy and Industrial Technology Development Organization (NEDO).