Gaia-X Hackathon Report

Results and Lessons from the Gaia-X Hackathon #6

3-4 May 2023

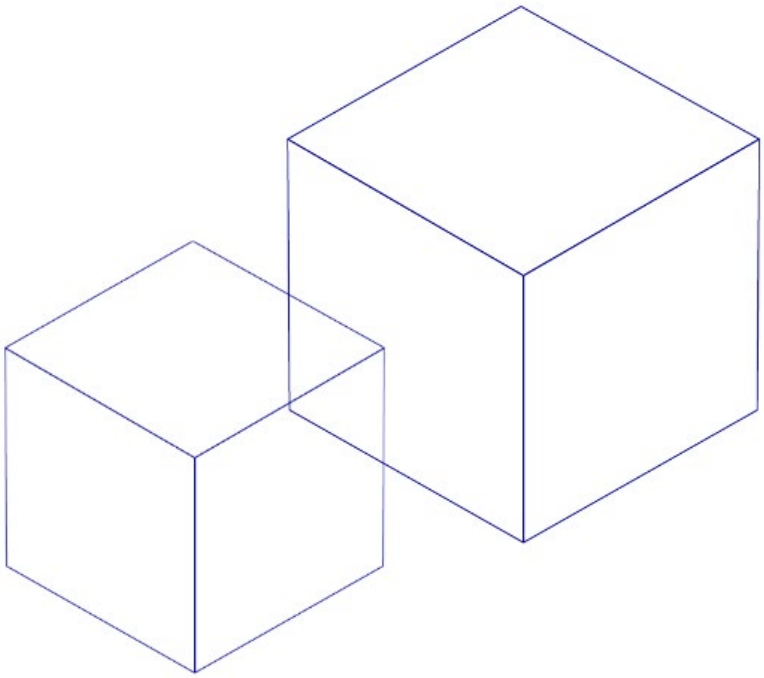

May 2023

Author: Cristina Pauna,

Program Manager for the Open Source Community

Contents

Executive summary

The Gaia-X Hackathon #6 was an on-site event, collocated with the Tech-X conference – the first conference dedicated to the technologies used to make the Gaia-X mission a reality. As in the previous events, it was a hands-on coding event, with teams mixed from different countries, companies and backgrounds, focused on real-life end-to-end scenarios which leverage the Gaia-X Open-Source Software and the Trust Framework.

The implementation of the Trust Framework 22.10 was made available before the hackathon by the Gaia-X Lab team. Besides updating to the latest specs, additional tools are now available to improve the user experience: a wizard to simplify signing the Gaia-X Verifiable Presentations, and a notarization service for the registration number which is used by the Compliance service in the background. All the hackathon teams tried out the new version, with direct support from our team on-site.

This hackathon was the first time we had a competition around it, awarding prizes in cash to the winning teams. Everyone was invested in fulfilling as much as possible of the criteria, with really great results and excellent anchoring in the real-life use of the software. Keep reading to find out who won.

The slides presented during the event can be found on our wiki page.

Sessions

Conditional Trust – ODRL

This session focused on the real-life scenario in the health space, where the regulation of a country (in this case Germany) states that a patient cannot get their test results directly from the laboratory, but it needs to go through a doctor. Initially, the data (the health test results) is held by the lab; as a first step, the patient needs to grant access to the doctor to see the results; then, after the doctor interprets the data, they grant access to the patient to see the results as well.

This flow ensures that both the privacy of the patient and the legal regulation is respected. The principles of SSI were highlighted in this session, and it exemplified how to deal with the complexity of it, where not only the data owner (the lab) is in charge of what happens to the data, but different stakeholders must be taken into consideration as well (the patient and the doctor).

Outcome & take-aways

The team showed a live demo with the access flow between the patient, lab and the doctor. Data was exchanged between the three participants only when the proper access was provided. Under the hood, Trust Services API was used for access control and the rules were described with ODRL.

One feedback from this team was that Self-descriptions for data resources are needed. While some basic information is described in the Trust Framework 22.10, the current implementation of the Gaia-X Compliance Service does not have the option to validate them, and such functionality is required for use-cases like these.

For any questions or follow-ups, the team can be reached through oss-community@list.gaia-x.eu.

Data Product Passport with European Manufacturing & Gaia-X standards

In this session members from seven companies teamed up and worked on implementing a digital product passport (DPP) based on Gaia-X specifications for Trusted Data Transactions, EU DPP specifications and the information available in an Asset Administration Shell (AAS). The use-case focused on calculating the carbon footprint in manufacturing, with an example of producing a steel roller. The carbon footprint of a single product can be calculated by adding the carbon footprint at each step: production of the raw material, transportation, product production.

This information was stored in a Gaia-X compliant Verifiable Presentations (VP) at each manufacturing step, thus making the source of the data transparent and reliable. Each company uses the Gaia-X Clearing House to verify that that all the information required by Gaia-X Compliance is present in the VPs for themselves as participant and for the services they provide.

Outcome & take-aways

The team successfully created the software to implement the described use-case. The implementation was based around an existing dataspace built upon the IDS architecture, with the open-source implementation of the TSG connector with AAS data app support, and IJssel technology.

A live demo was done to show how a user can calculate the total carbon footprint for the end-product by aggregating the VPs generated at each manufacturing step and all was done using Gaia-X compliant VPs. The end-to-end approach of this team paid off and they were voted on **the third place** in the hacking competition. The results can be found at the TSG gitlab (https://gitlab.com/tno-tsg/gaia-x-hackathon). For any questions or follow-ups, the team can be reached via Maarten Kollenstart (maarten.kollenstart@tno.nl).

Decentralized Compliance Service Hacking

This idea for this session rose from the various deployment scenarios of the Gaia-X Compliance service, described in this article. With the Clearing Houses the Compliance Service evolved from a centralized deployment to a distributed one. But further steps are needed in order to reach a decentralized deployment, and this hackathon was the first one in that direction.

The goal of this hacking session was to move from the centralized engine written in TypeScript to a private decentralized model, based on a blockchain infrastructure.

Outcome & take-aways

The team first brainstormed on whether the Registry – which is an essential component used by the Compliance Service – should also be decentralized. For this 2-days exercise the decision was to do so, and the docker images provided by the Gaia-X Lab for the Registry and Compliance Service were successfully deployed on a decentralized cloud infrastructure based on Zero-OS and ThreeFold chain.

While there are still many aspects to clarify until the Compliance Service is fully decentralized, this session showed that an integration between the software provided by the Association and a blockchain infrastructure is already possible from the technical point of view. For any questions or follow-ups, the team can be reached through oss-community@list.gaia-x.eu.

DID as a Service - Hands On

At the core of Gaia-X compliance sits the inconspicuous Self-Description, a file containing claims about participants and services, digitally signed by the issuer of those claims using a Decentralized Identifier. Someone who wants to verify the identity of the issuer would need to check the DID Document integrity, cypher and publick key correctness against predefined policies, privacy on URIs against allow-lists, and rules against specific dataspaces constraints.

While seemingly straight-forward, security best practices dictate that keys should rotate. Matching that with using Decentralized identifiers (DID) – which generally would require static files to be stored on a web server – can create one spiny issue, that his team tackled during the hackathon.

The participants worked together on a DID as a Service approach and explored the potential usage of a key management solution in combination with DID Documents, key rotation and Verifiable Credentials signing.

Outcome & take-aways

The outcome of this session was an up and running solution for creating did:web based DIDs plus the related DID Document, which are resolvable over the Universal Resolver. With this method, a user that wants to verify a VC (which contains a DID URI to identify who signed the VC), can access the DID document and verify that the proof generated with that identity is still valid, even after a new public key was generated.

The code is available here, and for any questions or follow-ups, the team can be reached through oss-community@list.gaia-x.eu or this matrix room.

Gaia-X compliant Trust and IAM framework

This team focused on the integration between Gaia-X and i4Trust frameworks, with focus on identity and access management. The main goal was being able to introduce new participants to a dataspace by using the Gaia-X services as trust anchor.

The onboarding service and its IAM-Framework already had the functionality to validate and persist Gaia-X compliant Self-Descriptions for legal participants in an NGSI-LD compliant broker. During the hackathon, an EBSI-compliant Trusted Issuers Registry based on those Self-Descriptions was implemented. By using this implementation as a participant registry inside a dataspace, the onboarded company and its users can participate inside the dataspace.

Outcome & take-aways

The team has successfully shown a live demo at the end of the 2 days of hacking, where a legal representative of a company – Animal Goods Org – can issue the legal participant SD, can store it in a FIWARE wallet, and from there call the Gaia-X Compliance Service to get the compliance proof.

After these steps, the legal representative can issue NaturalPerson Credentials for themselves and onboard the company in the dataspace, where Animal Goods Org is added as a trusted participant.

A detailed description of this session can be found here. For any questions or follow-ups, the team can be reached via Stefan Wiedemann (stefan.wiedemann@fiware.org) and Tim Smyth (tim.smyth@fiware.org).

Implementing Verifiable Credentials in a real live use case

In this session, the team had three use-cases, each in a different domain: automotive, academic, and manufacturing. Their challenge for the hackathon was to show how the same technology can be used for each of these industries, while having Gaia-X compliant Self-Descriptions at its core.

Outcome & take-aways

One of the nicest things about this session was to point out how technology has evolved in the past – for example going from letters to mobile phones to communicate remotely – and where we are with Gaia-X in that journey. It may be only an emerging technology today, but the potential is limitless.

The team successfully showed in a live demo how to create participant Self-Descriptions, how to join a data ecosystem, publish a service in a catalogue as a provider, and discover that service as a consumer. Not only that, but they went a bit the extra mile to also show that access to the data is denied when using the wrong credentials.

This team was voted on **the first place** in the hacking competition. The code can be found here, and for any questions or follow-ups, the team can be reached at maharshi@smartsensesolutions.com or via this matrix room.

Self description generation via infrastructure discoverability features

This session focused on the scenario of a cloud infrastructure service provider that needs to create compliant SDs for their services. It was a continuation of the work done in Hackathon #4 on using the underlying infrastructure discoverability features to generate Self-Descriptions. The rationale behind it is that while some of the self-description attributes like legal person and company number only exist in small numbers and don't change very often, this is not the case for technical properties of on-demand infrastructure such as clouds and k8s clusters. Clouds tend to evolve more quickly, the k8s clusters even are created and scaled on-demand. Properly describing them requires that the creation of these self-descriptions is automated.

In this hackathon, the goal was to extend the functionality to include not only OpenStack IaaS platform, but also Kubernetes clusters, and in the process, to also extend the vocabulary used to describe infrastructure resources.

Outcome & take-aways

The team had successfully updated the scripts to generate the newest version of Self-Descriptions for the OpenStack infrastructure. The relevant information needed in an SD was also retrieved for the K8S cluster, however more work is needed to convert the yaml files into JSON-LD format, in order to get the Gaia-X compliance stamp. An Apache Airflow Pipeline was created to generate the SDs and automatically call the Gaia-X compliance service to complete the use-case.

The work done by this team also resulted in a recommendation for the Service Characteristics WG, to extend the vocabulary for IaaS services and even go further to specify the ontology for OpenStack and Kubernetes, as this would ensure a uniform way of describing different infrastructures.

This team was voted on **the second place** in the hacking competition. The code is available at https://github.com/SovereignCloudStack/gx-self-description-generator. For any questions or follow-ups, the team can be reached at info@osb-alliance.com.

Feedback from the participants

Most of the respondents to our survey found the hackathon very engaging and valuable, but also very challenging. Given the limited time on-site during the two days, the onboarding sessions were done virtually the week before. This seems to have been appreciated by the participants, as they had time to get familiar with the challenge and set up the dev environment so that they are ready to go on the first day of hacking.

Some participants would have preferred to have the hacking separate from the Tech-X conference, as hacking made it difficult for them to participate in the workshops and talks going on in parallel. However, the overall satisfaction was high, and more than 80% of respondents indicated they would attend again in the future.

Conclusions

There were many new aspects to this instance of the Hackathon: it was the first on-site Hackathon, first time collocated with a bigger conference, first time we ran it as a competition with prizes in cash. The level of contribution and collaboration was higher than ever. People from about 40 different companies and institutions joined, around 170 Self-Descriptions created during the two days – more than double than in our previous hacking events –, and a clear view how Gaia-X is today the solution for the data economy of tomorrow.