

After Brussels: I Came Home With Requirements, Not a Product

A personal Brussels field note on civic tech, sub:stance, participation receipts, and the infrastructure a room needs after the demo is over.

2026.07.08 15:50 DENNIS HEDEGREEN ANALYSE V1.0

[HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/AFTER-BRUSSELS-I-CAME-HOME-WITH-REQUIREMENTS-NOT-A-PRODUCT/](https://hedegreenresearch.com/articles/after-brussels-i-came-home-with-requirements-not-a-product/)

Brussels did not give me a product.

It gave me requirements.

That is less dramatic than a launch and more useful than a victory lap.

I went to Brussels with more projects than a person should probably carry into one room.

The setting was the first Civic Tech Hackathon organised by the European Commission, in Brussels, on 22 and 23 June 2026, connected to the wider Democracy Shield initiative.

P4P was in the bag.

24 Doors was in the bag.

Windowsill was in the bag.

Work Doctrine was in the bag too, though that one travelled less like a tool and more like a sentence I kept finding myself saying to people:

I am tired of the AI conversation being only about AI taking our jobs. It is also taking our hours. The question is who receives those hours.

That sentence travelled better than I expected.

Maybe because it does not require a dashboard to understand. A worker can hear it one way. A union can hear it another way. A policy person can hear it as a measurement problem. A citizen can hear the simpler question underneath it:

if technology saves time, does life receive any of it back?

That was one of the first things Brussels taught me.

The thing I hoped would travel is not always the thing that travels.

The second thing was more personal.

In the early mingling, more than one person read me as some version of the anarchist.

It was not meant coldly. At least that is not how it landed. It was closer to a social shorthand for someone arriving from outside the institutional path, talking about AI, work, power, saved hours, and society without first asking where the safe edge of the conversation was.

That mattered because I did not feel like I had entered the room as a normal professional person.

I have not had a normal job.

I have not spent years learning the ordinary rhythm of offices, committees, workshops, team roles, status games, polite intervention, and how people take space in rooms like that.

So part of the field test was not only whether my projects could survive the room.

It was whether I could.

That showed up in small ways.

After the third presentation, I went outside to smoke. I started talking with someone there, and the conversation was good enough that one cigarette became three. When I came back in, the last presentation was already underway.

I am still annoyed by that.

Not because anyone did anything wrong.

Because I missed something I had come there to see.

But it also told me something honest about the setting. I cannot yet sit in a room with around a hundred people for that long without my attention, body, and ADHD starting to fight the format. I am still learning that part.

Four projects in the room

Our team battle cry said more than we probably meant it to.

We chose Yes, but no.

The YES was loud.

The but no was quiet.

It was funny because it was honest. We were there. We were building. But the room also carried the tired, sideways energy of people trying to make something real inside a format that wanted enthusiasm on schedule.

I thought P4P might become the main path.

It did not.

Within a few minutes, it was clear that the theme and the room were asking for something else. That is not a failure of P4P. It is what happens when an idea meets a real setting. The setting answers back.

24 Doors was closer to the room in another way. It was about access, language, and the first gate of participation. It looked like a hackathon object because it had a surface. You could point at it. You could show it. You could say:

this is a door.

Windowsill mattered as proof that I could build public tools that actually exist outside slides.

But Work Doctrine travelled best in conversation.

Not because it was finished.

Because it named a pressure people could already feel.

That became one of the quiet lessons of the trip. A room does not only test whether a project is good. It tests which part of the project can survive contact with other people.

It also tests the person carrying the project.

I do not naturally take word-power in a room. I do not naturally become the person who says: this is the direction, this is the story, this is what the team should do next.

My instinct is different.

Build the thing.

Show the thing.

Let the thing be inspected.

That can be useful.

It is also not enough.

Brussels made that clear. Prototypes do not speak for themselves. Someone has to translate why the prototype matters while there is still time for the room to understand it.

The interface was not the engine

The challenge, as I understood it, had a simple heart:

citizens should be able to read society better.

That is a good problem.

It is also a hard problem to hold inside a hackathon room, because the words around civic technology can become thick very quickly. Participation. Trust. Democracy. Engagement. Dialogue. Inclusion. AI. Transparency.

All true enough.

All easy to say.

Harder to build.

The room naturally rewarded what rooms like that reward: visible interface, pitch clarity, demo readiness, a name, a screen, a thing people can understand fast.

That is not a moral failure. It is a format pressure.

But it showed me the distinction I had been trying to name.

The team's project was called sub: stance.

It deserves to be named properly.

In two days, the team conceived, built, and presented a participation platform around a conversational interface. As I understand the public framing, the aim was to help people move beyond surface-level concerns and engage more deeply with the issues that mattered to them.

That is real work.

It is especially real work in a hackathon room, under time pressure, with people who have just met each other, different skills, different assumptions, and a deadline that turns every unfinished decision into a presentation problem.

My question was not whether sub: stance was good work.

It was.

My question was about the layer underneath any such surface.

A chatbot performs listening.

A civic instrument must prove what it did with the input.

Those are not the same thing.

The more I listened, the more I felt that the real object was not chat. It was directed listening.

That listening took most of the first day.

I asked some questions. I tried to understand the team direction. I respected the people around the table. They were sharp, patient, funny, and serious in different ways.

The distinction I kept circling was concrete: a conversational surface can help people enter a problem, but the civic engine still has to account for what happens between input and memory.

Question.

Answer.

Generated follow-up.

Answer.

Generated follow-up.

Answer.

Summary.

Confirmation.

Only then pattern.

Only then group meaning.

Only then decision memory.

That is different from a friendly conversational surface. It is less magical and more accountable. It does not ask the user to believe that the system listened. It has to show the route from voice to pattern.

The night prototype

I built a parallel prototype overnight.

That sentence sounds more heroic than it felt.

Mostly it felt like pressure, coffee, confusion, and the ordinary stubbornness of needing to see whether an idea has a motor.

The prototype was parallel to sub:stance, not a rival to it.

Same problem.

Second angle.

It was not the best-looking thing in the room. It was not the cleanest demo. It did not have the obvious pitch advantage of a polished app surface.

Someone else could make the visible surface faster and cleaner than I could in that format.

That is an honest thing to say.

A hackathon is not only a test of ideas. It is a test of presentation speed, team rhythm, shared language, sleep, confidence, and whether the right person can explain the right thing at the right minute.

But it had the part I cared about:

Q -> A -> generated Q -> A -> generated Q -> A

Not chat.

Directed listening.

The point was not that my version was the answer. The point was that the room made the missing requirement visible. If a democratic system asks people for input, it cannot treat the input as content and then vanish into a summary.

It needs a return path.

Participation needs a receipt.

The person should be able to see that their input was received, how it was understood, what pattern it joined or resisted, whether it changed anything, and what happens next.

Without that, participation becomes another polite extraction layer.

Prototype the room with the room

At one point I had a backup idea in my head.

Put a QR code on the screen.

Ask everyone in the room to answer a low-risk question for two minutes.

Use something shared and harmless: the food, the buffet, the service, the room itself.

Then show the pattern back a few minutes later.

Not as a grand civic claim.

As a demonstration of behavior.

Prototype the room with the room.

After coming home, I rebuilt that idea as a small Hackathon Food Pulse demo inside DeepAsk, using synthetic data. That distinction matters. It was not a record of real participants in the room. It was a reconstruction of the loop I wanted to test there.

That phrase stayed with me because it points to something civic technology often forgets. A participation tool should be able to demonstrate its own participation loop on the people who are present. If it cannot show a receipt inside the room, why should anyone trust that it will create one outside the room?

The same thought followed me into the food and drink details.

The table was also infrastructure

The food detail was not only about taste.

The venue had Coca-Cola products around the hackathon, while wine and beer were offered at the closing dinner. I drank Fritz-Kola instead, partly because I did not want the wine, and partly because the place we ate actually had better goods on the menu.

It stayed with me because Europe talks a lot about strengthening its own markets, capacities, and public infrastructure. That does not mean Europe has to get rid of American brands. That would be too simple. But it does mean the small defaults are worth noticing.

If a European civic-tech room is trying to imagine better democratic infrastructure, and the drinks on the table still arrive through the most obvious global default, then the table is part of the lesson too.

Not because Coca-Cola is the enemy.

Because procurement is also a language.

Because the menu is also a map.

At minimum, there could have been two variants, so the room could see what people actually chose. Better still, there could have been one soft drink from each EU country, not as decoration, but as a tiny market map on the table. Then the room would not only talk about European plurality. It could taste it, count it, and notice which choices people made when the default was not already decided for them.

Because a better European market does not only appear in speeches about competitiveness. It appears, or fails to appear, in the boring choices around a room: what gets bought, what gets offered, which local producers are visible, which habits repeat because nobody treated them as design.

That is why the Fritz-Kola mattered to me. It was not a grand political act. It was just the small relief of seeing that another choice was possible on the same table.

The missing role

The other thing Brussels made visible was a missing role.

Hackathon rooms have mentors, organizers, teams, juries, themes, slides, and deadlines.

But a civic-tech room also needs someone moving across the room asking a different kind of question.

What is the motor behind the interface?

What data is being processed?

Does this already exist?

Who owns the idea after the weekend?

What happens to the input?

What are the AI and data risks?

Where do teams overlap?

What would a participant receive back?

Call that person a floor reader.

Or a civic-tech preflight steward.

The title matters less than the function. Someone has to read the room before the room turns into demos.

I am saying this partly because I think I needed that role too.

Not as a special exemption.

As a more honest fit.

I am not trained in professional group work. But Brussels showed me that I can still contribute if the role is framed correctly: listen hard, find the motor, test the pressure, build an artefact, and return it to the room before the shape disappears.

The record was hard to find

Afterwards, I went looking for the public record of something I had lived through.

That was stranger than I expected.

I could find my own traces. I could find private notes. I could find messages, captures, names, and the small afterlife of contact threads.

And there is a public trace. A participant's LinkedIn post names the event, the team, substance, the organisers, and the people involved.

That matters.

But it also sharpens the problem.

The most complete public trace I could work from lives in one person's feed on one closed platform. It is not an institutional memory surface. It is not a durable public archive. It is not something a citizen can reliably find later without already knowing the right name, account, or platform path.

A European Commission civic-tech event whose clearest public memory lives in a participant's LinkedIn post is not a failure of that participant.

It is a sign of the same missing layer.

That matters because it is the same problem in another form.

The room happened.

But where did the room go?

Who can return to it?

What did it change?

Which ideas survived?

Which inputs became decisions?

Which prototype became a continuation?

Which concern was answered?

Which disagreement was preserved?

The searchable web did not answer that for me.

And that is not only a documentation problem.

It is a civic infrastructure problem.

The people were real

That is also why I do not want this to read as if the room was only a system failure or a format problem.

The people were real.

The team was real.

The days mattered because people gave them attention, energy, humor, patience, and their own angle on the problem. I am grateful for that.

So thank you to Alin Gabriel Gramescu, Cruz G., Katja Hermes, David Canat, Hendrik Nahr, and Duarte Silva for the days in the room, the conversations, and the team effort.

I am naming them with permission.

That gratitude is not the same thing as endorsement.

It is simpler than that.

It is a way of saying that the experience mattered because people were actually there.

Requirements, not product

So no, Brussels did not give me a finished product.

It gave me requirements.

Voice.

Pattern.

Receipt.

Memory.

DeepAsk is one name for the first layer: listen one person deeper.

Mind Hive is one name for the second: make group meaning readable without pretending everyone agrees.

A participation receipt is the third: show people what happened to their input.

Lineage is the fourth: let the conversation survive contact with time.

If civic technology wants people to participate, it has to do more than collect their voices.

It has to remember what it did with them.

RELATION MEMORY

FOLLOWS [Hedegreen Research Is Going to Brussels](#) · Returns after the Brussels field test that the earlier article announced.

NEARBY [I Shipped 24 Doors Before the Hackathon](#) · Shares the Brussels civic-tech room and the boundary that the event was not validation.

NEARBY [The Feed Was Not the Food](#) · Connects the Brussels lesson to the broader argument that conversation needs memory and return paths.

NEARBY [The Website Needs Memory Now](#) · Shares the claim that public work needs relation, memory, and traceable return.

CANONICAL URL: [HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/AFTER-BRUSSELS-I-CAME-HOME-WITH-REQUIREMENTS-NOT-A-PRODUCT/](https://hedegreenresearch.com/articles/after-brussels-i-came-home-with-requirements-not-a-product/)

PDF VIEW GENERATED: 2026-07-11T15:35:45.704Z