


Tiles: Own your AI with open models and decentralized protocols

Ankesh Bharti • 2026-09-28 • English

Full presentation and discussion

Edited transcript of the full presentation and Q&A. Revised on 29 September 2026 against the Soniox transcript, Zoom captions, and presentation material; the Chinese version was checked paragraph by paragraph against the English. Repetitions and broken sentences have been edited for reading, while examples, qualifications, and the speaker's claims are retained. Timestamps begin at 00:03:05.2 in the source recording. Setup and waiting are omitted. Selected unclear audio excerpts were retranscribed with Soniox. Unresolved wording is marked. This is not a word-by-word manual listening audit.

Introduction



The slide features the Tiles logo (a stylized 'T' made of geometric shapes) in the top left. The title 'Tiles: Own your AI with open models and decentralized protocols' is centered. Below the title, it says 'Ankesh Bharti @feynon, Founder & CEO, Tiles Privacy'. To the right is an image of a laptop displaying the Tiles logo. The bottom left has the 'INDIA FOSS 2026' logo and the date 'September 2026, Bengaluru'. The bottom right has the Tiles logo and '© 2026 Tiles Privacy'. The slide number '1 of 30' is in the top right.

Tiles: Own your AI with open models and decentralized protocols

Ankesh Bharti @feynon
Founder & CEO, Tiles Privacy

INDIA FOSS 2026 September 2026, Bengaluru

© 2026 Tiles Privacy

1 of 30

Presentation frame • 00:00:24

00:00:00 • Ankesh Bharti

Thanks for inviting me to speak. I'm Ankesh, founder and CEO of Tiles Privacy. Today I want to talk about what it means to own your AI: not just the model, but also the data, identity, and infrastructure around it.

00:00:28 • Ankesh Bharti

And I think two things make that possible now: recent developments in open models, and adoption of decentralized protocols.



- CEO & Founder, Tiles Privacy
- Co-founder User & Agents community
- Technologist based in Pune
- Ambition is to build a decentralized stack that expands user agency in our digital spaces

Presentation frame • 00:01:04

00:00:38 • Ankesh Bharti

A little more about me: apart from founding Tiles Privacy, I am a co-founder and steward of User & Agents, a community where we discuss user-agent systems that give people more agency, control, and choice. Tiles was born out of those discussions. I like to call myself a technologist, and I am based in Pune, India.

00:01:06 • Ankesh Bharti

My ambition is to build a decentralized stack that expands user agency in our digital spaces.

We need to reclaim
ownership of our digital lives
in the age of AI through
open models, and
decentralized protocols

Presentation frame • 00:01:24

00:01:12 • Ankesh Bharti

The thesis of this talk is pretty simple: AI is becoming part of how we think, work, research, and communicate. And as that happens, whoever controls the AI layer starts controlling a surprisingly large part of our digital lives. I think we now have the technology to build that relationship differently: run intelligence locally, and use open protocols where we need a network.

We all use AI agents now




Presentation frame • 00:01:44

00:01:40 • Ankesh Bharti

Most of us in this chat probably already use several of these AI agents. They are becoming much more than chatbots: they read code, browse the web, operate tools, work with files, handle everyday tasks, and increasingly act on our behalf. Naturally, we give them more and more context about ourselves.

I like ChatGPT



Presentation frame • 00:02:24

00:02:05 • Ankesh Bharti

And I actually like ChatGPT. What keeps me using it is not just the model. It's the product around the model, the consistency, the memory, and being able to pick up the same context across different surfaces. It starts to feel less like using a chatbot and more like using software that actually knows you. But the better that gets, the more one-sided the relationship becomes.



Presentation frame • 00:02:44

00:02:31 • Ankesh Bharti

The system knows more and more about me, but all of that context ultimately lives inside an account, an infrastructure I do not control. If I leave the platform, a growing part of that accumulated digital memory stays behind. I like the relationship. I just do not like the dependency. My conclusion was not that I should stop using AI. I need healthier boundaries. I want to keep the same convenience without depending on blind trust.



Presentation frame • 00:03:04


00:02:58 • Ankesh Bharti

I want my identity and context to stay with me across every app I use, with trust that I can verify computationally. I still want the relationship; I just want it to be on my terms.

00:03:11 • Ankesh Bharti

This became my preference: private and local. And the good news is, today's local AI tools are already technically excellent. They're naturally great at private, personal use.

Frontier is open-source today




© 2026 Tiles Privacy

Presentation frame • 00:03:24

00:03:23 • Ankesh Bharti

Frontier AI is also no longer locked behind proprietary APIs. For many practical workloads, frontier-level capabilities are already available in open weight models, such as GLM series, as shown in this chart by Artificial Analysis.

And there is compute to run them locally



© 2026 Tiles Privacy


Presentation frame • 00:03:44

00:03:39 • Ankesh Bharti

The compute available in consumer devices now makes it possible to run even frontier models locally. Technologies such as UltraFusion combine two Apple silicon Max chips into a single Ultra chip, enabling medium-sized local models to run on-device. We are also finding smarter ways to run larger models at the edge. Apple's upcoming Siri in iOS 27 is powered by Foundation Model 3, a distilled version of Gemini, and runs a 20-billion-parameter mixture-of-experts model directly on the iPhone.

20B models are now pocket-sized

11 of 30




Presentation frame • 00:04:24

00:04:14 • Ankesh Bharti

It loads expert weights from fast flash storage instead of keeping the entire model in memory, and by loading them once per prompt rather than once per token, it largely compensates for the higher latency of flash storage.

But collaboration breaks the private deal

12 of 30




© 2026 Tiles Privacy

Presentation frame • 00:04:44

00:04:31 • Ankesh Bharti

But local inference alone does not give us a complete product. As soon as we want identity, multiple devices, sync, or collaboration, we often recreate exactly the architecture we were trying to escape. So the model may be local, while the product relationship becomes centralized again.

Users and keys at the center



Presentation frame • 00:05:04

00:04:51 • Ankesh Bharti

What if we put users and their keys at the center and built outwards? The keys belong to users on their own devices, not to the provider. They establish users' identities and what they control. Then come their devices, models, and everything else. Ownership should not stop at the model layer. As an AI-use disclosure, the project I am about to introduce has manually reviewed code on its main branch, with core systems largely handwritten.


AI disclosure

Presentation frame • 00:06:04

00:05:40 • Ankesh Bharti

As I was saying, our main branch contains manually reviewed code, with core systems mostly handwritten. Our canary branch is where we experiment more aggressively with AI-generated code, primarily for client-side features and UI/UX work above the API layer. This lets us prototype quickly, test ideas with users and the team, and learn what works. Useful features gradually enter the main branch through small, carefully reviewed PRs, rather than wholesale merges of experimental code.

Tiles and the live demo




Presentation frame • 00:06:44

00:06:15 • Ankesh Bharti

To address this problem, we built Tiles, a local-first collaborative AI assistant. Here, local-first means something specific: the product should keep working even if our team stops working on it. Your data, identity, and core experience do not depend on Tiles continuing to exist as a company or service. We want to provide a familiar AI experience while changing the underlying architecture, so the user remains at the center of control. It's software built to outlive us.

App and CLI



© 2026 Tiles Privacy

Presentation frame • 00:07:04

00:06:50 • Ankesh Bharti

Before I go any further, I'd like to show a brief demo of what the app looks like.

00:07:36 • Ankesh Bharti


This is a prerecorded demo. I can give a live demo at the end if anyone is curious. Here I am using the app to find the dates for Mozilla Festival 2026 and save them in my calendar.

00:07:58 • Ankesh Bharti

This will use plugins. Actually, I think a live demo would be better, so I will share my screen again.

00:08:37 • Ankesh Bharti

I'll copy the prompt here. Let's keep the live demo casual. It looks much like ChatGPT or another chat app. I will send this prompt: "Find me dates for the festival conference," and then save them in my calendar. It is using an on-device model. There are two main parts to what we have built: the AI part and the peer-to-peer, decentralized networking part.




Presentation frame • 00:09:24

00:09:12 • Ankesh Bharti

First, how is the AI part different from something like Ollama? With Tiles, I took an opinionated approach: support a particular model and handle its configuration at the llama.cpp inference layer, the Pi harness layer, and in the prompt templates. Currently, Gemma 4 is the default model series, configured to work out of the box without setup on your end.

00:09:39 • Ankesh Bharti


You can also change the model by modifying the from field in the model file. It currently supports any Hugging Face endpoint, and we will support local paths and other options later. That is our opinionated model implementation. We also have plugins following the Agent Plugins standard.



Presentation frame • 00:10:24

00:10:05 • Ankesh Bharti

It is the same standard used by ChatGPT and Cursor, so compatible plugins should work here too. For example, Exa is a web-search plugin combining MCP and a skill. You can add other plugins by specifying a ZIP path or URL. The calendar automation plugin is a CLI with a skill explaining how to use it. Returning to the prompt, you can see that the model reasoned through it, first used Exa to find the dates, and then called the calendar tool.



Presentation frame • 00:10:59

00:10:56 • Ankesh Bharti

It worked and saved the event in my calendar. You can also share chats. Sharing saves the chat to AT Protocol. It doesn't seem to be working correctly right now... Okay, this time it works. We are not storing the shared chat in a database of our own.

00:11:26 • Ankesh Bharti

It is stored on the protocol, which I will explain later in the slides. You can see the source it is fetched from. On AT Protocol, apps share the same database, owned by the user. Users can self-host it or have a provider host it for them, while retaining ownership of their data across apps. This is the public record of that chat. Going back, these are other chats I shared earlier.

00:12:00 • Ankesh Bharti

You can share privately as well. Right now, data on the protocol is public, so we put the key in the URL and store the chat transcript as an encrypted blob. On the PDS, you see encrypted data. Even though the record is public, its contents cannot be read without the key.

00:12:25 • Ankesh Bharti

The key is in the URL, so you need to keep that URL secure. We haven't implemented all the decentralized networking features in the UI yet. Most supported features are available through the CLI, and we are gradually bringing them into the graphical interface.

00:12:52 • Ankesh Bharti

You can sync chats across devices without sending the chats to the cloud. We use client-to-client data transfer and peer-to-peer syncing. We have a relay, which I will explain in the slides. You can sync chats across your devices using a local account based on a decentralized identifier, which I will also discuss later.


00:13:20 • Ankesh Bharti

That is the main overview. By sharing this token, you can use a local model from a remote device. You can also keep a Mac awake for a specified duration if you want it to serve as an always-on inference device. Here is the AT Protocol account I am logged into. That is mostly what we currently have in the app.

00:13:43 • Ankesh Bharti

Now I will return to the slides and explain how these things work. I wanted to show the app first so you could picture what it does. Let me share the slides again.

Technical architecture



Presentation frame • 00:14:24

00:14:18 • Ankesh Bharti

The technical architecture has four modules. The core is written in Rust and includes the plugin system, model-file parser, encrypted SQLite, and local filesystem. Peer-to-peer networking is based on iroh. Zero-trust authentication and authorization use DID and UCAN. The core also includes keychain support, public-key infrastructure, syncing, and remote inference. The second module is inference, which communicates with the core over HTTP. It contains the llama.cpp server and an Open Responses API translation layer. This is the standard OpenAI currently uses, but adoption across the industry is limited because most apps still use the completions API.

00:15:13 • Ankesh Bharti

In the previous generation of Tiles, GPT-OSS was the default model. We implemented the Open Responses API to make it work at its best, and we still use that. On the client side, we have a CLI and UI. The web UI is based on Tilekit APIs: REST endpoints exposing all of Tiles' features.

Developers can embed those capabilities in their apps, which I will discuss later. We also embed the Pi harness as a Bun binary. That is the architecture. Now, how do user-owned keys work with a Tiles account? DID stands for decentralized identifiers and provides the identity primitive. UCAN stands for user-controlled authorization network and provides the authorization primitive. Together, they make up your Tiles account, generated locally and secured through your device's keychain, rather than created and controlled on a centralized Tiles server.

DID, UCAN and user identity

1. User owned keys with Tiles Account

17 of 30

DID
Decentralized Identifiers


 **UCAN**
User Controlled Authorization Network

 © 2026 Tiles Privacy

Presentation frame • 00:16:04

DID is who you are

18 of 30



© 2026 Tiles Privacy


Presentation frame • 00:16:44

00:16:21 • Ankesh Bharti

Basically, instead of using SSO or OAuth, which is centralized in nature, you have DID and UCANs, which work peer-to-peer. So DID is basically who you are. It's a portable identifier controlled by your keys. Services can verify you, but the services themselves do not have to be the source of your identity.

UCAN is what you can do

19 of 30




© 2026 Tiles Privacy

Presentation frame • 00:17:04

00:16:46 • Ankesh Bharti

If DID answers “Who are you?”, UCAN answers “What are you allowed to do?” Tiles lets you grant scoped authorization for syncing chats across devices and with other users. UCAN uses capability-based authorization rather than the access control lists, or ACLs, common in centralized systems. Instead of a central server checking a permissions list, you present a signed capability granting specific authority, potentially with an expiry.

Local, peer-to-peer and open social networks




Presentation frame • 00:17:44

00:17:14 • Ankesh Bharti

Instead of giving another device a permanent password, I can grant exactly the authority it needs. That's a much cleaner fit for local-first and peer-to-peer systems. Second is local-first AI with Tiles network, which works across local, online, and social layers. You get local devices, ambient compute, private P2P networking, and an open social layer without requiring a central Tiles server to sit in the middle. The basic rule is simple: keep ownership close to the user, and introduce network only where the network adds value.

Iroh based P2P sync through public relays



© 2026 Tiles Privacy

Presentation frame • 00:18:24

00:17:47 • Ankesh Bharti

You have the devices you actually use, your phone, laptop, and what I'd call ambient AI compute. That could be a home lab, a workstation in the office, or any computer you control that can sit in the background and provide AI inference when you need it. These devices can talk directly over your local network, so this can work without the internet at all. To connect across networks, Tiles uses end-to-end encrypted peer-to-peer networking with public relays built with iroh.

00:18:15 • Ankesh Bharti

The relays are used to discover IP addresses, while the actual data transfer happens directly between clients using UDP packets.

00:18:25 • Ankesh Bharti

The data never reaches a relay; the relay only sees the IP address, allowing devices to be discovered. [Editor's note: This is the speaker's description. In iroh generally, relays can forward end-to-end encrypted traffic during connection establishment or when a direct connection is unavailable. They cannot read its plaintext. This note does not establish the relay policy of this Tiles build.] These relays can also work on your local Wi-Fi or local network. With your own local network, you don't need the hosted relay. For social and collaborative features, we use AT Protocol and the user's own personal data server.

Collaborative features on AT Protocol

OPEN WEB



OPEN SOCIAL WEB



Presentation frame • 00:19:04

00:18:54 • Ankesh Bharti

The analogy here is the open web. Your website works across different browsers because the browser does not own the underlying web. AT Protocol brings a similar idea to social layer. Your data lives in your own repository, which you can think of a bit like a GitHub repository for your digital life. Different applications can read from and write to that same repository, rather than each app keeping your identity and data trapped in its own database.

3. Opinionated AI stack

23 of 30



© 2026 Tiles Privacy

Presentation frame • 00:19:44

00:19:21 • Ankesh Bharti

Your identity and content can work across applications, so switching apps does not mean starting your digital life over. Third is our opinionated stack, with Gemma 4 as the default series. We handle model-specific tuning across llama.cpp, the Pi harness, and quantizations, so you don't have to configure everything yourself. For extensibility, we support the Agent Plugins standard and compatible plugins from clients such as Cursor or ChatGPT.

Server

```
GET    /server/start
GET    /server/stop
GET    /server/ping
```

Agent

```
GET    /agent/start
GET    /agent/state
POST   /agent/prompt
GET    /agent/end_session
GET    /agent/reload
```

Session

```
POST   /session/new
GET    /session/list
GET    /session/{session_id}/chats
POST   /session/chat
```

Account

```
POST   /account/create
GET    /account/status
```

ATproto

```
POST   /atproto/login
POST   /atproto/logout
GET    /atproto/status
POST   /atproto/share-session/{session_id}
```

Presentation frame • 00:20:04

00:19:54 • Ankesh Bharti

Finally, you can embed Tiles in your own product with Tilekit. It is a set of REST API endpoints for developers who want the underlying runtime and capabilities in their own applications. They can build interfaces and products on the same infrastructure. It is similar to how Codex lets you embed its harness through Codex App Server. It's a similar design.

00:20:27 • Ankesh Bharti

Before I wrap up, there is one more thing I would like to share. We will announce a major partnership with Mozilla Foundation next month, and we are looking for organizations and customers to pilot what we are building and become design partners.



26 of 30

We're announcing a major partnership with the Mozilla Foundation soon and are looking for organizations to help pilot and test what we're building.

 © 2026 Tiles Privacy

Presentation frame • 00:20:44

Advisors

28 of 30



Boris Mann
[@bmann.ca](https://twitter.com/bmann.ca)

COO, Ink & Switch
Project Lead, AT Community Fund
Co-founder & CEO, Fission Labs



Dietrich Ayala
[@autonome](https://twitter.com/autonome)

Co-founder, Users & Agents
Founder, Web Transitions
Ecosystem Lead and Advisor,
Protocol Labs
Developer Advocate, Mozilla



Gordon Brander
[@gordon](https://twitter.com/gordon)

Research Fellow, Future of Life
Foundation
Research Fellow, Cosmos Institute
Founder and CEO, Subconscious
Web Platform Designer, Google
Design Engineer, Mozilla

 © 2026 Tiles Privacy

Presentation frame • 00:21:04

Over the past year, we have grown into a team of five. We will work full-time starting next month, once we receive the Mozilla Foundation grant. I am grateful to Boris Mann, Dietrich Ayala, and Gordon Brander for advice and guidance from their experience in decentralized systems, browsers, and the open web. Tiles is a product of the open-source community. User & Agents and OSDC, our university-based open-source community, have helped a great deal.



Partners

29 of 30



Sol PBC
solpbc.org



Jeremie Miller @jeremie.com
Founder, Sol PBC · Creator, Jabber/XMPP · Founding board member, Bluesky

Current sponsors

7

Past sponsors

17



Boris Mann
@bmann.ca



Luke Hubbard
@lukeinth



Dietrich Ayala
@autonome



Ethan Frey
@ethanfrey



Hugo Duprez
@HugoDuprez



Utkarsh Saxena
@saxenauts



Dev Doshi
@devdoshi



Xi Zhang
@aeftm



Josiah Witt
@josiahwitt.com



Yusuf Olokoba
@OlokobaYusuf



Curran Dwyer
@currandwyer



Rashid Aziz
@razberry



Gavin Owens
@goblinboats



Alex Komoroske
@komorama



Chad Kohalyk
@chadkoh



Anish Lakhwara
@chickensoupwithrice



Nick Perez
@nickthesick.com



Steve Manuel
@nilalice



Shobhit Srivastava
@SinisterLight



Incognitothief
@whereditgo.diamonds



AS Anonymous sponsor

© 2026 Tiles Privacy

Presentation frame • 00:21:44

00:21:21 • Ankesh Bharti

Darkshapes designed our logo and helped turn our ideas into Tiles' foundational visual identity. Special thanks to Solstone and its founder, Jeremie Miller, for being a close partner and sponsor, and to everyone else who sponsored Tiles and sustained it over the past year.

Talk recap



Presentation frame • 00:22:04

00:21:46 • Ankesh Bharti

To recap, I love ChatGPT, but I do not want to compromise on convenience or control. Local AI and open protocols let us keep that convenience while making our data, identity, and memory portable and truly ours. I'd love for the audience to try Tiles on Mac or Linux and tell me what you think. If you'd like to contribute or chat about local-first AI, please join our new Discord server through this QR. Or for anything else, mail me at this address.

00:22:14 • Ankesh Bharti

So that was my talk. Yeah. If you have questions about Tiles, want to see the CLI, or want to discuss trade-offs or other technical details, I would be happy to take those questions now.

Discussion and Q&A

00:22:46 • Ankesh Bharti (reading a question and answering)

“When agents on two devices write their own memories, which may be duplicated or contradictory, how do you plan to merge them?” Right now, chat syncing is controlled by users, and we only share chat transcripts, or conversation threads. You cannot yet resume a chat on another device, although you can sync it there.

00:23:17 • Ankesh Bharti

We haven't yet worked through those merging issues. This is still an early project, currently in alpha. The second iteration of Tiles has only been under development for six or seven months. We will probably consider this over the coming months.

00:23:49 • Ankesh Bharti

One interesting thing we are working on is making every chat a sandbox. Currently, sharing a chat only shares the transcript. Soon, each chat will be a microVM sandbox based on `microsandbox.dev`. Suppose a researcher installs Python packages and shares a chat with a friend. The friend would not need to install those packages again, because they would share the whole computing state. It would still be lightweight because it is based on a microVM.

00:24:41 • Ankesh Bharti

I am excited about that. With microVMs, you can take snapshots, pause, and resume. I think we will have a way for agents to collaborate in each other's computing environments, like Grok bot [name unclear], while also addressing merging issues.

00:25:07 • Ankesh Bharti

I see a link: `centaur.run`, "Multiplayer self-hosted secure agents for Slack." Okay.

00:25:16 • Ankesh Bharti

We are also building group chats over the coming six months, so this seems similar.

00:25:27 • Ankesh Bharti

My goal is to have my whole digital life in a form where I control the account, model, data, and everything around them.

Agent collaboration and shared environments

00:25:54 • Audience (chat question)

Do you plan to create some structure to enable direct collaboration between agents, for example using A2A protocol to let them talk to each other, or letting them access a shared kanban?

00:25:58 • Ankesh Bharti

We plan to support that, although we have not chosen a protocol yet. We can remain protocol-agnostic. The plan is for each chat to be a microVM. You could invite an agent into a conversation thread and collaborate with it, or invite a friend and collaborate with them.

00:26:26 • Ankesh Bharti

On the shared kanban question, I recently saw an LM Studio product that seems to integrate a canvas agents can access. I have not tried it, but it looks similar to this idea. LM Studio released it in Bionic [app name unclear], their new local AI app.

00:26:51 • Ankesh Bharti

For the first twelve to eighteen months, the product's users will mainly be students. We will begin by working closely with institutions in India to identify their use cases.

Local models and hardware limits

00:27:09 • Audience

Hi, thanks for the presentation. I am curious about the computing side of locally hosted language models. I have tried several on my MacBook M4 with 24 GB of memory, probably including Gemma. Even after a tiny task, my computer gets very hot. How could this become more feasible on a mid-range computer?

00:27:53 • Ankesh Bharti

There are smaller models now, such as Gemma E2B. The key question is whether the local model is good at tool calling. If it is, I think it can handle most use cases. As for thermal throttling, these are still early days. I think the recent iPhone 19, or iPhone 18, has twice the NPU capacity, and so does [chip/model name unclear] Pro Max, helping with prefill. [The speaker is uncertain about the device names; the wording is retained rather than treated as a verified specification.]

00:28:31 • Ankesh Bharti

We are all tinkerers here. I think hardware may need another year or two to catch up, but we will eventually get there. Now is a good time to experiment, participate, and influence the direction this takes.

00:28:53 • Audience

Okay. Thank you.

00:28:59 • Ankesh Bharti

Frankly, my own device also gets hot when I use local models. I don't worry too much, because I think better hardware will eventually address it. I even hope the AI bubble bursts so we get cheaper hardware, as happened during the internet bust.

00:29:29 • Audience

I personally think it depends on the scenario. Parts of a deterministic task can be handled by scripts. You can redesign the workflow and delegate some capabilities to a locally hosted language model. That seems very useful.

00:30:07 • Ankesh Bharti

Integrating DID and UCAN for decentralized authorization and authentication should also help with protocols such as A2A. Right now, even when agents communicate through A2A, a centralized server somewhere still mediates the interaction. Neither the users nor the agents have that freedom or agency.

Closing remarks

00:31:23 • Ankesh Bharti

It has been nice talking to you all. I hope to meet this community in person one day, perhaps at an event in China.

00:31:55 • Ankesh Bharti

I also want to apologize for being a little tired. Last week was hectic: I gave a talk and had a booth at IndiaFOSS, India's biggest digital commons and open-source conference. I was tired but did my best here. Feel free to contact me on social media. I will put my personal website in the chat; its footer has all my social links.

00:32:28 • Ankesh Bharti

I will also share the link to the product we are building so you can try it. That is everything from me. Thank you all for having me.