

State of Synapse 2025

What we're doing, where we're at and the future

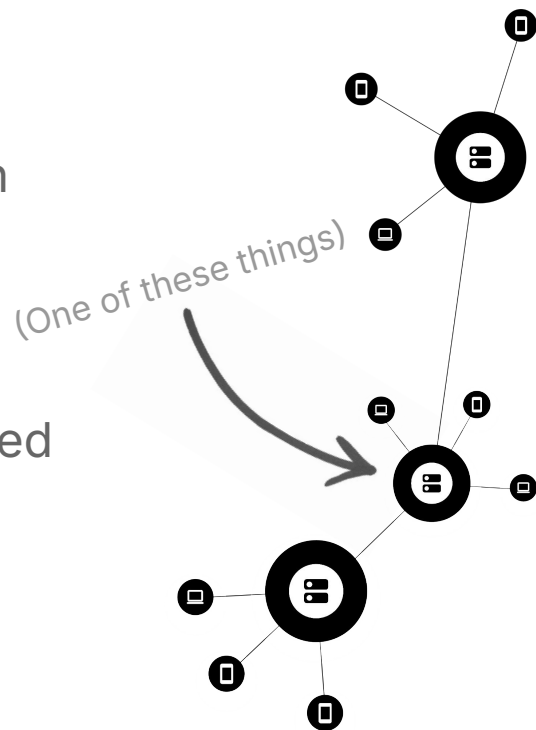
Erik Johnston

Staff Engineer, Element
@erikj:jki.re


What even is Synapse?


- A Matrix *homeserver* implementation, written in python
- It's the oldest, most featureful, most battle-tested and most widely used in the ecosystem


~10,000 active deployments (that have opted into phone home stats!)





The Evolution of Synapse...

Original prototype
back in 2014 



Reference implementation
demonstrating Matrix 




More and more features added
to support the ecosystem 



...

Synapse's goal is to
provide a **great experience**
out of the box

To provide a great UX...



For **users** this means:

- Being **reliable**
- Be **fast** and not slow and janky
- Having all the expected **features**

For **server operators** this means:

- Making it **easy to install** and get started with
- **Easy to maintain** and keep updated



The use cases we're focusing on now...

1. Fully supported out of the box:

Small and medium sized messaging deployments



2. Should work with some tinkering:

Larger deployments

Non-messaging use cases

3. Requires significant extra tooling and expertise:

Massive deployments that require huge scale

How are we doing?

Things going well:

- Easy to install
- Easy to keep updated
- Supporting clients with new features
- Reliability

Needs improvement:

- Resource consumption
- Performance

What we've been working on...

Native sliding sync

Used to get rooms on launch and getting updates

On the previous API this can take a while, initial launch on Matthew's account takes over 15 minutes...

With sliding sync that takes
less than a second

A **1000x speed up!** 

What we've been working on...

Rust support 🦀

We are *not* rewriting Synapse in Rust...

... but we are using it as a tool to optimize hot paths

We use PyO3 and integration has been surprisingly straight forwards!

Time to send a message in
a room with 5k local users:

~3s to ~500ms

A **6x speed up!** 🚀

What we've been working on...

... and everything else

Migration to OAuth2.0 with MAS...

... to support state of the art security like... MFA

Authenticated media...

... protects servers from their users treating them as free hosting

And performance improvements...

And bug fixes, and security updates, and etc etc etc...

What we want to do next!

Continue supporting new Matrix features

Better threading, etc

Further algorithmic changes for performance

API changes, state res, etc

Better Rust support

async/await, direct database access, etc



What we want to do next!

And continuing with all the day-to-day maintenance, bug fixes, security updates, performance tweaks, etc



Synapse provides a **great experience** out of the box

A big thank you to the community for
the support over the years ♥

Thank you!