

Large scale multi-tenancy hosting for TI-Messenger.

10.2025

Patrick Maier // @pmaier:element.io

Agenda.

1. TI-Messenger and the small hosts challenge
2. Synapse Pro for small hosts
3. Synapse Pro and ESS
4. What customers can expect

TI-Messenger and the small hosts challenge.

What is TI-Messenger?

- The secure & sovereign messenger for German healthcare
- Based on Matrix
- Specified and driven by **Gematik**
- Implemented and provided by industry vendors
- Two main use cases / products
 - TI-M ePA
 - For insured people (provided by their health insurance company)
 - Typically medium - very large size deployments
 - **TI-M Pro (this is the one to look at today!)**
 - For professionals (doctors, hospitals, insurance clerks, etc.)
 - Typically (many!) small - medium size deployments
- Most vendors chose **Synapse** at the heart of their TI-M products



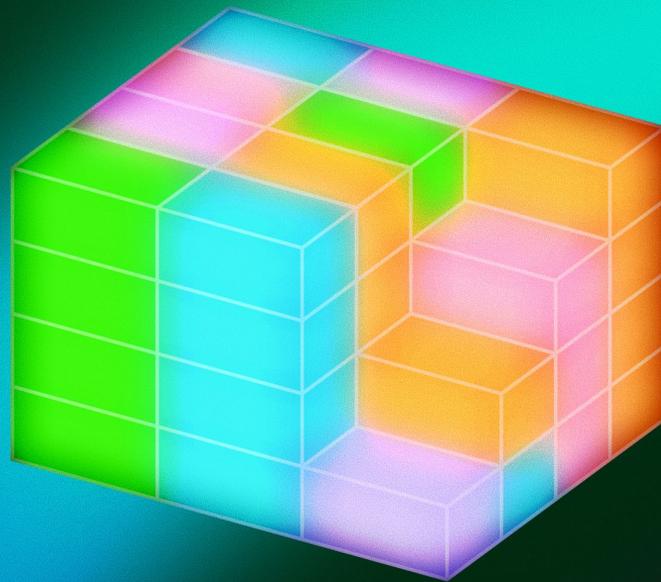
- Synapse is targeted at **community use cases** (small-/mid-scale)
- Synapse supports workers for scaleouts but
 - Config requires a lot of (manual) trial-and-error to find appropriate values
 - The number of workers needs to be set manually and adjustments require a restart of the entire system
 - Some workers cannot be scaled horizontally (exclusive database access)
- **Synapse has a relatively high resource footprint at scale**
 - Running large hosts is inefficient
 - **Running many small hosts is inefficient**
- Synapse can't be re-deployed without downtime
 - Some workers use static sharding for scaling (e.g., for persisting new messages in a room)
 - The number of workers needs to be configured statically
 - Config updates require downtime

- **Synapse Pro** is Element's commercial implementation of Synapse targeted at **professional use**
- Synapse Pro ships as part of **ESS Pro** which is being delivered as a **helm chart**
- Its mission is to **address the limitations of Synapse FOSS** so that it can
 - **Scale** to massive size effortlessly and automatically
 - Save **resources** and reduce operational **cost**
 - Ease operations and improve operational **stability**
- Synapse Pro is available in **two flavours**
 - Synapse Pro for **large hosts** (> 100 users up into the millions)
 - **Synapse Pro for small hosts (< 100 users)**

Synapse **PRO** for small hosts

Synapse PRO for small hosts

- Synapse Pro introduces **multi-tenancy** for better small host efficiency
- Main goal is to **reduce memory consumption** per tenant and **personnel cost for operations**
- Basic idea
 - Run multiple distinct tenants within one Python process ("shard")
 - Fully-featured Matrix servers
 - Segregate tenant data on db schema level
 - Tenant management API
 - Kubernetes controller to integrate with CD-tooling and GitOps processes

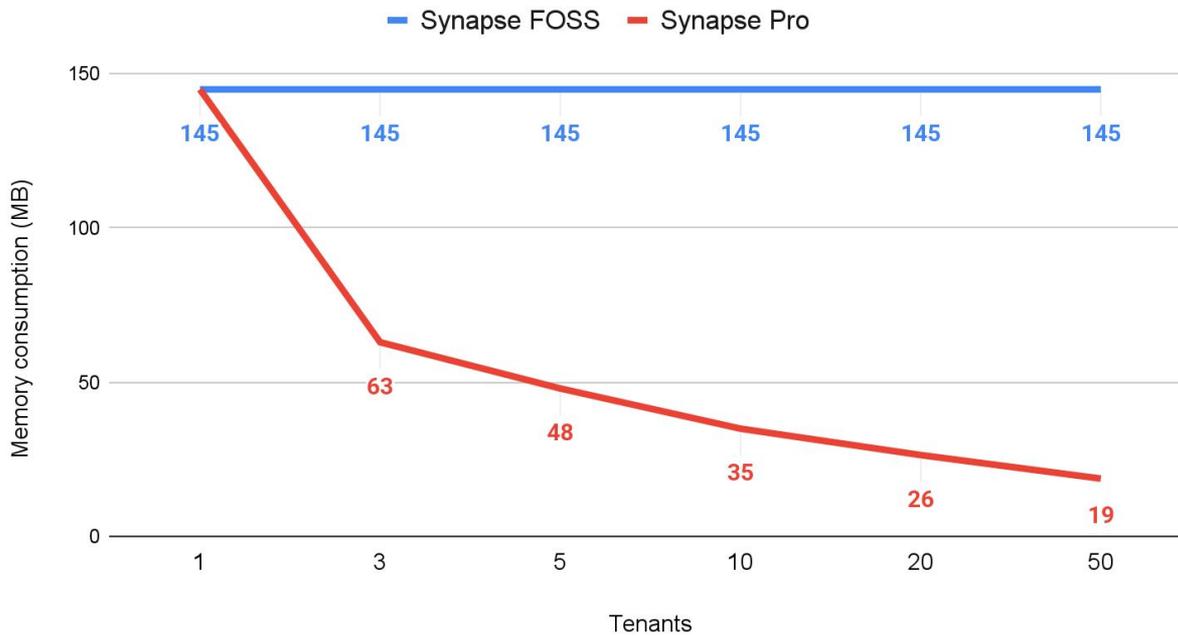


A journey:

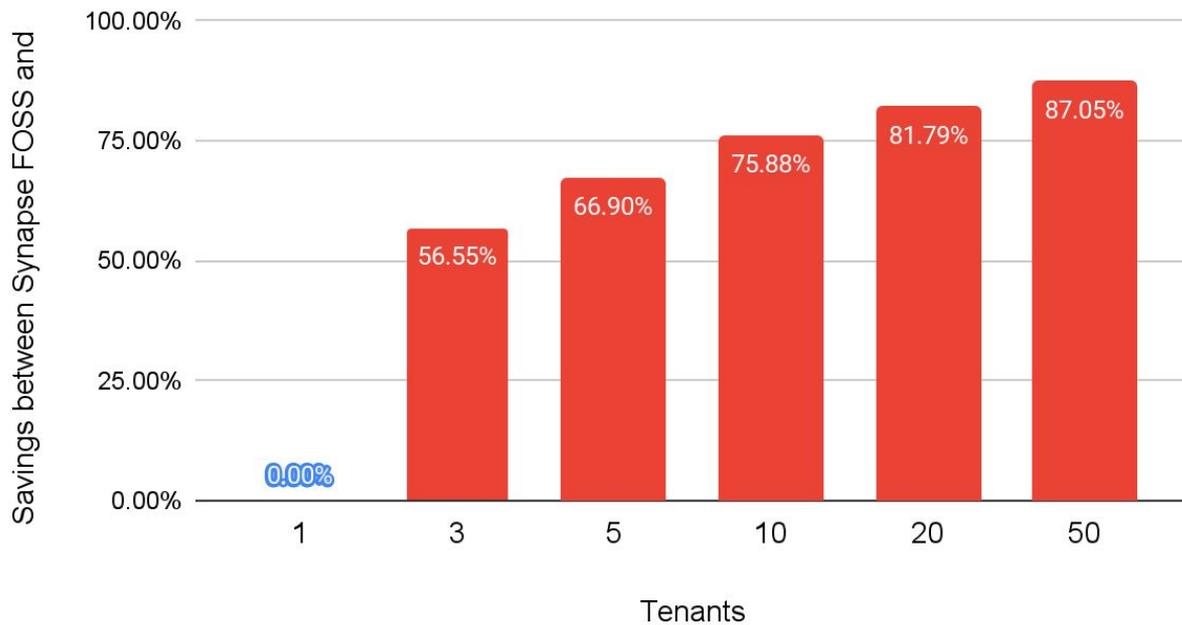
- We have built a **modified version of Synapse**
 - Initially prototype, then cloud PoC, now **ready for production**
 - Deployed as one **shard**
 - Can inhabit up to **50 tenants** each (limited due to Python being single-threaded)
 - Comes with a **Kubernetes controller** to manage multiple shards and their tenant allocations
- We have conducted **resource usage tests** to illustrate potential savings
 - 5 users per tenant
 - 1-5 DMs per user (encryption enabled)
 - 1-5 public rooms per user
 - Federated rooms between the hosts
 - On average between 60-120 seconds, each user sends a message into one of their rooms at random (equals 2 events/second; comparison: matrix.org has 35 events/s)

- Comparison of the **memory consumption** of **Synapse FOSS** and **Synapse Pro** in a scenario over **five years** to illustrate the **savings**
- Terminology
 - Tenant = An individual Synapse host with an own domain and DB
 - Shard = A Synapse main process that hosts a number of tenants
- Using our load testing suite we have measured **per-tenant memory usage** dependent on the number of tenants

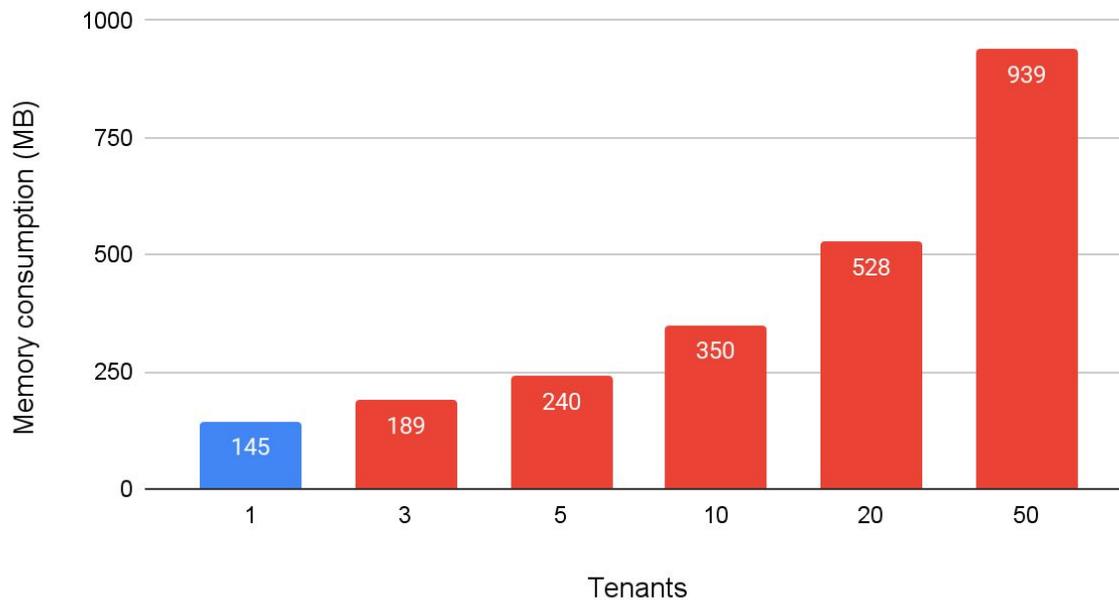
Per-tenant memory consumption



Per tenant savings with Pro (relative)



Memory consumption per Pro shard

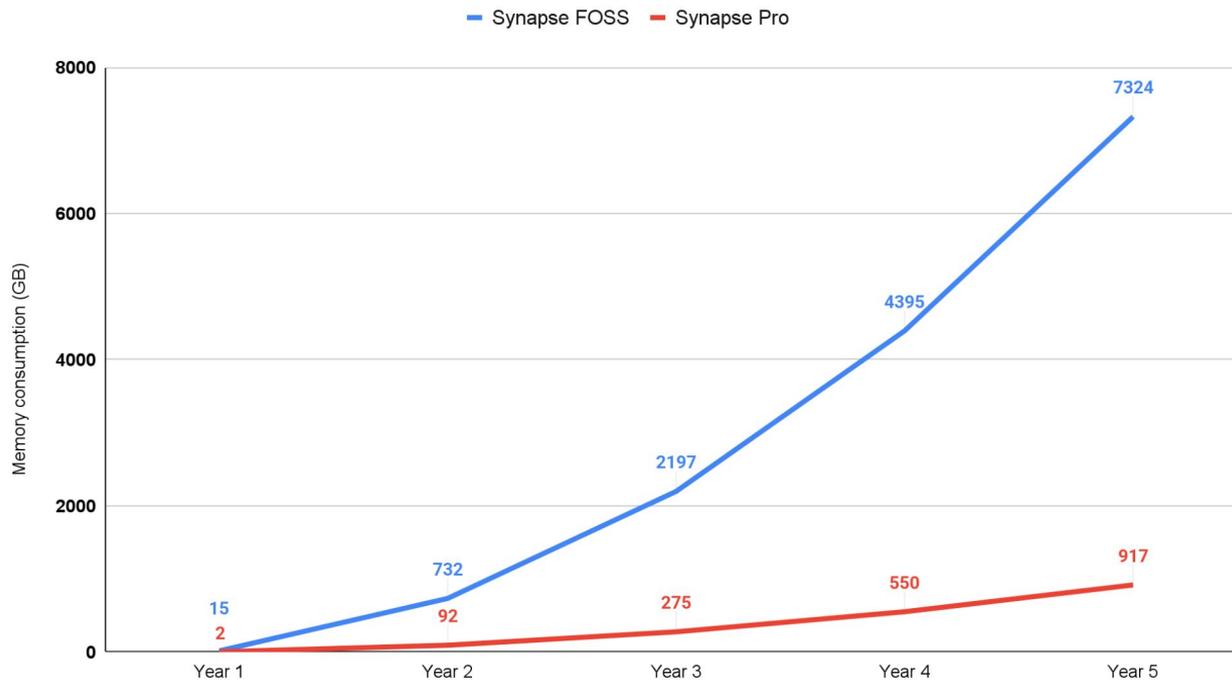




- We expect to be able to run up to **50 tenants in one shard**
- We calculate memory consumption over five years with the following assumption on the **number of tenants**
 - Year 1: 100
 - Year 2: 5.000
 - Year 3: 15.000
 - Year 4: 30.000
 - Year 5: 50.000



Average yearly memory consumption



Synapse Pro and ESS.

Synapse Pro and ESS.

- **Synapse Pro for small hosts (multi-tenancy)**
 - Kubernetes controller deploys Synapse shards (each can host up to 50 tenants)
 - Kubernetes controller dynamically manages shards and tenant allocations
 - **Up to 90% infrastructure cost savings, operational cost savings due to automation**
- **Synapse Pro for small hosts provides different modes of operation**
 - Standalone (Synapse-only)
 - ESS Pro
 - ESS TI-M
- **Fleet management with ESS**
 - ESS can easily be plugged into CD tooling (like **ArgoCD**) to manage whole fleets of Matrix servers automatically
 - Large hosts and small hosts can be **combined**
 - **Switch** from sharded small host to individual large host (and vice versa) easily

- **ESS TI-M** is a **special version of ESS Pro** focused on the requirements of **TI-Messenger Pro** and ePA as specified by the German National Digital Health Agency **Gematik**

Components

- Synapse Pro
(with TI-M modules)
- SBG
(with TI-M modules)
- Federation List Service (FLS)
- Signal
(Push Gateway)

Core properties

- Comes with the **same commitments and services as ESS Pro** but ships different components and features, compliant with the TI-M Pro and ePA specifications
- Can be deployed in **Pro or ePA mode**
- Provides a modern, scalable and lightweight **Messenger Proxy**
- Part of an **accredited TI-M product with T-Systems** for BARMER health insurance and others

What customers can expect from ESS with Synapse Pro.

What customers can expect.

90% reduced resource footprint and operational cost reductions for small hosts

- No need to worry about Synapse internals, load balancing, etc.
- Deployment and Matrix consulting
- **Long-term support** (LTS) (to prevent **Gematik re-accreditation risks**)
- TI-M Pro/ePA-compliant backend with Messenger Proxy (optional)
- Commercial license for Synapse without copyleft (to prevent **AGPL risks**)
- ISO 27001, ISO 5230 and Cyber Resilience Act compliance
- True L3 Support (from the code owners)
- Security notices in advance

What's next?

- Synapse Pro for small hosts is **ready for integration and production use**, first successful PoCs are running
- Synapse Pro for small hosts is **not only for TI-Messenger!**
Let us know your small hosts use case!
- If you haven't yet, start your PoC tomorrow!

Questions?