

Sentient White Paper

[Website](#) | [YouTube](#) | [LinkedIn](#)

“Generative AI is out. The age of **Proactive AI is here.”**

When most people think of AI, they generally think of it as a tool or a simple productivity booster that accepts instructions, or an input “prompt”, processes a response and provides an output. It started with simple text-in, text-out scenarios. Then people realised that these models can handle a lot more than text - they moved on to images, audio and even actions as outputs (something we call “AI Agents” today).

However, all of these applications still had one flaw - in most cases, these **“agents” were not acting autonomously** - they still relied on human input to be told what to do. Sure, you could say that they were autonomous in a sense, but it wasn't *true* autonomy, in every sense of the word - i.e. they can't anticipate the user's needs proactively and act on them without being given said “prompt”.

A lot of people say that current generative model architectures are still very primitive and nowhere near what we require to make fully autonomous agents that can think and reason like a human, without human intervention.

We at **Existence** believe that this is not an architecture problem but rather a **context problem**. Most people who complain about AI not performing right are **generally just prompting it wrong** - garbage in, garbage out. And the same goes for agents; they can't reason and make decisions on their own simply due to **a lack of context**.

A ton of agent startups are being built today, but most of these products still require large amounts of human intervention. AI-powered email clients still need users to sit and perform a sizeable fraction of the work, not even coming close to their promised “superhuman” levels of productivity. Most of them are just glorified email classifiers and labellers.

Also, most agent solutions are scattered across several apps. Want to automate your email? Download tool A. Want to automate your calendar? Download tool B. There is no single tool that manages and maintains context across your digital life seamlessly, bridging the gap between all your tools and acting as a central interface to manage every aspect of your human life.

How big of a problem is this?

According to a study conducted by McKinsey, "The average person spends about **28% of their workday** on email-related tasks, which translates to approximately **2.5 hours per day** on average."

And it's not just email - notifications are equally time-consuming. Studies have shown that "The average person checks their phone about **58 times a day**. Many of these checks are triggered by notifications from social media apps, messaging platforms, news outlets, and even games. In terms of time, it adds up to around **2-3 hours a day** spent interacting with notifications."

Can you imagine how all those wasted hours add up? 4-5 hours wasted per day per person is billions of hours of human life wasted - days, months, years wasted. Years that you could have spent doing something productive.

That is how our vision for Sentient was born.

We envision an **ever-present, personal AI** that lives across all your devices, streaming in context from every possible source, categorizing and storing it neatly. It will use all this contextual information to learn about you, organise your to-dos and tasks for you and help reduce cognitive load.

And it won't just stop at simple organization. It will start handling various aspects of the task for you, completely autonomously.

Consider this scenario - a new email has just popped up in your inbox informing you about a presentation you need to give next week. Before you even open and read the email, Sentient has already read

the email and started working on it. It has access to your other emails, your cloud storage like Google Drive, knowledge bases like Notion and more. It can also access your Google Docs, Google Slides and Google Calendar while having the ability to perform deep research on any topic using the internet. Sentient swiftly starts working on the presentation - it lays out talking points for you in a Google Doc, creates a basic outline of the presentation in Google Slides and sets up reminders and other events in your Google Calendar to help you complete the presentation on time. It will also save short-term memories about this presentation in its knowledge base, which can help it with planning other events in the coming week.

All that, without you saying a single word.

Of course, you can still interact with Sentient like you would with any AI companion. We offer low-latency text-based chats that happen in a single, unified interface without having to manage multiple chats for different topics. You can also seamlessly switch over to a voice call.

In the future, Sentient will be able to stream in context from many more sources, like the microphones in your device, which will allow it to truly be present in every aspect of your life.

But what about privacy?

Privacy is a very natural concern when talking about data collection and processing on this scale. Sentient will be fully customizable. Users will choose what data goes in and what actions Sentient will be able to perform and the level of autonomy it has.

Everything is transparent and being built open-source. Users will have full control over their data, and the highest standards of security and data privacy will be maintained. No user data will be used to train any of the models we use.

Users who want 100% control over everything will be able to host Sentient on their hardware.

We envision a world where everyone has their own personal, proactive AI companions that handle boring, monotonous tasks for them, ensuring that they have the time to focus on what matters most.