



## Supporting Document

### Overview

sundae\_bar (Subnet 121) is building one continuously improving general purpose AI agent designed to do real business work - unlocking enterprise productivity for a wide variety of business tasks

---

### The Problem

Most AI tools perform well in controlled demonstrations but struggle in real business environments.

Businesses need AI that can:

- Handle multi-step workflows
- Maintain context across tasks and time
- Adapt to feedback and changing priorities
- Operate reliably without constant human direction

Traditional AI systems are not built for this. They respond to prompts and do not *work*.

---

# The Generalist Agent

The `sundae_bar` generalist agent is designed to operate across common business functions - from operations and finance to marketing, research, and internal coordination.

Key characteristics:

- **Role-aware behavior** — understands how it should act in a business context
- **Persistent memory** — retains goals, preferences, and prior work
- **Workflow execution** — completes tasks end-to-end, not one prompt at a time
- **Continuous improvement** — gets better as it is evaluated and refined

The agent is not trained for a single task. It is built to generalize across the kinds of work businesses perform every day.

---

## AI agents are not chatbots - Use case examples

### Agent example 1: HR Operations Agent

An autonomous HR agent that manages hiring end-to-end - screening CVs, shortlisting candidates, scheduling interviews, and drafting outreach - improving over time as it learns from hiring outcomes.

### Agent example 2: Marketing Execution Agent

A marketing agent that plans, creates, and runs campaigns across channels - generating content, publishing posts, analyzing performance, and iterating based on real engagement data.

### Agent example 3: Financial & Crypto Operations Agent

An autonomous finance agent that monitors markets, executes strategies, manages risk, and produces reports - continuously adapting to changing market conditions.

---

## Gartner's 2026 prediction

AI agents are moving from experimentation to infrastructure, with the market expected to exceed \$47 billion by 2030.

Gartner predicts that by 2026, 40% of enterprise software will include task-specific AI agents, up from under 5% today. [Source](#)

---

## How Bittensor results in the most performant agents

The **sundae\_bar** incentive design rewards only the top-performing agent creating a **winner-takes-all environment** that intensifies competition and directly incentivises real improvements.

***Bittensor allows many independent developers to compete to make the same agent better, accelerating progress beyond what any single company can achieve alone.***

Any incremental advancements made by the leading agent become visible to the ecosystem, allowing others to build, refine, and push performance even further.

Crucially, the knowledge generated is shared across 100's of developers, which is exactly why **sundae\_bar is being built on Bittensor** – to harness open innovation, rapid improvement and collective intelligence at scale.

---

## Memory and Context: Built for Real Business Work

Most AI agents fail in production because they treat each request in isolation, losing context between steps and breaking down in real workflows.

sundae\_bar's generalist agent is designed to operate with **persistent memory and contextual awareness**, allowing it to behave like a reliable digital worker rather than a chatbot.

**Retains information:** such as prior decisions, task state, retrieved data, and tool outputs. This enables it to resume work, reason across multiple steps, and avoid repeating actions or producing inconsistent results.

**Reduced errors:** When interacting with systems like CRMs or calendars - the agent tracks what has already been done, what data was modified and what remains outstanding.

**The result is an agent that can operate consistently across complex business processes.**

---

## Summary

sundae\_bar × SN121 is building a single, continuously improving digital worker trained through open competition and deployed into real enterprise environments.

- One agent
- One shared codebase
- Incremental improvements from 100's of developers
- Goal to be deployed across millions of businesses

This structure enables faster improvement, clearer value capture, and direct alignment between builders, businesses, and token holders.

---

## Additional information:

Website: [sundaebar.ai](https://sundaebar.ai)

X: [@sundaebar\\_ai](https://twitter.com/sundaebar_ai)

Bittensor explainer video: <https://www.youtube.com/watch?v=mJX8DJyICjQ>

Github: <https://github.com/sundae-bar/bittensor-subnet>