

The company brain.

Using AI as an intelligence layer—not another chat window.

MINI-AGI

MEMORY

CAPABILITIES

CONTROL

Coffee, technology, and better questions.



A chat window answers. A company brain coordinates.

CHAT

Prompt → response

- Context expires
- One task at a time
- People reconnect every step

COMPANY BRAIN

Context → action → learning

- ✦ Memory persists
- ✦ Capabilities connect
- ✦ Verified work improves the system

The interface is not the architecture.

The company is already a mini-AGI.

It senses, models reality, pursues goals, acts, and receives feedback.

SIGNALS

Customers • market • operations

WORLD MODELS

What is true now?

GOALS

Strategy • constraints • intent

CAPABILITIES

Build • sell • support • operate

INTERFACES

Chat • voice • app • API

AI does not replace the objective function.

Three conversations describe different layers of the same system.

01

PERSONAL STACK**Tommy Eastman**

Persistent agent • model choice • tools • reusable skills

02

COMPANY BRAIN**TFTC / Conner Brown**

Durable memory • retrieval • knowledge relationships

03

ORGANIZATIONAL LAYER**Dorsey ♦ Botha**


World models • capabilities • task-specific interfaces

Individual intelligence compounds into organizational intelligence.

A company brain is memory that can be retrieved, related, and used.



How intelligence moves through the company.




Verification is what turns activity into learning.

Without it, the brain accumulates confidence—not knowledge.


Shared context changes why hierarchy exists.

PYRAMID



Information moves
up and down.


CIRCLE



People direct from
the edge.


A THESIS — NOT A LAW

Products become capabilities. Meetings become experiments.



Stable interfaces still matter where safety, repetition, or regulation demands them.

The newsletter becomes a company-brain loop.



Human role: choose the story, shape the argument, own the truth.

A company brain needs governance and context hygiene.

OWNERSHIP

Where memory, models, and tools live

LEAST PRIVILEGE

Only the access required for the task

APPROVAL GATES

Humans authorize consequential actions

VERIFICATION

Read back the real external state

CONTEXT HYGIENE

Curate, prune, and version examples

ROLLBACK

Logs, backups, and reversible operations

$$\text{AUTONOMY} = \text{CAPABILITY} \times \text{TRUST} \times \text{CONTROL}$$

A 30-day company-brain experiment.

WEEK 1

MAP

Choose one weekly workflow. Define inputs, decisions, risks.

WEEK 2

REMEMBER

Create durable sources, definitions, examples, and ownership.

WEEK 3

RETRIEVE

Add selective recall and read-only capabilities.

WEEK 4

OPERATE

Automate one reversible step. Verify and write back.

Measure coordination cost, cycle time, and correction rate.

Design how intelligence flows—not only what AI says.

- 01** **CONTEXT** What does the company know?
- 02** **CAPABILITIES** What can it safely do?
- 03** **LEARNING** How does verified work compound?
- 04** **CONTROL** Who remains accountable?

What repeated workflow should become your first company-brain loop?

Sources: TFTC #796 · TFTC/Conner Brown company-brain segment · Long Strange Trip with Jack Dorsey · Sequoia “From Hierarchy to Intelligence”