

The AI Native

Operating at the Speed of Information

AI



Armin Catovic

Director Data & AI

Redesigning the firm

From AI tools to the redesign of work

One question changes the design of work

How can agents solve this problem?

Assist

Act with approval

**Autonomous
business function**

**Keep with a
person**

AI use is widespread. Company impact is rarer

44%

report AI scaling across
their enterprise

6%

meet the survey's AI
high-performer definition

McKinsey Global Survey 2026 · self-reported results

Workflow redesign is associated with value

Nearly 3 in 4

AI high performers report
fundamentally redesigning workflows

About 1 in 4

other respondents report
the same level of redesign

The unit of change is the workflow, not the prompt.

The company becomes legible

Trusted context, shared standards, agent action

A small core, with ownership across the company

Central Data & AI

Data infrastructure

Production templates and quality standard

World model and data agent runtime

Every business team

Domain context and definitions

Data emitted to the shared standard

Agent tasks, acceptance and review

Leadership owns the reviewed decisions that define company intent.

The company world model

Tier 1

Reviewed company intent

Identity, customers, strategy, positioning, operations, hiring

Tier 2

Systems of record

Financial measures, CRM records, product and operating data

Tier 3

Observed signals

Customer calls, conversations and other unstructured evidence

How we do it at Funnel: Reef

The screenshot displays the Reef AI interface, which is divided into a main workspace and a right-hand sidebar.

Main Workspace:



- Header:** Shows the Reef logo, a task title "Funnel Customer Feedback Analysis", and status information: "Idle", "Created 11 min ago", and "Updated 2 min ago".
- Left Sidebar:** Contains navigation options: "Tasks", "Explorations", "Projects", "Skills", and "Semantics".
- Task Content:**
 - A prompt box at the top says: "what do funnel customers think about our measurement product, and how does it compare against the competition, prepare a google slide deck for me".
 - Below the prompt, it indicates a "Thought trace" with 49 tools, 3 subagents, and 2 notes.
 - A section titled "Open your Google Slides deck" contains text: "The 14-slide deck covers customer feedback from April–September 2026, comparisons with Rockerbox, Northbeam, Triple Whale, Measured and Fospha, and recommended priorities." and a "Main takeaway" paragraph.
 - Another paragraph states: "The analysis screened 319 customer-voice claims and 136 competitive claims separately; the sets may overlap. This is qualitative evidence—not a representative satisfaction survey or an accuracy ranking."
 - A note mentions: "Linked sources and caveats are in the speaker notes. [Preview the deck.](#)"
 - A "Files" section shows two items:
 - "Funnel Measure deck preview" (HTML, 1.8 MB, 3 min ago)
 - "Funnel Measure: Customer Voice & Com..." (PPTX, 4.0 MB, 3 min ago)
- Bottom:** A text input field says "Ask a follow-up..." and a footer note reads "Reef can make mistakes. Double-check important responses."

Right Sidebar (Slide Deck Preview):


- Title:** "Funnel Measure deck preview" with the file name "Funnel_Measure_Customer_Voice_Preview.html".
- Slide Title:** "Funnel Measure: Customer Voice and Competitive Positioning | October 2026".
- Controls:** Includes a "Dark theme" toggle, "Previous" and "Next" navigation buttons, and a dropdown menu currently showing "1. Funnel Measure".
- Slide Content:** The preview shows a slide with a dark green background, a lighthouse image, and the text "Funnel Measure", "Customer voice, selection criteria and competition", "April–September 2026", and the Funnel logo.
- Footer:** Indicates "1 / 14 Funnel Measure" and includes a "Slide text" button.
- Disclaimer:** A note at the bottom states: "Preview rendered with LibreOffice. PowerPoint and Google Slides may display some details differently."

Swimming in the Reef

DATA SOURCES



USERS & AGENTS



Voices on Reef



Fredrik Skantzé · 4:31 PM

@pellemade and I hanging out with Reef on a Friday afternoon in the office doing analysis for the commercial strategy that would have required a whole team of people to help with in the past. This thing is on fire 🔥 I'm so impressed with what the team has achieved. And it will just get better! 🙌



👍 37 🔥 49 🤝 32 🍷 12 🦀 9

👤 6 🙌 1 🗨️

3 replies · Sep 12th at 12:28 PM

Feedback from the organization:

- It's been aiding a lot in individual, team and company-wide learning and development. Helping to dig deep in areas that was too inaccessible or timely to get together previously (Manager Sales EMEA)
- Reef is my religion 🙏! It really helps to understand the business better when designing new experiences for our users. It doesn't hurt that the team working on it are the most absolute AWESOME people in the world 🌟 (Sr Product Designer)
- Amazing work by everyone involved 🙌 Reef is such a game changer, working with it on a daily basis 🔥 (Chief Data Science Officer)
- Most of my user research starts with Reef. Without it, my insights would be underwater and words too shallow! 🌊 (AI/UX Context Designer)
- This is such a game changer! What Reef enables is on another level. Super impressed 🤖 (Director Partners)

The deliberate agent loop

- | | | |
|----|-------------------------|--|
| 01 | Goal and mandate | A person sets the outcome, boundaries and approval points. |
| 02 | Trusted context | Reef retrieves relevant, sourced company knowledge. |
| 03 | Action | The agent uses approved tools to complete bounded work. |
| 04 | Check | It tests the result and records supporting evidence. |
| 05 | Learn | It proposes context corrections to a domain owner. |

Tier 1 changes require human review.

A positioning decision travels through the company

Decide

Leadership approves a change to target-customer positioning.

Propagate

Reef identifies affected material and drafts sourced updates.

Verify

Owners approve the changes. Reef checks for remaining drift.

Illustrative workflow

Every team gains two agent-facing responsibilities

Maintain domain context

- Keep definitions current.
- Publish data to the shared standard.
- Resolve contradictions.

Orchestrate agents

- Set goals and limits.
- Review consequential actions.
- Measure the finished work.

Slides, documents, sheets, become sourced outputs of the company world model.

The tensions

Leverage brings new costs and responsibilities

Token spend is an input. Completed work is the outcome

Cost per verified outcome

Model calls + tools + retries + review

Successfully completed work

1 in 5


say AI operating costs
constrain use

The duality of AI spend

Funnel AI Spend

Engineering & Internal AI Spend

Jan-Sep 2026



More leverage changes the path to expertise

The job bundle changes.

Agents handle more first drafts and routine execution.

People carry more responsibility for judgment and review.

How will a junior become the next expert?

More capability creates more dependency

Data location

Know where sensitive context is processed.

Model choice

Route tasks by measured quality and cost.

Continuity

Retain a workable path when a provider fails.

Portability

Keep Reef's tools and context usable across models.

An agent needs a mandate, not just access

**An agent needs a
mandate, not just access.**

Identity

Permissions

Approval

Traceability

Clear boundaries make useful autonomy possible.

The compounding loop

Every completed task can improve the next one

The AI-native scorecard

Verified outcomes

Useful work completed to the agreed quality bar

Cycle time

Time from a request to an approved result

Full cost

Model, tools, retries and human review

Corrections

Overrides, errors and escalations

Context health

Ownership, freshness and resolved conflicts

Reef engagement: 50 to 230 weekly users · 100 to 320 weekly tasks · 3K to 10K tool calls per week

Completed work builds reusable company knowledge

Better context

An owner resolves a missing definition or contradiction.

Reusable procedures

A successful approach becomes a shared workflow.

Better checks

A failure becomes a test that catches the mistake next time.

Evidence from agent research

42.4

Base agent

59.5

With ACE

+17.1 percentage points

Same model. Reusable lessons from execution feedback.

The loop compounds when later tasks reuse reviewed lessons.

Source: [Zhang et al., ACE, ICLR 2026, Table 1](#). Average AppWorld task and subtask completion score.

The AI Native

Design the company so agents can do useful work,
and leave us better informed after every task.