

• I — OPEN

DESIGN + AI SUMMIT

Recall versus inference

Building AI that actually knows you

Stuart Frisby · mrfrisby.com

• I — OPEN

MONDAY

• I — OPEN

Booking.com → Deliveroo → Preply

• I — OPEN

COSTA

Recall, not inference.

Interactive prototyping

AI drawing interfaces

Figma Make



claude.ai / Claude artefacts visual design and structure - Claude

How Claude artefacts look, and why

Chat

Chat

Share

Switch to dark theme

How Claude artefacts look, and why

An artefact is the thing Claude makes for you, shown beside the conversation rather than inside it. Its appearance follows from a handful of practical decisions: keep the work apart from the talk, keep it in one self-contained file, and let it adapt to wherever you open it.

The two-pane layout

Choose a kind of artefact. The conversation sits on the left; the finished thing sits on the right.

CodeDocumentAppDiagram

Show how a request reaches a website.

Here's the path. Each step hands the request to the next one until the page comes back.

Request path

Diagram

Request pathSVG

CopyPublish

Browser

sends the request

DNS

finds the address

Server

builds the response

Your screen

shows the page

This window is an illustration, not a screenshot. The real controls differ between the web app, desktop and mobile. The tip calculator and the theme switch above do work, because artefacts

An asset generator.
Generating assets.

**Design is what happens
between your ears.**

Abductive → Deductive

**Capability without context
is just a very fast amnesiac.**

Not memory. Synthesis.

The Force Multiplier Principle

**A weak designer with good
AI habits is still a weak designer.**

Notebooks. Slack threads.
Meeting notes.

**The working life generates its
own record. Most of it disappears.**

●●●

Search...

#K

COSTA |

Briefing

Meeting Prep

Journal

Tasks10

Notes

URLs

Slack

Images

Calendar

Transcripts

People

Tags

Profile

Sun 20 Sep 12:30

Briefing

play

Subscribe

Regenerate

Good morning, Alex.

Today

Two 1-1s today — Maya at 10 and Jordan at 15:30. The onboarding review with Lucas is at 14:00; worth reading his proposal before that (it's on your task list, 1h estimate).

Four tasks flagged for today. The two high-priority ones are linked to meetings: Maya's design review feedback (45m, before 10am) and Lucas's onboarding flow review (1h, before 2pm). Marcus's API error state response can wait until this afternoon.

People

Maya Chen — last spoke Tuesday. She's presenting the notification redesign at 10:00. Her last 1-1 flagged scope creep concerns; worth having a clear position ready before the meeting.

Jordan Park — 16:00 hiring panel. No open threads from last week; straightforward.

Focus

The Q2 principles deck is the highest-leverage output this week. It's blocking the headcount ask and shapes the growth squad's Q2 design direction. Two hours of clear thinking this morning beats three rushed hours this afternoon.

Calendar

Tue 26 Mar today

● 09:00

Design standup

● 10:00

Maya Chen 1-1

● 11:00

Focus time

now

● 13:00

Leadership sync

● 14:30

Onboarding review

● 15:30

Jordan Kim 1-1

● 16:30

Hiring panel

Meeting Prep →

Tasks

Today: 2 · All: 11

☐

H

Review Maya's redesign proposal

due:today

☐

H

Send Q2 principles deck

due:today

☐

M

Draft headcount rationale

due:fri

☐

M

Follow up DS contract renewal

due:fri

☐

L

Respond to Marcus

View all tasks →

ASSISTANT

YOU

Can you draft a Q2 design strategy doc?

COSTA

Done. It covers goals, team structure, and four strategic bets for Q2. I've kept it to one page of signal.

Q2 Design Strategy.md

Draft · 847 words

I pulled from your profile knowledge agent (career context and working style), your Q1 strategy doc, and Maya's 1-1 last week (she flagged the component audit as a Q2 dependency).

Can you add a section on the DS hiring plan?

**A chief of staff, for one
person, every day.**

Not a startup.
Not a pitch.

• UNDEFINED

One palette. Twenty-four hours.



- V – INTRODUCING COSTA

SL DYN EASE STARLIGHT

$$\Delta H(h) = 40 - \frac{h - 7}{14} \times 10$$

$$t = \begin{cases} \max\left(0, 1 - \frac{h}{7}\right) & h < 7 \\ \max\left(0, \min\left(1, \frac{h - 7}{11}\right)\right) & h \geq 7 \end{cases}$$

$$H_o = (H_a + \Delta H(h)) \bmod 360^\circ, \quad S_o = S_a, \quad L_o = L_a + t^2 (0.88 -$$

Full 24-hour extension of SL dyn ease and the active production formula for COSTA. The ease parameter t is bidirectional: it runs forward from 0 at dawn ($h=7$) to 1 at $h=18$ in the daytime half, and in reverse – from 1 at midnight down to 0 at dawn – in the overnight half. This ensures the opposite colour is highly luminous during the small hours as well as the afternoon. The hue drift ($+40^\circ$ at dawn, decreasing linearly to $+30^\circ$ at dusk) continues unclamped into overnight hours, reaching $+45^\circ$ at midnight. Continuity at $h=7$ is guaranteed by construction: both branches evaluate to $t=0$ at dawn.

Shape of work.
Not shape of a chat window.

ONE LAYER VERSUS THREE

GENERIC AI TOOL

- **Interactive layer**
Ask a question · get an answer

- **Automatic layer**
If it exists, it lives in their product

- **Corpus**
Exists, but isn't yours to take with you

Your context, held in someone else's product.

COSTA

- **Interactive layer**
Assistant · full corpus as context

- **Automatic layer**
Brief · meeting prep · task extraction

- **Corpus**
1-1 notes · transcripts · tasks · calendar

Context accumulates. Both layers draw from it.

The model is the same.
The corpus is the difference.

AGENT DIRECTORY STRUCTURE

context/agents/

profile/ Profile Knowledge Agent

- agent.json { "name": "Profile Knowledge Agent" }
- career.md Role, background, professional trajectory
- working-style.md Decision defaults, what energises and drains

organisation/ Organisation Knowledge Agent

- agent.json { "name": "Organisation Knowledge Agent" }
- people.md Team structure, reporting lines, key relationships

decisions/ Decisions Knowledge Agent

- agent.json { "name": "Decisions Knowledge Agent" }
- made.md Settled calls — the why, not just the what
- open.md Live tradeoffs, not yet resolved

- VI — WHAT IT TAUGHT ME

On demand

Chat interface

- 01 System prompt includes agent index + required-reads rules
- 02 Model reads task, identifies which agents apply
- 03 Model calls read_file for each relevant file
- 04 Content enters conversation context mid-loop

Model-directed • correct for general questions where relevance varies

Pre-loaded

Brief • meeting prep

- 01 Server identifies which agents are always needed for this task
- 02 Agent files are read and injected into the system prompt
- 03 Model receives full agent context on the first turn
- 04 No read_file round-trip required

Server-directed • correct when the same agents are always needed

• VI — WHAT IT TAUGHT ME

EVERYTHING IT'S ALLOWED TO KNOW

 Briefing

 Meeting Prep

 Journal

 Tasks

 Notes

 URLs

 Slack

 Images

 Calendar

 Transcripts

 People

 Tags

PEOPLE



Maya Chen

Head of Product Design · Meridian

Direct report · 8 mutual connections


SYNTHESIS

Maya joined 18 months ago and leads the design system work. Working relationship is strong — 1-1s are consistent and she's candid about blockers.

1-1 HISTORY

- Mar 18 Component audit · team capacity
- Mar 04 Q2 OKR planning · roadmap
- Feb 18 Team structure · hiring plan

- VI - WHAT IT TAUGHT ME



NOW PLAYING

Your Morning Brief

COSTA · Sun 20 Sep



0:47

2:14



My own voice.
Not another boss.

Assistant

CONTEXT LOADED
Meeting Prep · Design review · Maya Chen

YOU
Can you draft a quick agenda for the design review at 2pm?

COSTA
Based on your last session with Maya and the open component audit thread, here's a draft:

1. Component audit status — 10 min
2. Q2 design OKR sign-off — 15 min
3. Notification redesign feedback — 20 min

Ask COSTA...

Derived Actions.

Plain markdown.
My machine. No lock-in.

It's my context.
I wanted to own it.

**Every conversation
started cold.**

Averaged. Flattened.

coach

CONTENT

`get_dlm_writing`

`get_dlm_links`

`get_design_worldview`

`get_stated_positions`

`get_taste`

`get_frameworks`

INSTRUCTION

`get_writing_style`

`get_voice_anti_patterns`

`get_ai_instructions`

CONTEXT

`get_career`

`get_working_style`

`get_personal_context`

mrfrisby.com/ask

12 / 13

prompts won by context


“That isn’t a framework I use or have written about — I don’t want to walk you through something I’d be inventing on the spot and attributing to myself.”

• VII — A PORTABLE LAYER

• FULL REPORT

Can Personal Context Servers Improve AI Outputs?

[Link at the end](#)



**Recall failing gracefully
looks like honesty.**

What stops happening.

**It's a relief to know my bad
memory isn't making me look
like a bad person.**

206h

estimated saved this quarter

What energises you.
What drains you.

**COSTA doesn't
run itself.**

Proposed Actions

Right some of the time.
Wrong most of the time.

Bounded. Unbounded.

AI-native means expanded reach. Not new tools.

• X — ZOOMING OUT

Taste. Judgment.

**“This changes everything”
vs. a smaller, more useful idea**

• X — ZOOMING OUT

1. Collect your raw materials.

2. Consider building your own MCP server.

**3. Don't let design get
recast as asset generation.**

**Design is what happens
between your ears.**

**Recall didn't need to be
invented. It needed to be
taken seriously.**

Thank you.



SLIDES & LINKS

**mrfrisby.com/
design-ai-summit**

The deck, the eval write-up, /ask, coach, and the newsletter — all at that one link.