

MODULE 6 • EXTENSION

From VibeFlow. *To Agentic Programming.*

When the cast grows, the discipline scales. The pyramid does not grow a fourth floor. What changes is the cast, the audience, and the leverage.

THE CLAIM

You can outsource your thinking. You cannot outsource your understanding.



Where this *module sits.*

Documents 1–5 walked the pyramid: Consigliere → Agent → Product Developer → Cross-cutting. The pyramid has three levels and three levels only. This module does not add a fourth.

WHAT MODULE 6 ADDS

A horizontal extension above Level 3 for those who have shipped real products with VibeFlow 3.0 and now need to scale.

WHO IT IS FOR

Strong Level 3 builders. EON's internal team migrating to VibeFlow 4. AI Founder Programme Door 3 graduates.

THE WARNING

Skip Level 3 and you will scale weaknesses across many workers, expose them to other agents, then fork them into many domains. Master Level 3 first.



From director. *To showrunner.*

VibeFlow 3.0 was the director with one writer (CoWork) and one actor (Code). Every take watched. Every scene reviewed. VibeFlow 4 is the showrunner of a series — many actors at once, the writers' room intact, the showrunner's judgement still deciding what makes the cut.

WHAT DOES NOT CHANGE

The Brain plans. The Hands execute. They never collapse into each other. The discipline of separation is what made Level 3 work; it is what makes Module 6 work.

WHAT DOES CHANGE

The Hands are now plural. A Dispatcher manages the queue. A Reviewer gates what merges. A Front Desk lets other agents consume what you have built.



Vibe coding is over.


Agentic engineering begins.

At Sequoia AI Ascent in February 2026, Andrej Karpathy retired the term he had introduced one year earlier and replaced it with one that describes the new default workflow.

THE REFRAME

"agentic — because the new default is that you are not writing the code directly 99% of the time. You are orchestrating agents who do, and acting as oversight."

— Andrej Karpathy, Sequoia AI Ascent, February 2026



THE LINE THAT ANCHORS THIS MODULE

*You can outsource your
thinking.*



*You cannot outsource your
understanding.*

— Andrej Karpathy, Sequoia AI Ascent, February 2026

The agents do the production. The director does the deciding. Outsource the thinking; retain the understanding. Without it, the agents drift; with it, they ship.

Three signals.

All three required.

Two-out-of-three is a signal to stay at Level 3. Three-out-of-three is the green light. Escalating before all three are real produces complexity without proportionate gain.

SIGNAL 1 • HANDS CANNOT KEEP UP

You are waiting on the worker more than the worker is waiting on you. The bottleneck has moved off planning and onto execution.

SIGNAL 2 • AN AGENT NEEDS TO CONSUME

A specific named partner or customer has asked, in plain words: 'Can my agent just call this directly?' Not hypothetical. Real.

SIGNAL 3 • WORKFLOW WANTS TO FORK

Stable for ninety days. Two or more domains could use the same workflow with their own data. The opportunity is many forks, not one system.



VibeFlow 4.

Eight roles, three new.

The original five (Brain, Hands, Vault, Publisher, Memory) carry forward. Three new roles enable the three pillars.



BRAIN

CoWork. Plans. Decomposes. Specs. Reviews tricky diffs.

DISPATCHER

NEW

Hands work units to fresh workers. Tracks the queue.

HANDS × MANY

Several Code instances in parallel, each scoped to one unit.

REVIEWER

NEW

Reads each diff against its spec. Approves or flags.

FRONT DESK

NEW

MCP server, CLI, API. The door other agents use.

VAULT

GitHub. Saves every approved diff. Enables rollback.

PUBLISHER

Vercel. Pushes to staging then production.

MEMORY

Supabase. What the system remembers across runs.

Three pillars.

Most internal to most strategic.

Adopt them in order. Pillar 1 changes how your team works today. Pillar 2 changes how customers reach you. Pillar 3 changes what businesses can be built on top.

PILLAR 1 • ORCHESTRATION

Many workers, one showrunner. Decompose sprints into self-contained units. Run them in parallel. Review before merge.

PILLAR 2 • THE BACK DOOR

When the customer is partly an AI agent. Expose your capabilities as API + MCP + CLI. Roll out in three stages.

PILLAR 3 • THE RECIPE LIBRARY

One mature workflow forks into many. Each recipe is spec + prompts + tools + acceptance test + memory schema.



Many workers.

One showrunner.

Several Code instances run in parallel. Each reads only its own unit's spec and the files in scope. None sees the whole system. The whole-system view lives only with the Brain.

RULE 1 • UNITS STAY SMALL

Thirty-minute scope. One file. Named acceptance test. A worker that goes off the rails fails small.

RULE 2 • WORKERS DON'T SEE THE WHOLE

Each worker sees its own brick. The Brain sees the wall. This is what stops 'helpful' sprawl.

RULE 3 • REVIEWER ≠ BRAIN

The Reviewer asks 'does this match its spec.' The Brain asks 'is the spec right.' Two different jobs. Never collapsed.



Five components.

All five required.

A unit with all five parts is one a worker can ship and a Reviewer can approve in a single short cycle. Miss any one and the worker has slack to drift.

- 1 SCOPE**
Which files. Read and modify. Usually one. Never more than three.
- 2 SPEC**
Plain English. Concrete, ten-year-old precision. The change in a paragraph.
- 3 ACCEPTANCE TEST**
How worker and Reviewer know it is correct. Defined before the work.
- 4 NON-GOALS**
What the worker must not do. The fence around 'helpful' drift.
- 5 TIME BUDGET**
About thirty minutes. Larger means it has not been decomposed enough.



RADAR's lead-scoring stage.

*Three weeks →
one weekend.*

RADAR is EON's 263-agent lead-intelligence system, running on Hetzner at a few dollars a month. The lead-scoring stage had months of patches. The Brain decomposed the cleanup into six units. Dozens of workers ran in parallel. Reviewer flagged two for respect. Done by Monday morning.

6

WORK UNITS

4 hr

HUMAN ATTENTION

2

REVIEWER FLAGS

WHAT THE HUMAN ACTUALLY DID

Wrote the spec. Decomposed it into six units. Reviewed two flagged diffs. Wrote one final-pass spec. Signed off. Outsourced thinking, retained understanding.



When your *customer is an agent.*

Most products today are still built for humans to navigate. Pillar 2 asks the harder question: what changes when the user is partly an AI assistant working on someone's behalf?

THE FRONT DOOR

Browser. Login. Forms. Dashboards. Built for humans. Still serves most of your customers. Stays.

THE BACK DOOR (HEADLESS 360)

API + MCP server + CLI. Same kitchen, new entrance. The agent reads a tool list and calls directly.

THE BAR FOR BUILDING IT

A specific named partner has asked. Not hypothetical. If nobody has asked yet, stay at Level 3.




Why are you still *telling me what to do?*

Throughout his Sequoia talk, Karpathy returned to one frustration: software products are still publishing instructions written for humans, even though the actual user is increasingly an AI agent acting on the human's behalf.

THE FRUSTRATION

"Why are people still telling me what to do? What's the thing I should copy-paste to my agent?"

— Andrej Karpathy, paraphrased, Sequoia AI Ascent, February 2026



No browser required.

The API is the UI.

At TDX 2026 in March, Salesforce announced “Headless 360” — a top-to-bottom rebuild of its platform around a single architectural premise: the agent is the user.

"Why should you ever log into Salesforce again?"

— Parker Harris, co-founder, Salesforce

"No browser required. Our API is the UI."

— Headless 360 launch positioning

"We just use the best, and often we will offer all three."

— Platform engineering on API/CLI/MCP

THE LESSON EON IMPORTS

Don't bet on one protocol. Expose API + CLI + MCP. Let consuming agents pick.



THE HOLLYWOOD REFRAME

Have your agent.

Talk to my agent.

THE OLD HOLLYWOOD RITUAL

'Have your people talk to my people.' The principals don't negotiate the details. The agents — entertainment agents — work them out between themselves. The principals see only the outcome: a yes, a no, or a counter-offer.



Headless rollout.

Skip none.

Every successful headless platform has rolled out in this order. Teams that try to skip from Stage 1 to Stage 3 detonate in support and lose trust with their first customers.

STAGE 1 • INTERNAL USE ONLY

Your own team consumes the back door. At least 90 days. Find the documentation gaps before customers do.

STAGE 2 • TRUSTED PARTNERS

5–10 named partners. 60–90 days. Each has skin in the game. High-bandwidth feedback fixes the API.

STAGE 3 • PUBLIC LAUNCH

Documentation, pricing, runbook. Easy if Stages 1 & 2 went well. Impossible if they didn't.



Sarah opens *zero browsers.*

Genesis 3 is EON's XR training authoring environment — 30-day delivery, integrated with EON's spatial platform. Sarah, training director at an offshore-platform operator, asks her own AI assistant for a safety training. Her assistant calls Genesis 3's MCP server. The two agents do the work.

1. SARAH SPEAKS PLAIN ENGLISH

'Offshore platform safety training, three modules, certification, end of June. ~200 field workers, six rigs.'

2. AGENTS HANDSHAKE

Sarah's assistant calls Genesis 3's MCP. Project created, environment attached, certification set, deployed to staging.

3. SARAH DECIDES THE THREE THINGS

Pass threshold, retake policy, voiceover localisation. Five minutes of human judgement. Training goes live.



Duplicating RADAR.

Many businesses.

When a workflow is mature in one domain, you fork it cleanly into many. The Builder is not writing nineteen agents. They are forking nineteen recipes and composing five of them into a working system.

WHY 'COPY' IS THE WRONG WORD

Agents are not files you copy. They are recipes you fork. Same recipe, different ingredients, different dish.

WHAT FORKING ENABLES

Same RADAR recipe behind many businesses. Same Genesis 3 recipe behind welding, surgery, pipelines, kitchens.

WHO RUNS THIS PILLAR

AI Founder Programme Door 3 graduates. Venture Builders. Non-programmer founders with strong domain expertise.



Five components.

Each on one page.

A recipe is forkable when each of the five is small enough to fit on one page, named clearly, and stored in a versioned location separate from the others.

1**SPEC**

What this agent does in plain English. Inputs, outputs, behaviour.

2**PROMPT SET**

System prompt + templated user prompts. Where most of the domain work shows up.

3**TOOL LIST**

External tools the agent can reach. Where Pillar 2's back doors get consumed.

4**ACCEPTANCE TEST**

Fixture inputs and expected outputs. How a fork is verified to still work.

5**MEMORY SCHEMA**

The slice of memory this agent owns. Often shared, but sliced per agent.



RADAR forks *for welding.*

A Venture Builder graduate — former vocational trainer with a decade of welding-certification work — wants to build a business in South-east Asia. They cannot afford to build infrastructure from scratch. They fork instead.

19

RECIPES FORKED

6 days


TOTAL FORK TIME

3 mo.

TIME TO LIVE

WHAT THE BUILDER WROTE

One thin custom layer for the welding-specific workflow logic. That is all. The other 95% of the system is forked recipes.



Four traps.

All tempting.

Each trap has a structural counter-measure. Each is named because every team will be tempted by it.

TRAP 1

BACK DOOR BEFORE FRONT DOOR

Front door must be in production 90 days before back-door work begins.

TRAP 2

UNITS TOO LARGE

Thirty-minute time budget. A unit that exceeds it gets split until it doesn't.

TRAP 3


COPY INSTEAD OF FORK

Speak in recipe language always. Vocabulary discipline reinforces architectural discipline.

TRAP 4

SKIP THE REVIEW

No diff merges without the queue. The few minutes save you the months.



Non-programmers. *And their partners.*

The agentic era does not eliminate the need for technical skills. It relocates them. Be honest about which pillar you can run alone.

PILLARS 1 & 2 • TECHNICAL PARTNER NEEDED


Recovering from hallucinated functions, debugging diffs, diagnosing MCP errors — these need a software-engineering instinct.

PILLAR 3 • THE NON-PROGRAMMER'S HOME

Reading specs, picking recipes, retuning prompts, defining test fixtures — judgement and domain expertise more than code.

THE PRACTICAL PATTERN

Builder runs Pillar 3 daily. When growth hits Pillars 1 or 2, bring in EON staff, a Distribution Engine partner, or a contractor.



Three lenses.

One workflow.

Pick one workflow you have already shipped at Level 3. Look at it through three lenses. One concrete deliverable per lens. Pick the most pressing and ship it. Then come back for the next.

LENS 1 • ORCHESTRATION


Where is the Brain becoming the bottleneck? Write a six-unit decomposition for one piece of work. Don't run it yet.

LENS 2 • THE HEADLESS BACK DOOR

Has anyone asked, in plain words? If yes, draft the MCP tool list. If no, write down what the named asker would want.

LENS 3 • THE FORK

Could this work in a second domain? Walk through the five-step fork for one agent. Estimate the fork time.



THE LINE THAT ANCHORS THIS MODULE

Outsource the thinking. Retain the understanding.

That is the discipline of agentic programming.

MODULE 6 • END OF CURRICULUM SERIES

EON AI VENTURES • GLOBAL VIRTUAL CAMPUS • MAY 2026

