

**AI PLAYBOOK**

---

# Vibe Coding for Marketers Part 2

June 2026



[theCMA.ca](https://www.thecma.ca)

Content partially generated by artificial intelligence, refined by human expertise.

This Playbook is part of the CMA's AI Mastery Series, empowering marketers to Implement AI in ways that earn regulatory confidence, maintain strong brand reputation, and foster consumer trust

# What you'll learn

---

A quick reminder: vibe coding is the practice of using conversational AI to generate functional code through plain-language prompts – no programming knowledge required. You describe what you need, AI builds it.

Part 1 was about adoption, learning the tools, building quick wins and proving the concept works. Part 2 is about embedding: moving vibe coding from something you do individually to something your team does by design. This playbook helps close that gap.

## **What you'll learn:**

- Shift vibe coding from individual quick wins to embedded team practice
- Navigate IT, legal and compliance as your builds become more complex and higher stakes
- Upgrade your AI toolkit, when and why to move from ChatGPT to Claude for advanced analytical work
- Scale CRM workflows
- Build collaborative AI workspaces where teams build, review and maintain code together
- Apply a three-tier escalation framework so you know exactly when to iterate versus when to involve technical teams

# The scaling gap: From adoption to embedding

---

Most marketing teams have reached the same crossroads. Individual marketers are saving time with AI tools. Quick wins are being built. But the practice lives with one or two people, depends on their individual prompting knowledge and disappears when they're out of office. That's adoption. Embedding looks different.

## **What adoption looks like:**

One marketer builds a landing page in 10 minutes. The rest of the team still waits days for developer support. Prompts aren't documented. Builds aren't reproducible. Wins aren't scalable.

## **What embedding looks like:**

The team has shared prompt libraries. Builds go through a lightweight review process. New team members can pick up and extend existing workflows. AI-generated code is maintained, not just deployed and forgotten. The organizations seeing results are the ones that restructure how work gets done.

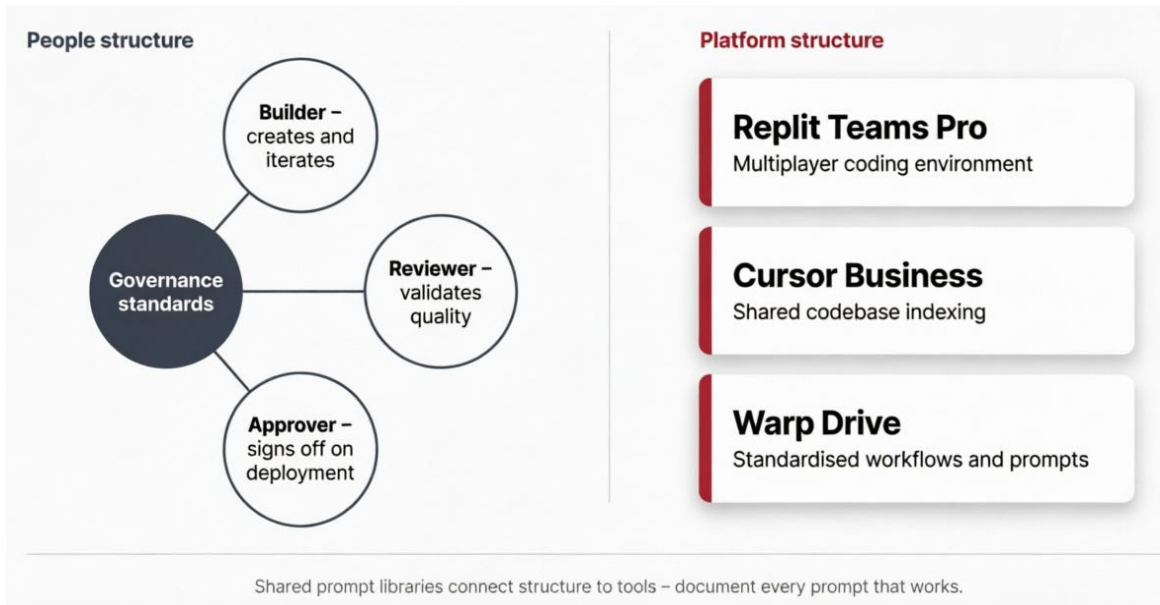
The path to embedding requires three shifts: from individual experimentation to standardized practice, from shadow AI to governed deployment and from personal productivity to team capability. This playbook maps each one.

# Building your collaborative engine

Scaling vibe coding requires both the right structure and the right tools. Structure without tools creates bureaucracy. Tools without structure create inconsistency. Shared prompt libraries are what connect the two. Document every prompt that produces reliable output, version it and make it available to the whole team.

## The people structure: Hub and spoke

Establish a small group of people responsible for governance standards, tool access and shared documentation paired with a community of practice where line-of-business marketers share what they're building, what's working and what's failing.



# Your governance framework: Guardrails for speed

Governance isn't a blocker. Done right, it's what allows your team to build faster with confidence. Follow a three-tier review protocol.

## Tier 1 – Build independently

Landing pages / Email templates / Internal dashboards / Analytics scripts / Prototypes

Go

## Tier 2 – Requires review

Personal data / API integrations / Email automation / CRM modifications / Automated decisions

Review

## Tier 3 – Mandatory sign-off


Production integrations / Security scripts / Payment processing / CASL or PIPEDA workflows

Stop

### Before deploying, ask:

*Does this collect personal data? / Connect to production systems? / Send communications? / Make automated decisions?*

**Yes to any = involve appropriate teams first.**



---

When the  
problem gets  
complex,  
upgrade the  
tool.

# Choosing the right AI for the task

The major LLMs operate at broadly similar reasoning levels. The right choice depends on which fits your workflow, your data environment and the specific task in front of you. The practical distinction is fit.

## Choosing the right AI for the task

The best tool is often the one your organisation already has licensed.

ChatGPT	Claude	Gemini	GitHub Copilot
<b>Best for</b> Fast iterative tasks, one-off builds, generalist marketing use	<b>Best for</b> Extended reasoning, multi-source data analysis, complex conditional logic	<b>Best for</b> Teams running campaigns through Google Ads, GA4 and Google Workspace	<b>Best for</b> Developer collaboration and teams in Microsoft 365 environments
<b>Strength</b> Widest plugin ecosystem, most familiar interface, versatile beyond coding	<b>Strength</b> Strong analytical coherence, long context windows, multi-step conversations	<b>Strength</b> Native Google ecosystem integration, real-time data access within Workspace	<b>Strength</b> IDE integration, widely enterprise-licensed, strong for structured code tasks
<b>Access</b> Widely adopted, free and paid tiers available	<b>Access</b> Subscription required, Pro plan available	<b>Access</b> Check your Google Workspace plan	<b>Access</b> Check your Microsoft 365 licence

**The practical rule:** use what your team already has. If quality issues persist on a specific task, try a different model.

### When to try a different model:

- Current tool producing inconsistent outputs on complex logic – try Claude or Gemini;
- Need native Google Workspace integration – try Gemini;
- Need the widest plugin and integration ecosystem – try ChatGPT;
- Running extended multi-source analytical work – try Claude; and
- Cost is a constraint – compare free tiers across all three before committing.

# Advanced capability: A/B testing at scale

---

**Before:** Your team runs single-variable A/B tests manually, one hypothesis at a time. Statistical significance takes weeks and results describe what happened, not why.

**After vibe coding (prompt for advanced multi-source A/B testing):**

"Using the three attached data sources (GA4 export, CRM segment data, ad platform performance), generate a structured A/B testing hypothesis for our Q3 lead generation campaign. Identify the highest-impact variable to test based on cross-source patterns. Design a synthetic control methodology to measure lift in the absence of a clean control group. Calculate the minimum sample size required for 95 per cent statistical confidence. Output a testing brief including hypothesis, variant specifications, success metrics and a week-by-week measurement plan."

**What Claude does differently:** Causal AI decision intelligence enables Claude to identify true cause-and-effect relationships rather than surface-level correlations. Synthetic control methods allow measurement of campaign impact when a clean control group is impossible to maintain such as during a major brand campaign running across all channels simultaneously.

How does Claude connect to your? MCP servers connect AI tools directly to live marketing data sources, your CRM, ad platforms and analytics environments – as a real-time data bridge. This enables zero-export analysis where Claude reasons directly within your data environment, removing the security risk of exporting sensitive customer data to an external interface.

Self-improving test systems build a memory of every hypothesis tested. If a past test showed emotional copy underperformed for a specific segment, Claude applies that learning automatically to future experiment design. Multi-source hypothesis generation in minutes. Causal inference replaces correlation. Testing velocity increases. Learning compounds across every experiment.

# Enhanced CRM workflows: Automating CASL compliance

---

Manual consent management fails at scale. Once your CRM reaches a few thousand active contacts, tracking express versus implied consent windows manually creates both compliance exposure and operational drag. The fix is automation built into your CRM workflows before you scale, not retrofitted after.

## **Build these four automated workflows in your CRM before scaling:**

- Flag express versus implied consent status per contact at point of entry;
- Automatically suppress contacts at the 6-month implied consent window;
- Hard-remove contacts at the 2-year window unless renewed express consent is obtained; and
- Log every consent interaction with timestamp and source for audit purposes.

When configuring agentic CRM workflows in platforms like Salesforce or HubSpot, set explicit guardrails to prevent AI agents from making unauthorized decisions such as sending unapproved discounts, accessing data outside their defined scope or triggering communications to suppressed contacts.

# Building resilience: When to iterate vs escalate

At scale, your team needs to have a clear protocol for what to do when it fails and learn from failures to build better. Set your own trigger thresholds based on your organization's size and risk tolerance – what constitutes a high-stakes build will look very different at a startup versus an enterprise. Define your thresholds before you scale, not after your first incident. Common factors to consider: volume of customer records affected, ad spend at risk and whether the build touches a regulated workflow.

## Your code isn't working. Which path do you take?

### Iterate independently

- Error is visible and describable
- Fix involves logic or formatting
- 2-4 iterations progressing
- Internal, non-critical build



### Escalate to MarOps

- Impacts over \$5,000 in ad spend
- Affects more than 10,000 records
- 4+ iterations unresolved
- Unexpected staging behaviour



### Escalate to IT

- Server or permission errors
- Fails consistently in production
- Regulated data or API auth
- CASL or PIPEDA implication

Start left. Escalate right only when criteria are met.

Document every failure. Note the error, the build context, what was tried and how it was resolved. Leading teams in 2026 use agentic observability supervisor agents running new scripts in sandboxed environments against historical data before deployment.

# Measuring what matters: Operational metrics

---

Only 19 per cent of marketers track AI-specific KPIs despite increasing adoption. Without measurement, embedding looks identical to adoption from the outside. These four metrics make the difference visible to your team, your stakeholders and your CMO. Establish a baseline before you scale. Improvement without a baseline is just a story.

## Team adoption rate



% of team actively using vibe coding weekly. Track active users not licence holders. Under 30% means you're still in adoption not embedding.

↑ Target: rising

## Speed to market



Days from campaign brief to deployed asset. Compare every AI-assisted build to your pre-vibe-coding baseline. Routine builds should move from days to minutes.

↓ Target: declining

## Campaign cost reduction



Cost difference between AI-assisted builds and traditional production. Include reduced developer dependency and agency scope in your calculation.

↓ Target: declining

## Workflows automated



Count of manual recurring processes replaced by AI-generated automation. Track cumulative quarter over quartet. Growth here is the clearest signal of embedding.

↑ Target: rising

# Measuring what matters: Governance metrics

Operational metrics tell you how fast the team is moving. Governance metrics tell you whether the practice is becoming more disciplined as it scales. These two sit at the intersection of marketing performance and organisational maturity. Share them with MarOps leads, IT partners and senior leadership as evidence that embedding is working, not just growing.

*These two metrics tell you if the team is getting better, not just bigger.*

## Complexity-adjusted throughput (CAT)

### CAT

How many custom marketing integrations can your team maintain without increasing headcount? This is your organisational leverage score. A rising CAT without growing headcount means vibe coding is creating genuine capacity not just speed. Track quarter over quarter and present alongside headcount data.

↑ Target: rising

## Escalation rate



What percentage of builds require Tier 2 or Tier 3 review? A declining rate signals growing team competency and stronger upfront governance. Break down by tier to identify where the team needs the most support. New team members typically start with higher escalation rates that decline as prompting skills develop.

↓ Target: declining

## Review cadence

### MONTHLY

- Team adoption rate
- Speed to market

### QUARTERLY

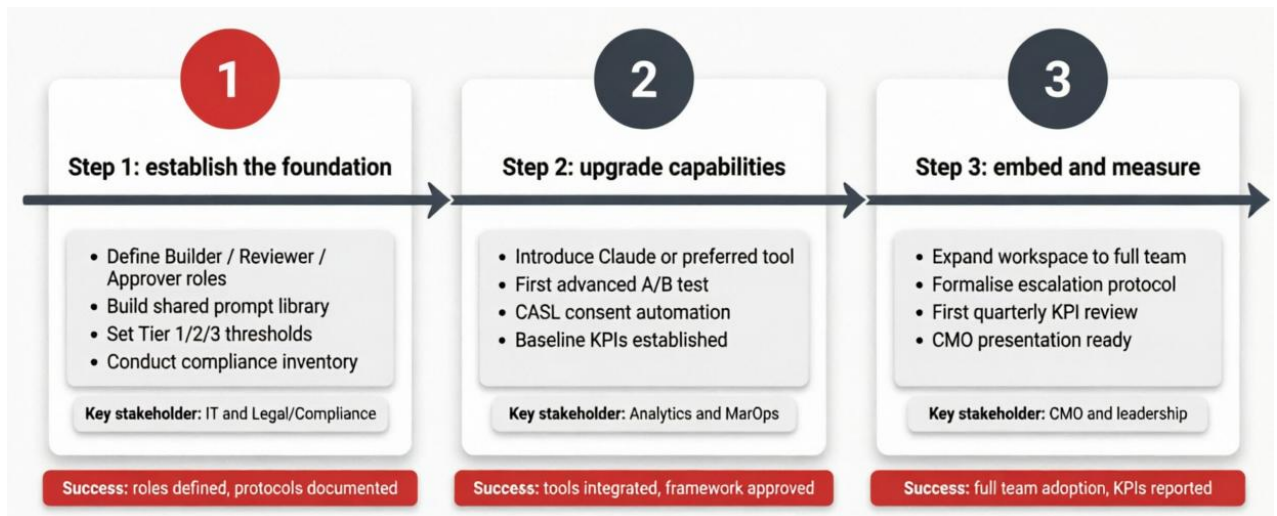
- Campaign cost reduction
- Workflows automated
- CAT
- Escalation rate

Present alongside campaign performance metrics with the same rigour.

# 12

# Your 3-steps caling action plan

Each step builds directly on the last. Move at the pace that reflects your organization's readiness. Don't advance to step 2 until step 1 success criteria are met.



# Recommended reading and references

---

## **CMA resources**

- [CMA Guide on AI for Marketers](#)
- [Setting the Stage on Artificial Intelligence: A CMA Primer on AI for Marketers](#)
- [CMA Accountability Checklists for AI in Marketing](#)
- [CMA Mastery Series: AI Playbooks](#)
- [CMA Generative AI Readiness Survey](#)
- [Canadian Marketing Code of Ethics and Standards](#)

## **External resources**

- ShareDuo: [Claude Artifacts Complete Guide 2026](#)
- Improvado: [Model Context Protocol \(MCP\) for Marketing Analytics](#)
- McKinsey: [Where AI will create value - and where it won't](#)

# The CMA

---

This playbook is developed with guidance by the [CMA AI Committee](#) and is part of the CMA's comprehensive AI initiative designed to empower Canadian marketers with the knowledge, skills, and ethical frameworks needed to implement AI responsibly and effectively.

The CMA is the voice of marketing in Canada, and our purpose is to champion marketing's powerful impact. We are the catalyst to help Canada's marketers thrive today, while building the marketing mindset and environment of tomorrow.

We provide opportunities for our members from coast to coast to develop professionally, to contribute to marketing thought leadership, to build strong networks, and to strengthen the regulatory climate for business success. Our Chartered Marketer (CM) designation signifies that recipients are highly qualified and up to date with best practices, as reflected in the Canadian Marketing Code of Ethics and Standards. We represent virtually all of Canada's major business sectors, and all marketing disciplines, channels and technologies. We advocate with government stakeholders, while also providing Canadian consumers with information to help them better understand their rights and obligations.

For more information, visit [thecma.ca](https://thecma.ca).



---

Build  
together.  
Scale further.