

Becoming AI Native.

A four-page map of what to learn, who to follow, and what *AI-native* actually looks like at thirty days. Built for executives, mid-level managers, and senior engineers reading the same field at three different altitudes.

3

TRACKS

4

CREATORS

12

CURATED VIDEOS

30

DAY RAMP

INSIDE THIS GUIDE

CLICK ANY ROW · JUMPS IN PDF

- 01 Three tracks. Pick the one that matches your role.** p. 02
Executive Management · Mid-Level Manager · Senior Engineer
- 02 Four creators worth a sustained follow.** p. 03
The strategist · the operator · the agent engineer · the principal AI engineer
- 03 The 30-day starter plan.** p. 03
Watch & frame → try one thing → plug in → teach it back
- 04 Cross-track mindset.** p. 03
Five principles paired with five signals worth watching in 2026
- 05 Best videos, ranked by category.** p. 04
Twelve hand-picked, current 2026 videos across all four channels. Every title is clickable.

"Model capability is no longer the bottleneck. The harness around it is. This guide is a four-week ramp to think that way."

THE THESIS

01 Three tracks. Pick the one that matches your role.

Cross-pollinate after week 4.

TRACK A · ~2 HRS / WEEK

Executive Management

Set the narrative, fund the bets, hire AI-first leaders. You don't need to write prompts. You need to spot which ones are working.

30-DAY OUTCOMES

- Can explain LLMs vs. agents vs. *harnesses* to a board without flinching
- Holds a defensible POV on build vs. buy, including which parts of the harness belong in-house
- Names the top three risks (exfiltration, lock-in, talent drain) for your firm specifically

90-DAY OUTCOMES

- Funded one measurable initiative with a written kill criterion
- Replaced one personal strategy ritual with an AI-assisted version

Primary follows Nate · Mark

TRACK B · ~4 HRS / WEEK

Mid-Level Manager

Choose the tools, ship the pilots, raise the floor of your team. The fluency multiplier. You set the cultural default.

30-DAY OUTCOMES

- Ships one team workflow that uses AI end-to-end (notes → actions, drafts, screens)
- Evaluated 2-3 tools against a written rubric, not a vibes-based shortlist
- Can spot weak prompting in a teammate's output and coach a fix

90-DAY OUTCOMES

- Half the team has a personal AI workflow they would refuse to give up
- Ships every AI pilot with a written eval, not vibes, for "is this actually helping?"

Primary follows Mark · Nate · Dan

TRACK C · ~6-8 HRS / WEEK + HANDS-ON

Senior Engineer

Build the systems, set the patterns, mentor the next layer. Code is your medium. The harness is your new abstraction.

30-DAY OUTCOMES

- Built and deployed at least one agent with real tool-calling and sub-agents
- Has wired up at least one MCP server (table stakes in 2026, not aspirational)
- Uses parallel delegation and task queues as a default working pattern

90-DAY OUTCOMES

- Treats the *harness* (skills, hooks, sub-agent routing, context windowing, evals) as a first-class engineering surface, not just prompts
- Authored an internal AI-engineering playbook the team actually uses
- Debugs agent failures by reading traces, not guessing

Primary follows Cole · Dan

02 Four creators worth a sustained follow.

Click any name. Links resolve in your browser.

THE STRATEGIST

01

Nate B Jones
youtube.com/@NateBJones ↗
Daily takes on the AI market and how product shifts land inside companies. Unusually good at separating signal from launch theater. Watch for the framing.
Start with: his pieces on Claude Code, Cowork, and where the harness is heading.

THE OPERATOR

02

Mark Kashef
youtube.com/@Mark_Kashef ↗
Workflow-first AI for people who don't write code. Mark shows the exact prompts and stitches behind real productivity wins. The kind a manager hands to a team on Monday.
Start with: his MCP-for-beginners primer and the "2026 AI Playbook."

THE AGENT ENGINEER

03

Cole Medin
youtube.com/@ColeMedin ↗
Full-stack agent builds: Pydantic AI, LangGraph, n8n, RAG, MCP servers, local models. Code-heavy and aimed at engineers shipping agentic systems in production, not demoware.
Start with: his AI Agents Masterclass playlist and MCP walkthrough.

THE PRINCIPAL AI ENGINEER

04

IndyDevDan
youtube.com/@indydevdan ↗
Patterns for living inside Claude Code: sub-agents, hooks, task systems, prompt engineering as a discipline. The closest thing to a senior peer reviewing your agentic workflow.
Start with: his task-system, hooks, and Agent Teams videos, all on page 4.

03 The 30-day starter plan.

Same shape for every track. Adjust the verbs.

WEEK 01

Watch & frame.
Pick your track. One or two videos a night from the primary creator. Write down three things you want to be able to do by day 30.

WEEK 02

Try one thing.
Build the smallest working version of one outcome. A prompt library, a Zapier+AI flow, a single agent. Cap at two hours. Feel it, don't finish it.

WEEK 03

Plug in.
Add an MCP server, a sub-agent, or someone else's workflow. Replace one recurring 30-minute task with an AI-first version that runs without you.

WEEK 04

Teach it back.
Show your work to one peer. Your team, a code review, a fellow exec. Teaching is the test that tells you whether you actually got it.

04 Cross-track mindset.

Principles you adopt · signals you watch.

04A · ADOPT THESE

Five principles of AI-native work.

01

Reach for an agent before a search bar.

02

Default to delegation. Humans review, models draft.

03

Write prompts like specs, not wishes.

04

Read traces, not vibes, when something breaks.

05

Trade the 30-minute task for the 30-day capability.

04B · WATCH THESE IN 2026

Five signals reshaping the field right now.

01

Harness engineering. The system around the model now defines reliability. The model alone doesn't.

02

Spec-driven development. Written specs become the primary artifact. Agents implement.

03

Long-horizon autonomous agents. Work that spans hours, days, and multiple context windows.

04

Evals as first-class engineering. Observability plus offline plus online evals beat vibes-based reviews.

05

Task queues replacing chat. The chat UI is sunseting. The queue is the new interface.

05 Best videos, ranked by category.

Validated May 2026 · every title links to YouTube.

EXECUTIVE LENS

01

Read the market.

For leaders who need a credible POV on where the field is moving.

① Task Queues Are Replacing Chat Interfaces (Claude Cowork Demo)

NATE B JONES · 32 MIN · JAN 2026

The clearest "chat is sunseting, task queues are the new UI" argument out there. Show this to any executive still living inside a chat window.

② Anthropic Didn't Build a New Browser. They Did Something Smarter.

NATE B JONES · STRATEGY READ

A clinic in how to read AI product moves for what they actually signal, not what the launch deck says.

MANAGER · OPERATOR

02

Operator playbooks.

For people running teams and rolling out AI without writing code.

① The 2026 AI Playbook: Focus, Scale, or Die

MARK KASHEF · STRATEGIC FRAMING

A practitioner's "what to actually do this year" framing for managers who need to pick three bets and ship them.

② Model Context Protocol (MCP) Explained for Beginners

MARK KASHEF · PRIMER · FLIGHT-BOOKING DEMO

The MCP primer non-engineers can actually follow. Hand to any manager who has heard "MCP" twelve times and never wanted to ask.

CROSS-CUTTING · THE FOUNDATION

03

Harness engineering.

Everyone watches these. The new mental model of how AI systems actually work.

① Your AI Stack Isn't Ready for Claude Mythos

NATE B JONES · THE HARNESS THESIS

The "model is no longer the bottleneck, the harness is" argument, framed for execs and ICs alike.

② Claude Code Task System: ANTI-HYPE Agentic Coding (Advanced)

INDYDEV DAN · PATTERN IN PRACTICE

What a real harness pattern looks like hands-on. The task-queue model, with code, not slides.

ENGINEER · ORCHESTRATION

04

Sub-agents & multi-agent.

For engineers ready to stop building "one big agent" and start orchestrating teams.

① One Agent Is NOT ENOUGH: Agentic Coding BEYOND Claude Code

INDYDEV DAN · THE CEILING-OF-ONE THESIS

Why every senior engineer hits a wall with single-agent setups, and what the parallel pattern unlocks.

② Claude Code's Agent Teams Are Insane. Multiple Agents Coding in Real Time.

INDYDEV DAN · AGENT TEAMS WALKTHROUGH

What Anthropic's Agent Teams actually unlocks in practice. The first reference video on production multi-agent coding.

ENGINEER · COMPOSITION

05

MCP, hooks & skills.

For engineers composing the harness. Tools, guardrails, and reusable skills.

① The MCP Integration EVERYONE Is Sleeping On

COLE MEDIN · MCP + CUSTOM AGENTS

How to actually wire MCP into a custom agent properly. Cole's reference walkthrough for production use.

② I'm HOOKED on Claude Code Hooks: Advanced Agentic Coding

INDYDEV DAN · HOOKS AS GUARDRAILS

Hooks treated as a first-class harness primitive. Guardrails, not afterthoughts. Hooks save you from rogue agents.

ENGINEER · PRODUCTION

06

Production agent builds.

For engineers shipping real systems. Pydantic AI, RAG, autonomous builders.

① AI Agents Masterclass (full playlist)

COLE MEDIN · END-TO-END BUILDS

Cole's most-cited series. Full agent builds across n8n, Pydantic AI, and the OpenAI Agents SDK. The 101 to 401 path.

② I'm Building an AI Dark Factory That Ships Its Own Code

COLE MEDIN · PUBLIC EXPERIMENT · APR 2026

The state of the art in agent-built systems, narrated as a public experiment. Worth watching for where the field is going next.

"The competitive edge has shifted to everything around the model. Orchestration, feedback loops, sandboxes, tooling. That is the work."

SLIPSPACE / AI

END · V1 · MAY 2026

Published by **SlipSpace AI Technologies LLC**. We build custom agentic AI for regulated workflows. Tribal councils, federal contractors, healthcare compliance, legal filings. The Harness pattern runs inference on customer-owned hardware behind a private VPN the customer owns.

[SLIPSPACE.AI/VIDEOS](https://slipspace.ai/videos)[SLIPSPACE.AI](https://slipspace.ai) · [BRIEF US](#)[TIKTOK · @SLIPSPACE.AI](https://tiktok.com/@slipspace.ai)