STEP-BY-STEP GUIDE

Build a 9-Agent Viral Content System

Research, ideas, hooks and scripts on autopilot.

Set up a team of 9 AI agents inside Claude Code that researches what is going viral, writes and scores hooks, and hands you a finished short-form script.
Written for creators with zero coding experience.

Intermediate · 2 hours setup · Claude Code · Claude Pro + free APIs

@learnwithseb

One system at a time. Built, not theorised.

Contents

01 The Game Plan
02 What You Need
03 How the 9 Agents Fit Together
04 Phase 1: Research (Agents 1 and 2)
05 Phase 2: Ideation (Agents 3 and 4)
06 Phase 3: Writing (Agents 5, 6 and 7)
07 Phase 4 and 5: Quality Check and Delivery
08 Create Your Project and Rulebook
09 Build the Dashboard
10 Add Your Keys and Your Voice
11 Fill Your Reference Library
12 Run Your First Brief
13 Tune the Agents and Close the Loop
14 Validation
15 Keep Building
16 Quick Reference

The Game Plan

You type one short brief. A team of 9 AI agents does the rest: it researches what is working on YouTube, TikTok and Instagram right now, turns that into ideas, writes hooks, writes the script and checks everything before it reaches you.

The secret is not one clever prompt. It is **pairs**. Every agent that creates something is followed by an agent that judges it. If the work doesn't pass, it goes back for a rewrite. Nothing moves forward until it earns its place.

Who this is for

- Creators who post short-form video and are tired of guessing what to make
- Business owners who want content built on real data, not hunches
- Anyone curious about how AI agents actually work together in practice

What you end up with

- A local dashboard (a web page on your own computer) that shows each agent working in real time
- A finished script with timestamps, an alternate hook, title options, thumbnail text and a production checklist
- A library of every idea and hook the system has ever generated, with scores
- A learning file that gets smarter each time you log how a video actually performed

NOTE

This guide is based on the content pipeline I built and use for LearnWithSeb. There is no video to follow. Everything here comes from the system's own master prompt and the settings for each agent, rewritten in plain English so you can build your own version.

HONEST TAKE

No system guarantees a viral video. What this does is stack the odds: it only builds on ideas that have already been proven to travel, and it refuses to hand you weak hooks. You still film, edit and post.

What You Need

The only thing you must pay for is Claude. Everything else is free or optional. Start with the free tools and add paid ones later if you want richer TikTok and Instagram data.

Tool	What it does	Cost
Claude Pro	Runs Claude Code, the engine behind all 9 agents	About $20 a month
VS Code	A free editor where you open your project and chat with Claude Code	Free
Node.js	Lets your computer run the dashboard	Free
YouTube Data API	Pulls YouTube view counts and channel sizes	Free
ScrapeCreators	TikTok and Instagram data	Optional, pay as you go
Apify	Grabs video transcripts for the very top performers	Optional, free monthly credit

Tool	What it does	Cost
Notion	Sends finished scripts to a Notion database	Optional, free

Three words you will see a lot

- **Agent:** one AI worker with one job and its own instructions. Think of it as a team member with a job description.
- **API key:** a password that lets a tool fetch data from a service like YouTube. You paste it in once and forget about it.
- **CLAUDE.md:** a text file in your project folder that Claude Code reads first. It is the rulebook for the whole system.

BUDGET TIP

Set a monthly spending cap on any paid API before you connect it. The system is built to use free sources first and only spend credits on transcripts of the top 3 to 5 videos.

How the 9 Agents Fit Together

The agents work in five phases, one after another. Each agent has a name, a persona (a classical philosopher, which makes them easy to remember) and a single job.

Phase	Agent	Job in one line
Research	1. Trend Scanner (The Watcher)	Finds what is going viral right now, with numbers
Research	2. Research Scorer (The Logician)	Checks the research is real, then summarises it
Ideation	3. Idea Engine (The Visionary)	Turns the research into 10 packaged ideas
Ideation	4. Idea Scorer (The Cynic)	Scores the top 3 and picks one winner
Writing	5. Hook Writer (The Orator)	Writes 8 opening hooks for the winning idea
Writing	6. Hook Scorer (The Judge)	Scores every hook, picks a primary and a backup

Phase	Agent	Job in one line
Writing	7. Script Writer (The Playwright)	Writes the full script around the best hook
Quality	8. Script Scorer (The Examiner)	The final check before anything reaches you
Delivery	9. Final Assembly (The Emperor)	Packages it all and records what was learned

The rules every agent follows

- ✓ Nothing moves forward until it passes its scorer (every dimension scores 7 or more out of 10)
- ✓ A failed piece goes back to its writer with specific feedback. Maximum 3 attempts, then it is flagged for you to review
- ✓ Your voice guide is law. Every writing agent reads it before writing a single word
- ✓ Your reference examples (hooks, scripts, calls to action) are studied, not ignored
- ✓ A missing optional API key never stops the run. The system works around it and tells you
- × Scorers are not allowed to give everything a high mark. Each gate has an anti-grade-inflation rule

WHY PAIRS WORK

An AI asked to write and judge its own work will usually say it is great. Splitting creation and judgement into separate agents, each with strict scoring rules, is what stops the quality drifting.

Phase 1: Research (Agents 1 and 2)

Agent 1: Trend Scanner (The Watcher)

This agent finds what is working for your topic **right now**. Its big idea is the **outlier multiplier**: a video's views divided by that channel's usual views. A video with 300K views on a channel that normally gets 20K is a 15x outlier. That proves the idea travels. A huge channel getting big views proves nothing.

- Starts with free sources: a sweep of Reddit, X, YouTube and the web for what people are saying
- Searches YouTube for recent videos and ranks them by outlier multiplier, not raw views
- Gives extra weight to the last 7 days. A wave that broke two months ago is not a wave
- Captures the exact first two sentences and the first frame of the top 3 to 5 videos
- Pulls TikTok and Instagram numbers, including saves and shares per view, if you have a key

- Mines comments and sorts real phrases into frustrations, desires, objections and identity statements

PROMPT: COPY & PASTE INTO CLAUDE CODE

You are a short-form content research analyst. Find what is going viral RIGHT NOW for this topic and prove why with hard numbers. Rank videos by outlier multiplier (video views divided by the channel's median views), not raw views. Weight the last 7 days most heavily. Capture the verbatim first two sentences and first frame of the top 3 to 5 outliers. Mine comments for real audience language sorted into frustrations, desires, objections and identity statements. Output a Research Brief with: top 10 outliers, best hooks found, winning title templates, audience language, top 3 pain points, format recommendation and gaps. Every claim needs a number.

Agent 2: Research Scorer (The Logician)

This agent checks the research is strong enough to build on. It scores five things out of 10: data quality, outlier evidence, freshness, hook strength and audience insight. It must also quote three real data points from the brief. If it can't, the research is treated as made up.

When the research passes, this agent squeezes it into a **Research Digest** of about 400 words. Every later agent reads the digest instead of the full research. That keeps them focused and saves a lot of processing.

PLAIN ENGLISH

Agent 2 is your fact-checker. It stops the system confidently writing a script based on numbers that don't exist.

Phase 2: Ideation (Agents 3 and 4)

Agent 3: Idea Engine (The Visionary)

This agent generates **10 ideas**, but not by random brainstorming. Each idea must come from one of six named methods:

Method	What it means
Outlier remix	Take a proven viral video and change one thing: the audience, the tool or the stakes. At least 3 of the 10 ideas use this

Method	What it means
Contrarian	Disagree with what everyone else is saying
Identity ladder	"There are 3 types of X. Which one are you?" Best for videos of 60 seconds or more
Demo with visible payoff	A result you can see in a single frame
Curiosity gap	A question the viewer can't answer and can't leave unanswered
Stakes and urgency	What changed this week and why it matters now

Every idea also needs an **emotional driver** (surprise, recognition, aspiration, FOMO, righteous anger or identity), a **borrowed proof** (the real outlier video that shows the pattern works) and **packaging** (a title plus what the first frame looks like). No emotion, no proof or no packaging means the idea is dropped.

Agent 4: Idea Scorer (The Cynic)

This agent scores the top 3 ideas out of 50, across five dimensions:

- **Angle clarity:** can the whole argument be said in one sentence?
- **Evidence base:** is there real data to back it up?
- **Creator fit:** would this feel natural for you to make?
- **Packaging power:** can you picture the title and first frame stopping a stranger's thumb?
- **Share trigger:** who exactly would send this to whom, and why?

Scores must spread by at least 4 points so the gate actually chooses. It then picks one winner and predicts how it will perform: **Base**, **Solid** or **Outlier-potential**. That prediction is checked against real results later.

Phase 3: Writing (Agents 5, 6 and 7)

Agent 5: Hook Writer (The Orator)

The hook is the first 3 seconds. If it fails, nothing else matters. This agent writes **8 hooks**, one for each style: bold claim, pain call-out, contrarian, reveal, story drop, identity ladder, costly mistake and receipts (opening on the result).

Each hook is a set of three parts that work together:

- **Spoken line:** 8 to 12 words. Any longer and the viewer has already scrolled
- **First frame:** what is on screen at the very first second

- **On-screen text:** 6 words maximum, adding something the spoken line doesn't say

BANNED OPENINGS

No hook may open with "Did you know", a greeting, or the topic's name. No yes/no questions that can be answered with "no". No vague phrases like "many people" when the research has an exact number.

Agent 6: Hook Scorer (The Judge)

This agent scores every hook out of 50 on scroll-stop power, specificity, emotional charge, voice match and platform fit. It uses written examples of what a 4 and a 9 look like, so a score means the same thing on every run. At least two hooks must score below 35. If everything passes, the judge is being too soft.

It also runs the **stranger test**: would someone who has never heard of you care? Then it picks a **primary** hook and a **backup** that uses a different style and a different emotion, so you have a genuinely different second option.

Agent 7: Script Writer (The Playwright)

This agent writes the full script as a retention plan, not a fill-in-the-blanks template.

- ✓ Word budget based on speaking speed: about 2.3 to 2.5 words per second, so 140 to 150 words for 60 seconds
- ✓ A re-hook every 10 to 12 seconds (a question, a twist, a number tease) to keep people watching
- ✓ The best payoff is held back until the final third
- ✓ One honest, balanced moment. Credibility keeps viewers
- ✓ Call to action at around 85% of the way through, not as the last words
- ✓ A loopable ending that connects back to the first line, so the video rewatches smoothly

It writes a first draft, diagnoses where attention would drift, then does a read-aloud pass to cut anything that trips off the tongue. It also adds on-screen text for every beat, because most people start watching on mute.

Phase 4 and 5: Quality Check and Delivery

Agent 8: Script Scorer (The Examiner)

The last gate before the script reaches you. It scores seven dimensions: voice match, retention, value, shareability, call-to-action strength, filler and **hook-payoff integrity**. That last one asks: does the script actually deliver what the hook promised? If not, the video is bait and it fails,

whatever the other scores say.

- **Beat map:** walks the script in 10-second chunks and marks each as fresh or coasting. Two coasting chunks in a row is a fail
- **Pacing maths:** word count divided by seconds must land between 2.1 and 2.7. Outside that, the script can't be spoken naturally
- **Send test:** would you send this to one specific person? If it can't name who, shareability is capped
- **Spoken flow check:** flags any sentence too long for one breath

Agent 9: Final Assembly (The Emperor)

This agent packages everything so you can go straight to filming. The finished output includes:

- The final script with timestamps, visual cues and on-screen text
- The backup hook as an alternative opening
- A production checklist of every screen recording, B-roll clip and visual you need, with animations marked as self-producible
- A publishing kit: captions for each platform, 3 to 5 niche hashtags (never #fyp), a first-frame spec and a suggested posting window
- 3 ranked title options and thumbnail text ideas
- The key research insight, so you understand why the script takes this angle

The learning loop

Agent 9 also writes a short entry into a learnings file: which idea method, hook style and scores were used, plus the predicted tier. After you post, tell Claude Code to log the results. It fetches the real numbers, compares them with the prediction and notes what it got right or wrong. Agents 1, 3, 4 and 5 read this file at the start of every run, so the system learns what works for **your** audience.

STEP 01

Create Your Project and Rulebook

Install VS Code and Node.js (choose the LTS version). Create a new folder on your computer called **content-pipeline**. Open VS Code, go to **File** → **Open Folder** and select it. Click the **Claude** icon in the left sidebar and sign in with your Claude Pro account.

Now ask Claude Code to write the rulebook for you. This prompt creates a **CLAUDE.md** file describing the whole system.

PROMPT: COPY & PASTE INTO CLAUDE CODE

Create a CLAUDE.md file in this folder for a 9-agent short-form content pipeline. I give it a Content Brief (Topic, Platform, Length, optional Angle) and it produces a research-backed script for TikTok, YouTube Shorts and Instagram Reels.

The agents run in order: 1 Trend Scanner, 2 Research Scorer, 3 Idea Engine, 4 Idea Scorer, 5 Hook Writer, 6 Hook Scorer, 7 Script Writer, 8 Script Scorer, 9 Final Assembly. Every scorer scores out of 10 per dimension and passes only if every dimension is 7 or more. Failed work goes back to its writer with specific feedback, maximum 3 attempts, then flag it for me.

Store each agent's instructions in its own editable file at dashboard/config/agents/agent-1.md to agent-9.md. Those files override anything else. Before every run, load my voice guide (dashboard/config/voice-guide.md), my references (dashboard/references/hooks.md, scripts.md, ctas.md), my API keys (dashboard/config/api-keys.json) and past learnings (dashboard/learnings.md). If an optional API key is missing, work around it and tell me. Save every run's research, ideas, hooks, scores and final script as JSON in dashboard/runs/[date-topic-slug]/.

TIP

Read the CLAUDE.md it creates. It is plain text. If something doesn't match how you work, tell Claude Code in normal words and it will change it.

STEP 02

Build the Dashboard

The dashboard is where you manage settings and watch the agents work. It runs only on your computer, at an address your browser opens like a normal website. Paste this into Claude Code:

PROMPT: COPY & PASTE INTO CLAUDE CODE

Read CLAUDE.md. Then build a local dashboard for this pipeline at http://localhost:3000 using Node.js with a simple dark theme. Create these pages, with a navigation bar on each:

1. Pipeline: the 9 agents as cards showing status (grey waiting, blue running, green passed, red failed, amber resubmitted), score and attempts, updating live during a run. Below it, a collapsible How to Run card, the latest script and a run history.
2. Settings: fields for YouTube, ScrapeCreators, Apify and Notion API keys, saved to dashboard/config/api-keys.json and masked after saving, with a Test Connection button for each. Add an Agent Configuration section listing all 9 agents, where I can expand and edit each agent's spec file.
3. Voice Guide: a large editor saved to dashboard/config/voice-guide.md, pre-filled with a starter template.
4. References: three tabs for Hook, Script and CTA examples, saved to dashboard/references/.
5. Outputs: every finished script as a collapsible list with copy and export buttons, showing spoken text only.
6. Ideas and 7. Hooks: every idea and hook from every run, grouped by run, with their scores.

Also create a .gitignore that excludes api-keys.json and the runs folder, and a start.bat file I can double-click to launch the dashboard. Then start it.

Claude Code will create the files and start the server. Open your browser and go to **localhost:3000**. From now on, double-click **start.bat** to open it.

COMMON MISTAKE

The dashboard only works while its window is open. If the page won't load, double-click start.bat again. Closing that black window switches the dashboard off.

STEP 03

Add Your Keys and Your Voice

Open the **Settings** page. Add your **YouTube Data API** key at minimum. To get one, go to Google Cloud Console, create a project, search for **YouTube Data API v3**, click **Enable**, then **Credentials** → **Create credentials** → **API key**. Paste it in and click **Test connection**.

Add ScrapeCreators and Apify keys later if you want TikTok, Instagram and transcript data.

Next, open the **Voice Guide** page. This is the most important thing you will write. Every writing agent treats it as law. Fill in each heading of the starter template:

- Who you are, in one sentence
- How you sound (casual, calm, excited, honest)
- Words and phrases you actually use
- Words you never use
- How you sign off
- 3 to 5 sentences written exactly the way you talk
- What makes your content different

PROMPT: COPY & PASTE INTO CLAUDE

Here are transcripts of 3 of my best-performing videos. Study how I talk and write me a voice guide with these headings: Who I am, How I sound, Words I use, Words I never use, My sign-off style, Example sentences that sound like me, What makes my content different. Use my real phrases, not a polished version of them.

TIP

Don't have transcripts to hand? Record yourself explaining your niche for 3 minutes on your phone, get it transcribed, and paste that in instead.

STEP 04

Fill Your Reference Library

Open the **References** page. It has three tabs. The agents study these to copy your patterns, so real examples beat perfect ones.

Tab	What to paste	Who uses it
Hook examples	10 to 20 of your best openings, or hooks from creators you admire	Agents 3, 5 and 6
Script examples	2 or 3 full scripts that performed well	Agents 3, 7 and 8
CTA examples	Your go-to calls to action and sign-offs	Agent 7

NOTE

If you use hooks from other creators, the agents learn the pattern, not the words. The ban list and voice match scoring stop them copying lines directly.

STEP 05

Run Your First Brief

Open the Claude Code panel in VS Code and paste a brief. Change the topic to suit your niche.

PROMPT: COPY & PASTE INTO CLAUDE CODE

Run the content pipeline with this brief:

Topic: 5 free AI tools that replace expensive software
Platform: all
Length: 60 seconds

Keep the **Pipeline** page open in your browser. You will see each agent card turn blue while it works and green when it passes. Amber means a scorer sent the work back for a rewrite. That is the system doing its job.

A full run takes about 10 to 20 minutes. When it finishes, open the **Outputs** page for your script. Check the **Ideas** and **Hooks** pages too. The runners-up are often worth filming as well.

COMMON MISTAKE

Don't skip the voice guide and references before your first run. Without them, the scripts sound like generic AI and the voice match scores will keep failing.

STEP 06

Tune the Agents and Close the Loop

After a few runs you will notice patterns. Maybe the hooks are too long, or the scripts are too salesy. Go to **Settings** → **Agent Configuration**, expand the agent and click **Edit agent**. Change its instructions in plain English and save. The next run uses the new version.

Once you have posted a video, come back a few days later and tell Claude Code how it did:

PROMPT: COPY & PASTE INTO CLAUDE CODE

Log results for run [paste the run name from the Outputs page]. Fetch the video's current stats, compare them with my channel's usual views and the predicted tier, and write one sentence about what the prediction got right or wrong.

TIP

Change one agent at a time. If you edit three at once and the scripts improve, you won't know which change did it.

Validation

Run through this list after your first full run.

- ✓ The dashboard opens at localhost:3000 when you double-click start.bat
- ✓ Test connection shows green for your YouTube key
- ✓ All 9 agent cards turn green during a run
- ✓ The research quotes real videos with view counts and outlier multipliers
- ✓ The script reads like you when you say it out loud
- ✓ The Outputs page shows the script, alternate hook, titles and production checklist
- ✓ A new entry appears in the learnings file after each run
- × Research full of round numbers with no sources means Agent 2 isn't doing its spot-check
- × Every hook scoring 45 or more means the Hook Scorer is being too generous

Keep Building

Once the core system runs smoothly, here is where to take it next.

- **Turn on Notion delivery** in Settings so every finished script lands in your content calendar
- **Ask for derivative cuts:** a 30-second version of a winning 60-second script, using the backup hook rather than the primary
- **Generate a script from a runner-up idea:** each idea on the Ideas page has a Generate Script button with a ready-made prompt
- **Add a news check** to the Idea Engine for fast-moving niches like AI, so urgency ideas are genuinely fresh

- **Review your learnings file monthly** and update the agents that keep predicting wrong

> **TIP**
>
> The system is only as good as what you feed it. A sharper voice guide and better reference examples will improve every agent at once.

Quick Reference

Agent	Scores out of	Passes when	Key output
1 Trend Scanner	Not scored	Sent to Agent 2	Research Brief with outlier multipliers
2 Research Scorer	5 dimensions	All 7+ and 3 real data points quoted	400-word Research Digest
3 Idea Engine	Not scored	Sent to Agent 4	10 ideas, each with method, emotion, proof, packaging
4 Idea Scorer	50	One idea scores 7+ on all 5	Winning idea and predicted tier
5 Hook Writer	Not scored	Sent to Agent 6	8 hooks (spoken line, first frame, on-screen text)
6 Hook Scorer	50	All 5 dimensions 7+	Primary and backup hook
7 Script Writer	Not scored	Sent to Agent 8	Timed script with overlays and titles
8 Script Scorer	7 dimensions	All 7+, clean beat map, pacing 2.1 to 2.7 words/sec	Scores and rewrite notes
9 Final Assembly	Not scored	Always runs last	Final script, publishing kit, learnings entry

BuildWithSeb | The 9-Agent Content System

Built by @learnwithseb · One system at a time. Built, not theorised.

© 2026 Learnwithseb. All rights reserved.