

The Content Factory Workbook

One recording in. Weeks of content out. You do the 1% only a human can do — **AI agents run the other 99%.**

DIGIMARCON · DIGITAL MARKETING,
MEDIA & ADVERTISING CONFERENCES

AI Agent Workshop — The Content Factory

The DigiMarCon edition — for attendees across the
circuit

Las Vegas · Minneapolis · Toronto · and every stop where
we meet

With Dennis Yu — keynote speaker on the DigiMarCon
circuit



dennisyu.com/digimarcon

Scan for the agents, the tap-to-copy
prompt, this workbook, and your audit
scorecard.



Dennis Yu on the DigiMarCon stage, 2025 — every keynote is Stage 1: Produce. The agents make the other 19 assets.

Dennis Yu

CEO, BlitzMetrics · Search engine engineer · Creator of the Dollar-a-Day
strategy

Print edition for DigiMarCon attendees
digimarcon.com · the conference circuit

The buttons now click themselves.

A message from Dennis Yu

Welcome to the Content Factory. Whether you're a business owner, a marketing director told to "do more with less," or a freelancer carrying five clients with the bandwidth of one — this workbook teaches the system that makes you **Googleable**. And in 2026, it makes the AI engines describe you correctly, too — because ChatGPT, Perplexity, and Google's AI read the same signals.

I'm a search engine engineer. When I built Yahoo!'s analytics over 20 years ago, my job was to protect search results from people trying to trick them. That's still my job today: protecting you from SEO "experts" who sell fake signals. When Google catches them — and it always does — you're the one who gets penalized.

The Content Factory is four stages — **Produce, Process, Post, Promote** — on a foundation we call **Plumbing**, with a feedback loop we call **MAA** (Metrics → Analysis → Action) running at every stage. None of that has changed. What changed is *who does the work*.

The first edition of this workbook taught your VAs and account managers which buttons to click in Descript and ChatGPT. This edition is different, because the world is different: **the buttons now click themselves**. AI agents — not chatbots you prompt, but persistent workers with job descriptions, memory, and a definition of done — run the factory floor. You bring the raw ingredients only a human can bring: your stories, your job sites, your handshakes. And here's what the best owners I know have figured out: anyone can talk to agents, but it's often a **young adult managing the agents** who makes the factory hum — meet three of them on page 16.

Losing weight isn't one activity — it's diet, exercise, and sleep working together. SEO is the same. It isn't one trick; it's the compound result of running this process, every week, with the machines doing the heavy lifting.

Before you leave this conference you'll have the same agent team I use running on your own Claude account (the free tier works), plus a path to a personal audit showing exactly where Google stands on you today — we run them for the rooms we speak to, and yours is one QR scan away.

Let's get started.

Dennis Yu

CEO, BlitzMetrics · dennisyu.com · 612-707-8045

How to use this workbook

Pages 4–11 are the **strategy** — read them like I'm explaining it to you over coffee. Pages 12–18 are the **proof and the setup**: the audit system, your agent team, the DigiMarCon scrapbook, and the 60-second install. Pages 19–22 are **worksheets** — bring a pen. Everything in here lives digitally at dennisyu.com/digimarcon.

What changed since the first edition — and what never will

THE OLD WORKBOOK SAID (2023–24)	THE AI AGENT EDITION SAYS (2026)
Hire a young adult and a team of VAs to run your factory.	Deploy an agent team — 10 skills running on Claude. One human (you, your kid, your VA) approves and QAs. The org chart got smaller; the output got bigger.
Paste your transcript into ChatGPT and prompt it, revision by revision.	A persistent agent runs the whole pipeline — transcript → canonical article → clips → captions → posts → ads — then grades its own work before you ever see it.
Use Descript to transcribe and cut clips; track assets in a spreadsheet Content Library.	Agents transcribe, cut, caption, and file every asset in your Content Library automatically. The spreadsheet maintains itself.
Our custom GPTs: Jennifer (article grader), Brandon (blog helper), Ethan (mentions), Olivia (Knowledge Panel), Christopher (site auditor), Adrian (page builder).	All seven grew up and moved into one agent running 10 skills — including a QA skill that makes every other skill sharper each run. Meet the team on page 13.
Boost manually; check your dashboards weekly if you remember.	The agent runs MAA on a schedule and brings you three decisions, not thirty charts. You spend dollars only on proven winners.
Post to Twitter.	It's called X now. (Some things change for no good reason. Post there anyway.)


Why agents, not prompts

Prompting is *asking a stranger for a favor* — you get one answer and start over tomorrow. An agent is an *employee*: it has a job description (a skill file), a memory of your business, and a definition of done it loops against until the work passes. The newest models — Anthropic's Mythos-class Claude, the engine we run our audits on — hold entire businesses in their head and don't get tired. When the marginal cost of completeness is near zero, "good enough" is no longer a strategy. Do the whole thing.

THREE RULES THAT DID NOT CHANGE

- 1. Zero times infinity is still zero.** Agents amplify what you feed them. Great raw content gets multiplied; garbage gets multiplied too — into more garbage. Bill Gates said AI is a multiplier of what you already have. Stage 1 is still yours.
- 2. Your brand is what others say about you.** Not what you say about yourself. The factory exists to collect, amplify, and compound other people's words — mentions, reviews, interviews, photos together.
- 3. Google rewards proof and punishes tricks.** E-E-A-T — Experience, Expertise, Authority, Trust — is the rulebook. Authentic, geotagged, story-driven content is timeless. Fake signals eventually get caught. Every Google update since 2012 says the same thing louder.

Five stages. One loop. A clear human/agent split.



↳ MAA — Metrics → Analysis → Action — runs at every stage. Measure, decide, adjust. Forever.

STAGE	THE HUMAN DOES (THE 1%)	THE AGENT DOES (THE 99%)
0 Plumbing	Grants access (Google, site, socials). Once.	Audits the foundation, finds the leaks, writes the fix list, drafts schema and pages.
1 Produce	Lives an interesting professional life on camera: 1-minute videos, geotagged photos, podcasts, handshakes.	Tells you exactly <i>what</i> to record (PAA questions, topic gaps) so no take is wasted.
2 Process	Approves. Adds the story only you know.	Transcribes, writes the canonical article, cuts clips, drafts platform-native posts, grades itself, files everything.
3 Post	Hits publish (or just says "ship it").	Publishes to your site, queues cross-posts, keeps every bio and headshot consistent, maintains schema.
4 Promote	Sets the budget. Says thank you like a human.	Finds boost candidates, runs the \$1/day tests, kills the bottom 90%, harvests and scores mentions.

GCT — the strategy filter

Before any piece of content exists, answer three things: **Goals** (what outcome?), **Content** (what piece earns it?), **Targeting** (who must see it?). Your agent asks these questions for you — refuse to skip them.

The 1% rule

Every entrepreneur needs a factory that handles 99% of the work — editing, distributing, boosting, optimizing — so you spend your 1% on the thing machines can't do: **being the proof**. Robots can't eat tacos with a client or shake hands in the expo hall.

Set up the pipes before you pour the water

In three decades I've watched smart owners make the same basic mistakes: location and service pages that don't link to their Google Business Profile, sites that load slow, no retargeting, data scattered across ten tools that don't talk to each other. Digital Plumbing is connecting all of it so every tool feeds the next — your call tracking feeds your analytics, your analytics feed your ad platforms, and you can finally make decisions from one picture.

You don't have to do this yourself anymore. This used to be a young adult's week of work from our Digital Plumbing Course. Now it's the first thing your agent audits — it produces this checklist, pre-filled with what you're missing, and drafts most of the fixes itself.

The Plumbing checklist

- Google Business Profile created, claimed, and linked from your site
- Google Analytics 4 installed (via Google Tag Manager)
- Search Console verified — and its monthly report actually read
- Conversion tracking: calls, forms, bookings wired as events
- Remarketing audiences built (site visitors, video viewers)
- Site loads fast on a phone; service pages link to GBP
- One consistent name, headshot, and bio on every profile
- Person + Organization schema (JSON-LD) on your site

Why Search Console outshines every SEO tool

It's Google's own data about your site — better than any third-party tool. Every month it shows your top queries, your top pages, and your leaks (one client found 876 visitors a month landing on a 404). Your agent reads this report and turns it into MAA: more of what's working, fix what's broken.

If you want to be Googleable, use Google's tools.

Agent shortcut

After you install the agents (p. 14), say: **"run knowledge-panel-entity-seo on me"** — the entity audit covers most of this page and hands you the fix list, ranked.

KEY TAKEAWAYS

Plumbing is the prerequisite, not the strategy. Do it once, properly. · The agent audits and drafts; a human grants access and approves. · A fast site + claimed GBP + verified Search Console + consistent identity = the foundation every later stage compounds on.

The only stage agents can't do for you

Produce is the first stage — raw, unprocessed proof that you do what you say you do, in the places you say you do it. Photos of Frank fixing the water heater near the ballpark. A one-minute video answering the question every customer asks. The podcast where a peer interviews you. Lunch with your team, geotagged.

This is where E-E-A-T's extra E — **Experience** — lives. The AI doesn't have your stories. It wasn't there when the basement flooded at 2 am. When you tell those stories on camera, Google can recognize content that no machine generated and no competitor can copy. Your phone geotags every photo automatically — device, place, time. That's a signature of trust, and it's free.

Quality matters more than ever, because everything downstream multiplies it. When the original is good, four stages of amplification make it great. When the original is poor, zero multiplied by infinity is still zero.

The one-minute video — 4 components

Hook · 0–3s Earn the next 57 seconds.

Pain / pleasure · 3–15s Name the problem they feel.

Solution · 15–50s Teach like a neighbor, not a brochure.

CTA · 50–60s One next step. Just one.

Start with your **WHY video** — the story of why you do this. Ordinary and relatable beats dramatic.

What to record (your agent makes this list)

- 50 People-Also-Ask questions, answered 1 minute each — about an hour of talking, a quarter's worth of content
- One 15-minute interview or customer story (record one this week)
- Geotagged job-site photos & team moments, weekly
- Shout-outs to other local businesses — elevate others, and Google connects you to your city

Posts that work: broetry + humble service

Every sentence gets its own paragraph — it reads faster on a phone. Hook first. And never "Look, I'm famous!" — make it about *them*: "Most homeowners in your city overpay for this. Here's how to tell..." On LinkedIn, ten engagements in the first hour is the tipping point, even for a "nobody" account.



Dylan Haugen & Dennis Yu on the DigiMarCon Minneapolis stage — the Topic Wheel, live. A talk like this is raw proof: Stage 1 in the wild.

KEY TAKEAWAYS

Raw proof is the feedstock; one hour of honest talking feeds the factory for months. · Geotagged, story-driven, imperfect-on-purpose beats polished and fake. · Agents tell you *what* to record so nothing is wasted — but the face and the handshake are yours.

One recording in. Twenty-plus assets out.

This is the stage that used to eat your VA's week — Descript for transcripts, ChatGPT for drafts, Canva for graphics, a spreadsheet to track it all. Now it's one command: **"use content-factory on this recording."** The agent transcribes, finds the angles, and builds the whole pyramid — pillar content at the top, micro-content beneath, distribution at the base.

What comes off the line from one 15-minute interview

The canonical article

One definitive page on your site that owns the topic — the page Google and AI engines quote. Everything else points here.

6–10 captioned clips

Short verticals cut at the strongest moments, captioned, ready for Reels / Shorts / TikTok.

Platform-native posts

LinkedIn broetry, an X thread, Facebook + Instagram posts, a GBP update, an email blurb, ad variants.

That's 20+ assets, every one linking back to the canonical article — that's how authority concentrates instead of scattering.

The Content Library: your greatest hits

Kids replay *Finding Nemo* until the disc wears out. Your market works the same way. The Library is where every asset, mention, and review lives, scored — so you stop feeding the newspaper-publisher treadmill and put 80% of your effort behind what's already proven. The agent maintains it automatically now; it used to be the most-skipped chore in the factory.

QA: Jennifer grew up

Our article grader used to be a custom GPT you had to remember to use. Now grading is built into the pipeline: the agent grades structure, voice, E-E-A-T elements, and links — A through D — and loops until it earns the A *before* you see the draft. Then a second skill, **recursive-self-improvement-qa**, sharpens the factory itself after every run.

Your voice is a guardrail, not a casualty

The #1 fear in this room: "AI content sounds like a corporate robot." That's prompt-era thinking. Your agent is trained on *your* transcripts, *your* stories, *your* phrases — and instructed to quote you verbatim where it matters. If a draft doesn't sound like you explaining something to a neighbor over coffee, reject it; the agent learns. Generic fluff is a configuration error, not a law of nature.

KEY TAKEAWAYS

Process is now a command, not a job description. · Everything points at ONE canonical article per topic. · The Library compounds your greatest hits; the QA loop compounds the factory itself.

Your entity home, then everywhere else

Posting isn't spraying links. It's building a constellation of **properties** that all point to one center: your **entity home** — a website on your own name. Not a sales page; a *facts* page. Who you are, what you do, your story, your proof, what people say about you, where you've appeared. It's the one URL Google treats as the authoritative source about you — the page a Knowledge Panel gets built from.

Then every platform gets its native version of your content: YouTube for the pillar video, Facebook and Instagram for discovery, LinkedIn for the professional crowd, X for speed, your GBP for local intent — all linking back to the canonical article. On Facebook, boost from a personal profile in professional mode; the reach beats a page several times over.

Three objects per network

Google validates you by triangulating: a **personal profile**, a **public-figure page**, and a **company page** on each network — all using the same name, headshot, and bio, all linking to your entity home. Inconsistency is why panels don't trigger.

The agent enforces consistency — it keeps a master bio and flags every profile that drifts.

The entity-home blueprint

- Hero + authority photo, name as the domain
- Stats bar (years, jobs done, markets, reviews)
- My Story · What I Do (3 cards)
- Featured interview + What People Are Saying
- As Seen On · Connect
- Person schema (JSON-LD) with sameAs links to every profile

Say **"run personal-brand-website-agent on me"** and the agent drafts this whole structure from your proof library.

Don't have a site on your name yet?

That's a Week-1 fix, not a someday project. We build these in a day now — see the worked examples linked at dennisyu.com/wichita. Your name as the domain. Facts, not hype. The factory needs a home to point at.

KEY TAKEAWAYS

One entity home on your name; everything else is a spoke. · Same name, face, and bio on every property — drift kills panels. · Agents publish and cross-post; you approve.

Dollar a Day + the Thank You Machine

The Dollar-a-Day philosophy: nobody — not you, not me, not the algorithm — knows in advance which piece of content will win. So stop guessing and start auditioning. Put \$1/day for 7 days behind each candidate post. That's seven dollars for the truth about a piece of content — the cheapest market research that exists.

The Dollar-a-Day recipe

- | | |
|---------------------|---|
| 1 • Audition | Boost each candidate at \$1/day × 7 days. Organic winners and "lighthouse" content (you with marquee people and brands) go first. |
| 2 • Kill | After 7 days, kill the bottom 90%. No mercy, no sunk-cost stories. Expect failure 90% of the time — there's no penalty for shots on goal. |
| 3 • Scale | Winners get \$30 over the next 30 days. Real unicorns get real budget and become remarketing audiences. |
| 4 • Repeat | It's a system, not a stunt — the same recipe works on Facebook, Instagram, X, YouTube... even Amazon ads to make your book a bestseller. |

Why so little money? Because the dollar isn't the strategy — the **selection** is. You're paying the platform to tell you what your market actually responds to, then concentrating real dollars only on proof. Most businesses do the opposite: big budgets on guesses.

The Thank You Machine

It's not what you say about yourself — it's what others say about you. So engineer gratitude: when someone says something good, thank them publicly and ask, "Can I quote you on that?" Turn the best mentions into banners, listicles, and posts that elevate *them*. Then boost those for a dollar a day. Influence is a currency; this machine mints it.

Agent + human split

Agent: harvests every mention (*positive-mentions-harvester*), scores each on the 30-point authority scale — Who said it, What they said, Where it lives — picks boost candidates, watches the tests, reports the winners.
You: set the budget, approve the spend, and deliver the human thank-you. A gift, a call, a handshake. That part doesn't delegate.

KEY TAKEAWAYS

- \$1/day × 7 buys the truth; kill the bottom 90%; scale the winners. · Lighthouse content (you + people with authority) boosts first.
- Gratitude is a growth channel — automate the harvesting, never the thanking.

Metrics → Analysis → Action

MAA is the heartbeat of the factory — the chapter the old workbook mentioned in passing and never taught. Here's the whole discipline: look at the numbers (**Metrics**), figure out what they're telling you (**Analysis**), then change something (**Action**). Most businesses collect dashboards and change nothing. The loop only counts if it ends in an action.

STAGE	METRICS THAT MATTER	WHAT "GOOD" LOOKS LIKE	WHEN IT'S BAD, CHANGE...
Produce	Raw assets captured per week; % geotagged; PAA questions answered	3+ assets/week, every week — boring consistency	Your calendar, not your standards. Schedule one recording hour weekly.
Process	Assets shipped per recording; article QA grade; days from record → publish	20+ assets per pillar; A-grade; under a week	The skill's guardrails — tighten voice and structure rules, re-run.
Post	Search Console queries & clicks; engagement in first hour; profile consistency	Canonical pages climbing for named queries; 10 engagements in hour one on LinkedIn	One of G, C, or T — usually Targeting (right content, wrong audience).
Promote	Cost per engaged view; winners found per 10 tests; cost per lead on scaled winners	~1 winner per 10 auditions; CPL dropping as winners compound	Nothing — kill and re-audition. The system IS the action.

When a video is "hot," you'll know within the \$7 audition — engagement above your account's norm, comments from strangers, shares. That's your Standard of Excellence: not perfection, but *beats the baseline*. Pin it, repurpose it, scale it.

Agents turned MAA from a meeting into a pulse

The old way: a monthly report nobody read. The new way: your agent checks Search Console, ad results, and post performance on a schedule, applies the table above, and brings you **three decisions** — "kill these two, scale this one, record more on topic X." You decide in five minutes. That's the whole meeting.

KEY TAKEAWAYS

A metric without an action is decoration. · Each stage has its own scoreboard — don't grade Produce on leads or Promote on likes. · Let the agent watch the numbers; you make the three decisions.

What Google (and AI search) actually wants

The Knowledge Graph is Google's brain — not pages linked to pages, but *entities*: you, your company, your city, every event and interview, each one an object with citations. The more your objects connect — to each other and to objects Google already trusts — the more confident Google is about who you are. Confidence becomes a **Knowledge Panel**: that box on the right when someone searches your name. It's verification you can't buy with a blue-check subscription — you earn it by linking your world together.

Find your KG MID (your entity's ID)

- 1 Google your name. Look for results with the small icons — those are Knowledge Graph objects.
- 2 Open the panel, tap the : three dots, choose **Share**.
- 3 The link looks like `g.co/kgs/...` — that string is your **Knowledge Graph ID**: a social security number for everything on the internet.

No KG MID? You're invisible. Five KG MIDs? You're scattered — Google thinks you're five different people. Both are fixable; your audit says which one you are.

Why you don't have a Panel (yet)

Google hasn't connected enough objects about you — because they aren't linked. Same name, role, headshot, and bio everywhere; Person schema on your entity home with *sameAs*s pointing to every profile; third-party corroboration (press, podcasts, speaker pages) pointing back. The agent's **knowledge-panel-entity-seo** skill runs this exact plumbing.

Why it matters now: Google's Search Profiles require ~100K followers — or a Knowledge Panel. The Panel is the small-business side door into being a recognized source.

2026 reality: the answer engines read the same graph

SEO grew a sibling: **AI search visibility**. ChatGPT, Perplexity, and Google's AI answers describe you based on your entity signals and your canonical articles — the exact outputs of this factory. Ask each engine "Who is [your name]?" today. If the answer is wrong, thin, or about somebody else, that's not trivia — that's how your next customer, lender, or journalist meets you. The **ai-search-visibility** skill measures it and fixes it.

KEY TAKEAWAYS

Entities, not keywords: you're building an object Google can trust. · One KG MID, claimed — not zero, not five. · The same signals now decide what AI engines say about you. E-E-A-T is the rulebook for both.

THE AUDIT

We audit the rooms we speak to. Get your scorecard.

We run the rooms we speak to through the same audit system we use for clients: **Mythos-class Claude agents** (Anthropic's newest model tier — we run Claude Fable 5) working the BlitzMetrics Knowledge Graph Explorer, Ahrefs, and live search, the way an analyst would — except it boils the ocean on every name and never gets tired. We recently did this for a room of 20 dealmakers in Denver (see dennisyu.com/dealcon): **zero of twenty had a Knowledge Panel**. Million-dollar operators, invisible to the machine. DigiMarCon, let's see if you can beat that.

What the agents check on your name

- ❑ **The search test** — who actually owns page one of your name? You, or a stranger with your name?
- ❑ **Knowledge Graph** — your KG MID count: invisible (0), scattered (2+), or clean (1, claimed)
- ❑ **Domain reality** — your site's rating, traffic, and the keywords your name ranks for (often: none)
- ❑ **Proof vs. visibility gap** — reviews, results, press you've earned that Google can't see
- ❑ **AI answers** — what ChatGPT and Google's AI say when asked who you are

The four tiers (find yours)

- Panel ✓** Google answers your name with *you*. Rare air — now compound it.
- Object** You exist in the graph, unclaimed and under-connected. Closest to the prize.
- Buried** Somebody else owns your name — a footballer, a folk singer, an obituary. Fixable, with focus.
- Invisible** No entity at all. Google's database doesn't know you exist. Greenfield — fastest gains.

Your scorecard

A score out of 100, the verdict in one sentence, your three biggest gaps, the exact content you already own that should be repurposed first, and a **90-day plan** your new agents can execute. Request yours at dennisyu.com/digimarcon — or just grab Dennis in the hallway.

Why we score: a plan beats a pep talk. You shouldn't leave a conference with inspiration; you should leave with a diagnosis, a to-do list, and the staff to do it. The next four pages hand you the staff.

KEY TAKEAWAYS

The audit is the same system, just pointed at you. · Your tier tells you the first move: claim, consolidate, displace, or build. · The 90-day plan is written to be executed by the agents you're about to install.

The 10-skill agent team you're taking home

One agent, ten job descriptions — the **Local Service Spotlight skill library**. Run them in order for a full brand build, or jump straight to the one today's job needs. Each skill knows its inputs, its outputs, and its definition of done.

#	SKILL	WHAT IT DOES
1	personal-brand-strategist	Names your buy box, passion, and one-sentence differentiation by mining your own history and bio. Start here.
2	positive-mentions-harvester	Finds and scores every good thing the market says about you into a ranked proof library.
3	reputation-gap-analyzer	Compares the reputation you want vs. the one you have → a dated 30-day plan.
4	personal-brand-website-agent	Drafts and builds your entity home — a site on your own name, structured to earn a Panel.
5	knowledge-panel-entity-seo	The Panel plumbing: entity audit, Person schema, third-party signals, the KG MID claim.
6	ai-search-visibility	Makes ChatGPT, Perplexity, and Google AI describe you the way you want.
7	dollar-a-day-strategist	Runs the audition→kill→scale promotion system from page 9.
8	content-factory	Turns one recording into the 20+ assets from page 7, all pointing at one canonical article. Today's workhorse.
9	definitive-article-writer	Writes the ONE canonical page that owns a topic — what Google and AI quote.
10	recursive-self-improvement-qa	Makes every skill document, QA, and sharpen itself each run. The factory improves the factory.

The 11th file: boil-the-ocean.md

Operating principles for persistent agents: loop until the definition of done passes, use memory so every run compounds, document everything, no "good enough."
The marginal cost of completeness is near zero — so do the whole thing, every time.

How to give an order


Say it plainly: "run **personal-brand-strategist on me**" · "use **content-factory on this interview: [link]**" · "run **dollar-a-day-strategist on my last 30 posts.**" The agent asks for what it needs, then works. Don't micromanage; inspect the output and demand the A.

We don't just teach this at DigiMarCon. We do it there.

Every photo on these two pages is **Stage 1: Produce** happening in real time at DigiMarCon stops — speaking, podcasting, interviewing, and putting young adults on the big stages. Worked like a factory, one conference produces a quarter of content: the keynote becomes the canonical article, the hallway interviews become clips, the meals become the relationships that no robot will ever replace.



The podcast booth, DigiMarCon Las Vegas 2025 — Liana Ling brings a studio to every conference; Dennis brings the stories. One booth session = a month of clips once the agents get it.



DigiMarCon Las Vegas 2025, Luxor — Dylan Haugen teaching the 4 Stages of the Content Factory, the exact framework in your hands right now. We put young adults on the big stages.

“It is a true pleasure and treat working with Dennis Yu.”

Jedediah Jenk, DigiMarCon's event coordinator and site manager, who calls Dennis one of the circuit's most requested speakers — because he arrives prepared, tunes every talk to the room, and stays to network with attendees afterward. (That last part isn't politeness. It's the Thank You Machine from page 9, running in person.)

The hallway is the factory floor

The keynote is one asset. The conference is dozens: the podcast cut you review at a side table, the meal where a partnership starts, the question in the hallway that becomes next week's canonical article. Here's what that looks like across the circuit, over the years.



DigiMarCon Toronto 2026 — reviewing podcast cuts with Liana Ling between sessions. Stage 2: Process happens anywhere.



Las Vegas, DigiMarCon week 2025 — In-N-Out with Marko Sipila after the sessions. Robots can't do this part; relationships are the moat.



Toronto, 2018 — same message, eight years earlier: authentic beats polished. The tools changed; the physics didn't.

Run THIS conference through the factory


- Record one 1-minute video per session you attend: "the one thing I'm stealing from this talk"
- Get 3 hallway interviews on your phone — speakers love a thoughtful question (it's how the Topic Wheel grows)
- Geotag photos with every person you meet; post one tonight and tag them (Thank You Machine, day one)
- Tonight: **"use content-factory on today's recordings"** — fly home with the assets already drafted

KEY TAKEAWAY

A conference ticket is a content budget. Most people spend it on a notebook they never reopen; the factory turns it into a quarter of authority. The next two pages introduce the people and the install that make it automatic.

Yes, agents. But put a young adult in charge of them.

Everyone in this room can talk to agents by Friday afternoon — that's the point of this workbook. But watch what we do in our own shop: the owner doesn't babysit the machines. A **young adult runs the factory** — managing the agents, QA-ing the output, owning the weekly rhythm — while the owner spends the 1% being the proof: on stage, on podcasts, on job sites. The young adult brings hustle, accountability, and a face your customers actually meet. The agents give that young adult senior-grade output on every task. Together, that's a marketing department for less than a part-time hire.




Cam Hazzard, Dylan Haugen, and Dennis Yu on stage at the DSDT AI Summit in Detroit — googling "Who is Cam Hazzard?" live for the room. We practice what we preach: speaking, podcasting, and putting young adults on the stage.



Cam Hazzard

Pro dunker out of Texas — one of 24 athletes Shaq picked for the DunkMan League on TNT this summer. Runs his own factory: grew Instagram from 2K to ~20K in months, interviews Dennis and local owners, and ships the assets the same week.

camhazzard.com



Dylan Haugen

The world's youngest pro dunker (DunkMan League, 100M+ views) and host of the 70-episode Dunk Talk podcast. Co-founded Local Service Spotlight with Dennis — he runs the agents that run client factories.

dylan-haugen.com



Marko Sipila

Started an agency in high school; seven figures by 21. Founder of HVACQuote.ai — instant quoting used by 300+ home-service companies — and CoatingLaunch. Runs backcountry marathons at altitude when he isn't shipping.

markosipila.com

Want one on your team? (Or want your kid to become one?)

This is the only pitch in the whole workbook, so here it is, gently: our **AI Builder program** trains young adults to run agent-powered Content Factories for real businesses — reputation first, proof always, hype never. Cam, Dylan, and Marko are what it looks like when it works: athletes who build agents and grow companies. Read *"How Young Adults Can Build a Thriving Digital Marketing Agency by Mastering Reputation and Trust"* — linked at dennisyu.com/wichita — or just grab Dennis at the conference.

Put the agents on your own Claude



dennisyu.com/digimarcon

Scan with your phone camera — opens this page so you can copy the prompt.

- 1 Open Claude** — the Claude app or claude.ai. **Free works.** On a work laptop, the browser is fine.
- 2 Paste the prompt** — at dennisyu.com/digimarcon, tap *Copy the prompt*, paste it as your first message, send. Claude becomes your Content Factory system. (The full prompt is printed on the next page, so this workbook works even with a dead phone battery.)
- 3 Give an order** — like **"use content-factory on this interview."** The agent asks a few questions and gets to work.

■ Make it permanent (Pro / Max)

Create a **Project** in Claude and paste the prompt into the Project's **instructions** — every chat inside it already knows the method. On Claude Cowork? Add one line: *"...and save this to memory so every future chat knows it."*

Then try it — say any of these

```
run personal-brand-strategist on me
run positive-mentions-harvester on me
use content-factory on this talk: [paste link]
run dollar-a-day-strategist on my last 30 posts
```

Then verify it worked

Claude's first reply will start with **✔ Content Factory installed — 10 skills loaded.** From then on, ask **"what skills do you have?"** in any chat — you should get all 10, numbered. Didn't? Just re-paste the prompt. ⚠ Don't drag the .zip bundle into a chat — chat attachments can't open zips. For the permanent install (Claude's Skills uploader), scan this QR for click-by-click steps.



The prompt, printed in full

You are my Content Factory and Personal Brand system, built on Dennis Yu's "Local Service Spotlight" method. You run a 10-skill library, in order, to turn one recording into weeks of content, build a personal brand that earns a Google Knowledge Panel, and turn proof into customers.

THE 10 SKILLS

1. **personal-brand-strategist** – name my buy box, passion, and one-sentence differentiation by mining my own AI history and bio. Start here.
2. **positive-mentions-harvester** – find and score every good thing the market says about me into a ranked proof library.
3. **reputation-gap-analyzer** – compare the reputation I want vs. the one I have, then give a dated 30-day plan.
4. **personal-brand-website-agent** – draft and build my entity home: a site on my own name, structured to earn a Knowledge Panel.
5. **knowledge-panel-entity-seo** – the plumbing for a Google Knowledge Panel: entity audit, Person schema (JSON-LD), third-party signals, and the KGMID claim.
6. **ai-search-visibility** – make ChatGPT, Perplexity, and Google AI describe me the way I want for customers.
7. **dollar-a-day-strategist** – turn proven content into customers with \$1/day boosts: pick winners, boost 7 days, kill the bottom 90%, scale the unicorns.
8. **content-factory** – turn one recording into a blog post, clips, reels, posts, and ads, all pointing at one canonical article.
9. **definitive-article-writer** – write the ONE canonical page that owns a topic, the one Google and AI quote.
10. **recursive-self-improvement-qa** – make any skill document, QA, and sharpen itself each run.

HOW TO WORK WITH ME

- When I say "run [skill] on me," follow that skill: ask for the inputs it needs, then produce its output.
- If I don't name one, start with personal-brand-strategist. For a recording, talk, or interview, use content-factory.
- Be concrete and honest. Separate fact from aspiration. Quote me where you can – my voice, not corporate robot voice.
- Remember my business, city, buy box, and proof so later skills build on earlier ones.
- If I ask "what skills do you have?" – anytime, in any chat – reply with the numbered list of all 10 skills above, each with its one-line job.

START NOW, in this order:

- 1) Confirm the install. Your first reply must begin with exactly: "✅ Content Factory installed – 10 skills loaded. Ask me 'what skills do you have?' anytime to verify."
- 2) Then ask me (1) my name, my business, and my city, (2) the kind of customers I want more of, and (3) optionally, a link to one recording – a podcast, talk, or customer story – to run through the factory. Then run personal-brand-strategist.

Tip: this exact text, tap-to-copy, lives at dennisyu.com/digimarcon. Free Claude accounts work; Pro/Max users should paste it into a Project's instructions to make it permanent.



This is what Stage 1 looks like

Cam Hazzard interviewing Dennis on Shaq's DunkMan League – fifteen honest minutes in a gym, one phone, no crew. Hand a recording like this to the prompt on this page and say "**use content-factory on this interview**" – the agents turn it into the article, the clips, and the posts while you're driving home. Record first. Polish never.

Your buy box & one-sentence differentiation

Fill this in by hand, then read it to your agent — it's the input personal-brand-strategist needs. Specific beats clever.

WHO I SERVE (THE CUSTOMER I WANT MORE OF)

THE EXPENSIVE PROBLEM I SOLVE FOR THEM

PROOF I CAN POINT TO (JOBS, RESULTS, REVIEWS, YEARS, NAMES)

MY ONE SENTENCE — "I'M THE ONLY ____ IN ____ WHO ____, PROVEN BY ____."

PROOF INVENTORY — MY 5 BEST MENTIONS (WHAT WAS SAID · WHO SAID IT · WHERE IT LIVES)

WHAT WAS SAID	WHO SAID IT	WHERE IT LIVES

Score each later on the 30-point scale (Who / What / Where, 10 points each). Anything 20+ is a boost candidate.

What you know × who you know

Nine topics you can talk about for an hour each — the intersection of your expertise and your passions. Then, for each topic, one person with authority you could co-create with (a "lighthouse"). Introverts: start with topics, add people. Extroverts: start with people, map topics.

TOPIC 1	TOPIC 2	TOPIC 3
TOPIC 4	CENTER Your name + your mission	TOPIC 5
TOPIC 6	TOPIC 7	TOPIC 8 / 9

LIGHTHOUSES — ONE PERSON OF AUTHORITY PER TOPIC (PEERS, CLIENTS, LOCAL FIGURES, PODCAST HOSTS)

TOPIC #	PERSON	WHY THEM / HOW WE CONNECT

A shared mission attracts people bigger than you — their authority compounds yours. That's content you can't fake.

Your first nine one-minute videos

Three videos per row. WHY = stories and values (make them feel). HOW = teach what you'd do if your product didn't exist. WHAT = the offer: price, features, testimonials. This grid is an evergreen funnel — record once, let the factory and Dollar-a-Day do the rest.

WHY · 1	WHY · 2	WHY · 3
HOW · 1	HOW · 2	HOW · 3
WHAT · 1	WHAT · 2	WHAT · 3

SCRIPT ONE VIDEO RIGHT NOW (PICK ANY CELL ABOVE)

HOOK · 0-3S	PAIN / PLEASURE · 3-15S
SOLUTION · 15-50S	CALL TO ACTION · 50-60S

Don't polish. Authentic and imperfect outranks scripted and sterile — that's not a pep talk, it's how E-E-A-T scoring works.

PAA starter + your weekly rhythm

10 QUESTIONS MY CUSTOMERS ACTUALLY ASK (YOUR FIRST PAA VIDEOS)

Your agent will pull the other 40 from live search data — say "find the People Also Ask questions for [your service] in [your city]."

THE WEEKLY HEARTBEAT (≈ ONE OWNER-HOUR TOTAL)

Mon	MAA: agent brings 3 decisions; you make them. (5 min)
Tue	Produce: record 1 interview or 3 one-minute videos. (30 min)
Wed	Process: agent builds the 20+ assets; grades itself.
Thu	Post: you approve; agent publishes everywhere. (10 min)
Fri	Promote: \$1/day auditions launch; one human thank-you. (10 min)
Monthly	Full MAA review: kill, scale, and re-aim the Topic Wheel.

THIS MONTH'S MAA LOG

METRIC I WATCHED	WHAT IT TOLD ME	WHAT I CHANGED

The 90-day standard

12 weekly cycles = 12 recordings → 240+ assets → 12 canonical articles → a tested ad library of proven winners → an entity home Google can finally trust. That's not a heroic quarter; that's one owner-hour a week with the factory running. Your audit's 90-day plan maps it to your specific gaps.

Conference notes — capture it while it's hot

WHAT I HEARD THAT I DON'T WANT TO FORGET

PEOPLE I MET (AND THE PROOF WE COULD CO-CREATE)

COMMANDS I WANT TO RUN ON MY AGENTS THIS WEEK

The one action I'll take Monday morning

Everything in this workbook, live:



dennisyu.com/digimarcon	This workbook, the agents, the tap-to-copy prompt, and your audit request.
dennisyu.com/dealcon	See a full cohort audited — 20 dealmakers, every scorecard public.
blitzmetrics.com/build-agents	How we build agents and keep them current — the living hub behind this system.
The AI Builder program	How young adults learn to run agent-powered factories (page 14) — article and details linked at dennisyu.com/wichita .

Built by agents. QA'd by humans.

This edition was tuned up the way we're teaching you to work: a Mythos-class Claude agent restructured the original 319-page workbook against its editorial blueprint, rewrote it for the agent era, and generated this print file — then a human read every word. One recording in, weeks of content out. The workbook practices what it preaches.

Thank you, DigiMarCon

To the **DigiMarCon** team — **Jedediah Jenk** and the crew who build these stages around the world — and to every attendee who stops us in the hallway with a question. Those questions become the next canonical articles. The Thank You Machine starts with actually saying it.

Dennis Yu

CEO, BlitzMetrics · Local Service Spotlight
dennisyu.com · 612-707-8045

 Kept current — June 2026 edition, updated for persistent agents (Claude Fable 5, OpenAI, Google).
The Content Factory Workbook · AI Agent Edition