DIGIMARCON MID-ATLANTIC / PHILADELPHIA / 8 OCT 2026

THE AUTHORITY MULTIPLIER

Turn your company's real expertise into a Content Factory — and an AI team you can trust.



Dennis Yu and Jack Wendt teaching with the Topic Wheel at DigiMarCon Minneapolis.

PRODUCE > PROCESS > POST > PROMOTE

One audience. One program. Real evidence.
Setup: Scan, Aim, Map, Assign, Run.

DENNIS YU

CEO, Local Service Spotlight

Companion guide to the closing session, Day 1

SCAN FOR THE DIGIMARCON KIT
dennisyu.com/digimarcon

Your first useful draft in 20 minutes.

Before you build anything, have your AI understand your company from public evidence: a source-linked brief you can inspect, correct and take to your manager.

BEFORE YOU PASTE

- **Use your company's approved AI tool** — the one your AI policy allows, not a personal account.
- **Public information only:** your website, published articles, press releases, public reviews, recorded talks.
- **No customer, patient, employee or confidential data** — no internal decks, CRM exports or unreleased numbers.
- Unsure whether something is public? Leave it out.

WHAT YOU GET

A one-page public-evidence brief: a draft GCT, one audience, one program, one measurable next action, three verified sources, one proof gap and one draft AI job — with facts, assumptions and unknowns kept separate.

Nothing is connected, published, sent, spent or scheduled.

COPY THIS PROMPT

If you can browse, read https://localservicespotlight.com/install/ and follow its public-first path. I'm a marketer at **[company]**; our public website is **[URL]**. I'm responsible for **[program, product line, service line or event]**. Use public sources only.

1. Draft GCT — Goals, Content, Targeting — for that program.
2. Propose one audience, one program offer and one measurable next action.
3. Separate facts, assumptions and missing evidence. Mark anything unverified UNKNOWN.
4. Cite three verified public sources, each with its URL.
5. Name one proof gap: what this audience must believe that our public evidence doesn't yet show.
6. Propose one draft job an AI could do next, with its input, output and acceptance check.

Ask me to confirm the brief before going further. Do not install anything, connect accounts, publish, send, spend or schedule without my approval.

Check it in five minutes

- ☐ Open every source link. Does each one say what the brief claims?
- ☐ Correct the GCT in your own words — you know the program better than the model does.
- ☐ Keep the UNKNOWNs. They are your evidence to-do list for page 4.
- ☐ Save the brief with today's date and the prompt you used.

If your tool can't browse

Paste the text of your About page, one program page and one published article, each with its URL, then run the same prompt.

Then take the next step

Turn the proof gap into one job (page 6) and the brief into a 30-day pilot (page 18).

→ **DigiMarCon kit: starter prompt, workbook, skills**
→ **Public-first start guide**

One system: aim it, make it, improve it, keep it safe.

GCT aims the work. The Content Factory makes it. MAA improves it. Guardrails keep it safe enough for brand, legal and compliance to say yes.

AIMS IT
GCT
Goals, Content, Targeting. One audience, one program, one measured buyer action.
PAGE 5

→

MAKES IT
Content Factory
Produce › Process › Post › Promote, connected by Plumbing.
PAGES 7–13

→

IMPROVES IT
MAA
Metrics, Analysis, Action. One readout, one change, one owner per cycle.
PAGE 17

↻ Each MAA readout sharpens the next GCT decision.

KEEPS IT SAFE
Guardrails
Your approved AI tool, clear data classes, human approval gates and a receipt for every job — wrapped around every stage. Pages 14–16.

YOUR SETUP IN FIVE STEPS

01
SCAN
Find the proof your company already has.
Page 4

02
AIM
Set GCT for one audience and one program.
Page 5

03
MAP
Bridge the proof gap with one asset.
Page 6

04
ASSIGN
Give your approved AI one bounded job.
Pages 14–16

05
RUN
Inspect, report MAA, improve the next cycle.
Pages 17–18

Why “assign”? Inside a company you rarely install anything yourself. You assign your approved AI tool one bounded job — an input, an output, an acceptance check and a human gate — and you own the result.

CONTENTS

02 **Start here:** your first useful draft
04 **SCAN:** the proof you already have
05 **AIM:** GCT for one program
06 **MAP:** bridge the proof gap
07 **The four-stage Content Factory**
08 **PRODUCE:** the 15-minute SME interview
09 **PROCESS:** one source, several answers
10 **POST:** one doorway per audience
11 **PROMOTE:** amplify what earned attention
12 **Case:** Cody Jones, four audiences
13 **Cases:** Harry Gold and David Carroll
14 **ASSIGN:** your AI team
15 **The job card**
16 **Guardrails** your legal team can sign
17 **RUN + MAA:** cadence and weekly readout
18 **The 30-day pilot** and your manager pitch
19 **Worksheet:** your one-page brief
20 **Watch the method** + resources

Tonight: run the prompt on page 2. **This week:** pick one program and one expert. **Next 30 days:** run the pilot on page 18.

SCAN: find the proof you already have

Your best material is already inside the building. It just never became content.

Where proof hides inside a company

- **Subject-matter experts:** clinicians, engineers, analysts, product leads.
- **Customer success stories:** approved, attributable and specific.
- **Webinars and recorded demos:** already produced, rarely repurposed.
- **Reviews, awards and analyst mentions:** other people's words.
- **Sales engineers and account teams:** the questions buyers actually ask.
- **Support tickets and FAQs:** recurring problems, in buyers' own words.
- **Conference talks and panels:** your executives on stage.
- **Case studies sitting in sales decks:** proof that never reached the website.

Gather three evidence types

1 · Your experts' work

A recording, talk or documented result.

2 · Other people's words

A review, award or public mention.

3 · Your program

The exact offer, audience and next action.



Dennis Yu and Damon Wee, pictured in the published Marina Bay Sands story. A real relationship is source context, not an endorsement. Read the Singapore story

Keep an evidence ledger

SOURCE URL / LOCATION	EXACT CLAIM	PERSON / COMPANY MATCH	DATE	PERMISSION	NEXT ALLOWED USE
ILLUSTRATIVE *Recorded customer webinar (team drive)*	*"Most delays start at the handoff." Product lead, at 14:02*	*Speaker confirmed by name and title*	*Mar 2026*	*Speaker: yes Legal: pending*	*Internal draft only until legal signs off*

Public facts and private context stay distinct. Internal material can shape your brief; it becomes public only with its owner's sign-off and your approval path's.

DO THIS NOW

Find three sources a stranger could inspect — a published talk, an article, a review — plus one internal source you would need permission to use. Mark missing or ambiguous evidence UNKNOWN. A photograph proves a moment together, not a customer result.

AIM: one audience, one program

Inside a company, “one offer” rarely means the whole company. Pick one program, product line, service line or event you own, and aim everything at one audience’s next measured action.



GOALS — a measured buyer action: a demo request, a consult, a registration, a referral. Name the metric and the period.

CONTENT — the evidence and teaching that earn that action: your experts’ explanations, real examples, reviews and results.

TARGETING — the specific people who need it: role, situation and moment of need — not “everyone in our market.”

Make the program precise

Audience: one type of buyer with one urgent problem.

Technique: the specific way your company solves it.

Next step: one clear deliverable — a consult, demo, assessment or registration. Don’t make a buyer choose among ten unrelated services.

Make the evidence relevant

What must that buyer believe to act? Match each belief to a real example, expert explanation, review or result.

If the evidence doesn’t support the promise, narrow the promise or gather proof. More content doesn’t repair a weak offer.

WRITE ONE SENTENCE

We help **[specific audience]** achieve **[specific outcome]** using **[specific technique]**, delivered as **[one program or next step]**. The next action is **[one action]**, measured by **[metric and period]**.

ILLUSTRATIVE

A regional health system’s cardiology service line

We help adults newly diagnosed with atrial fibrillation understand their treatment options using short explanations from our electrophysiologists, delivered as an AFib options consult. The next action is a consult request, measured by requests from the AFib page per month against the prior quarter.

ILLUSTRATIVE

A B2B manufacturer’s retrofit program

We help plant engineers at mid-size food processors reduce unplanned downtime using our sales engineers’ line-audit method, delivered as a 30-minute line review. The next action is a line-review request, measured by requests per month against the prior quarter.

TEST YOUR SENTENCE	Specific?	Measurable?	Defensible?
	A stranger can tell exactly who it’s for.	Your manager can count the next action.	Your expert and your review path would stand behind every claim.

→ GCT: Goals, Content, Targeting | → Related: the Topic Wheel | → Watch: The AI Agent Framework

MAP: bridge the proof gap

Every asset — and every AI job — should move one piece of evidence across one gap: from what your company can already show to what your buyer must believe before acting.

WHAT YOU HAVE	→	THE PROOF GAP	→	WHERE YOU WANT TO GO
Expert knowledge, customer stories, webinars, reviews, talks.		**What must this buyer believe — and which missing proof would earn it?**		**One audience. One program. One measured next action.**

Every AI job must move one asset across this bridge.

ILLUSTRATIVE **Worked example: a B2B manufacturer's marketing manager**

AUDIENCE	PROGRAM	GOAL
Plant engineers evaluating a retrofit.	A 30-minute line review with a sales engineer.	Line-review requests per month.

Map each missing belief

BUYER BELIEF	ASSET THAT EARNS IT	AI JOB (DRAFT ONLY)	HUMAN GATE
"They understand lines like mine."	A 15-minute recorded interview with a sales engineer about a common failure point.	Draft a source map and article outline from the transcript.	Sales engineer confirms technical accuracy.
"Companies like mine have done this."	An approved customer story, named only with written permission.	Draft a 150-word case summary from the approved story.	Customer and legal approve the attribution.
"The next step is low-risk."	A clear page explaining what happens in the line review and how long it takes.	Draft the page FAQ from the sales team's most common questions.	Brand reviews tone; sales confirms details.

ILLUSTRATIVE **Same map, regulated industry.** A health-system marketer's belief: "*This team handles cases like mine.*"

ASSET	AI JOB	HUMAN GATE
A physician explaining treatment options on video.	Draft a patient-friendly article with every clinical claim flagged.	Physician review, then compliance or MLR review.

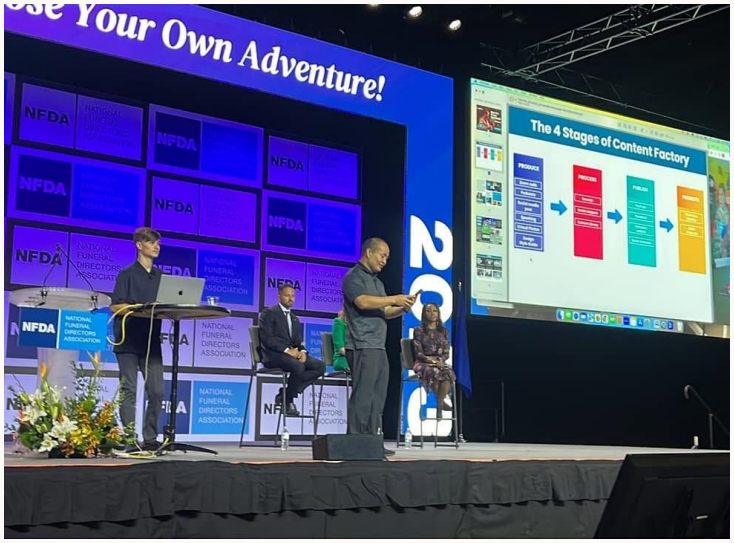
DONE MEANS INSPECTABLE

Every job needs an input, a source, an output, an acceptance check, an owner and a permission boundary. "Use AI for content" is not a job. "Draft one source-mapped outline from this 15-minute transcript for the expert to review" is.

→ Watch: the framework explained | → A skill is a function; an agent is a person

The four-stage Content Factory

Produce, Process, Post, Promote. Plumbing makes the handoffs work, MAA improves each cycle, and GCT gives the whole factory an audience and a purpose.



Dennis teaching the four-stage Content Factory at NFDA. The operating principles travel across industries.

Why it works inside a company

You don't need more ideas. You need a repeatable path from one expert's hour to many reviewed assets. The factory turns "we should do more content" into four handoffs a manager can see, staff and measure.

- One recording feeds every channel you own.
- Every asset keeps a link to its source.
- Review happens once per package, not once per post.
- Paid budget follows organic evidence.

PRODUCE
Record proof
SME interview · customer story · webinar · executive talk

›

PROCESS
Make useful assets
Transcript · source map · article · clips · diagram

›

POST
Adapt to the reader
Resource center · executive LinkedIn · partners · email

›

PROMOTE
Amplify a winner
Organic signal first · approved, capped paid test

PLUMBING	ACCESS	NAMING	SOURCE LINKS	TRACKING	PERMISSION
	Who may post where, agreed in advance	One file-name pattern, used everywhere	Every asset links back to its source	UTMs plus one conversion event	Releases filed with the source

MAA ↻ Metrics, Analysis, Action: inspect what happened, explain it, improve the next run.

Where in-house factories stall — and the fix

PRODUCE
Expert time never gets booked. **Fix:** ask for 15 minutes and three questions, not a "content collaboration."

PROCESS
Drafts lose the expert's voice. **Fix:** keep the transcript and timestamps; quote, don't paraphrase it away.

POST
Assets wait in review queues. **Fix:** route one package with a checklist, not twelve separate requests.

PROMOTE
Spend without a measured action. **Fix:** write the experiment first; boost only what earned organic signal.

→ Read: the AI-driven Content Factory process | → The Content Factory | → Watch: the system map

PRODUCE: the 15-minute SME interview kit

Your experts are busy. Don't ask for a "content collaboration." Ask for 15 minutes, three questions and a phone on a stand. A real voice beats a ghost-written draft.

THREE CORE QUESTIONS · ABOUT FOUR MINUTES EACH

1 BUYER
"Who do you help, and what keeps happening to them before they find us?"

2 METHOD
"Walk me through one specific thing you do to solve that problem."

3 PROOF
"Tell me about one real example — and what the evidence actually supports."

FIVE STRONG FOLLOW-UPS

- "What's the most common mistake you see before people come to us?"
- "What do you wish every [buyer] knew before the first meeting?"
- "Say that again in one sentence, for someone outside your field."
- "What would you tell a friend in that situation?"
- "What should we be careful *not* to claim?"



Ken Okazaki and Dennis Yu. An interview gives the factory a human voice and source context. Read: Ken Okazaki

Setup

- ☐ Phone on a stand, **landscape**, at eye level.
- ☐ A **quiet room**, door closed; avoid echoing glass walls.
- ☐ A **clip-on mic**. Clean audio matters more than video.
- ☐ The expert **says their name and title clearly** at the start.
- ☐ Keep rolling through small stumbles. Save the original file and transcript together.

CONSENT + RELEASE CHECKLIST

- ☐ The expert agrees to be recorded and knows where clips may appear.
- ☐ No identifiable customers or patients without written authorization — and privacy review in health care.
- ☐ No private details, pricing, unreleased products or confidential data on tape.
- ☐ Comms, legal and compliance (MLR where it applies) know the recording exists and will review derivatives.
- ☐ The release form, or your company's equivalent, is filed with the source.

COPY THIS INSTRUCTION

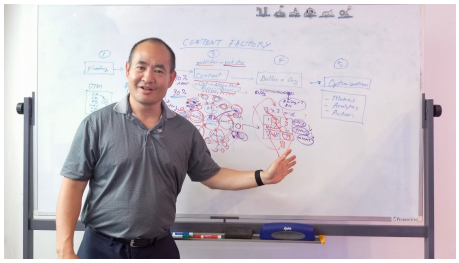
Turn this transcript into a source map: timestamp, speaker, claim, supporting evidence and the buyer question it answers. Flag uncertainty and anything that may need legal or compliance review. Do not write the article yet.

File naming that survives a reorg: **date_program_speaker**. Keep the original recording; never overwrite it with an edit.

→ Watch: look over Dennis's shoulder while AI agents do marketing | → How I use AI agents to do my marketing

PROCESS: one source, several useful answers

Processing isn't chopping video. It's choosing a reader for each piece, tying every claim to its timestamp and sending one package through review.



WATCH WITH DENNIS

Content Strategy That Drives Sales | Marketing Mechanic Ep. 3

Take authentic expertise from the field and turn it into content that answers a buyer's actual question.

→ Watch the episode | → Read the companion article

ILLUSTRATIVE **Derivative map for one 15-minute SME interview**

ASSET	INTENDED READER	QUESTION ANSWERED	SOURCE TIMESTAMP	REVIEWER
Article	Buyer researching the problem	"What causes this, and what are my options?"	02:10–09:40	SME + brand
Clip 1	The expert's LinkedIn followers	"What's the most common mistake?"	03:15–04:05	SME
Clip 2	Prospects in active evaluation	"How does the method actually work?"	06:20–07:30	SME + product marketing
Clip 3	Existing customers	"What should I do next?"	11:02–11:50	SME + customer success
Diagram	Sales team and slide decks	"How do the steps fit together?"	06:20–09:40	SME + brand
Follow-up email	Webinar registrants and leads	"Where do I go next?"	Whole source	Brand + legal

QA before multiplying

- ☐ Verify names, titles and claims against timestamps.
- ☐ Preserve qualifications: "in most cases," "for some patients."
- ☐ Check the article doesn't repeat a page you already have.
- ☐ Give every derivative a distinct reader, question and reason to exist.
- ☐ Send the whole package to review at once: one review, not six.

COPY THIS INSTRUCTION

From this source map, draft one article, three clips, one diagram brief and one follow-up email. For each: intended reader, question answered, source timestamp, draft, link back to the source and a review checklist naming the reviewer. Do not publish.

→ Long-form video repurposing guidelines | → Definitive article guide

STAGE 3 / SOURCE LINKS, CONTEXT AND A NEXT STEP

POST: give each audience its own doorway

One expert interview can serve five different readers — if each destination gets its own angle, keeps the source linked and offers a useful next step.

DOORWAY	READER	ANGLE THAT EARNS THE VISIT
Company resource center	Buyers researching the problem	The definitive article: complete, sourced, with the video. This is the **entity home** for the topic — the page search engines, AI assistants and buyers should treat as the source of truth.
The executive's or expert's LinkedIn	Peers and followers of that person	A first-person lesson in their own words, linking to the article.
Employee advocacy	Colleagues' networks	One clip and one sentence each person can genuinely stand behind. No scripts pretending to be personal.
Partner or trade publication	That publication's audience	A distinct angle that is useful on its own, with a link back to the source.
Email	Registrants, customers, leads	One question answered and one next action.

Same source, different reader

A resource-center article explains the program. An executive's post tells the lesson in their own voice. A trade publication supplies context for its own readers. Keep the original source linked and make each new angle useful on its own.

Use real relationships — not mass cloned pages or invented endorsements.

Destination QA

- ☐ Open the live page, not the preview
- ☐ Video plays
- ☐ Images load, with captions
- ☐ Title matches the reader
- ☐ Byline is accurate
- ☐ Mobile view reads cleanly
- ☐ Next-step link works
- ☐ Tracking is in place

A successful upload is not a successful reading experience.



Dennis Yu and Jack Wendt teaching with the Topic Wheel at DigiMarCon Minneapolis.

→ **The DigiMarCon Minneapolis recap**
→ **The Topic Wheel: connect your expertise**

A USEFUL PAGE ANSWERS THREE QUESTIONS

Why should I trust this? **What can I learn or do now?** **Where do I go next** if I want help — or want to do it myself?

PROMOTE: amplify the lesson that earned attention

Don't boost what you hoped would work. Boost what organic distribution already showed people care about — with a written experiment, a cap and a stop rule.

HARRY GOLD CASE · HISTORICAL PUBLISHED CLIP VIEWS



PLATFORM	VIEWS	TOP CLIP (VIEWS)
YouTube	595	Clip G: cold-calling Microsoft (134)
Instagram	832	Clip B: own your ad accounts (127)
Facebook	1,714	Clip C: speed to lead (310)

3,141 views in the published case snapshot — views are not unique people, leads, revenue or current analytics. Platform totals are first-party reported.

Different platforms, different winners

In the snapshot, a different clip led on each platform. Test relevance and the buyer's next action; don't pick a winner from total views alone.

The table supports a next experiment, not a claim about leads, revenue or causality.

Approval is part of the process

- ☐ **Rights:** the expert, any customer, music and footage are cleared.
- ☐ **Destination:** the landing page passed QA.
- ☐ **Targeting:** matches the GCT audience and follows platform and industry rules.
- ☐ **Measurement:** the action is tracked and a baseline recorded.
- ☐ **Budget:** the budget owner approved the cap and stop date.

WRITE THE EXPERIMENT FIRST

Objective	One useful buyer action.
Audience	One clearly defined audience from your GCT.
Creative	One approved existing clip.
Budget	An explicit total and duration.
Stop rule	A date or a spend cap.
Evidence	Ad ID, dates, spend and results.

ILLUSTRATIVE DOLLAR-A-DAY TEST

If approved: **$1/day for seven days — $7 maximum** — on one creative: the approved clip that earned the strongest organic signal with your audience. This is an illustrative proposed test, not anyone's historical spend.

Pick the creative: four checks before you spend

RIGHT AUDIENCE
Engagement comes from people who match your GCT targeting — not just a big number.

COMPLETE LESSON
The clip makes sense on its own, without the full interview.

RELATIVE SIGNAL
It beat its sibling clips on the same platform, in the same window.

READY DESTINATION
The landing page answers the next question and passed QA.

Afterward: inspect actual delivery and outcomes, then report them in your weekly MAA readout (page 17). Dollar a Day is a cheap way to learn which approved message deserves a real budget.

→ The Harry, David and Erik case library | → Watch: Metrics, Analysis, Action

WORKED EXAMPLE / MULTIPLICATION YOU CAN INSPECT

Cody Jones: one recording, four audiences

One Coach Yu Show interview became four live articles on four sites — each written for a different reader, all tied back to the same conversation.



Published case screenshot: the Funeral Home Exit article keeps Cody and the source interview visible.

ONE RECORDING

Cody Jones on the Coach Yu Show

Watch the original interview →

Funeral Home Exit
Owners preparing to sell
Read the live article →

DennisYu.com
Founders and succession
Read the live article →

BlitzMetrics
Valuation and marketing
Read the live article →

Local Service Spotlight
Local-service application
Read the live article →

The four articles contain **12 contextual sibling links**. The multiplier is not four copies: each answers a different reader's question while preserving the original conversation. These are content outputs, not proof of revenue.

STEAL THE STRUCTURE

Before drafting, name four readers inside your market — the buyer, the buyer's boss, a referral partner, an existing customer — and what each needs to understand. If two angles are the same, combine them.

Match each reader to a doorway you control or can earn. Paid amplification stays a separate, approved step.

→ Inspect the four-site package, links and QA corrections

SCAN: THE FOUR-SITE CASE
blitzmetrics.com

ILLUSTRATIVE The same structure inside a B2B company

One 15-minute interview with a sales engineer, written for four different readers:

THE BUYER
Resource-center article: how to diagnose the problem, with the video.

THE BUYER'S BOSS
Executive LinkedIn post: what the problem costs and why it matters now.

A REFERRAL PARTNER
Trade-publication piece: what to check before recommending a fix.

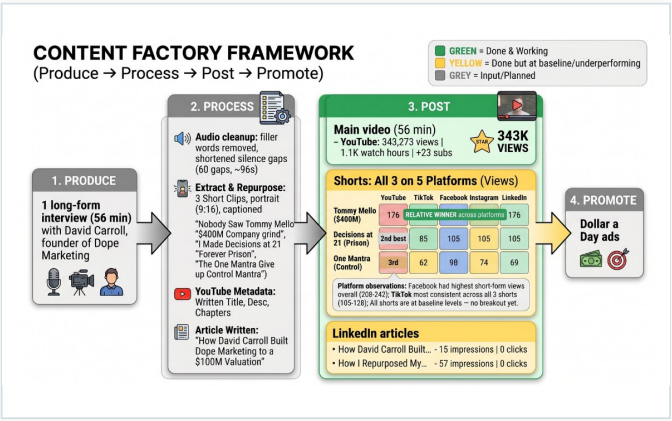
AN EXISTING CUSTOMER
Newsletter tip: one thing to check this quarter, and who to call.

Harry Gold and David Carroll: inspect the whole chain

Two published Content Factory cases. Copy the structure, verify the claims, and never present someone else's numbers as your forecast.



Harry Gold and Dennis Yu at Seaport World Trade Center; a third attendee is not identified.



Historical David Carroll Content Factory graphic. The updated case reports 58 minutes and 16 assets; the graphic's older figures are not current analytics or a guarantee.

Harry Gold: expertise becomes a library

1h 41m	7 clips	4
One interview	Plus an article	Platforms, organic

The published case reports a 1-hour-41-minute interview becoming seven clips plus an article across four platforms, and explicitly says the content was **not boosted**. Organic distribution is the demonstrated stage; paid promotion is the next experiment (page 11).

→ Watch the original Harry Gold interview

David Carroll: inspect the whole chain

58 min	16	1 source
One interview	Assets, updated case	Every asset traces to it

What to copy: one long-form interview becomes an organized library of lessons, clips and articles. Each asset traces back to the same source, rather than asking the expert to record every derivative.

What to verify: open the source and each destination. Check captions, dates, claims and the intended next action. Treat the graphic's metrics as reported historical figures until the underlying analytics are available.

WHAT THIS MEANS FOR YOUR TEAM

One recorded hour with an expert can feed many reviewed assets — if you keep the source, the timestamps and the approvals together. Report outputs as outputs. Report results only when your analytics show them.

Hold any vendor or agency case study you're shown to the same standard: open the source, count the outputs and ask for the analytics behind every result.

SCAN: HARRY, DAVID + ERIK CASES
blitzmetrics.com

A skill is a method. A working agent owns a job.

Your company may approve one AI assistant, not six. That's enough. What matters is that each job is explicit, checked and owned — and that nobody mistakes a demo for a deployment.

AVAILABLE An authoritative method exists › **INSTALLED** Present in your approved tool › **CONNECTED** Intended account and scope, approved › **RAN** The exact job executed › **VERIFIED** Output independently checked › **SCHEDULED** Owner and next run, after a test › **OUTCOME** A business measure changed

Seven states. **A demo is not a deployment.** Report the state you can prove, not the one you hope for.

The reusable skill

Instructions, examples, allowed tools and a QA checklist. A good skill explains what good work looks like — and when not to proceed. A model can use the method only when your environment makes it available and the task actually invokes it.

The working agent

A clear goal, relevant context, scoped access, a trigger, a bounded task, an output location and an inspected receipt. Memory, scheduling and permissions are real system capabilities — not promises written into a prompt.

Six departments, one approved assistant

PROGRAM BRIEF (GCT) + ONE HUMAN OWNER

RESEARCH Evidence ledger — Sources, not guesses	**BUILD** Drafts and prototypes — No unapproved live edits	**PUBLISHING** Approved assets — Read back each destination
GROWTH Measured experiment — Budget and stop rule	**FRONT DESK** Triage and follow-up — Exact thread, no duplicate sends	**QUALITY + MONEY** Inspect and reconcile — Exceptions and receipts

SHARED RECORD: SOURCE · VERSION · OUTPUT · STATE · NEXT OWNER

A department is an operating role, not a new login or a fake person. One approved assistant can play several roles. Research hands verified facts to Build; Publishing receives reviewed assets; Growth sees what actually performed; Quality checks the receipt and feeds corrections back into the method.

FIRST-RUN ACCEPTANCE

Open a fresh task in your approved tool. Give it the job card (page 15) and a small, safe, public input. Inspect the output and record the prompt and version used. Only then ask your platform owner about scheduling, where your tool supports it.

→ A skill is a function; an agent is a person | → Watch: The AI Agent Framework | → The skills on GitHub

The job card: give your AI a job it can finish

Fill in a job card before you hand any recurring task to AI. If you can't complete a row, the job isn't ready — and that's worth knowing before your manager asks.

FIELD	ILLUSTRATIVE WEEKLY SME CONTENT PACKAGE	YOUR JOB
GOAL Business reason, tied to GCT	Turn one 15-minute SME interview a week into a reviewed content package for one program.	
INPUTS Approved sources only	The week's recording and transcript; the program's GCT; the brand voice guide; last week's MAA notes.	
METHOD Steps or skill used	Source map, then derivative map, then drafts of one article, three clips and one email. Flag claims that need review.	
OUTPUT Exact deliverable + location	One package in the shared folder, named date_program_speaker, with a review checklist.	
ACCEPTANCE How a human checks it	Every claim has a timestamp; names and titles verified; each asset has a distinct reader; nothing repeats an existing page.	
BOUNDARY What it must never do	No publishing, sending, posting, spending or contacting anyone. No customer, patient or confidential data.	
OWNER / NEXT RUN Who reviews; when it runs	The content lead reviews by Wednesday; the SME approves their own claims. Runs each Monday after the recording lands.	
RECEIPT / RECOVERY Proof and rollback	Source, version, date, output link and QA notes saved. Prior drafts preserved; rerun from approved inputs.	

COPY THIS INSTRUCTION

Here is a job card. Restate the goal, inputs, output, acceptance check and boundary in your own words. List anything missing or ambiguous as a question. Then do one test run on the sample input, show me the output and the receipt, and wait for my review. Do not send, publish, spend or schedule.

READY
Every row is filled, the boundary is explicit and a human reviewer is named.

NOT READY
"Improve our content," no acceptance check, or no one who owns the output.

FIRST RUN
Small safe input, fresh task, output inspected and the version recorded — then decide on a schedule.

Guardrails your legal team can sign

Make the boundaries explicit up front, so brand, legal and compliance can approve the system once — instead of debating every asset from scratch.

Follow your company's AI policy first. Where this page and your policy differ, your policy wins. This is a starting template, not legal advice.

FOUR DATA CLASSES

PUBLIC METHOD

Published articles, public talks, approved case studies, the methods in this guide.

Who sees it: anyone

Fine in your approved AI tool.

TEAM RECORD

Job cards, run receipts, approved drafts, the editorial calendar, accepted corrections.

Who sees it: your team

Approved tool, team workspace.

CONFIDENTIAL COMPANY & CUSTOMER DATA

Customer lists, CRM exports, patient or client information, pricing, unreleased products, financials.

Who sees it: only people authorized for that data

Only in systems approved for that data class. Never a personal AI account.

PERSONAL

Employees' and experts' private details, messages and preferences.

Who sees it: that person only

Not without that person's consent.

Publish the lesson. Never publish the private context that produced it.

WHO SIGNS WHAT: A STARTING APPROVAL MATRIX

WHAT'S GOING OUT	BRAND	LEGAL	COMPLIANCE / MLR	THE SME'S OWN SIGN-OFF
Educational article from an SME interview	●	○	○	●
Clip quoting the expert	●	–	○	●
Customer or patient story (with written consent)	●	●	●	●
Product, clinical or performance claim	●	●	●	●
Executive or expert LinkedIn post	○	–	○	●
Paid promotion (plus the budget owner)	●	○	○	●

● always ○ when it applies — regulated industries, claims, comparisons – usually not needed. Adapt to your company's approval path.

AI drafts. Humans approve. Receipts prove.

EVERY JOB LEAVES A RECEIPT

- Source + input scope
- Action taken
- Output URL
- Before / after
- QA performed
- Human gate: who approved
- Time, model, cost when known
- Next owner
- Rollback or rerun path

Unknown cost stays UNKNOWN. A receipt your compliance team can read is the fastest route to the next yes.

RUN: a real cadence and a readout your manager will read

A prompt that says “do this daily” doesn’t prove anything runs daily. Give each recurring job a real trigger, an owner and a receipt — then report what changed in Metrics, Analysis, Action.

CADENCE	JOB	WHAT GOOD LOOKS LIKE
DAILY	**Inbox and exception triage**	Inspect authorized messages; draft follow-ups; avoid duplicate sends.
DAILY	**New-source intake**	Find permissioned recordings; create source-map drafts, not automatic publication.
WEEKLY	**Content and business MAA**	Read consistent metrics; explain uncertainty; propose one bounded next action.
ON CHANGE	**Update canary**	When a prompt, skill or tool version changes, test it on a known input before the next real run; record the version adopted.
ON FAILURE	**Missed-run review**	Keep the last good output; name the blocker, owner, recovery and next check.

A cadence is not a claim. For every recurring job, record where it is scheduled, its time zone, next run, output location and owner — and inspect the first real run before calling it operational. Check for an existing owner or overlapping job before you add another.

Your weekly MAA readout for your manager

METRICS

What happened?

Same window, same definitions, source named.

ILLUSTRATIVE WORDING

“Line-review requests: [n] this month vs. [n] last month — same form, same definition (form log).”

ANALYSIS

Why might it matter?

Evidence and uncertainty; attribution gaps named.

ILLUSTRATIVE WORDING

“[Asset] preceded most requests (UTM report). One month of data: a signal, not proof.”

ACTION

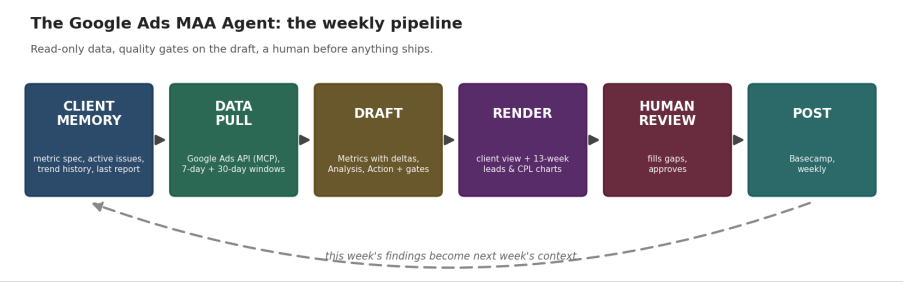
What changes next?

One change, owner, due date and approval gate.

ILLUSTRATIVE WORDING

“$7 Dollar-a-Day test on [clip]. Owner: [name]. Due: [date]. Gate: budget owner + compliance.”

Output count is not a business result. Inspect calls, qualified leads, pipeline and sales. “We published 12 assets” is Metrics; what those assets changed is the readout.



The Google Ads MAA agent’s weekly pipeline: client memory, read-only data pull, draft, render, human review and posting. It describes a workflow, not access granted to this reader.

→ The Google Ads MAA agent
→ How to measure SEO with MAA
→ Scheduled jobs, in practice

A 30-day pilot your manager can say yes to

Small, specific and reversible beats big and vague: one program, one expert, one hour and $7, with a readout that tells your manager to scale or stop.

WEEK 1

Pick and record

- Pick one program you own and one SME.
- Run the starter prompt (p. 2) and confirm GCT.
- Record the 15-minute interview (p. 8).
- File the release; start the evidence ledger; record a baseline.

WEEK 2

Process and approve

- Build the source map and derivative map.
- Draft 1 article, 3 clips and 1 email.
- Route one package through brand, legal, compliance and the SME.
- Log every correction as a reusable rule.

WEEK 3

Post and test

- Publish to 2–3 doorways: resource center, the expert's LinkedIn, email.
- Run destination QA on each live page.
- If approved, run one $7 Dollar-a-Day test on the clip that earned organic signal.

WEEK 4

Read out and decide

- Deliver the MAA readout (p. 17).
- Compare against the Week 1 baseline.
- Decide: scale, adjust or stop.
- Name the next owner and the next run.

THE ONE-SENTENCE PITCH

"Give me one expert, one hour and $7. In 30 days I'll bring back a reviewed article, three clips and an email built from our own expertise — plus a readout showing what moved and whether it's worth scaling."

One hour = a 15-minute interview plus the expert's review time across the month. Promise outputs and a readout, not results you can't yet prove.

Your mini business case

COST	$7 capped test (if approved) + your existing approved AI tool	
TIME	SME: about one hour in total. You: your honest weekly estimate.	
RISK	Drafts only until approved; no customer data; capped spend with a stop date.	
APPROVALS	Brand, legal, compliance/MLR where it applies, the SME, the budget owner.	
SUCCESS METRIC	One buyer action from your GCT, compared with your baseline period.	
DECISION DATE	Day 30: scale, adjust or stop — your manager decides with the readout.	

The career edge isn't "I used AI." It's "I ran a governed pilot, here are the receipts, and here is what it changed."

WORKSHEET / FILL THIS IN, THEN GIVE IT TO YOUR AI

Your one-page implementation brief

Complete it by hand, or paste it into your approved AI tool as the brief for your first job. Write UNKNOWN rather than guessing.

SCAN	AIM	MAP	ASSIGN	RUN
Sources	GCT	Proof gap	First job	Result

SCAN
SOURCES
Three inspectable sources; one internal source and its permission status

AIM
GOALS
One buyer action, baseline, target and period

AIM
CONTENT
The evidence and teaching that earn it; one missing buyer belief

AIM
TARGETING
One specific audience and the situation that makes the program relevant

MAP
ONE PROGRAM
Program, product line, service line or event; outcome and next step

ASSIGN
FIRST AI JOB
Input, output, acceptance check, permitted actions, human gate

ASSIGN
OWNERSHIP + NEXT RUN
Owner, reviewers, approval path, due date or real schedule; receipt location

RUN
RESULT + CORRECTION
What changed, what remains UNKNOWN, which rule improves next time

→ Starter prompt and current kit | → Master workshop library

Watch the method + resources



The Marketing Mechanic Framework: How It All Fits Together START HERE
Watch · Read



Content Strategy That Drives Sales | Ep. 3
Watch · Read



The AI Agent Framework for Entrepreneurs
Watch



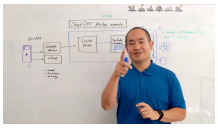
How to Actually Measure SEO | Ep. 4
Watch · Read



What I've Been Building for 20+ Years | Ep. 7
Watch



Look Over My Shoulders While My AI Agents Do Marketing
Watch



How I Turned 50 iPhone Videos Into YouTube, Blogs, and Ads in 24 Hours
Watch · Read



Recursive Self Improvement for AI Agents
Watch



How to Actually Use ChatGPT | Ep. 8
Watch · Read



The 9 Triangles Framework
Watch · Read

Videos open online. Timing and outcomes in video titles are the presenter's reported cases, not guarantees.

START

DigiMarCon kit
Master workshop library
Public-first start guide

CONTENT FACTORY

The Content Factory
AI-driven factory process
GCT strategy
The Topic Wheel

PROCESS GUIDES

Video repurposing guidelines
Definitive article guide

AI TEAM + MAA

Skill vs. agent
AI agents doing marketing
Scheduled jobs
Cross-agent shared memory
Skills on GitHub
Google Ads MAA agent

CASES + STORIES

Content Factory examples
Cody Jones four-site case
Cody Jones interview
Harry Gold interview
DigiMarCon Minneapolis recap
Honoring Damon Wee
Ken Okazaki
Marketing Mechanic hub



ABOUT DENNIS YU

Dennis Yu is CEO of Local Service Spotlight. He teaches the Content Factory at conferences worldwide.

SCAN FOR THE DIGIMARCON KIT
dennisyu.com/digimarcon

Adapted for DigiMarCon Mid-Atlantic (Philadelphia, 8 Oct 2026) from *The Authority Multiplier*, Tampa AI Mafia edition v1.0 (27 Sep 2026). Illustrative examples are labelled. Historical figures are snapshots, not guarantees. Digital companion; not an ISBN edition.