KSCPA ANNUAL TAX CONFERENCE · NOVEMBER 20, 2026

Claude for CPAs

From AI Assistant to Agentic Workflows

Session I · Understanding Claude · 2:30–3:20 PM CT

Marc Hamilton, CPA, CGMA

HOW THESE TWO SESSIONS RUN

Understand it first, then watch it work

- Session I (2:30): what Claude is, what it is good for, and where the professional lines sit.
- Session II (3:30): live demonstrations on a tax-season task, start to finish.
- Polling questions appear throughout. Please answer in the panel.
- Every client you see today is fictional. No real client data is used.
- A verified fact carries a number like [1]; a forward-looking view is labeled as judgment.

WHAT YOU'LL LEAVE HOLDING

Three ideas that outlast any one tool

The ladder

Answer, guide, act. Know which kind of tool is in front of you, and how much checking it needs.

The pre-flight

Five questions, thirty seconds, before any AI-assisted work leaves your hands.

The pattern

Prompt, instruction, skill. Write the procedure once, in plain English, and reuse it.

POLL 1

Where are you with AI today?

- ☐ I haven't tried it yet
- ☐ I've tried a free chatbot once or twice
- ☐ I use it weekly, but not on client work
- ☐ I use it on client work

Respond in the polling panel.

1

What is Claude?

THE PLAIN-LANGUAGE DECODER

Four words you will hear all day

Term	What it means
Large language model (LLM)	The engine underneath. It predicts and writes language.
Chatbot or assistant	An LLM you converse with. The label says nothing about quality.
Generative AI	Any AI that produces new content: text, images, code.
Agent, agentic	AI that takes a goal and carries out several steps on its own.

WHERE CLAUDE FITS

Three families of AI tools

General assistants

ChatGPT, Microsoft Copilot, Gemini, Claude. Broad, fast, inexpensive. Claude lives here.

Tax and accounting platforms

Built on licensed professional libraries: the Code, regulations, standards.

Embedded features

AI added to software you already run: tax, practice-management, and accounting systems.

Products are named as examples of a category, not endorsements.

CLAUDE, IN ONE SLIDE

A general assistant built by Anthropic

- Runs in a web browser, a desktop app, and on your phone.
- Individual plans: Free, Pro, Max and Ultra. Business plans: Team and Enterprise. [10]
- Your plan decides which features you get and what happens to your data. More on that in Section 4.
- It reads what you give it: text, PDFs, spreadsheets, images, and, on the desktop, folders you choose.

THE LADDER

Answer, guide, or act

	Chatbot	Copilot	Agent
What it does	Answers when asked	Guides you inside one app	Takes a goal and does several steps
Where your judgment sits	After you read the answer	Between every step	Only where you place a checkpoint
The risk lives in	The answer	The step you accept	The seams between steps

WHAT CHANGED IN SEPTEMBER

Chat and task are now one Claude

- On September 16, 2026, Anthropic combined its Cowork task mode and chat into one Claude. [8]
- Claude now works out what a task needs, from a quick answer to multi-step work, in one conversation. [8]
- By default it asks before taking an action. You can let it keep working and check in only when something needs a closer look. [8]
- Pro, Max and Ultra got it first; Enterprise since Oct. 23. Team has it now, on by default Dec. 1. Free: no date yet. [8][12]

KNOWLEDGE CHECK 1

A tool takes a goal and completes several steps before you see any of them. What is it?

- ☐ Chatbot
- ☐ Copilot
- ☐ Agent
- ☐ Depends on the vendor's label

Respond in the polling panel.

2

What can it do for a small firm?

TAX-SEASON WORK

Where it saves time in a tax practice

Client letters

Organizer cover letters, reminders, and plain-English explanations.

Missing-document follow-ups

Compare what arrived with what you expected; draft the request.

Notice summaries

Summarize an IRS or state notice with the identifiers removed.

Research starting points

A first map of an issue. Every citation gets opened and read.

Engagement letters

First drafts from your own template and terms.

Checklists and procedures

Turn how you do it into a written, repeatable checklist.

BEYOND TAX

Accounting, audit, and advisory work

Accounting

Variance explanations, reconciliation support, month-end checklists.

Audit

Request lists, workpaper indexes, first drafts of memos for review.

Advisory

Client newsletters, cash-flow explanations, meeting notes into action items.

Start where no client data is involved: internal drafts, public guidance, your own procedures.

DRIVING THE TOOL

A good prompt is an assignment, not a question

Part	What it does	Example
Role	Tells it who to be	"Act as a tax preparer at a small Kansas firm."
Context	Gives the situation	"The client is a married couple, first year with us."
Specificity	Says exactly what, in what form	"A numbered list, under 150 words."
Constraints	Sets the limits	"Plain English. Do not mention amounts. Flag anything you are unsure of."

THE DIFFERENCE IN PRACTICE

Same task, two prompts

Weak

"Write a letter about missing tax documents."

It has to guess the client, the tone, the length, and the list.

Strong

"Act as a tax preparer at a small Kansas firm. Draft a friendly email to a married couple asking for three documents: [list]. Plain English, numbered, under 150 words. Do not mention amounts."

Nothing left to guess. You review a draft instead of rewriting one.

POLL 2

Which task would you hand to Claude first?

- ☐ Client emails and letters
- ☐ A research starting point
- ☐ A checklist or procedure
- ☐ Cleaning up a spreadsheet

Respond in the polling panel.

3

Where does it go wrong?

THE CLAIM IT ALL RESTS ON

“

The most dangerous AI risk is not a wrong answer. It is a good one.

A wrong answer trips a control: a tie-out, a review note. A good-looking one walks right past it.

WHEN THE TOOL IS WRONG

Four ways it fails, and the check for each

Confident fabrication

A fluent, invented citation or figure, with no hedge. Check: open the source and read it.

Plausible but stale

Right-sounding, based on old law. Check: verify against current authority.

The subtly wrong number

A transposed input that still foots. Check: tie it back to the source document.

Hidden instructions

A file that tells the AI what to do. Check: treat files as evidence, not orders.

AN INDEPENDENT MEASUREMENT

17% to 33%

Stanford researchers tested leading legal research AI tools. They found the tools “each hallucinate between 17% and 33% of the time,” though they did better than a general chatbot. Claude was not tested. [9]

Judgment: grounding earns a tool a place in your practice. It does not earn your signature.

KNOWLEDGE CHECK 2

Claude cites a Treasury regulation in a research answer. What do you do first?

- ☐ Open the regulation and read the cited passage
- ☐ Trust it; it is a well-known model
- ☐ Ask Claude whether it is sure
- ☐ Paste it into a second AI tool

Respond in the polling panel.

4

The professional lines

THE RULES ALREADY FOUND YOU

Your existing duties apply to AI

- IRS guidance, Issue 2026-19 (June 24, 2026): “ethical obligations of competence, diligence, and confidentiality remain unchanged.” [1]
- Circular 230 §10.22, diligence: “verifying the accuracy of facts, citations, and calculations produced by AI.” [1][2]
- §10.35, competence: “understand both the law and the technology” you use. [1][2]
- §10.36, firm procedures: firms “must deploy internal policies and procedures” for AI. [1][2]
- SSTS §1.4, reliance on tools: “Use of a tool does not absolve the member of professional obligations.” [5]

Circular 230 is the rule. The June guidance explains how the IRS applies it to AI. More IRS guidance is expected in the coming months. [13]

WHERE DOES THE DATA GO?

Read the terms for the plan you are on

Individual plans: consumer terms

- The 2025 update covers Free, Pro and Max. You choose in Privacy Settings whether chats may be used to improve Claude. [6]
- If you allow it, data may be kept up to five years; if not, 30 days. [6]
- A privacy setting is not permission to share client data.

Business plans: Team, Enterprise

- Covered by Anthropic's Commercial Terms, not the consumer terms. [6]
- Your firm's contract governs. Read it before client data goes in.
- A contract is not a §7216 consent. Tax return information has its own rules. [3]

IRS: use "only secure, enterprise-approved AI" and "Never upload sensitive data to unsecured sites" [1]. Return information: IRC §7216 [3]. Client confidentiality: AICPA ET §1.700.040 [4].

WHERE IS THE LINE?

AI-assisted versus AI-performed

AI-assisted (defensible)	AI-performed (exposed)
A professional directed the tool	The tool decided; a professional forwarded
Output was checked against the source	No independent check against the source
A named person owns the result	Responsibility pushed to the tool or 'the team'
The file shows what a human checked	The signature asserts more than the file supports

THE PRE-FLIGHT CHECKLIST

Five questions before AI work leaves your hands

1. Is any client-confidential data going into an unapproved tool? If so, stop.
2. Am I competent to use this tool and judge its output?
3. Has a competent person checked the output against the source?
4. Did I verify the confident parts: the numbers and the citations?
5. Do I own this result, regardless of the tool?

KNOWLEDGE CHECK 3

On an individual Pro plan, when may Anthropic use your chats to improve Claude?

- ☐ Never
- ☐ Only if you allow it in Privacy Settings
- ☐ Always
- ☐ Only for uploaded files

Respond in the polling panel.

POLL 3

Does your firm have a written AI-use policy today?

- ☐ Yes, written and current
- ☐ A draft, or informal
- ☐ No
- ☐ I don't know

Respond in the polling panel.

NEXT: SESSION II

Programming business processes in English

Prompt

A one-time request

Instruction

Standing rules for every task

Skill

A written procedure, saved and reused

Workflow

Skills chained, with checkpoints

10-minute break

Session II · Claude in Practice · begins at 3:30 PM CT

KSCPA ANNUAL TAX CONFERENCE · NOVEMBER 20, 2026

Claude in Practice

Session II · From AI Assistant to Agentic Workflows

3:30–4:20 PM CT

Marc Hamilton, CPA, CGMA

THE IDEA

“

AI is beginning to let professionals program business processes in English.

Describe the procedure, the rules, the exceptions, the controls, and the output you expect, in plain language instead of code. (Judgment)

THREE LEVELS

Prompt, instruction, skill

	What it is	How long it lasts	Example
Prompt	A one-time request	One conversation	"Draft an email asking for these three documents."
Instruction	Standing rules for how work is done	Every task you attach it to	"Never change a source file. Never send anything."
Skill	A written procedure Claude follows	Saved, reused, improved	"Run the client intake review for Brown."

ANATOMY OF A SKILL

The same parts as a good SOP

Trigger

When to use it

Inputs

What it reads

Steps

What it does, in order

Rules

What it must and must not do

Exceptions

What to do when it doesn't fit

Checkpoints

Where it stops for a person

Output

What it produces, and where

In Claude

A folder with a SKILL.md file of instructions [7]

POLL 4

How are your firm's procedures documented today?

- ☐ In writing, and current
- ☐ Partly written down
- ☐ Mostly in people's heads
- ☐ Not sure

Respond in the polling panel.

BEFORE ANY PROMPT

Three moves come first

1 · The folder

One folder per client. Turn off connectors and web access the task doesn't need. [11]

2 · The standing rules

Preserve originals. Never send. Cite the source. Never invent a number.

3 · The method

The skill: the procedure in a file, not in my memory.

WHERE THE HUMAN STANDS

Three checkpoints

Checkpoint	Where it sits	What you inspect
Intent	At the start	Did it understand the job: client, year, deliverables?
Correctness	Right after the riskiest step	Did it read the documents correctly, before building on them?
Gate	In front of anything that can't be undone	Nothing is sent until a named person says go

DEMO SCENARIO · FICTIONAL CLIENT

Tax-season intake for the Smith family

In the folder

- Seven documents that arrived this season
- Last year's document summary
- This year's client organizer

What we'll ask for

- An inventory of every document
- A missing-items list, with the reason for each
- A draft email to the clients, not sent

Watch for: where I stop it, what I check, and what it refuses to do.

LIVE DEMO 1

The Smith intake, in plain English

1. Attach the demo folder (fictional clients only)
2. Paste the standing rules
3. Give the task in one paragraph
4. Intent checkpoint: it restates the job
5. Correctness checkpoint: inventory before anything else
6. Finish: the review and a draft email, saved

All client data is fictional.

WHAT YOU JUST SAW

Four things worth noticing

It read every file

Including last year's summary and the client's organizer.

It caught a misfile

Another client's W-2. It stopped reading and carried no figures.

It showed its reasons

Every missing item traced to a prior-year line or an organizer answer.

It drafted, it didn't send

The gate stayed closed. A named preparer sends.

KNOWLEDGE CHECK 4

Which checkpoint sits in front of sending the client email?

- ☐ Intent
- ☐ Correctness
- ☐ Gate
- ☐ None is needed for a draft

Respond in the polling panel.

LIVE DEMO 2

Turn what we just did into a skill

1. Ask Claude to write the procedure as a skill
2. Compare it with the firm's reviewed version
3. Install the reviewed skill

All client data is fictional.

LIVE DEMO 3

Same skill, a second client

1. New task, same demo folder (fictional clients only)
2. One line: "Run the client intake review for Brown."
3. Check that the same checkpoints appear

All client data is fictional.

WHAT CHANGED

From a conversation to a procedure

Demo 1

- I wrote the procedure into the chat
- It worked once, for one client
- It lived in my head and my typing

Demo 3

- One line named the client
- The skill carried the steps, rules, and checkpoints
- It lives in a file anyone in the firm can review

ONE SENTENCE IS THE CONTROL

Loosen one rule and the answer goes wrong

The rule as written

- "Never infer that a document is not needed."
- Result: the brokerage 1099 and the child-care statement are flagged as missing.

The rule loosened

- "If last year's document is missing, assume the client no longer needs it."
- Result: both drop off the list. Confident, tidy, and wrong.

POLL 5

Which process would you write as a skill first?

- ☐ Client document intake
- ☐ Engagement letters
- ☐ Monthly bookkeeping close
- ☐ Notice responses

Respond in the polling panel.

WHERE THIS IS GOING

From one skill to an agentic workflow

One skill

A procedure you run

Chained skills

One's output feeds the next

Scheduled and connected

Runs on a calendar, reaches your systems

A checkpoint at every seam

More autonomy, more oversight

Two failures to design against: an early error built upon (compounding), and a slightly different job done well (drift).

A SAFE WAY TO BEGIN

Your first four workflows, in this order

1. A reversible, low-stakes task you can check completely, such as sorting and renaming files.
2. A data task with one clear correctness check, such as pulling figures from documents.
3. A drafting task with an intent check, such as an internal memo.
4. Only then, something that touches a client, with the gate closed until you say go.

THE NEXT 30 DAYS

A practice plan

Week 1

Try it on non-client work. Check your plan's privacy setting.

Week 2

Write one procedure as standing rules and a skill.

Week 3

Run it on fictional data. Note every correction you make.

Week 4

Review with a colleague. Decide whether it is ready for a controlled pilot.

KNOWLEDGE CHECK 5

When you let an agent run a task, what are you accountable for?

- ☐ Only what it reports as an error
- ☐ What you let it reach, where it stops for you, and the final review
- ☐ Nothing; the vendor is responsible
- ☐ Only errors that reach a filing

Respond in the polling panel.

POLL 6

After today, what is your next step?

- ☐ Try it on non-client work this week
- ☐ Write or update our AI-use policy
- ☐ Write one procedure as a skill
- ☐ Wait and watch

Respond in the polling panel.

AI can do the steps.

A named human still owns the result.

GO DEEPER IN 2027

The LearnVoyage AI series for CPAs

The AI-Driven CPA

Strategy, governance, and ethics

The Authoritative AI Toolkit

Benchmarking the AI landscape

Mastering Workflow Automation

AI agents and efficiency

The AI Practice Lab

Deployment and implementation

Marc Hamilton, CPA, CGMA · marc.hamilton@learnvoyage.com · 205-799-8285

Questions

Ask in the panel

REFERENCES

Sources cited in this program

[1] Internal Revenue Service. (2026, June 24). Introductory Guidelines for Responsible AI Use in Federal Tax Practice (Issue No. 2026-19). https://content.govdelivery.com/accounts/USIRS/bulletins/41d6e70 ; PDF: https://www.irs.gov/pub/opr-taxpros/issue-number-2026-19-introductory-guidelines-for-responsible-ai-use-in-federal-tax-practice.pdf

[2] U.S. Department of the Treasury. Treasury Department Circular No. 230 (Rev. 6-2014), §10.22 Diligence as to accuracy, §10.35 Competence, §10.36 Procedures to ensure compliance. https://www.irs.gov/pub/irs-pdf/pcir230.pdf

[3] IRC §7216 (criminal) and §6713 (civil penalty); Treas. Reg. §§301.7216-2 and 301.7216-3 (consent). https://codes.findlaw.com/cfr/title-26-internal-revenue/cfr-sect-26-301-7216-3/

[4] AICPA Code of Professional Conduct, ET §1.700.001 and §1.700.040, Disclosing Information to a Third-Party Service Provider. Text checked against the Code as updated through July 2026. https://www.aicpa-cima.com/resources/landing/aicpa-ethics

[5] AICPA. Statements on Standards for Tax Services (effective Jan. 1, 2024), §1.3 Data Protection and §1.4 Reliance on Tools. https://www.aicpa-cima.com/resources/download/revised-statements-on-standards-for-tax-services-no-1-4-1-1-2024 ; quoted in Holets, D. J. (2025, Sept. 30), The Tax Adviser.

[6] Anthropic. (2025, Aug. 28). Updates to Consumer Terms and Privacy Policy. https://www.anthropic.com/news/updates-to-our-consumer-terms

[7] Anthropic. (2025, Oct. 16; updated Dec. 18, 2025). Equipping agents for the real world with Agent Skills. https://www.anthropic.com/engineering/equipping-agents-for-the-real-world-with-agent-skills

REFERENCES (CONT.)

Sources cited in this program

[8] Anthropic. (2026, Sept. 16; updated Oct. 23). Claude Cowork and chat are now one Claude. https://claude.com/resources/articles/cowork-is-now-claude

[9] Magesh, V., Surani, F., Dahl, M., Suzgun, M., Manning, C. D., & Ho, D. E. (2025). Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools. Journal of Empirical Legal Studies, 22(2), 216–242. https://doi.org/10.1111/jels.12413

[10] Anthropic. Claude plans and pricing. https://claude.com/pricing (accessed Nov. 13, 2026).

[11] Anthropic. Get started with Claude Cowork. https://support.claude.com/en/articles/13345190-get-started-with-claude-cowork

[12] Anthropic. Set up and manage the new unified Claude experience for your Team organization. https://support.claude.com/en/articles/18342806-set-up-and-manage-the-new-unified-claude-experience-for-your-team-organization

[13] Waggoner, M. (2026, Nov. 9). IRS official signals additional guidance on AI, practitioner fees. Journal of Accountancy. https://www.journalofaccountancy.com/news/2026/nov/irs-official-signals-additional-guidance-on-ai-practitioner-fees/

[14] Coverstone, M. B. (2026, Oct. 31). AI risks to the Sec. 7525 practitioner privilege. The Tax Adviser. https://www.thetaxadviser.com/issues/2026/oct/ai-risks-to-the-sec-7525-practitioner-privilege/

To keep up: AICPA Artificial Intelligence (AI) Tax Resource Center. https://www.aicpa-cima.com/resources/landing/artificial-intelligence-ai-tax-resource-center

All sources checked Nov. 13, 2026. Framings drawn from Marc Hamilton, LearnVoyage AI series Sessions 1–4 (2026) and The Next Standard. Judgments are labeled as such.