

BMSS Family of Companies Presents:

AI Ready Train. Adopt. Thrive.



TUESDAY
AUGUST 18



8:30 A.M. TO 9:30 A.M. CDT

AI



Taylor DeCelle
BMSS Advisors & CPAs



Brian Jackson
Abacus Technologies



Micah Murphy
BMSS Advisors & CPAs



Using AI Tools Daily

Creating a baseline for safe, everyday use

■ OPEN VS. CLOSED SYSTEMS

- **Open systems (free)** — allow prompts and data to be introduced into the model and used for learning and enhancement of results.
- **Closed systems (paid*)** — provide a private platform where data introduced is not used for learning within the model.

■ SECURE CONFIGURATIONS

- As with any technology, there is inherent risk associated with it.
- AI platforms have security settings and configurations that must be reviewed and configured to minimize your risk.

TAKEAWAY

Know which system you're using and what happens to your data — then lock down the settings before you rely on it.

Using AI Tools Daily

Establishing policy and shared ownership

■ AI POLICY AND GOVERNANCE

- Organizations should already have an AI Policy in place.
- **Authorized systems** — for business use.
- **Assessment process** — for evaluating new tools.
- **Professional use** — and responsibility.

■ AI ADVISORY OR STEERING COMMITTEE


- Establish a committee of stakeholders and users to help evaluate and share their experience with AI within the organization.

TAKEAWAY

***Policy sets the guardrails.
A standing advisory
committee keeps them
current as new tools and
use cases emerge.***

Integrated AI Governance Framework

FROM TRUSTED DATA TO ACCOUNTABLE DECISIONS



Data Governance

“Can the system trust its inputs?”

■ THE PILLARS, NOW EXTENDED

- Classic pillars still apply — quality, lineage, privacy.
- “Data” now also means runtime context — retrieved documents, embeddings, conversation history, agent memory, tool responses.

■ NEW QUESTIONS TO ASK

- Was this input authorized for this purpose?
- Was it current?
- Can we trace exactly what was used?
- Could it have been manipulated?

KEY LINE

Context lineage is now as important as data lineage. If you can't reconstruct the context, you can't explain the decision.

Model Governance

“Can the system trust its intelligence?”

■ THE GOVERNED SURFACE IS WIDER

- Same four concerns still apply — bias, explainability, drift, stability.
- But governance must now also cover — prompts, tool-selection logic, agent planning, guardrails, model routing.

■ THE BIG MISTAKE

- Treating approval as a one-time event.
- The model can stay frozen while the data, prompts, or retrieval index around it shifts.

KEY LINE

Shift the mindset from static approval to continuous assurance. Not “was it valid at launch” — “is the whole system still behaving acceptably right now.”

Risk Management

“How much exposure will we accept?”

■ NOT EVERY USE CASE DESERVES THE SAME SCRUTINY

- **Low-risk:** Summarizing a long document — if it's wrong, a human catches it before anything happens.
- **High-risk:** Auto-sending a payment or auto-editing a production system — if it's wrong, the damage is already done.

■ RISK TIERING (IMPACT · REVERSIBILITY · AUTONOMY · SENSITIVITY · NOVELTY)

...should directly set:

- Validation depth
- Approval authority
- Monitoring frequency
- Permitted autonomy

KEY LINE

Risk appetite is only real if it's enforced at runtime. A policy statement that isn't encoded into the system has no teeth.

Decision Governance

“Who or what is authorized to act?”

■ THE LEAST MATURE PART OF MOST PROGRAMS

Organizations govern data and models tightly — but let final decision logic scatter across code, spreadsheets, and human convention.

■ ALIGN HUMAN REVIEW TO CONSEQUENCE

Human-in-the-loop shouldn't be a blanket rule bolted onto every interaction — that kills automation economics.

■ CRITICAL DISTINCTIONS

- A model output isn't a decision.
- A probability isn't an approval.
- An agent's ability to call a tool isn't permission to act.

AUTOMATE

Low-risk, reversible



GUARDRAIL

Moderate-risk




REQUIRE APPROVAL

High-risk, irreversible,

Policy → Controls → Evidence → Accountability

Policy defines intent — but intent alone doesn't control behavior.



■ ACCOUNTABILITY, DIVIDED

Tech owns platforms · Data teams own products · Risk owns independent challenge
— *but the business owns the decision and its outcome.*

Agents don't just predict — they act.

Agents combine reasoning with execution, which opens new failure modes.

■ NEW FAILURE MODES

- Goal hijacking
- Prompt injection
- Wrong tool selection
- Excessive permissions
- Compounding multi-step errors

KEY LINE

***Agents combine reasoning
with execution — they
don't just predict, they act.***

**AGENTIC AI TURNS GOVERNANCE FROM A DOCUMENTATION DISCIPLINE
INTO A RUNTIME ARCHITECTURE.**

MEET CLAUDE

Claude, in plain terms

Cut through the hype: what Claude actually does, and how to put it to work.

Taylor DeCelle & Micah Murphy · AI Advisory · August 18, 2026



What we'll cover

1

How AI reads your words

Tokens and the context window, in plain terms

2

Three ways to work with AI

Chat, Cowork, and Code, plus when to use each

3

What else it can do

See your files, build deliverables

4

Skills and plugins

Bundle the right tools by team

→ Then: a live demo on real numbers

How AI reads your words: tokens

A PERSON READS

“Will we make payroll?”

Sees words, grammar, and meaning all at once.

CLAUDE READS

Will

we

make

pay

roll

?

→ split into tokens (word-chunks)

→ then turned into numbers it can compute on.

Size sense: ~100 tokens ≈ a short email · ~2,000 tokens ≈ a 3–4-page document

The more advanced or longer the task, the more tokens it uses.

The context window: the AI's working memory

What it is

The context window is how much the AI can hold in mind at once: your prompt plus everything in the conversation. Go past it, and the earliest details start to drop off.

Think of it like a desk: only so much fits before older pages get pushed off.

How we manage it

- **Projects: set custom instructions and key info once, per client**
- **Point it at files on your computer or in a Project, no re-uploading**
- **Split big engagements into focused sessions**
- **Fresh chat when the task changes**

Three ways to work with AI

Same AI, three doors in. Pick the one that fits the task.

Chat

Ask, draft, and analyze in a conversation.

EXAMPLE

Summarize a report; draft a client email.

WHO IT'S FOR

Everyone

Cowork

Assign multi-step work on your files and tools.

EXAMPLE

Build a 13-week cash forecast from a file.

WHO IT'S FOR

Business and ops users

Code

Write, review, and debug software.

EXAMPLE

A script that merges monthly sales files.

WHO IT'S FOR

Developers and IT

Consumption is measured in tokens: bigger, more complex tasks simply use more.

BMSS Family of Companies Presents: AI Ready – Train. Adopt. Thrive.

Pick the right horsepower

Four model tiers. You can switch anytime.

Haiku

Fast & light

WHEN TO USE

Quick lookups, sorting invoices, tagging emails.

Sonnet

Daily default

WHEN TO USE

Drafting deliverables and analyzing reports.

Opus

Deep reasoning

WHEN TO USE

Complex analysis and high-stakes judgment calls.

Fable

Most capable

WHEN TO USE

Long, multi-step work it can run on its own.

Rule of thumb: start on Sonnet. Step up to Opus or Fable for complex, multi-step work; drop to Haiku when speed and volume matter most.

BMSS Family of Companies Presents: AI Ready – Train. Adopt. Thrive.

It reads more than text

Drop files and images in as-is. Claude reads them.

Documents

PDFs, Word, contracts, scanned pages, even handwriting.

EXAMPLE

Pull the key clauses from a 30-page contract.

Spreadsheets

Excel and CSV: it reads the numbers, not just the words.

EXAMPLE

Explain what changed in this month's P&L.

Images & photos

Screenshots, charts, a photo of a receipt or whiteboard.

EXAMPLE

Snap a receipt; it captures the totals.

No retyping, no reformatting. It works from what you already have.

BMSS Family of Companies Presents: AI Ready – Train. Adopt. Thrive.

It builds things you can use

Ask, and Claude builds a real, editable deliverable right in the chat, then refines it with you.

Documents & letters

Drafts you edit line by line.

Tables & checklists

Structured and ready to use.

Calculators & dashboards

Small interactive tools, live.

Charts & one-pagers

Visuals built from your numbers.

You edit it in place and keep it. No starting from a blank page.

BMSS Family of Companies Presents: AI Ready – Train. Adopt. Thrive.

It can run on a schedule

Set it once. Claude does the work; no one has to kick it off.

Mon 7am

A cash, sales, and pipeline brief waiting in your inbox

Daily

Overdue invoices flagged and follow-up emails drafted

Month-end

Books reconciled and a close summary drafted

It runs unattended. The answer is ready before anyone asks.

Plugin vs Connector vs Skill

Connector

The live link to one app. It handles the login and permissions so Claude can read or act there.

EXAMPLE

QuickBooks, Outlook, Slack

Skill

The how: plain-English instructions for one specific workflow.

EXAMPLE

Cash-flow forecast, meeting recap

Plugin

The bundle: a set of skills plus the connectors they use, in one click.

EXAMPLE

Claude for Small Business

Connector = where your data lives. Skill = what to do with it. Plugin = the bundle that brings them together.

Skills and plugins: bundle it by team

A skill is one workflow in plain text. A plugin bundles the skills and connectors a team actually uses, shared in one click.

Accounting plugin	Sales plugin
SKILLS Cash-flow forecast · Month-end close · Invoice follow-up CONNECTORS Outlook · Excel · QuickBooks	SKILLS Lead triage · Proposal drafting · Pipeline recap CONNECTORS HubSpot · Outlook · Slack

Most people don't need every skill, just the handful their role uses, packaged together.

BMSS Family of Companies Presents: AI Ready – Train. Adopt. Thrive.

Powerful, but you stay in control

You approve the important steps

Anything that touches money or customers waits for your OK.

Review it thoroughly

You are accountable for everything AI produces. Review it in full, every time.

Read-only by default

It can't act in a system until you connect it and grant access.

You own the output

You are responsible for whatever it submits under your name.

Now let's put it to work.

[Over to the live demo](#)

Questions?

Micah Murphy

Manager, AI Advisory Lead

BMSS Advisors & CPAs

micahmurphy@bmss.com

205.982.5587

Taylor DeCelle

AI Advisory Consultant

BMSS Advisors & CPAs

tdecelle@bmss.com

205.271.5439

Brian Jackson

Chief Executive Officer

Abacus Technologies

bjackson@abacustechnologies.com

205.443.5915

*If you have a question that you'd like to ask, please use the **Q&A Button** at the bottom of your screen.*

CPE certificates will be issued approximately two weeks after the presentation.



BMSS Family of Companies Presents:

AI Ready Train. Adopt. Thrive.



TUESDAY
OCTOBER 06



8:30 A.M. TO 12:30 P.M. CDT



Taylor DeCelle
BMSS Advisors & CPAs



Brian Jackson
Abacus Technologies



Micah Murphy
BMSS Advisors & CPAs

