

SWTA LEADERSHIP LAB · BRECKENRIDGE

Work Smarter. Market Better. In 90 Minutes.

Ninety minutes. Hands on keyboard. No experience needed.
You leave with an assistant you built yourself.

Crys Black · Tuesday, September 15, 2026 · Breckenridge, CO



WHAT YOU LEAVE WITH

01

A prompt formula for any AI tool

02

A custom assistant for your agency

03

A roadmap guide to take home

ChatGPT, Copilot, or Gemini.

Whichever your agency already uses.

THE PROMISE

By the end you own three things.

Not notes. Not inspiration. Three things that actually run.

01

A formula for better answers

Five things every good prompt has.
It works in any AI tool your agency uses, today and next year.

02

An assistant you built yourself

Your instructions, your files, your voice. Built for one real task at your agency.

03

A kit you take back home

Your roadmap from the Friction Finder plus the reference guide to redo this on Monday.

Zero prerequisites. If you have never opened an AI tool before, today is a fine day to start.



YOUR GUIDE FOR THE SESSION

Crys Black

Fractional CMO and AI Strategist

I started as a programmer. Then two English degrees, excellent preparation for talking to machines.

And yes, I used to drive a bus.

ALSO

Lead instructor, Women in AI Labs

Treasurer, Executive Board of Rocky Mountain AI

Contributor, Colorado AI News

You know your riders and your community.
I just know where the switches are.



WHAT I DO ALL DAY

Operationalize AI

For startups, agencies, and public-sector teams.

Cut through the hype

Making AI approachable and useful, not just impressive.

crys@crysblack.com

[linkedin.com/in/crystalblack](https://www.linkedin.com/in/crystalblack)

HOW THIS WORKS

The contract for the next 90 minutes.

Four rules. Three are for you, one is mine.

RULE 01

Every section has two levels

A starting move and a stretch move. Take the one that fits you today. Both count as finishing.

RULE 02

Questions go in the parking lot

Ask any time. If it pulls us off the build, it goes on the wall and we answer it before you leave.

RULE 03

You cannot break anything

Nothing you type today touches a live agency account. Wrong turns teach you the shape of the tool.

MY SIDE OF IT

I stay hands on keyboard with you

No lecture block runs longer than fifteen minutes. Raise a hand and I come to you.

Bring whichever AI your agency uses. ChatGPT, Copilot, or Gemini all work today.



HANDS UP

Where are you with AI right now?

Raise your hand for the one that is true. There is no wrong hand to raise.

Every day

You are my co-teachers.

Take the stretch move in every block, and float to whoever is stuck near you.

Now and then

This is built for you.

You know the chat box. Today we turn that into a system that remembers your agency.

Barely, or never

You have the best session ahead.

No habits to unlearn. Sit near a hand that went up first and copy shamelessly.

Why I ask. One room, three starting points. I need to know who to check on first.




RUN OF SHOW

Here is the plan.

Welcome, plus the mental models	ACT ONE
Prompting that works: GRACE Live demo, the framework, prompts you keep forever.	ACT TWO
Build your own custom assistant 30 minutes. Hands on keyboard.	YOU BUILD
Share-out. Screens up, show your build.	
What happens next The 30-day habit, the take-home kit, parking lot answers.	CLOSE
Done. You leave with working tools An assistant, a formula, and a guide to run it all again at home.	CLOSE

The Path we walk this session.

We come back to this at every block change, so you always know where you are.



Step one takes five minutes In your Friction Finder. **Then we spend the session on two through five.**

HANDS ON

Find Your Friction

Three things that eat your week. Capture them now, score them later.

Step 01

Open the Friction Finder.

Scan



or type

swta2026.crysblack.com

Step 02

Add three friction points.

Real ones, with names. Steal from the sparks if you are stuck.

Hold here

Step 03

Stop at the scoring screen.

Do not score yet. Do not pick the winner. We come back to it at build time.

Done early? Add a fourth. There is always a fourth.



ACT ONE

How This Works

Fifteen minutes of mental models. Then you build for the rest of the session.

IN THIS ACT

Chatbot to agent

The three layers

Custom assistants

Your data, one slide

Step one, Friction, happens in the Friction Finder



Chatbot, automation, agent.

Three different things and people use the words interchangeably. Here is the map.

CHATBOT

Stays in its lane

Answers what you ask. It will not go off and do anything on its own.

IN YOUR WORLD

A rider FAQ bot on your website.

AUTOMATION

Runs the steps you set

A trigger fires, your prompt runs. Wording varies, shape is the same every time.

IN YOUR WORLD

A weekly service alert draft. A social post queue.

AGENT

Decides what to do

You give it a goal, not a recipe. It picks the steps, uses tools, adapts.

IN YOUR WORLD

AI reading your ridership data and writing the report.



Today we build in the first column. You stay in charge of all three.



Three layers. You only work in one.

So you always know exactly where you are standing.

1

The Model

The engine

The AI brain itself. Claude Opus, GPT, Gemini. Somebody else trains it, somebody else pays the electricity bill.

You never touch this layer.

2

The App

Where you log in

ChatGPT, Copilot, or Gemini. The workspace that connects you to the engine.

Whichever your agency uses.

3

What You Build

Your custom tools

Custom assistants, saved prompts, templates. Built by you, for your agency, reusable forever.

This is where every minute of today goes.

Layer three is the only layer you own. Good news: it is also the only one that knows anything about your agency.



Instructions set the voice. Files hold the facts.

INSTRUCTIONS

The personality

Who your assistant is, how it answers, and what it never says.

- Written once, in plain English
- Holds your tone, rules, and standards
- Every chat starts already briefed

You stop re-explaining your agency in every chat.

KNOWLEDGE

The source material

The files your assistant reads before it answers anyone.

- Your FAQs, schedules, templates
- Your brand voice doc and past examples
- Update the file, the assistant updates

You stop pasting the same background every single time.

THE WHOLE EQUATION

**Instructions give the voice.
Knowledge gives the facts.**

Consistent answers, at a speed that feels unreal.

Same recipe, three names

Instructions plus your files. Every platform sells it under a different name.

ChatGPT

Projects & custom GPTs

A Project holds ongoing work. A custom GPT is the shareable assistant your whole team can open.

In your world

An Alert Writer GPT pinned in everyone's sidebar.

Microsoft Copilot

Agents

Built with Create an agent. The government-tenant friendly option, and it can read your M365 files.

In your world

An agent that checks your comms SharePoint before it answers.

Google Gemini

Gems

Write the instructions, attach the files, save the Gem. It waits in your sidebar after that.

In your world

A Board Report Gem that knows your ridership numbers.

Build in whichever one your agency already pays for. The recipe transfers; you now know how to cook it.



MENTAL MODEL FOUR OF FOUR

One slide on data safety.

Exactly what you would tell a new hire on day one. Same rules apply to AI.

Green

Go, no thinking needed

Brainstorms. Outlines. Drafts. Your own notes. Anything already on your public website.

Yellow

Anonymize first

Work in progress. Use [Name] placeholders. Strip personal details out before you paste anything in.

Red

Keep a human on it

Rider personal info. Personnel records. Anything confidential or bound for the public without review.

Read your agency's AI policy. Write down what it says. Then stop worrying and go build.



ACT TWO

Prompting That Works

Same model. Same account. Same seven seconds.
The prompt is the entire difference.

IN THIS ACT

The gap, live

GRACE

Prompts you keep

Then you build

**The Path: steps two, three,
and four.**



THE FRAMEWORK

GRACE: five things a good prompt has.

Use it when you are creating something you would put your name on.

G

Goal

What am I trying to achieve?

"Be ready for Tuesday's board meeting."

R

Role

Who should it be while it answers me?

"A transit comms pro who knows riders."

A

Action

What exactly should come out the other end?

"A one-page brief I can read in four minutes."

C

Context

What raw material am I handing over?

The export, the transcript, the actual file.

E

Example

What does good look like, in your hands?

Two pieces your team was proud of.

The two filled circles are the ones nobody does. Next slide is entirely about them.

One habit on top: say who the output is for.

Same ask. Two prompts. Seven seconds apart.

Read them both. Next slide, I run them live.

PROMPT A · WHAT MOST PEOPLE TYPE

Write a service alert.

WHAT YOU GET BACK

- A stiff, generic announcement
- Confident tone, zero specifics
- Nothing you would actually post

Elapsed time: seven seconds.

PROMPT B · THE SAME ASK, IN GRACE

Role, plus goal, plus the real detour info and a past alert pasted in.

WHAT YOU GET BACK

- Your route numbers and dates, correct
- Your agency's voice, not a robot's
- Short enough to actually get read

Elapsed time: also seven seconds.

Same model. Same seven seconds. The only variable is the prompt.



Watch me run it - RTD, live.

A real transit agency with public board docs. Public data, green zone.

ROUND 1 · THE WEAK ASK

Tell me about RTD.

ROUND 2 · THE SAME ASK, IN GRACE

- Goal: brief me for a partnership meeting
- Role: transit comms director prepping a brief
- Context: their board packet PDF, attached

Example: one past brief we were proud of

WHY ROUND 2 WINS

Watch for these three moments in the output.

THE LETTERS DOING THE WORK

- R gave it a job, not a topic
- C made it use their words, not its memory
- E set the bar for what finished looks like

Total elapsed: about fourteen seconds.

The model did not get smarter. The prompt did.



SAME ASK, BUILT UP ONE LETTER AT A TIME

Watch a bad prompt become a good one.

WHERE MOST PEOPLE STOP

Write a press release about our new route.

You get generic fluff. You rewrite most of it. You blame the AI.

WHY IT UNDERPERFORMS

It has a Goal and nothing else. The AI has to guess the reader, the format, the standard, and the agency. It guesses average.

THE RULE OF THUMB

If a new hire could not do the task from your prompt, the AI cannot either. Same missing information, same shrug.

THE SAME ASK, IN FULL GRACE

- G** A press release our local paper will actually run, announcing the new Route 12 service.
- R** You are a transit communicator who cares about riders and about being worth reading.
- A** Give me a headline, a lead, two quotes to attribute, and what to leave out.
- C** Pasted below: the route details, launch dates, and our style guide.
- E** Also pasted: the two releases of ours that got picked up last year. Match their shape.

Ninety extra seconds of typing. A release you can send.



The service alert writer.

Give this one to your assistant in the build block and never write it from scratch again.

GOAL A rider alert about [DISRUPTION] that answers questions before they get asked.

ROLE You are a transit communicator who is clear first and clever second.

ACTION Give me the what, when, where, the rider impact, and what NOT to include.

CONTEXT [Paste: the detour details, dates, routes, and your style guide]


EXAMPLE [Paste: your best past alert. Match the structure exactly.]

Copy it any time: swta2026.crysblack.com/prompts



YOU ARE HERE ON THE PATH

Steps three and four. Build then Connect.



WHAT CHANGES RIGHT NOW

I stop talking. You open your AI tool, and for 30 minutes you build.



LIVE DEMO · TEN MINUTES, WATCH ONLY

One assistant, built in front of you.

Then you do these exact three moves for your own agency.

1

Start the assistant

ABOUT THREE MINUTES

A custom GPT, Gem, or Copilot agent for a fake transit agency. I name it and write three lines about its job.

2

Add instructions

ABOUT FOUR MINUTES

The service alert prompt from earlier, pasted in as standing instructions. Watch it become the default.

3

Add the files

ABOUT THREE MINUTES

The rider FAQ and style guide, uploaded once. Now every answer uses your words, not its guesses.

Call out anything you want me to slow down on. That is what the demo is for.



Two ways to get your real data in.

One is elegant. One is boring and works every single time. Both are correct.

PATH A · UPLOAD ONCE

The assistant holds your files

Upload the FAQ, the style guide, the schedule PDF once. Every chat after that can read them.

Today we use: one or two files you brought along.

Heads up: if uploads are locked down at your agency, Path B always works.

PATH B · EXPORT AND DROP

You paste it in the chat

Copy the text of the doc, paste it into the chat, ask your question. Old fashioned and unbeatable.

Works in every AI tool ever made. No permissions, no waiting.

If you are new today, start here. Nobody can tell the difference in the output.

YOUR RAW MATERIAL, EITHER WAY

Rider FAQs · Route schedules · Alert templates · Style guide · Board reports · Press releases



Your agency. Your assistant.

I am circulating the whole time. Raise a hand the second something looks wrong.

1 Pick your friction

Look at your Friction Finder. Finish scoring friction. Pick one of those friction points!

2 Start your assistant

A GPT, Gem, or Copilot agent. Name it and write its job in three plain sentences.

3 Give it your words

GRACE goes in the instructions. Add a file or paste in the background it needs.

4 Run it twice

Two real tasks. Fix the biggest miss. That second run is where it becomes yours.

STARTING MOVE, IF YOU ARE NEW

Steps one and two only. A named assistant with a job is a complete win.

STRETCH MOVE

Give it a second task and tighten the instructions.



THREE DEMOS FROM THE ROOM

Tell us what you built.

Three volunteers, two minutes each. Everyone else shares with the table next to you.

01

What did you build?

One sentence on what it does. Not how, not the struggle.

02

What surprised you?

The good, the weird, and the “huh, that actually worked.”

03

When will you first use it?

Name the actual day. Tools used in the first week are the ones that survive.

GROUND RULE FOR DEMOS

Broken counts. Half finished counts. Ugly definitely counts. Nobody apologizes for their tool in this room.



WHAT HAPPENS NEXT

Keep the muscle warm.

This sticks when you use it. It fades in about ten days if you do not. Here is the plan and the kit.

THE 30-DAY HABIT

One task a week. Four by October.

- Small and finished beats big and someday
- Tell one coworker about each one, even the ugly ones
- Teaching a teammate is what cements it
- Week four, open your Friction Finder roadmap

YOU ARE LEAVING WITH

Four things, all of them real.

- The reference guide, every prompt from today
- The template builder that writes your next ones
- A working assistant built for your agency
- A friction roadmap you can rerun any time



LET'S KEEP TALKING

What to call me for

Today was the free sample. Here is what the next call looks like.

01

A workshop for your agency

This session, sized to your team. Half day or full day, hands on keyboard, everyone leaves with something running.

02

An AI Roadmap Sprint

We run the Friction Finder across your whole team, set your data guardrails, and ship your first three builds together.

03

An advisor on call

Monthly 1:1 time for you or your comms team. Bring what you are building, leave unstuck.

Book 15 minutes:
app.usemotion.com/meet/crys-black. Or send your roadmap from the tool.





Let's keep building.

Crys Black

Fractional CMO and AI Strategist · Colorado

crys@crysblack.com

linkedin.com/in/crystalblack

AVAILABLE TO YOU

These slides, the Friction Finder, both available at swta2026.crysblack.com.



BEFORE YOU GO

Trade one idea with your neighbor.

A room full of transit people who build things. That is the network you came for.

THE 30 DAY HABIT

One task a week. Four by October.

Thank you for a wonderful session.
A room of people who build things now.

Nine words, in plain English.

Custom assistant

A chatbot you set up once with instructions and files. A GPT, a Gem, a Copilot agent.

Instructions

The standing orders your assistant follows. Who it is, how it answers, what it avoids.

Knowledge

The files you upload for your assistant to read first. FAQs, schedules, style guides.

Prompt

What you type to the AI. The better the prompt, the better the answer.

Model

The AI brain under the hood. GPT, Claude, Gemini. Your app is a window onto one.

Context window

How much the AI can hold at once. Long chats get foggy. Start fresh when that happens.

Hallucination

A confident wrong answer. The reason you paste in real data instead of trusting recall.

Token

Roughly a word chunk. Only matters because it is the unit behind usage limits.

GRACE

Goal, Role, Action, Context, Example. The five things a prompt needs when the output matters.

