wundertre

REV UP YOUR BUSINESS WORKSHOP

Leverage AI for 10x output

Ryker Taylor · Wundertre · September 17, 2026

Customer-exclusive session · This session is being recorded

Tonight

Section 1	Three AI concepts that carry forward
Section 2	AI-empowered humans > AI agents
Section 3	Prompts & prompt engineering
Q&A	The last 30 minutes are yours

One worksheet. One live build. Three things to do Monday. Bring a pen.

10x doesn't come from the tool. It comes from the operator.

LIVE POLL

Honestly — how do you use AI today?

- A Never
- B Occasionally, for drafts
- C Daily, for real work
- D We've automated things with it

The 10x claim, honestly

- 10x is a real number and a specific one. Not "your business grows 10x." **"The task that took an hour takes six minutes."**
- The tasks where that's true aren't exotic: first drafts, summaries, structuring messy information, first responses, analysis nobody was going to do anyway.
- The catch: it only shows up for people who know how to operate the thing. The same tool in unskilled hands produces 1x output that *looks* like work.

- **So the question isn't "should we use AI." It's "who on our team is the operator, and how do we make them good."**

01

Three AI concepts that carry forward

True last year. True tonight. True in eighteen months.

Why concepts, not tools

- Every AI tool that will exist next year is a recombination of the same few capabilities.
- Learn the capabilities and a new tool takes a week to evaluate instead of a quarter.
- Still no tool list tonight. But by Section 3 you'll be able to make any tool on the list actually work — which is worth more.

Concept 1 — Capabilities, not products

Generation | **Extraction & structuring** | Analysis & reasoning | **Retrieval** | Agency

THE TWO THAT MATTER MOST FOR 10X

Extraction & structuring is the sleeper.

Almost nobody uses it, and it's where the most hours hide. Every call, email thread, and PDF that someone reads and re-types into a system.

You have a person whose job is partly being a human OCR machine. That job is over — and they'd be thrilled.

Retrieval is what makes the other four trustworthy.

An AI grounded in your actual documents — pricing, process, past proposals — is a different animal from one working off general internet knowledge.

LIVE EXERCISE · 2 MINUTES

Sort these — which capability is each one?

- Following up with a lead who went quiet ----------
- Writing the monthly newsletter ----------
- Figuring out why three customers churned ----------
- Getting every sales call into the CRM ----------
- Answering "what are your hours?" at 11pm ----------

Most people give AI a task and no context, get generic output, and conclude AI is generic. It isn't. It was working with nothing.

Concept 2 — Context is the moat: the four kinds

In order of how much they change the output.

1 · Identity context

Who you are, who you serve, how you sound. Brand promise, voice, personas.

Changes tone and framing.

2 · Reference context

Your actual material. Past proposals, best emails, pricing sheet, process docs.

Changes accuracy. This is the big one.

3 · Situation context

What's happening right now. This lead, this complaint, this goal.

Changes relevance.

4 · Example context

Two or three samples of what "good" looks like for this task.

Changes quality more than any other single thing.

As models converge, context becomes the only difference between you and a competitor using the same tool.

Concept 3 — The delegation ladder, operating manual

1	**Assist**	Human approves every output. *Start every use case here. No exceptions.*
2	**Automate**	Human spot-checks a sample. Only once you know its failure modes.
3	**Delegate**	Human reviews exceptions only. High-volume, low-stakes, reversible work only.

BEFORE YOU MOVE ANYTHING UP A RUNG

Have we watched it fail? · Is the failure visible? · Is it reversible? · Who owns the standard?

You don't promote a task because the tool got better. You promote it because you've seen it fail and you know what failure looks like.

Where the 10x actually lives

- Everyone wants to get to rung 3. **Almost all the real productivity is happening on rungs 1 and 2 — with a human in the loop.**
- That's not a limitation. That's the whole point.
- The person on rung 1 catches what an agent on rung 3 would have compounded.

- **Which raises the question: who's doing the handing over?**

LIVE POLL

Who on your team is best at using AI right now?

- A The owner
- B One specific person
- C Nobody, really
- D Everyone, about equally

02

AI-empowered humans > AI agents

The most valuable thing you can build this year isn't an agent

What agents are good for — and why they underperform in a business your size

Agents are real, they're improving fast, and they're excellent at high-volume, well-defined, low-stakes, reversible work. Rung 3 work.

The problem: the hype sells rung 3 to businesses that haven't done rung 1. Four reasons it goes wrong:

1 · They compound errors.

A human makes a mistake and notices. An agent makes a mistake and builds the next five steps on it.

2 · They don't have your context.

They don't know this customer is difficult, this quote is a favor, you never discount in Q4.

3 · Nobody owns the output.

"The AI did it" is not an answer a customer accepts.

4 · They're expensive to get right.

Setup, supervision, cleanup. The demo hides all of it.

The most valuable thing you can build this year isn't an AI agent. It's an AI-empowered person.

The AI-empowered human, defined

A person doing the same job, with the same judgment, at 3–10x the throughput — because AI handles the first 80% of every task and they handle the 20% that needed a human anyway.

THE 80%

Drafting, gathering, structuring, summarizing, first-passing.

THE 20%

Deciding, relating, judging, approving. **It was always the valuable part.**

- Why this beats the agent: the human catches the error, has the context, and owns the output. Three of the four failure modes gone — for the price of keeping one person in the loop.
- **The honest cost:** it doesn't remove headcount. It changes what headcount does.

Before → after

NAME / ROLE

BEFORE

AFTER

What a 10x operator does differently

1 · Starts by deciding what the AI does first.

Not "should I use AI" — "what's the 80% here."

2 · Feeds it context before the task.

A brand doc, two examples, the actual situation — before a single instruction.

3 · Treats the first output as a draft, never an answer.

Edits, pushes back, asks for three versions. The second prompt is where the quality is.

4 · Builds a library.

Every prompt that works gets saved. After a month, nothing starts from zero.

5 · Reviews with a standard.

Knows what "good" looks like and rejects what misses it.

6 · Knows what not to hand over.

Pricing, relationships, crisis, anything under their own name.

None of that is technical. Every bit of it is learnable in a month.

Building the role — the owner's job

- **Pick one person, not the whole team.** One operator who gets good, documents what works, then teaches. Everyone at once = no one.
- **Pick the right one.** Not the most technical. The most *curious*, with the highest standards for their own work. Taste and patience beat technical skill every time.
- **Give them one specific workflow to 10x.** Not "use AI more." The monthly report. The lead follow-up. The proposal process. With a before-and-after measure.
- **Give them 30 days and protect the time.** Fails as a side project. Works as Tuesday afternoons, blocked, for a month.

- **Then they build the library and train the second person. That's how one operator becomes a team.**

Where agents *do* belong

- Once a task has lived on rungs 1 and 2 long enough — every failure seen, high-volume, low-stakes, reversible — *then* it's agent work.

Speed-to-lead first response | Appointment reminders | Tier-1 support | Data entry from structured sources

- **An agent built by someone who's been the operator for that task will work. An agent built to skip being the operator won't.**

Agents are a graduation, not a starting point. Build the person first.

03

Prompts & prompt engineering

A method, a template, and one prompt you build yourself

A prompt is not a question. It's a brief.

- The same thing you'd hand a new contractor on day one.
- *Give a freelancer a Post-it that says "write a newsletter" and you'd get garbage — and you'd know it was your fault. That's what most people do to AI, then blame the AI.*
- Everything in a good brief is context plus instruction. That's the whole method.

- **"Prompt engineering" is a grand name for learning to brief well. Anyone who's managed a person can do it.**

Anatomy of a prompt that works — six parts

1 · Role

Who the AI is being.

"You are a senior copywriter for a home services company."

2 · Context

The four kinds. Who you are, your reference material, the situation, examples of good.

Where most of the words in a great prompt go — and where most people put zero.

3 · Task

What you want, specifically. Not "write an email."

"Write a follow-up to a lead who got a quote four days ago and hasn't responded."

4 · Format

Length, structure, tone, what to include and leave out.

"Under 120 words, plain text, one clear next step."

5 · Constraints

The rules.

"No discount. No apology. Don't use the word 'just.'"

6 · Examples

Two or three of what good looks like.

The single highest-leverage part. One good example beats three paragraphs of description.

Before / after

THE PROMPT MOST PEOPLE WRITE	THE PROMPT THAT WORKS
"Write a follow-up email to a lead."	"You are a senior copywriter for [company]. We're a family-owned HVAC company — direct, warm, never pushy. [paste voice notes]. This lead got a quote for a full system replacement four days ago and hasn't responded. Write a follow-up under 120 words, plain text, one clear next step. Don't offer a discount, don't apologize. Here are two follow-ups we've sent that worked: [paste]."
Generic, could be from anyone	*Sounds like you, addresses the real situation, ready to send after a light edit*

Three techniques that separate good from great

1

Iterate, don't restart.

The first output is a draft. *"Shorter." "Less formal." "Three versions with different openings."* The second and third prompt is where 10x lives. Most people take output #1 or give up.

2

Make it show its work.

"Before you write it, list the three things this lead is probably worried about." Think first, then produce. Catches a bad assumption before it becomes a bad email.

3

Ask it to critique itself.

"Review that against the brief and tell me what's weak." Then fix what it found. The cheapest quality pass available. Almost nobody does it.

The mistakes that make AI look bad

- No context
- Vague task
- No examples
- Accepting the first draft
- One giant prompt instead of a conversation
- Asking for things it can't know — real-time facts, your unstated preferences

The prompt library — the compounding play

- **Every prompt that works gets saved** with a name and a note on what it's for. *Follow-up-after-quote. Monthly-report-summary. Review-response-negative.*
- **Standardize the context block.** Identity context — brand, voice, personas — is the same in every prompt. Write it once, well, paste it every time. This is the Brand Kit's job.
- **The library is what you hand the second operator.** It's how the skill transfers.
- **Review it quarterly.** Prompts drift as models change. Same cadence as the framework review.

LIVE EXERCISE · 3 MINUTES

Build one prompt

Pick one real task from this week that took longer than it should have.

1 · Role
"You are a..."

2 · Context
Who we are / what I'll paste in / the situation right now

3 · Task
Exactly what I want

4 · Format
Length, tone, structure

5 · Constraints
What it must NOT do

6 · Examples
What "good" looks like

Then: my first follow-up prompt. And the name I'll save it under.

LIVE DEMONSTRATION

Live build

One task from the room, six parts, run it, iterate once.

Brief it like a person. Give it what it needs. Don't take the first answer. Save what works.

Three things

1 **The concepts don't expire.**

Capabilities not tools. Context is the moat. The ladder. Everything new plugs into those three.

2 **Build the operator, not the agent.**

One person, one workflow, thirty days, protected time. The agents come later, out of what they learn.

3 **A prompt is a brief.**

Six parts. Context does the heavy lifting. Never take the first draft. Save what works.

Do these three things Monday

1 **Name your operator.**

One person. Curious, high standards, not necessarily technical. Tell them Monday.

2 **Pick their one workflow.**

The thing that eats the most hours. Measure it before you touch it.

3 **Write your context block.**

Brand, voice, who you serve — one page, written once, pasted into every prompt from now on. Tonight's worksheet is the start.

Questions

Thank you

Recording + slides + worksheet	Within 48 hours
Three of my own prompts	Included with the recording
Questions after tonight	care@wundertre.com

10x doesn't come from the tool. It comes from the operator.