



AI × CRE · SESSION 06

FRIDAY, AUGUST 21, 2026 · 10:00 AM ET · MICROSOFT TEAMS

From data dump to deal flow

Building a prospect list you'll actually work

Filter inside the platform you already pay for, let AI do the research, and score every row against what you actually want.

Ivy Marguerite

Senior Territory Manager, EyeQ Monitoring · Co-Lead, AI × CRE

We've been stacking this all year.

Every session has added one layer. This one puts them together on the task most of us do worst: finding who to talk to next.

S3

Frame · Focus · Finish

Michelle's prompting framework. Still the backbone of every prompt today.

S4

Automating the busywork

Kaitlyn on where the hours actually go, and how to get them back.

S5

Copilot in the flow of work

Ashley showed us AI inside the tools we already pay for.

S6

Prospecting, end to end

Today. Source to spreadsheet, with the compliance check baked in.

The data dump is not a prospect list.



1 You find a page with 200 companies on it.

Great. That's a page, not a pipeline.

2 You copy it into Excel.

Now it's a mess in Excel. Congratulations.

3 You mean to research each one.

You research four. On Thursday. Maybe.

4 You call whoever's at the top.

Which is alphabetical, not strategic.

The gap isn't finding data. It's everything between finding it and knowing who to call first.

What you leave with.



A 60-second permission check

Four questions you run on any page before you point AI at it. Written down, so you never have to guess.



A working four-prompt sequence

Vet, pull, enrich, qualify. Copy them, swap in your own filters, run them Monday.



A scored list, not a pile

One page in, a ranked spreadsheet out, with the rows that failed still visible so you know you checked.

You stay the expert. The AI does the typing.



PART ONE

You are already logged in.

Six minutes on what that actually permits. Then we build something.

Licensed to look. Not licensed to take.

Your subscription buys you access. It almost never buys you the right to reproduce the database behind it.

1

Reading is what you pay for

Pull up a record, read it, summarize it, draft from it. That is the product working as sold.

2

Harvesting is a different act

Systematically capturing many records into your own file is what these agreements are written to stop.

3

The line is volume and intent

One account you are actively working, in the ordinary course, is not the same as walking the whole result set.

The test that actually works: am I reading, or am I building a copy of their database?

Two checks, once per platform.

Do this once for each tool you use, write down the answer, and never think about it again.

ROBOTS.TXT

Add /robots.txt to any domain.

It is a public file saying which automated visitors are welcome where. Look for two things: Disallow lines covering the pages you use, and AI agents named directly. Some sites now list ClaudeBot and GPTBot by name. That tells you the posture without reading a word of legalese.

THE TERMS

Search the page for six words.

Robot. Scrape. Crawl. Automated. Data mining. Artificial intelligence. Ctrl-F, six searches, one minute. If none of them hit, you are probably fine. If one does, read that clause properly before you go further.

THE ONE THAT WILL CATCH SOMEBODY

At least one very large CRE platform revised its terms in July 2026 to bar sharing passcodes with “any third-party AI services, including, without limitation, artificial intelligence agents,” and separately bars exposing the product to outside AI tools. If that is your platform, none of what follows applies to it. Run the six-word search on your own subscription before you run any of this.

Signing in raises the stakes.

Same capability. Much bigger blast radius. Three things worth knowing before you turn it loose.

Everything on screen gets captured

Claude screenshots the tab to decide what to do next, and it cannot filter sensitive content out of what it sees. Sessions save to a history that syncs across your devices.

A hijacked agent inherits your access

Prompt injection is malicious instructions hidden in page content. Anthropic's own testing has driven it very low and they still say in bold that it is not zero.

Client data is on that screen too

If your platform shows you deals, contacts, or anything under NDA, that is now in the conversation. Use a separate browser profile.

“Use a separate browser profile without access to sensitive accounts.”

Anthropic's own recommendation. Two minutes to set up, and it is the single best thing on this slide.



PART TWO

Let the platform do the filtering

You already pay for a very good filter engine. Use it, then hand the result to AI.

Claude in Chrome, in one slide.

A side panel in your own browser that can see and act on the page you are looking at. Paid Claude plans. Chrome only.

Manually approve

Claude pauses and asks before every action.

Use this the first few times on any platform.

Automatically approve

Claude keeps working and checks its own actions for safety.

The sensible default once you trust the workflow.

Skip all approvals

No pauses. Nothing checks its actions.

Never on a site you are logged into. Ever.

It will also ask before entering anything sensitive, even on a site you granted standing permission to. Read that prompt when it appears.

Filter in the platform. Not in the prompt.

This is the whole trick, and almost nobody does it.

WHAT PEOPLE TRY

“Find me every industrial owner with 50-plus properties.”

Now the model is guessing at criteria your platform already knows precisely. You get plausible names, some of them wrong, and no way to tell which.

WHAT ACTUALLY WORKS

Build the search in the platform. Then say “read this.”

The platform's filters are exact, indexed, and something you already pay for. Let it do the part it is good at. AI does the part it is good at, which is everything after the result set.

Set your filters in their interface: asset type, market, size, ownership. Get the result set on screen. Only then open the side panel.

Also the safer path: you are reading a screen you were entitled to see, not asking a model to reconstruct a database.

Demo one: read the result set.

Filters already set in the platform. Result list on screen. Now hand it over.

PROMPT

Read the results list on the page I have open and give me a table, one row per company. Use the columns this page actually displays. Only include what is visible on this screen. Do not open individual records, do not navigate anywhere, and do not fill in anything the page does not show. If a field is blank here, leave it blank. Tell me how many rows you found and whether that matches the result count shown on the page.

“Only what is visible”

Stops it inventing a column your platform did not give you.

“Do not open individual records”

Keeps you on one screen you were already looking at.

“Tell me how many rows”

Makes it check its own work before you have to.

Check it before you fall in love with it.



Row count matches the page

If the platform said 34 results and you got 27, ask which seven and why.



Blanks are actually blank

A field the platform never showed should be empty, not plausibly filled in.



Spot-check three at random

Pick three rows, find them on screen, confirm. Thirty seconds, every time.



Columns are the ones you asked for

A helpful extra column is still a column you now have to verify.

Thirty seconds of checking here saves an hour of calling the wrong people later.



PART THREE

Company research

The part that used to eat your evenings. This is where the hour actually gets paid back.

Your platform knows who exists. It does not know who to call.

WHAT THE PLATFORM GAVE YOU

- Company or entity name
- Address and market
- Square footage on record
- Owner of record, maybe a contact
- Sale or lease history

All true. None of it tells you whether to pick up the phone.

WHAT ACTUALLY DECIDES IT

- How big are they really, across the whole portfolio
- What changed in the last ninety days
- Whether they have a reason to care right now
- Who owns the problem you solve
- What they already have in place

None of that is in your subscription. All of it is public, and it used to take you twenty minutes a company.

Four layers, in this order.

Each one answers a different question, and each has a source you can point at afterwards.

1 The filings

For any public company, the 10-K. Property counts, total square footage, markets, and a risk-factors section that often names your problem in their own words. The SEC requires it, so it is accurate and it is dated.

2 Their own words

Investor relations decks, the portfolio page, the newsroom. This is where strategy lives: what they are buying, where they are expanding, what they are calling a priority this year.

3 What changed

News, permits, acquisitions, incidents, filings. This is the trigger, and it is the difference between an opener that lands and one that gets deleted.

4 Who

Leadership pages and announcements. You are looking for the person who owns the function you sell into, not the person with the biggest title.

Run them as separate passes, not one giant question. You get better answers and you can see which layer failed.

Demo two: the research pass.

Same conversation, table already in it. Layers one and two, in one prompt.

PROMPT · FIRMOGRAPHICS

For each company in this table, research and add: approximate total properties owned, approximate total square footage, primary asset types, and the markets they operate in. Check the 10-K first for any public company, then their investor relations and portfolio pages. For every figure give me the number, the source, and how confident you are. If you cannot find a reliable figure, write “not found” rather than estimating. I would rather have gaps than numbers I have to go back and check.

WATCH FOR

Round numbers with no source. “Approximately 200 properties” from nowhere is a guess wearing a suit.

CELEBRATE THIS

A “not found.” It refused to invent a number and told you so. That is the behaviour you want.

“Why now” is the whole opener.

Layer three, on its own pass. No trigger is a complete answer: a row without one is not ready, and manufacturing one is how you sound like everybody else.

PROMPT · TRIGGERS

For each company, find anything that changed in the last ninety days: acquisitions, new development or construction starts, expansion into a new market, leadership hires in operations or security, incidents in the news, or anything in a recent filing that touches loss prevention. Give me the date and the source for each. If there is nothing real, write “no trigger” rather than reaching.

Acquisition or portfolio growth	They just took on assets they have not standardized yet.
New development or construction	There is a site with no coverage on it right now.
Expansion into your market	They need local answers and they do not have a bench yet.
A hire in your function	Somebody new owns the problem and is looking to make a mark.
An incident in the news	Handle with care, and only if it is genuinely public.

Make the research cite itself.

The research is only worth what you can defend when somebody asks where the number came from.



Source on every figure

A link or a document name. If it cannot produce one, treat the number as unverified and mark it that way.



Confidence, stated plainly

High, medium, low. Low-confidence numbers are still useful, they just do not go in an email.



“Not found” beats a guess

A visible gap costs you nothing. A confident wrong number costs you the meeting.



Spot-check three at random

Every single run. Not the ones that look wrong, three at random. That is how you catch the pattern.



Date everything

A trigger from fourteen months ago is not a trigger. Make it tell you when.

You are still the expert. The AI just did two days of reading in four minutes.



PART FOUR

Make it actionable

A list you don't work is a hobby. Here's how it turns into Monday.

The platform filtered. Now you score.

Filters got you a result set. Scoring is what turns it into an order to work it in. Write these down before you look at a single name.

Size floor

Below this, not worth your time.

e.g. 50+ properties

Markets

Where you can actually service it.

e.g. national

Asset types

Where your offer genuinely fits.

e.g. industrial, retail

Hard disqualifiers

Kills a row on sight.

e.g. existing customer

What makes a 5

Describe your best-ever account.

the one you'd clone

Reachable

What a warm path actually looks like.

e.g. mutual contact

Three minutes. Right now, on paper. The two you have not written down before are the last two.

Most rows should die here.

A shorter list is the point, not a failure. Thirty names you'll never call is worse than six you will.

PROMPT

Apply these filters. Size floor: [YOURS]. Geography: [YOURS]. Asset types: [YOURS]. Hard disqualifiers: [YOURS]. For each company, mark whether it passes each filter, then score Fit from 1 to 5 and explain the score in one sentence. Put the companies that fail a hard filter at the bottom under a heading called Cut, with the reason. Do not quietly drop anything.

“Do not quietly drop anything” is the most important six words on this slide.

A model that silently discards eleven rows hands you a clean-looking list and no idea what it threw away. Making it show its cuts, with reasons, is how you catch a filter you set wrong. It is also how you find the row it was wrong about.

Two numbers, one order to work them in.



FIT

1 to 5

How well they match the filters you just wrote down.

×

REACH

1 to 5

How realistically you can get in front of the right person.

=

SCORE

/ 25

Sort descending. Start at the top on Monday.

Why two numbers and not one: a perfect-fit account you have no way into is not your best Monday. A decent-fit account where you already know somebody is. Multiplying them stops you from confusing the two.

The worksheet does this math and sorts for you.

What actually happens next.

1 Sort by score.

The order is now a decision, not the alphabet.

2 Take the top ten.

Not the top forty. Ten is a week you'll finish.

3 Write the opener from the row.

Use the trigger. If the row has no real trigger, it isn't ready.

4 Keep the Cut list.

In ninety days a cut row becomes a live one. You'll already know why you passed.

If you do nothing else from this hour: sort your existing list and call the top of it instead of the top of the alphabet.

Five things to carry out of this hour.

1 Check the platform once.

Sixty seconds: robots.txt and a six-word search of the terms. Write down the answer, move on.

2 Licensed to look, not to take.

Your subscription buys access. It does not buy the right to reproduce the database behind it.

3 Gaps beat guesses.

Make it say “not found.” A visible hole is safer than a confident wrong number.

4 Never let it drop rows quietly.

Make it show its cuts and the reason for each one.

5 You are responsible.

For everything the tool does on your behalf. That is in writing, in their terms and yours.

Not legal advice. I'm a territory manager, not your attorney. If your firm has counsel, this is a fifteen-minute conversation worth having with them.

Take it with you.

Everything from this hour goes up on the site straight after the session.

INTERACTIVE WORKSHEET

Prospect Build Sheet

Source-vetting log, your six filters, and a scoring table that sorts itself. CSV export.

FIVE PROMPTS

The prompt pack

Vet, pull, enrich, qualify, and the opener. Copy buttons, added to the site prompt library.

VETTED STARTING POINTS

Source bank

Public sources by category, flagged green, check-it, or don't. With the reason for each.

Three asks, same as every month: send us your prompts, tell us what you want covered next, and tell us how you're actually using this.



Questions.

And the one I'll get anyway: yes, I know you want to do this on LinkedIn. Their user agreement names browser extensions specifically. Read one profile you were already going to read and let AI help you write to that person. Don't build a list.

NEXT SESSION

Friday, September 18, 2026 · 10:00 AM ET · Microsoft Teams

Michelle Hamilton · Build Your Own AI Chief of Staff

Ivy Marguerite · imarguerite@eyeqmonitoring.com · aixcre site: sessions, prompt library, worksheets