

AVS 2026 · CLEVELAND

The Field Guide

AI in Asset Management & Valuation

Reference material for after the room empties: what AI does beyond chat, the guardrails, the economics of the buildout, residual values, a thirty-day plan, and every resource worth following.

NICK COSCIA · CAPTERIS

LOGAN CAMPBELL · IRONTRAX

What's inside

The session runs an hour. This runs as long as you want it to. Everything the hour could not hold lives here — plus the parts you will want again next week.

Beyond the chatbot

3–4

Skills, projects, agents and artifacts — what AI looks like once it stops being a chat window.

Standards & obligations

5–7

The AO-41 four-step framework, what never goes in a public tool, and the regulatory landscape.

The economics of AI

8–11

Inverted margins, GPU rents, token spend, credit spreads, power and water.

Residual values

12–13

Nvidia's residual guarantee, and why the real risk is overbuild rather than obsolescence.

Getting started

14–15

From rough idea to implementation-ready, and a thirty-day plan.

Keep going

16–18

Podcasts, newsletters, books, and who to follow in AI.

None of this is a compliance opinion, and none of it replaces your own judgment or your firm's policy. It was written by two people who work in this industry.

Skills and projects

Chat is level one. These are the four shapes AI takes once it stops being a text box.

Skill The playbook you hand a new hire.

Write your procedure down once — how you decode a serial number, how you evaluate a spec sheet — and the AI applies it consistently every time. No retraining, no re-explaining. You still check the output.

HOW YOU WRITE ONE

“Use this skill whenever the user describes equipment, an asset, a deal, or a transaction for collateral evaluation, valuation, credit memo preparation, OLV IP / FMV IP analysis, or secondary market assessment... Include everything relevant: borrower legal name, invoice detail, serial numbers, condition notes.”

Simply put: like a recipe card — anyone who follows it gets a consistent result. A ChatGPT “custom GPT” is the same idea packaged as its own walled-off assistant.

Project A war room for one topic.

A workspace that keeps every document, conversation and instruction about one subject together — so the AI already knows the background every time you walk in. No re-explaining, no starting from scratch.

BEFORE

Underwriting more exposure to an existing obligor meant hand-rebuilding their in-force book — hunting through the monthly portfolio report and keying every schedule into the eval. Tedious and error-prone at thousands of assets a year.

AFTER

The project holds the portfolio report, the eval template and the procedure. The analyst types “build the existing exposure for [customer]” and reviews the output instead of assembling it.

Agents and artifacts

Agent An analyst you delegate to.

Instead of answering one question, an agent takes a whole assignment — plans the steps, does the work, checks itself, and comes back with the finished result. You review the output, not the process.

A third-party appraisal PDF lands in your inbox. You drop it in and say “file this.”

STEP ONE

Your extraction skill invokes itself — values, comps and serials pulled into your master file.

STEP TWO

The PDF is renamed to your convention and moved to the appraisal folder on Box or SharePoint.

STEP THREE

You get the finished report back, not a checklist of steps to work through yourself.

Artifact The deliverable, not just the answer.

Instead of a paragraph describing the work, the AI builds the actual thing beside the chat — a document, spreadsheet, chart or working tool. You edit it, iterate in plain English, and reuse it next quarter with fresh data.

IN PRACTICE

Ask for OLV curves across your crane portfolio and get a live, interactive chart — not a description of one. Drop in a new appraisal, ask for a credit-committee one-pager, and hand off the finished document.

Chat is level one. Skills, projects, agents and artifacts are how AI goes from answering questions to doing the work.

AO-41: the four-step framework

The Appraisal Standards Board adopted Advisory Opinion 41 on April 23, 2026 — the first USPAP guidance to explicitly name generative AI, alongside machine learning and automated valuation models. It creates no new standards. It clarifies how the Ethics, Competency, Scope of Work and Record Keeping rules apply once a technology tool enters an assignment.

01 · EVALUATE THE TOOL

Understand what it does and where its limits are — enough to judge relevance, not enough to audit the algorithm.

02 · EVALUATE THE DATA

Know what the tool relied on, and whether it is current, reliable and appropriate for this assignment.

03 · PERFORM INDEPENDENT ANALYSIS

Tool output is information, not a conclusion. Your judgment still has to be layered on top.

04 · ASSESS CREDIBILITY OF OUTPUT

Decide it is accurate, relevant and appropriate before it becomes part of your assignment results.

Every use case in this guide fits inside this framework — as long as a human owns step four.

A tool cannot comply with USPAP. Only the appraiser can.

Source: The Appraisal Foundation, Advisory Opinion 41 (adopted Apr. 23, 2026); ISA; Dennis Badger & Associates AO-41 summary.

What never goes in a public tool

This is minimization, not a legal opinion. Some of these fields are necessary inside an authorized environment — the test is whether the task actually requires it.

DO NOT PASTE

Borrower financial statements · client names tied to assignment results · unreleased appraisal conclusions · loan documents · anything under an NDA · personal information · anything you would not email to a competitor.

GENERALLY FINE

Public specifications · manufacturer literature · published auction results · your own procedures and templates · anonymized or synthetic examples · general market questions with no client identifiers.

Four controls that actually work

USE ENTERPRISE TIERS

Not consumer chatbots. The interface looks the same; the contract does not.

GET TERMS IN WRITING

Confirm retention and training terms rather than assuming them.

WRITE ONE PAGE

A one-page AI use policy beats silence — and beats a fifty-page policy nobody reads.

CONSIDER OPEN WEIGHTS

Models you can run behind your own firewall, when the data cannot leave.

The question is never “is AI allowed?” It is “which environment, which data, and under whose agreement?”

The regulatory landscape

SR 11-7

Remains the anchor model-risk-management framework for banks and bank-affiliated lessors deploying AI models.

REGULATION B (ECOA)

The CFPB's final rule, effective July 21, 2026, eliminates disparate-impact liability under ECOA. ECOA covers business credit too, so the change reaches commercial equipment finance.

STATE DISCLOSURE LAWS

Ten states now require commercial financing disclosures: CA, CT, FL, GA, KS, MO, NY, TX, UT, VA. Texas MCA providers and brokers must also register with the OCCC by Dec. 31, 2026.

THE RULE OF THUMB

Disparate-impact risk has not gone away for commercial lenders. State fair-lending laws (e.g., NY, IL) and private ECOA suits keep the theory alive.

This one genuinely needs your counsel. Nothing here is a legal opinion.

Sources: Federal Register; CFPB; state commercial financing disclosure statutes.

Questions worth asking your own shop

- Do we have a written AI use policy, or do we have silence?
- Which environment is approved — and does everyone know which one they are in?
- Who owns this internally? Is there a name, or is it nobody's job?
- What do our client agreements and NDAs already say about third-party processing?
- If a file went to litigation tomorrow, would the workfile show how the tool was used?

Margins are inverted

In most businesses, profit margins are highest closest to the end customer. In AI it is inverted — margins rise the further you get from the end user, and fall as you approach the layer actually selling to customers.

Silicon & equipment

Nvidia, AMD, TSMC, Broadcom, Micron

+41%

FAR FROM THE CUSTOMER

Energy & grid

Constellation, Vistra, NextEra, Eaton, Vertiv

Compute & cloud

AWS, Azure, Google Cloud, Dell, CoreWeave

Models & apps

OpenAI, Anthropic — the layer you actually buy from

−59%

CLOSEST TO THE CUSTOMER

Source: Torsten Sløk, Apollo Chief Economist, Daily Spark.

**The layer you buy from is the least profitable layer in the stack.
That is unusual, and it shapes who can afford to keep building.**

For a room that finances equipment: the money in this buildout sits in silicon, power and physical infrastructure — all of it assets somebody has to value, finance and eventually remarket.

On the tape: compute and tokens

COMPUTE · APR 2026

SemiAnalysis: on-demand GPU rental capacity sold out across all GPU types, with new Blackwell clusters booked into June–July.

COMPUTE · AUG 2026

A100, H100 and B200 rents all rising at once — older chips holding value instead of depreciating.

Hardware rents are rising while the price per token keeps falling — the clearest sign the demand is real.

This is our territory. Rental rates for financed equipment — cranes, trucks, construction machinery — have always been tied to the health of the industries renting them, not set in a vacuum. GPU rents behaving the same way is a sign this asset class is starting to trade like any other piece of financed iron.

Token spend

\$2.07 → \$0.97 average cost per million tokens, late-May peak to Aug. 31, 2026

THE SIGNAL

A usage-weighted benchmark of what the active LLM market is effectively paying per million tokens.

WHAT IT MEANS

Provider price cuts, plus cheaper open-weight models pulling usage toward lower-cost options — often several times cheaper than closed models.

Token costs have roughly halved even as GPU rents rise — buyers are scaling usage while optimizing model mix and price per task.

Sources: SemiAnalysis (Apr 2026); Silicon Data LLM Token Expenditure Index; Bloomberg (Aug 2026).

On the tape: credit

S&P GLOBAL · MIDYEAR 2026

Hyperscalers and related issuers sold \$225B in bonds in the first half, and spreads on their debt have widened as supply keeps coming.

MORGAN STANLEY RESEARCH

Private credit projected to supply about \$800B of data-center financing through 2028 — the largest single source.

“SPREADS WIDENED”

Lenders want more yield over Treasuries to hold the risk. Borrowing just got more expensive.

“REAL YIELDS ARE UP”

Interest rates after inflation. When they rise, every project's hurdle rate gets tougher to clear.

Apollo: “the margin for error is narrowing.” If debt gets pricier, the weakest projects are the ones that stop.

“I don’t think anyone in ’24 or ’25 thought that the prices of old GPUs would still be going vertical in 2026.”

GAVIN BAKER · ATREIDES MANAGEMENT · AUGUST 2026

A room that underwrites residual value should sit with that for a second: two-generation-old chips are appreciating, not depreciating. The whole depreciation debate assumed old GPUs would be near-worthless in three years.

Sources: S&P Global via Fortune (Jul 2026); Reuters; Apollo (Aug 2026); Morgan Stanley (Jul 2025); Dallas Fed (Feb 2026); Invest Like the Best.

Power is the new bottleneck

THE DRAW

A single frontier AI campus can pull hundreds of MW; the largest top 1 GW, enough to power a city.

THE RESPONSE

Nuclear restarts, new gas turbines and utility rate cases are now tied directly to AI demand.

THE CONSEQUENCE

If power becomes the bottleneck instead of chips, the buildout timeline — and the equipment coming to market — shifts with it.

The water behind the power

1–5M gal

Per day a large hyperscale facility can use for cooling under peak conditions.

~210B gal

Indirect water used at power plants to generate that electricity in 2023 — about 12× the water used directly on site.

176 TWh

U.S. data-center electricity use in 2023, 4.4% of all U.S. power. LBNL projects 6.7–12% by 2028.

The figures are debated — estimates vary with how power-plant water is counted — but every methodology agrees on direction. Water, power, and increasingly noise from on-site generation and cooling equipment are all now gating factors on siting.

Texas is the outlier: roughly 7 GW operating today and nearly 100 GW planned — more than the next two states combined.

Virginia still leads on capacity already built, at about 17 GW. Virginia and Utah follow Texas on planned capacity, and even together they fall short of it.

TO TRACK IT YOURSELF

EIA Electricity Data Browser · IEA “Electricity” and data-center reports · Grid Strategies transmission reports · NextEra, Constellation and Vistra earnings calls · Torsten Sløk’s Daily Spark.

Sources: Apollo Daily Spark (Aug 2026); Cleanview (Sept 2026); LBNL (2024); EESI; EIA; IEA.

Nvidia backstops its own chips

\$500B+

On Aug. 10, 2026, Nvidia signed MOUs with Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs and KKR to mobilize third-party capital for AI infrastructure.

UP TO 25%

Nvidia may provide residual-value support for up to this share of an individual project, assessed deal by deal. Detailed terms are not yet public.

THE MECHANISM

The capital providers underwrite the customer, demand and cash flow. Nvidia's capped support sits on top, so its exposure is to what the hardware is worth second-hand.

Gavin Baker: Nvidia is becoming “the central bank of AI” — matching compute buyers with the lenders who fund them.

This should be familiar. A manufacturer standing behind residual value to make lender economics work is exactly what happens in equipment finance every day — it is just happening at a scale, and on an asset class, that did not exist five years ago.

WHY IT MATTERS TO THIS ROOM

If manufacturer support is what makes the paper financeable, then the credibility of that support becomes a valuation question. Someone has to form a view on whether the residual floor holds — and that is the same work you already do on iron.

Sources: Nvidia (Aug. 10, 2026); Jensen Huang on X (Aug 2026); Gavin Baker on All-In (Aug 2026).

Old iron, and the real risk

THE EQUIPMENT-FINANCE PARALLEL

Decades-old American trucks and machinery, fully depreciated at home, still find a second life exported overseas. Residual value and useful life are not the same question — secondary markets extend both.

A 2020 CHIP, RENTED THROUGH 2029

CoreWeave signed an A100 contract running to 2029 — nine years on a 2020-vintage chip. Prior-generation pricing is at or above historical levels, and expiring H100 capacity was rebooked at 95% of the original rate.

THE CAVEAT

One contract is not a fleet — and CoreWeave’s own 10-K still depreciates GPUs over six years. Treat the rebooking data as a signal, not a curve.

The bigger risk isn’t obsolescence — it’s overbuild.

If capacity gets built ahead of demand, the constraint holding rents up disappears and prior-generation pricing corrects fast. David Sacks calls that “dark GPUs” — in his words, a disaster for everyone. The chips do not have to stop working for the residual assumption to break — there just has to be too much of everything.

What to watch

UTILIZATION

Are the new facilities filling as fast as they are opening?

REBOOKING RATES

What share of the original rate does expiring capacity command?

ANNOUNCEMENT SLIPPAGE

Planned capacity that quietly moves right is the earliest signal there is.

Sources: CoreWeave Q2 2026 earnings call; David Sacks on All-In (Aug 2026).

From rough idea to implementation-ready

One of the least discussed uses of AI is also one of the most valuable: helping the expert arrive with the idea already half, three-quarters, or fully baked — which makes buy-in easier and implementation faster.

0%

Rough idea

Start with the asset manager's judgment: the problem, the opportunity, or the workflow worth improving. This part is still yours.

50%

Build the case

AI researches the context, pressure-tests the assumptions and organizes the business case.

75%

Make it tangible

Turn the idea into a draft process, a mockup, requirements and clear success criteria.

100%

Hand off to build

Give developers and software teams something concrete to react to — not a pie-in-the-sky request.

**The judgment at the start is the part that cannot be delegated.
Everything after it can be accelerated.**

Your first thirty days

WEEK 1 · GET ACCESS

Ask IT and compliance what is already allowed. Check whether your association offers a member assistant — ELFA members should ask about ELFie. Read AO-41 if you are credentialed. Subscribe to one daily email.

WEEK 2 · ONE TASK

Pick a repetitive document task. Do it the old way and the new way on the same input, then compare honestly. Take a free hour-long AI fluency or foundations course.

WEEKS 3-4 · WRITE IT

Take your most-repeated judgment call and write it out in plain English. Give it to the AI as a skill. Run it on five cases where you already know the answer. Show someone else in your shop.

Most people quit in week two, when it gets something confidently wrong. That is the correct reaction to the output and the wrong conclusion about the tool.

Where the industry research actually lives

ELFA STRATEGIC RESEARCH AGENDA

Pillar One is “Industry Modernization & AI,” published June 24, 2026 — evidence the trade body has institutionalized this rather than treating it as a fad.

ELFF HORIZON REPORT

Source of the widely quoted adoption figures: the 2024 edition found 42% of end-users use generative AI and 42% more plan to. Likely understated today.

ELFIE

ELFA’s member-exclusive AI assistant, launched Feb 2025. Ask your association contact.

MARKET DATA

Rouse Services · EquipmentWatch · Sandhills / Ritchie Bros. / IronPlanet · Iron Solutions.

A CAUTION ON VENDOR CLAIMS

Check before you repeat. If a vendor calls a data provider “AI-powered,” ask where the provider itself says so — and look at the provider’s own site before it goes in a memo.

What to listen to

ODD LOTS

Bloomberg. “What It Takes to Get a Data Center Financed” is the single best forty-five minutes for this room.

AI DAILY BRIEF

Nathaniel Whittemore. Daily, framed around business adoption rather than research.

HARD FORK

Roose & Newton, NYT. The lowest barrier to entry — start here if you are starting anywhere.

INVEST LIKE THE BEST

O'Shaughnessy. Source of the “going vertical” Gavin Baker quote.

DWARKESH PODCAST

One episode on compute prices bears directly on residual assumptions.

LEX FRIDMAN

Long-form. Episode #494 with Jensen Huang is the relevant one.

ALSO WORTH A LISTEN

All-In · BG2Pod · No Priors · The a16z Show (renamed from “a16z Podcast”). All confirmed actively publishing as of August 2026.

If you only do one thing: subscribe to one free daily email and listen to the Odd Lots data-center financing episode. Ninety minutes total, and it covers most of what this room needs.

What to read, where to learn

THE DAILY SPARK · FREE

Torsten Sløk, Apollo. The highest-relevance subscription on this list for this audience.

ONE USEFUL THING · FREE

Ethan Mollick, Wharton. If you print one URL from this guide, print this one.

THE NEURON · FREE

Five-minute daily email written for professionals, not builders.

POWER 2026 · FREE

Neel Somani. The best free resource here for valuing energy infrastructure.

ARTIFICIAL ANALYSIS · FREE

The vendor-neutral model and provider scoreboard.

STRATECHERY · PAID

Ben Thompson. The right register for modeling infrastructure depreciation.

ALSO

Equipment Finance News (free, industry-specific) · SemiAnalysis (paid research).

FREE COURSES — PICK ONE

Anthropic AI Fluency · OpenAI Academy · Google AI Essentials. An hour each, and enough to stop guessing.

What to read on the plane

CO-INTELLIGENCE

Ethan Mollick. Treat AI as a co-worker with a jagged skill profile — the right mental model for this room. Read this one first.

HUMAN + MACHINE

Daugherty & Wilson. The most operational book on the list; six hybrid role types across 1,500 organizations.

SUPERAGENCY

Hoffman & Beato. The optimistic case — read alongside the next one, not alone.

EMPIRE OF AI

Karen Hao. The counterweight: investigative reporting on compute, energy and labor inputs.

Who to follow in AI

A short list rather than a long one. Each of these people publishes reasoning you can check, not just conclusions — which is the only useful filter in a field this noisy.

ETHAN MOLLICK · @EMOLLICK

Wharton. The most useful account here for non-technical professionals — practical, tested, and honest about what does not work. Writes One Useful Thing.

GAVIN BAKER · @GAVINSBAKER

Atreides Management. Source of the “going vertical” quote. The clearest voice on compute economics and what they mean for residual value.

DYLAN PATEL · @DYLAN522P

SemiAnalysis. The deepest reporting anywhere on chips, capacity and data-center supply. Paid research, free feed.

BRAD GERSTNER · @ALTCAP

Altimeter. Note the handle is the firm’s, not his name. Also co-hosts BG2Pod.

NEEL SOMANI · @NEELSOMANI

Power markets and energy infrastructure — the constraint most people in finance are underweighting. Writes Power 2026.

TORSTEN SLØK

Apollo Chief Economist. No personal X account — the Daily Spark email is where his work actually lands.

ALSO WORTH FOLLOWING

Nathaniel Whittemore for daily adoption news · Ben Thompson for the strategy and depreciation lens · Dwarkesh Patel for long-form technical interviews · Karen Hao for the reporting that runs against the consensus.

TWO CAUTIONS

@AnthropicAI is Anthropic’s official account. Impersonation and lookalike handles are common in this space, so check the profile against the person’s own site or newsletter before you follow.

Follow the people who show their work. In a field moving this fast, sourcing is the only filter that holds up.

Trusted data in. Assisted reasoning. Human judgment out.

Start with one real task, in an approved environment, with data you trust.
The tool assists the analysis. The professional owns the conclusion.

STAY IN TOUCH

Nick Coscia

Capteris · ncoscia@capteriscapital.com

Logan Campbell

Irontrax · lcampbell@irontrax.com

DISCLOSURES

This guide is for educational and discussion purposes only. It is not a recommendation for or against financing, investing in, acquiring, disposing of, or otherwise transacting in AI infrastructure, GPUs, data centers, power assets, companies, securities or related technologies. Nothing presented constitutes investment, credit, legal, tax, accounting, valuation, appraisal or underwriting advice. Any transaction or policy decision should be based on independent diligence, current information, applicable standards, and advice from qualified professionals. Third-party information may include estimates, opinions or data that change over time. Views expressed are our own and do not represent the official positions of Capteris Capital or Irontrax.