

AI VOLS

Meeting #6

Building Agents

How agents actually work, and how to use them — no coding required.

LIVE DEMO

Stock Ticker Walkthrough

Spring 2026 · AI Club

College of Emerging & Collaborative Studies · University of Tennessee

April 28, 2026 · 6:30 PM

Agenda

Last Week: AI in Business · Two ways into AI businesses · the build-team collapse

- 01 | AI News Update** *What happened this week*
- 02 | Stock of the Week** *AI stocks worth watching*
- 03 | Model Standings** *LM Arena*
- 04 | Building Agents** *What an agent is — and how Claude Code works*
- 05 | Live Demo** *Stock Ticker Weekly Report*
- 06 | AI in Your Field** *What AI looks like in your field*
- 07 | Closing** *Next week, GroupMe, Website*

What's new in AI

GPT-5.5: Incremental, Not Insane

- Better at coding + reasoning, but not a huge leap
- Feels like refinement, not a breakthrough
- Models are stabilizing, not exploding

GPT Image Gen 2: Actually Useful

- Much stronger realism + consistency
- Handles text in images way better
- Real use: ads, slides, UI mockups

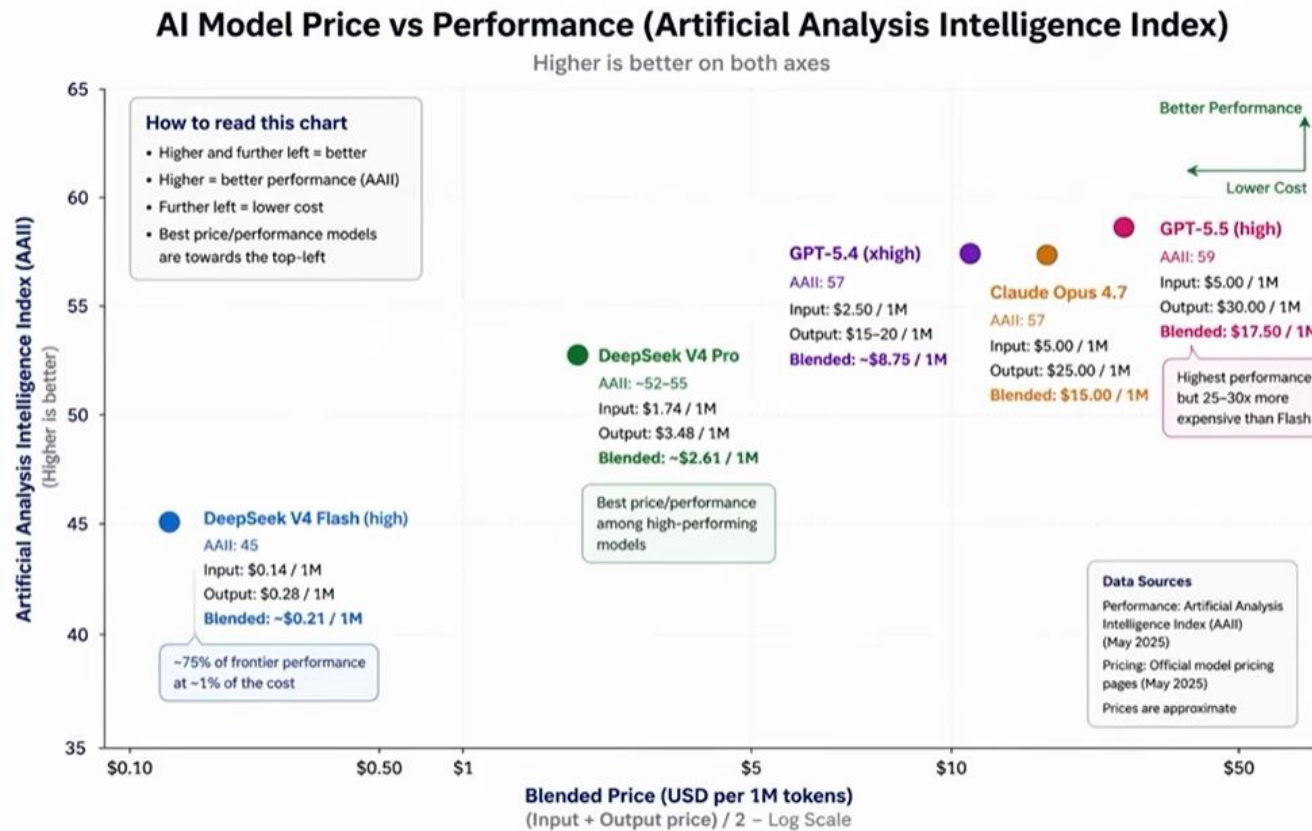
DeepSeek V4: Price Disruption

- Near top-tier performance
- ~1/6 the cost of leading models
- Forces price competition across AI

Musk vs Altman: OpenAI Lawsuit

- Musk claims OpenAI abandoned nonprofit mission
- Wants structural changes + massive damages
- OpenAI says it's competitive/personal
- Trial started — could impact AI company models

AI in Global Economy



AAII is a composite index of 10 challenging evaluations across reasoning, coding, mathematics, and real-world problem solving.
 Blended price = (Input price + Output price) / 2 per 1M tokens.

AI Global Power Shift

- chip bans → China forced to build its own
- DeepSeek: cheaper ≈ competitive → price war
- U.S. = compute | China = efficiency
- Distillation/scraping to catch up
- Risk: cheaper AI → global adoption → censorship/influence



The quiet engine powering the AI build-out.
Massive upside. Structural moat. Software floor.

STRONG BUY



Market Cap
\$2T+



AI Chip CAGR
~64%
(Analyst proj.)



AI Revenue Target
\$100B
by end of FY2027



Revenue Growth
+29% YoY
(Recent FY)

1 THE CONTRARIAN ANGLE



- Everyone talks about Nvidia as the AI chip king — and they're right.
- Broadcom is playing a different, arguably smarter game.
- Instead of competing for the GPU market, Broadcom **co-designs custom AI accelerator chips (XPU)s** directly with hyperscalers — Google, Meta, Apple.
- Once a hyperscaler builds their entire inference infrastructure around Broadcom's silicon, switching costs are enormous. That's a moat that compounds.

2 THE NUMBERS ARE INSANE

Custom AI chip revenue target

\$100B

Annual revenue by the end of FY2027

Revenue Growth (Analyst Projections)

FY2026 **64%**

FY2027 **50%**

Nvidia (FY2024–FY2028 CAGR)

~26%



Broadcom is projected to nearly double Nvidia's growth rate through 2028. That's elite territory.

Source: Wall Street Analyst Estimates, The Motley Fool (Apr 2025)

3 THE SOFTWARE FLOOR

vmware

High-margin recurring software business

- The VMware acquisition gives Broadcom a massive recurring enterprise software business.
- Generates billions in revenue every quarter.
- This isn't a pure-play chip bet that craters if AI capex slows — the software business provides real earnings ballast.
- Higher margins, lower volatility, durable cash flow.

4 WHY NOW?



"Anything-but-AI" sentiment led to a selloff in numerous AI-related stocks during Q1 2026.



Better entry point created after gains in 2025.



Broadcom recently joined the \$2 trillion market cap club, validating the thesis.



EARLY INNINGS

The AI chip revenue ramp is just getting started.

5 THE BULL CASE

- ✓ Hyperscaler lock-in creates durable, sticky revenue
- ✓ \$100B AI revenue target achievable by end of 2027
- ✓ VMware = high-margin software floor (~\$4B/quarter)
- ✓ Fastest projected CAGR among major chip stocks
- ✓ Strong balance sheet and cash generation
- ✓ Increasing AI capex visibility from hyperscalers

6 RISKS TO WATCH

- ⚠ Concentrated customer base — Google, Meta are huge %
- ⚠ Hyperscalers may internalize chip design over time
- ⚠ High valuation — premium priced for perfection
- ⚠ Macro / export restriction headwinds
- ⚠ Execution risk in continued AI chip ramp
- ⚠ Technological disruption or competition



THE TAKEAWAY

Broadcom is the quiet engine powering the AI build-out. Custom silicon lock-in with hyperscalers creates an enormous moat, while the software business provides real earnings ballast. With a \$100B AI revenue target in sight and growth that outpaces even Nvidia, this is a rare combination of explosive upside and structural defensibility.

LM Arena Leaderboard

Rank	Rank Spread	Model	Score	Votes	Price \$/M	Context
1	1 - 7	claude-opus-4-7-thinking Anthropic · Proprietary	1503 ±8	5,996	\$5 / \$25	1M
2	1 - 6	claude-opus-4-6-thinking Anthropic · Proprietary	1501 ±5	20,917	\$5 / \$25	1M
3	1 - 7	claude-opus-4-6 Anthropic · Proprietary	1496 ±5	22,200	\$5 / \$25	1M
4	1 - 8	claude-opus-4-7 Anthropic · Proprietary	1493 ±7	6,746	\$5 / \$25	1M
5	1 - 8	gemini-3.1-pro-preview Google · Proprietary	1493 ±5	26,125	\$2 / \$12	1M
6	2 - 12	muse-spark Meta · Proprietary	1489 ±7 Preliminary	7,931	N/A	N/A
7	1 - 14	gpt-5.5-high OpenAI · Proprietary	1488 ±10	3,587	\$5 / \$30	1.1M
8	4 - 12	gemini-3-pro Google · Proprietary	1486 ±4	41,400	\$2 / \$12	1M
9	6 - 17	grok-4.20-beta1 xAI · Proprietary	1481 ±5	15,438	N/A	N/A
10	6 - 17	gpt-5.4-high OpenAI · Proprietary	1479 ±6	14,328	\$2.50 / \$15	1.1M

[View all](#)

Chat → Agent

Chat

You ask. It answers.

- Lives in a chat box
- Can't touch your files
- Can't run code
- You copy answers out manually
- Limited flexibility

Agent


You set the goal. It does the work.

- Reads + writes files on your computer
- Runs commands, fetches data, browses the web
- Plans → does → checks → fixes
- Remembers your project across sessions
- Kicks off other agents to work in parallel

CONTEXT

Understanding the Context Window

Context is the model's working memory. Managing it is the single most important skill in Claude Code.



CONTEXT ROT

The more you load in, the more the model forgets.

By ~10K tokens, most models lose ~50% of recall. A full window isn't a smarter agent — it's a fuzzier one.

/clear

Wipe the slate

Drop the entire context. Use it between unrelated tasks so old chatter doesn't poison the new one.

/compact

Summarize, don't dump

Summarize the long session — keep the facts, lose the noise. Use it when one task drags on.

Mental model: the agent is choosing what to load. Your job is to help it choose well.

5 Things You'll See in the Demo

These five concepts are how every modern coding agent — Claude Code, Cursor, Codex — actually works.

01

CLAUDE.md

Project memory

Tells the agent who you are + how you work, every session.

02

Plan Mode

The safety brake

Agent shows you the plan first. You approve, then it acts.

03

MCPs + APIs

Universal adapter

Lets the agent reach outside the chat — files, sites, services.

04

Skills

A button to press

Reusable workflow packs — instructions, scripts, and references load when relevant.

05

Sub-agents

AI employees

Parallel workers that handle their own piece without bloating the main chat.

CLAUDE.md

"Project memory — the briefing the agent reads every session."

~/.claude/CLAUDE.md	global — every project you open
└─ my-project/CLAUDE.md	this project — tech stack, rules
└─ my-project/web/CLAUDE.md	this folder — frontend-only notes
└─ my-project/api/CLAUDE.md	this folder — backend-only notes

Files cascade: global → project → folder. Closer wins.

BEFORE

```
my-project/
└─ CLAUDE.md ← 2,143 lines
  ├ Project overview
  ├ Full brand voice guide
  ├ Audience personas
  ├ Database schema dump
  ├ API patterns + examples
  ├ Testing strategy
  ├ Content workflow
  ├ Client onboarding
  ├ Release process
  ├ Old decisions from 2024
  ├ Meeting notes from March
  └─ Things to never do
```

AFTER

```
my-project/
└─ CLAUDE.md ← 180 lines
  ▼ Only loaded when needed
  context/
  │ └─ brand/
  │   │ └─ voice.md
  │   │ └─ audience.md
  │   └─ positioning.md
  │
  └─ tech/
    │ └─ architecture.md
    │ └─ database-schema.md
    │ └─ api-patterns.md
    └─ testing.md
  workflows/
  │ └─ content-creation.md
  │ └─ client-onboarding.md
  └─ release-process.md
```

When you reach for it: The agent should know your stack, conventions, or rules without you re-typing them.

Plan Mode

"The agent shows you the blueprint before it builds."

WHAT A PLAN LOOKS LIKE

Here's what I'll do — confirm before I write code:

1. Read the relevant files to understand context
2. Make the change in <these specific places>
3. Add a test that covers the new behavior
4. Run lint + typecheck before handing back

Approve (Y)

Refine (N)

THE FLOW

- 1 You ask**
Describe what you want.
- 2 Agent plans**
Reads code, drafts a plan.
- 3 You approve / refine**
Catch gaps before code is written.
- 4 Agent executes**
Now it actually builds it.

Shift+Tab toggles Plan Mode in Claude Code.

When you reach for it: *You're kicking off real work and want to catch misunderstandings BEFORE any code gets written.*

MCPs + APIs

"The universal adapter — how the agent reaches outside the chat."



MCP = the standard. Plug it in once, every agent (Claude, Cursor, Codex) can use it.

When you reach for it: *The agent needs live data: stock prices, your database, GitHub, the browser, anything.*

Skills

"A reusable recipe — instructions and helpers the agent loads only when needed."

WHAT'S IN A SKILL

my-skill/

SKILL.md

the recipe — what to do, when to do it

scripts/

reusable helper code the agent can run

references/

docs, examples, formats to follow

HOW IT WORKS

1

You write the skill once

Save it in ~/.claude/skills or in your project.

2

Agent only sees descriptions

Short one-liners — costs almost no tokens.

3

Full skill loads when relevant

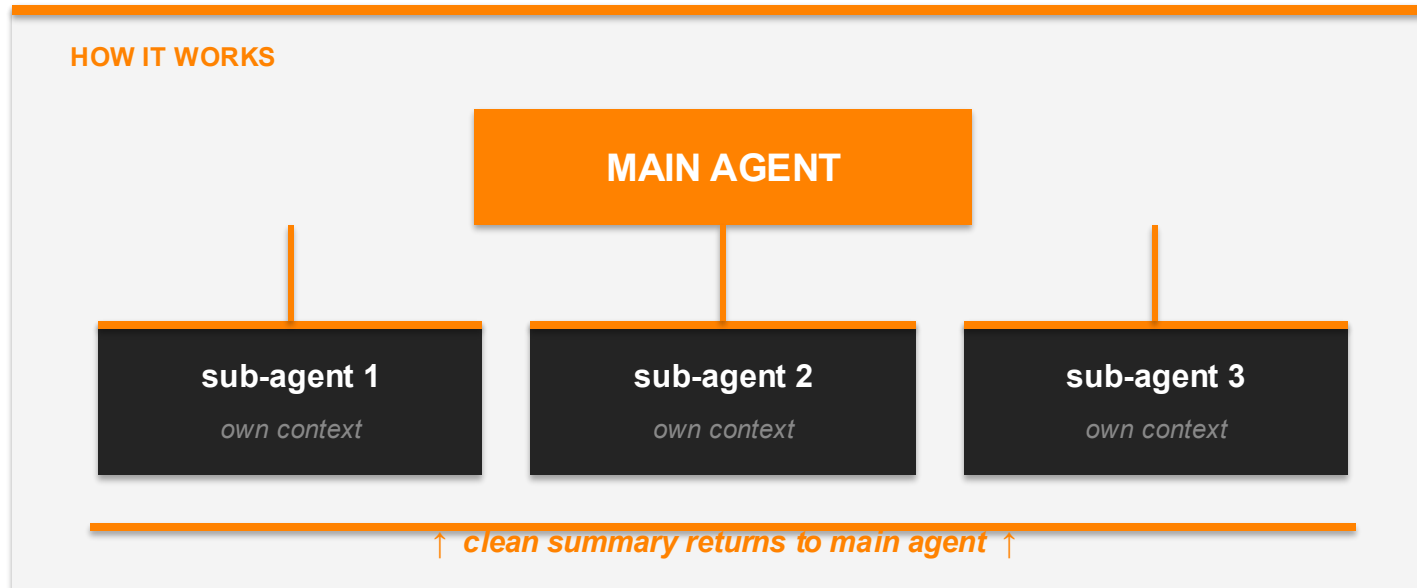
Instructions + scripts open only when triggered.

Progressive disclosure — full details load only when triggered.

When you reach for it: *You catch yourself explaining the same multi-step process more than twice.*

Sub-agents

"Dedicated workers — each handles one piece, then hands you back a clean summary."



ISOLATED CONTEXT

Each sub-agent gets its own clean window.

PARALLEL SPEED

All workers run at the same time, not in line.

CLEAN HANDOFF

Main thread only sees the result, not the mess.

When you reach for it: *A task naturally splits into independent pieces — or you want messy details kept out of the main thread.*

When to Reach for Which

If you remember one slide from tonight, this is it.

WHAT YOU NEED	REACH FOR
Agent should know my stack & rules	CLAUDE.md
Agent needs live data outside the chat	MCPs / APIs
I do this same flow over and over	Skill (slash command)
Work splits into parallel pieces	Sub-agents
Don't act until I approve the plan	Plan Mode

Stock Ticker Weekly Report

What we're going to build, live, in the next 15 minutes.

Pick tickers

AAPL, NVDA, TSLA

Fetch live data

prices, news, earnings

Summarize per ticker

what changed, why

Email me a report

Markdown, in my inbox

What does the agent need to pull this off?

Remember my preferences → **CLAUDE.md**

See data outside the chat → **MCPs + APIs**

Run the same flow every Sunday → **Skills**

Research 3 tickers in parallel → **Sub-agents**

Ask before sending the email → **Plan Mode**

Claude Code · Cursor · Codex

All three give you a coding agent. The difference is where it lives and how you drive it.

Claude Code

by Anthropic

CLI · Desktop · Plugin

- Terminal CLI, Mac/Win desktop, VS Code & JetBrains extensions
- Best primitives: skills, sub-agents, plan mode
- Plugin ecosystem

STUDENT DEAL

NA

Cursor

Visual editor

- VS Code, but AI-native
- Tab autocomplete on every line
- Model selection

STUDENT DEAL

FREE PRO FOR 1 YEAR
with .edu email

Codex (OpenAI)

by OpenAI

CLI · Desktop · ChatGPT

- Codex CLI + new desktop app
- Backed by GPT-5 family
- Great computer use

STUDENT DEAL

\$100 in Codex credits
with ChatGPT signup (US/CA)

MASTER THESE 15 CONCEPTS



Idea to Strategy
Roadmap



workflow



DECAY.md



Pause-Play



CLAUDE.md



Network



Hook



Terminal



Brain



AI in Healthcare

 AI can improve productivity up to 40% and reduce errors

AI is shifting healthcare from reactive → predictive

- **Used across:**
 - Patient care
 - Clinical workflows
 - Drug discovery
- **Improves:**
 - Patient outcomes
 - Efficiency
 - Speed of innovation

How AI is Used Today

- **Clinical AI Assistants**
 - Automate notes, reduce doctor burnout
 - Improve diagnosis accuracy
- **Predictive Health**
 - Detect diseases *years earlier*
 - Enable early intervention
- **Personalized Medicine**
 - Tailored treatments using genetics + lifestyle data
- **Drug Discovery**
 - AI can reduce timelines from **years** → **months**

AI in Healthcare- Examples

 AI can improve productivity up to 40% and reduce errors

AlphaFold

- A company spun out of DeepMind is using AI (AlphaFold) to **design entirely new drugs**
- These drugs are now **entering human clinical trials**
- AI predicts how proteins fold → lets scientists design targeted treatments faster (usually a 10 year process is now taking months)

UCSD AI Research- new type of Alzheimer's

- Modeled how proteins fold, interact, and change over time
- Identified patterns humans couldn't see
- Suggested new mechanisms behind how the disease starts and progresses

Clinical Co-Pilots (OpenEvidence)

- AI systems now:
 - Analyze patient data instantly
 - Suggest diagnoses
 - Reduce documentation burden

Personalized Medicine/ Wearable health (fall-risk prevention, real- time crisis prevention)

AlphaGenome

Medical Imaging

Tairex Agent Hospital (10k+ patients daily, 93+% accuracy)

AI in YOUR Field

What does AI look like in YOUR major or field of interest — and how are you going to use it for an edge?

If you're stuck, think:

Business

outreach, market research, deal memos

Engineering

CAD review, simulation, spec drafting

Pre-med

study summaries, case prep, clinical notes

Education

lesson plans, IEP drafts, grading rubrics

Architecture

site analysis, code lookups, mood boards

Law

case research, brief drafts, deposition prep

Finance

earnings reads, comp analysis, model checks


Design

wireframes, copy passes, design critique

New here?
Complete the **Sign Up Form**



https://groupme.com/join_group/11141080/zBZkxNij



https://www.instagram.com/aivols?igsh=bXlwYWtpeGJoa29j&utm_source=qr

Aivols.org