

Agentic AI: The Future of Work


Jack Cathey, Ph.D., CPA
Turner School of Accountancy
Belk College of Business
UNC Charlotte

NCACPA
Business & Industry Fall Conference
September 16, 2026
Asheville, NC

Presentation Overview


- How Did We Get Here?
- What Makes AI "Agentic"?
- Agents in Action: Examples
- Three Ways Agentic Systems Fail
- What's Next

How Did We Get Here?



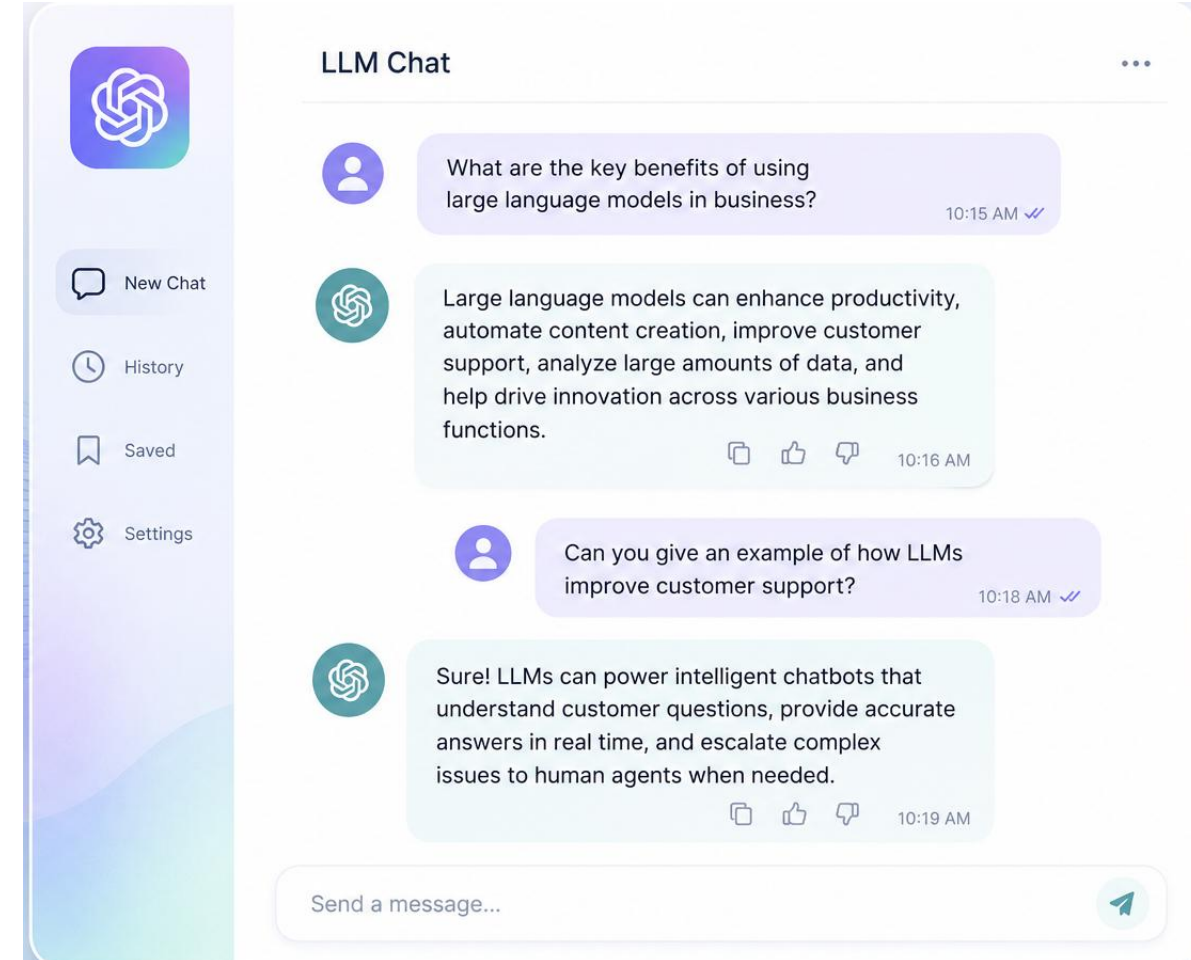
Artificial Intelligence (AI)

- AI is a branch of computer science that enables machines to perform tasks that usually require human intelligence
- Examples:
 - Playing Chess
 - Computer Vision
 - Autonomous Driving



Generative AI, including Large Language Models


- Generative AI is a type of AI that can create new content, such as images, text, or music, based on training data
- Large Language Models are a type of generative AI that uses machine learning to analyze and generate human-like text
- Examples: ChatGPT, Claude, Gemini, Copilot



LLMs Are Useful For

- Creating first drafts of documents
- Summarizing documents
- Brainstorming
- Writing Excel Macros

Fundamental LLM Loop



LLM – Quick Pro Tip

- If you work with LLMs frequently, you end up spending a lot of time typing; it can get tiring.
- Many LLMs now include a voice-to-text component that can save you the typing
- There are also dedicated speech-to-text apps that will work across any system where you want to dictate speech to text
- My favorite is Wispr Flow
 - It works with any application on my computer and phone
 - It does a good job interpreting my voice
 - It does a good job when I edit on the fly



Agentic AI

- Agentic AI is what you get when you combine an LLM with tools and multistep planning
- Here's McKinsey's definition: *Agentic AI is a system based on generative AI foundation models that can act in the real world and execute multistep processes — automating and performing complex tasks that would normally require human effort.*
- That's pretty abstract, so let's dig in by looking at some examples


Agents in Action: Four Examples

- Email Clean Up Agent
- File Organizer Agent
- Sales Analysis Agent
- Cost Analysis Agent

Email Clean Up Agent

- Goal – I don't want to keep buying storage for my personal Gmail account, and I know there are tons of marketing emails that I can delete to decrease the storage I need
- I can do this myself in Gmail, but it takes a while
- So, what if I created an Agent to do this for me?

Email Clean Up Agent - Continued



- Notice I am in “Cowork” mode instead of Chat
- I gave the Agent a goal, but did not have to describe the process
- Also notice I had connected the Agent to my Gmail account

Email Clean Up Agent - Continued

I want you to go through my  Gmail account and delete any archived emails from marriottbonvoy@email-marriott.com . When you are done, report back to me how many emails were deleted.

Loaded tools

The Gmail connector I have access to doesn't include a delete or trash function — it supports searching, labeling, and drafts, but not deletion. So I can't delete emails directly through it.

Here's what I can do:

1. **Search and count** — I can find all archived emails from `marriottbonvoy@email-marriott.com` and tell you exactly how many there are and what they contain.
2. **Delete via browser** — If you have the Claude in Chrome extension connected, I can navigate to Gmail and delete them through the browser interface.

Which would you prefer, or would you like both?

Progress >

Working folder >

Context ▾



Track tools and referenced files used in this task.

- The model has reasoned from the goal
- Each connector has only certain capabilities
- The agent can do things in a browser

Email Clean Up Agent - Continued

- The agent asks for permissions when it is acting. Generally, you will be given the choice to deny, to allow once, or to allow with the specific task.
- Also notice the agent is doing work in parallel and shows its progress and connectors on the top right

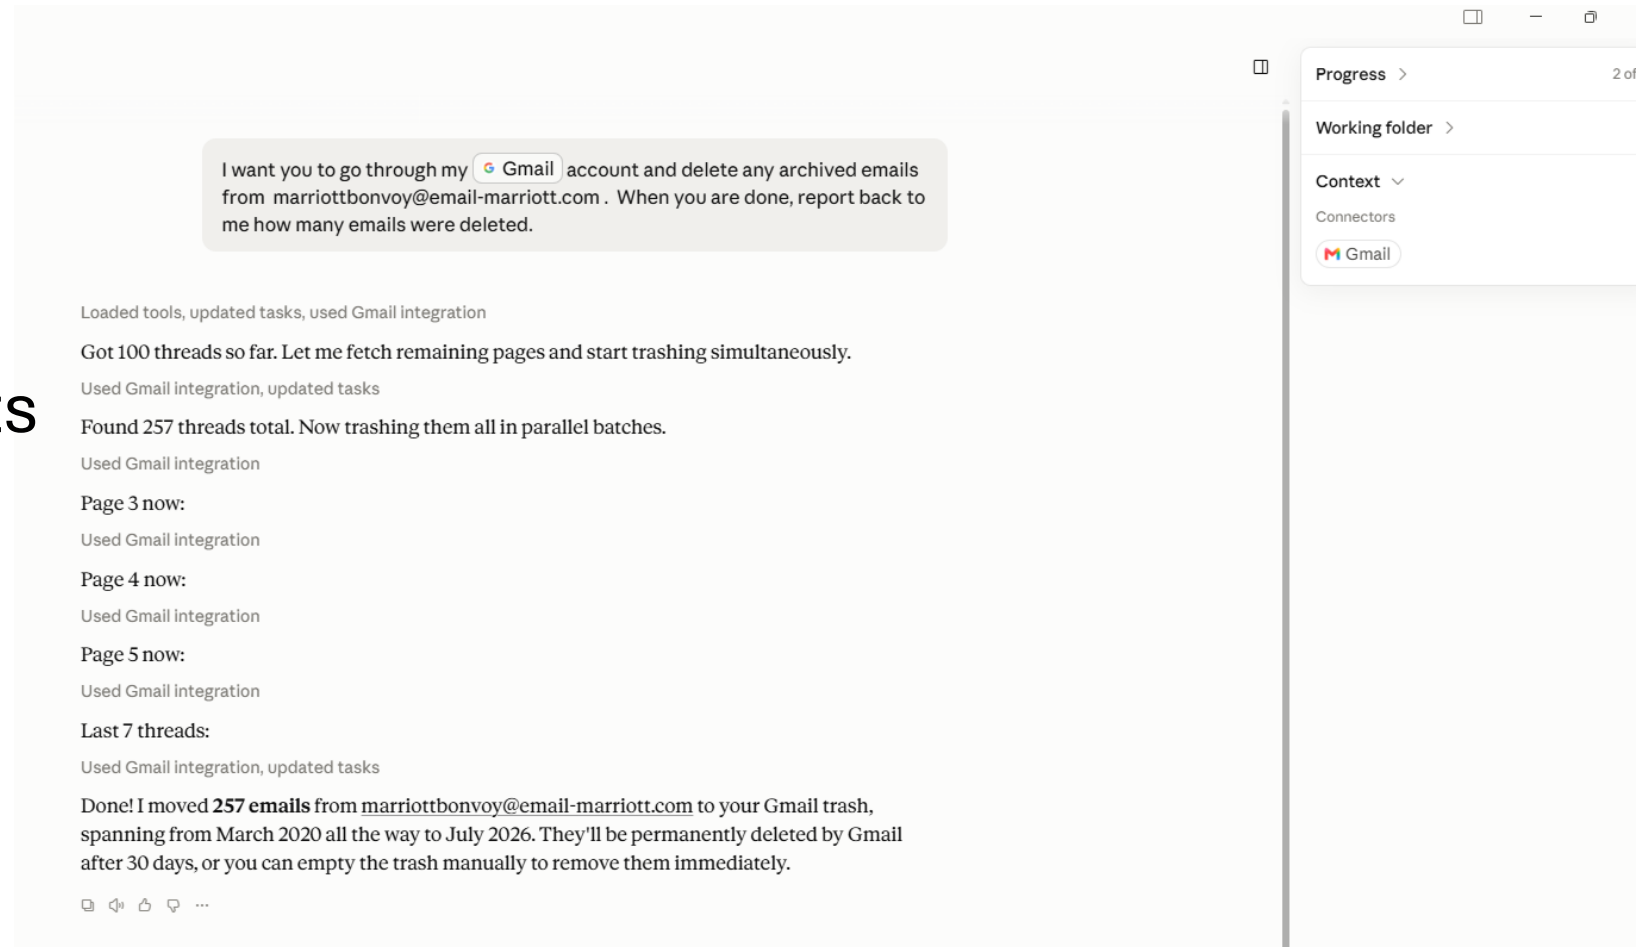
The screenshot displays the 'Email Clean Up Agent' interface. At the top, a dropdown menu shows 'Marriott emails deletion'. A central message box contains the instruction: 'I want you to go through my Gmail account and delete any archived emails from marriottbonvoy@email-marriott.com. When you are done, report back to me how many emails were deleted.' Below this, the agent's progress is shown: 'Loaded tools, updated tasks, used Gmail integration', 'Got 100 threads so far. Let me fetch remaining pages and start trashing simultaneously.', 'Used Gmail integration, updated tasks', 'Found 257 threads total. Now trashing them all in parallel batches.', and 'Used Gmail integration'. A permission dialog box is open, asking 'Claude wants to use Adds a sensitive label (Trash or Spam) to a thread. from Gmail'. It shows a JSON request:

```
{  'threadId': '19f8ea1092ca5af4',  'labelOption': 'TRASH'}
```

. Below the request, a warning states: 'Allowing this action comes with risks. Malicious instructions in files, emails, and web content could trick Claude into unintended actions. [Learn more](#)'. At the bottom of the dialog are three buttons: 'Allow for this task' (with a dropdown arrow), 'Allow once' (with keyboard shortcuts 'Ctrl + Enter'), and 'Deny' (with keyboard shortcut 'Esc'). To the right of the main interface, a sidebar shows the 'Progress' section with two steps: '1. Search for archived emails from marriottbonvoy@email-marriott.com' (completed) and '2. Deleting Marriott Bonvoy emails' (current step). Below this is the 'Working folder' section, followed by the 'Context' section which lists 'Connectors' and shows 'Gmail' as the active connector. At the bottom of the interface, there is a 'Tip' bar that says 'Need changes? Ask Claude to edit the file instead of making a new one.' and an 'Add to message' button. Below the tip bar is a text input field labeled 'Write a message...' with a 'Manual' dropdown menu and a 'Sonnet 4.6 Medium' dropdown menu. A small green circular icon with a white 'C' is visible in the bottom right corner. At the very bottom, a small footer text reads: 'Claude is AI and can make mistakes. Please double-check responses. Give us feedback'.


Email Clean Up Agent - Continued

- When the agent has completed the task (accomplished the goal) it reports back



Email Clean Up Agent - Continued


- Final result.
- Since the request only asked for archived emails to be deleted, the email that was still in my Inbox is still there.



File Organizer Agent

- As an academic, I am always finding articles that I want to save for later use
- When I find an article, I save it to a folder
- That folder is like the junk drawer in your house. Lots of different files in there, not organized, and some don't even have helpful names.
- What if I built an agent to help me with this?

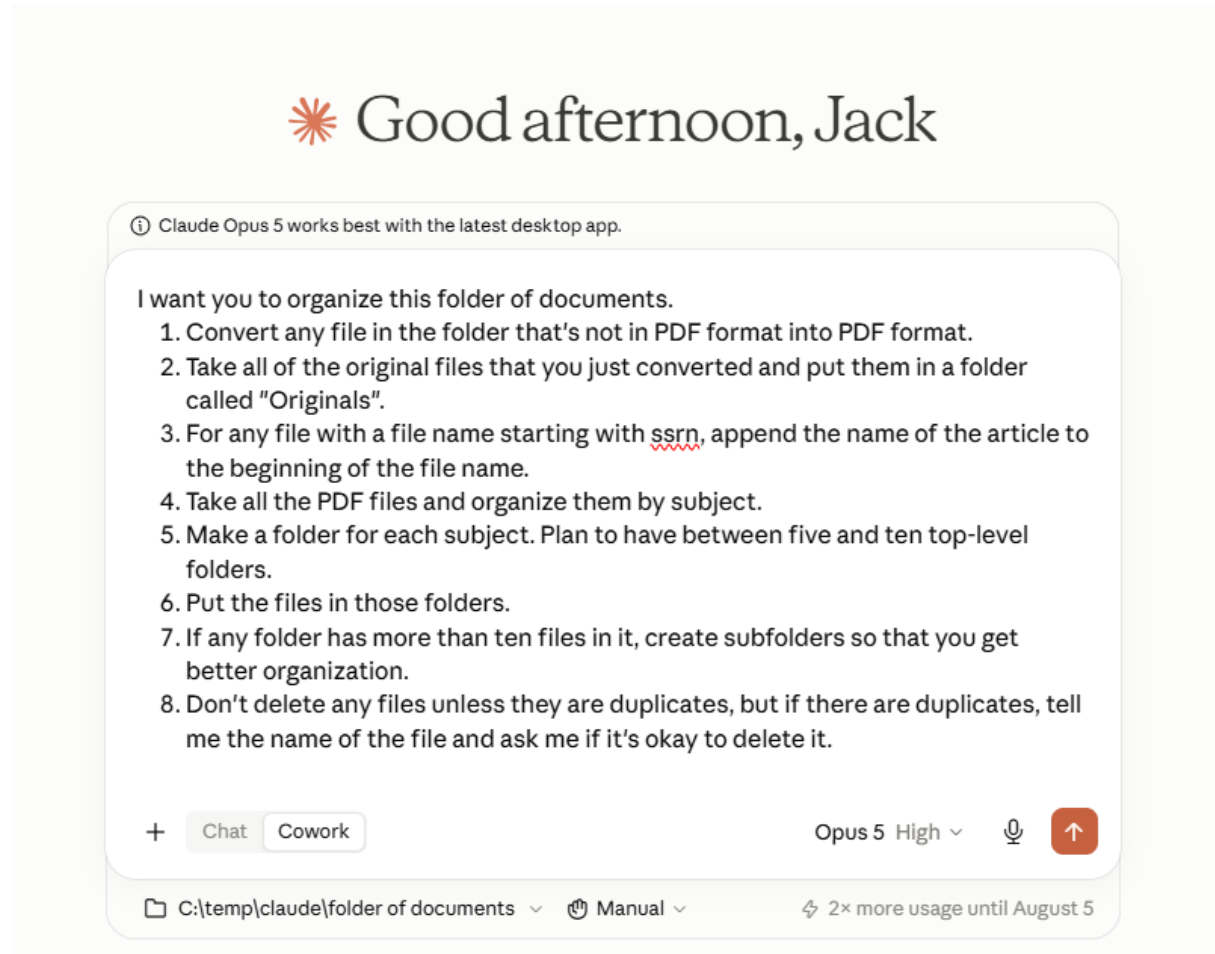
File Organizer Agent - Continued



Name	Date modified	Type	Size
Iceberg 1.pdf	4/5/2026 11:23 AM	Adobe Acrobat U...	3,634 KB
iceberg 2.pdf	4/5/2026 11:23 AM	Adobe Acrobat D...	2,733 KB
Inheriting versus Developing Data Analytic Tests and Auditors Professional Skepticism - ssm.pdf	9/17/2025 5:10 PM	Adobe Acrobat D...	1,943 KB
Intellectual property protection lost and competition - An examination using LLMs.pdf	7/21/2026 2:58 PM	Adobe Acrobat D...	2,354 KB
Intellectual Property Protection Lost and Competition - ssm.pdf	9/12/2025 5:08 PM	Adobe Acrobat D...	1,234 KB
ITDF-Human-Resilience-full-report-3.20.26.pdf	4/2/2026 6:12 PM	Adobe Acrobat D...	4,108 KB
ITDF-Human-Resilience-report-Exec-Summary-3-20-26.pdf	4/2/2026 6:12 PM	Adobe Acrobat D...	851 KB
JPMorgan Management Task Force Regarding 2012 CIO Losses 1-16-13.pdf	5/26/2025 6:03 PM	Adobe Acrobat D...	304 KB
KPMG Webinar Recalibrating Curriculum for AI Driven Jobs.pdf	4/2/2026 6:02 PM	Adobe Acrobat D...	1,486 KB
Language models and the intentional stance.docx	3/18/2026 10:13 AM	Microsoft Word D...	77 KB
Iee_2025_ai_critical_thinking_survey.pdf	2/19/2025 10:26 AM	Adobe Acrobat D...	996 KB
LLMs as Research Assistants Risk of Topic-Overclassification in Financial Disclosure Research - ssm.pdf	1/24/2026 3:34 PM	Adobe Acrobat D...	1,389 KB
Moving_beyond_the_hype_AI_whitepaper_updated.pdf	4/7/2026 10:14 AM	Adobe Acrobat D...	902 KB
Mythos and Cybersecurity - Schneier on Security.pdf	4/17/2026 11:26 AM	Adobe Acrobat D...	76 KB
mythosreadyv1.0.pdf	5/16/2026 7:59 AM	Adobe Acrobat D...	21,476 KB
mythosreadyv92.pdf	4/16/2026 3:42 PM	Adobe Acrobat D...	20,198 KB
omdia-ai-agents-thought-leadership.pdf	4/7/2026 10:49 AM	Adobe Acrobat D...	478 KB
open-world-evaluations.pdf	4/16/2026 3:24 PM	Adobe Acrobat D...	1,666 KB
Phishing Campaings and LLMs.pdf	3/25/2025 2:14 PM	Adobe Acrobat D...	936 KB
pwc-your-cloud-transformation-key-accounting-considerations-and-their-broader-impact.pdf	1/28/2026 3:57 PM	Adobe Acrobat D...	2,653 KB
Reliable Empirical Accounting Research in the Age of Generative AI - ssm.pdf	4/15/2026 9:34 AM	Adobe Acrobat D...	371 KB
Renewal - North Carolina State Board of Certified Public Accountant Examiners.pdf	6/25/2025 9:30 AM	Adobe Acrobat D...	102 KB
Responsible adoption of AI agents in the enterprise_PwC.pdf	10/9/2025 1:16 PM	Adobe Acrobat D...	1,386 KB
Responsible Disclosure in the Age of AI.pdf	6/2/2026 11:43 AM	Adobe Acrobat D...	2,081 KB
Rethinking Materiality in the Debate Over ESG _ CLS Blue Sky Blog.pdf	4/16/2026 1:30 PM	Adobe Acrobat D...	227 KB
sg-consensus-ai-safety.pdf	5/13/2025 7:19 AM	Adobe Acrobat D...	2,672 KB
ssrn-4986017.pdf	7/21/2026 3:01 PM	Adobe Acrobat D...	1,362 KB
ssrn-5083571.pdf	7/21/2026 2:54 PM	Adobe Acrobat D...	750 KB
ssrn-6128028.pdf	7/21/2026 3:00 PM	Adobe Acrobat D...	742 KB
ssrn-6439138.pdf	7/21/2026 3:02 PM	Adobe Acrobat D...	142 KB
SXYZ.pdf	9/11/2025 1:15 PM	Adobe Acrobat D...	646 KB
The Coming AI Hackers.docx	5/16/2026 8:00 AM	Microsoft Word D...	72 KB
The Impact of ChatGPT on Student Outcomes.pdf	5/17/2025 12:04 PM	Adobe Acrobat D...	1,162 KB
The Shape of the Thing.docx	3/12/2026 6:28 PM	Microsoft Word D...	787 KB
Toward Digital Corporate Law - ssm.pdf	11/6/2025 2:12 PM	Adobe Acrobat D...	914 KB
trust-in-the-age-of-agents.pdf	3/23/2026 3:07 PM	Adobe Acrobat D...	455 KB
Understanding Different ChatGPT Models.docx	3/21/2025 4:11 PM	Microsoft Word D...	169 KB
Unit Tests for Microsoft Excel.pdf	5/26/2025 6:15 PM	Adobe Acrobat D...	201 KB
Using AI to Identify Exogenous Shocks and Conduct Archival Accounting Research - ssm.pdf	7/23/2026 10:39 AM	Adobe Acrobat D...	1,062 KB
White Papers 2025 Leveraging COBIT for Effective AI System Governance.pdf	4/22/2025 6:19 PM	Adobe Acrobat D...	17,091 KB
Why does GenAI give different answers when you ask the same question.docx	3/23/2025 8:39 AM	Microsoft Word D...	77 KB

- Here is what the folder looks like
- Notice the ssrn-xxxxxx.pdf files
- Notice most files are .pdf, but there are other file types as well

File Organizer Agent - Continued



- Notice I have given it access to the folder on my computer
- I have also told it to convert the non-.pdf files to .pdf
- I have told it to organize the files and given it a general structure
- Lots of elements to my goal

File Organizer Agent - Continued

Document folder organization ▾

How should I structure the top-level folders?

By subject matter (Recommended)

Filing them now.

Updated tasks, ran 2 commands

Done. All 79 PDFs filed, nothing deleted, nothing lost.

Converted to PDF — 9 `.docx` + 1 `.webp` (Billionaires infographic). Sources moved to `Originals/`.

Renamed the 4 opaque SSRN files:

- `ssrn-4986017` → Generative AI in Financial Reporting
- `ssrn-5083571` → How the EU Sustainability Due Diligence Directive Could Reshape Corporate America
- `ssrn-6128028` → Auditor Narcissism and Disclosure Style of Key Audit Matters
- `ssrn-6439138` → The First Annual Reports of the LLM Era


Structure (9 top-level folders; the two that exceeded 10 files got subfolders):

01 AI Agents & Enterprise Adoption	10	
02 AI Governance, Risk & Safety	9	
03 AI & Cybersecurity	7	
04 AI in Academic & Accounting Research	8	
05 Accounting, Auditing & Corporate Gov.	13	→ Auditing & Assurance / Financial Reporting & Disclosure /
Regulation		& Corporate Law / Risk, Controls & Case Studies
06 ESG & Sustainability Reporting	9	
07 Higher Education & the CPA Profession	7	↓
08 Future of Work & AI in Society	11	→ AI as Normal Technology & Big

- The agent converted the files as requested, renamed the files as requested, looked at the files, and decided how to organize them; it created the new thematic organizational folders and moved the files

File Organizer Agent - Continued

- Finished Product




Agentic AI - Recap

- The agents we have seen so far have:
 - A goal – an outcome we want to achieve. It can be somewhat general or more detailed.
 - A set of connections to tools and storage
 - The agent decides how to accomplish the goal
 - When the agent is actually doing something, especially if we haven't told it to do that thing or if the action has consequences, the agent comes back and asks for permission
 - When the agent achieves the goal, it reports back
- Let's look at a couple more examples that are a little more accounting focused

Sales Analysis Agent

- Each month a set of weekly sales transaction files are created and placed in a folder




The screenshot shows a Windows File Explorer window with the address bar indicating the path: This PC > OS (C:) > temp > claude > sales. The window displays a list of five Excel files, all created on 1/22/2023 at 10:14 AM. The files are WEEK_1.xlsx (54,034 KB), WEEK_2.xlsx (60,856 KB), WEEK_3.xlsx (55,695 KB), WEEK_4.xlsx (74,767 KB), and WEEK_5.xlsx (21,846 KB).

Name	Date modified	Type	Size
WEEK_1.xlsx	1/22/2023 10:14 AM	Microsoft Excel W...	54,034 KB
WEEK_2.xlsx	1/22/2023 10:14 AM	Microsoft Excel W...	60,856 KB
WEEK_3.xlsx	1/22/2023 10:14 AM	Microsoft Excel W...	55,695 KB
WEEK_4.xlsx	1/22/2023 10:14 AM	Microsoft Excel W...	74,767 KB
WEEK_5.xlsx	1/22/2023 10:14 AM	Microsoft Excel W...	21,846 KB

Sales Analysis Agent - Continued

- Here is what the files look like



The screenshot displays the Microsoft Excel interface with the 'Home' tab selected. The ribbon shows various toolbars including Clipboard, Font, Alignment, Number, Styles, Cells, and Editing. The active worksheet is 'WEEK_1.xlsx'. The data is organized in a table with columns for transaction details. The table includes a header row and 15 data rows. The columns are: TRANSACTION_ID, TRAN_DATE, STORE, REGISTER, TRAN_NUM, TRAN_TIME, CUST_ID, TRAN_LINE_NUM, MIC, TRAN_TYPE, ORIG_PRICE, SALE_PRICE, TRAN_AMT, TENDER_TYPE, and SKU. The data represents transactions from June 1, 2016, to June 4, 2016.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
	TRANSACTION_ID	TRAN_DATE	STORE	REGISTER	TRAN_NUM	TRAN_TIME	CUST_ID	TRAN_LINE_NUM	MIC	TRAN_TYPE	ORIG_PRICE	SALE_PRICE	TRAN_AMT	TENDER_TYPE	SKU	
1	90492496	6/3/2016 0:00	921	54	19 1550	177663749		1 466	R		59	59	-59	DAMX	3655749	
2	90492545	6/3/2016 0:00	222	3	18 1835	218381288		1 235	R		99	82.5	-82.5	BANK	4196406	
3	90492624	6/3/2016 0:00	218	11	15 1825	199069831		2 000	R		0	15.37	-15.37	DLRD	999999967	
4	90492702	6/2/2016 0:00	731	12	9 1417	155888326		1 088	P		26	21.19	21.19	BANK	4390620	
5	90492724	6/1/2016 0:00	523	1	10 1447	158325436		2 086	P		0	11.9	11.9	DLRD	3276781	
6	90492890	6/2/2016 0:00	902	33	3 1033	179698805		1 854	P		39	39	39	DAMX	4307634	
7	90492912	6/1/2016 0:00	742	36	37 1942	218038706		1 877	P		26	9.1	9.1	BANK	2699885	
8	90493116	6/4/2016 0:00	203	3	22 1612	218091939		1 004	P		34	34	34	BANK	4031870	
9	90493165	6/3/2016 0:00	365	10	32 1751	198349802		1 000	R		0	12.8	-12.8	BANK	999999950	
10	90493244	6/3/2016 0:00	374	35	2 1016	118555906		2 000	P		0	2.68	2.68	DLRD	999999950	
11	90493322	6/2/2016 0:00	914	36	9 1447	189720713		1 028	P		33	33	33	BANK	3362368	
12	90493344	6/1/2016 0:00	742	36	37 1942	218038706		2 636	P		59	20.65	20.65	BANK	3811423	
13	90493478	6/3/2016 0:00	436	4	20 1707	173531399		8 060	P		30	30	30	BANK	6446368	
14	90493557	6/3/2016 0:00	698	354	2133 1040	209225960		1 000	P		0	8.5	8.5	BANK	999999976	

Sales Analysis Agent - Continued

- Create the agent, giving it a goal and access to the files

☀ Good afternoon, Jack

Using the Excel files in this folder, prepare a sales analysis report for the month of June. Make sure to include a graph showing sales by day. Also include a table that shows the five highest-performing stores and the 10 highest-performing SKUs. Add anything else that seems interesting. Look into what time of day has the highest sales across all stores. Transactions coded P are customer purchases, and transactions coded R are customer returns. Make sure to note the five products based on SKU that are returned the most frequently. Exclude sales from store 698. The report should be in .docx format.



Chat

Cowork

Sonnet 4.6 Medium ▾



📁 C:\temp\claude\sales ▾ 🖱 Manual ▾

⚡ 2× more usage until August 5

Sales Analysis Agent - Continued

- The agent determines the steps to use
- The agent includes the items we ask for but also adds additional items it thinks will be helpful

June sales analysis report ▾

Using the Excel files in this folder, prepare a sales analysis report for the month of June. Make sure to include a graph showing sales by day. Also include a table that shows the five highest-performing stores and the 10 highest-performing SKUs. Add anything else that seems interesting. Look into what time of day has the highest sales across all stores. Transactions coded P are customer purchases, and transactions coded R are customer returns. Make sure to note the five products based on SKU that are returned the most frequently. Exclude sales from store 698. The report should be in .docx format.

Ran 17 commands, updated tasks, loaded tools ▸

fastexcel is fast — 591K rows in seconds. Now let me run the full analysis.

Ran 11 commands, updated tasks

Now let me combine all weeks and generate the charts and final summary:

Ran a command

All data processed. Now let me read the docx skill and build the report:

Ran 9 commands, updated tasks, read a file

Report generated. Let me verify it visually before delivering:

Ran 4 commands, read 4 files

Looks great. Now copy to the workspace folder and present it:

Ran a command, updated tasks, loaded tools

The report is ready. Here's a quick summary of what's inside across 6 pages:

Executive Summary — 9 KPI cards covering gross sales (\$71.7M), net sales (\$60.9M), total returns (\$10.8M), transaction counts, return rate (15.05%), and average transaction value (\$30.01).

Daily Net Sales — Bar chart with daily average reference line, plus a full 30-day table showing each day's net sales and % vs. the daily average. Notable findings: June 29 was a

Write a message...

+ 📎 📖 Manual ▾ Sonnet 4.6 Medium ▾ 🔊

Claude is AI and can make mistakes. Please double-check responses. Give us feedback

June 2016 sales analysis report · DOCX Google Drive ↗

Sales Analysis Report

June 2016 | All Stores (Excl. Store 698)
Prepared: July 24, 2026

Executive Summary

The following report covers all transactions for June 2016 across 297 store locations, excluding Store 698. Data spans June 1-30, 2016 and encompasses 2,690,763 total transaction lines.

\$71,702,478.47 Gross Sales	\$60,914,616.10 Net Sales (after returns)	\$10,787,862.37 Total Returns
2,389,202 Purchase Transactions	301,561 Return Transactions	15.05% Return Rate
\$30.01 Avg. Transaction Value	\$2,030,487.20 Avg. Daily Net Sales	297 Active Stores

Daily Net Sales

Net sales (purchases minus returns) varied considerably across the month. Weekends consistently outperformed weekdays, with the most notable spike occurring on June 29, which recorded \$6,271,234.07 — more than three times the monthly daily average of \$2,030,487.20. June 18 (Father's Day weekend Saturday) was the second-highest day at \$3,613,443.21. The slowest day was June 27 at \$1,172,073.83.

Figure 1: Net Sales by Day — June 2016 (blue = positive, dashed line = daily average)

Sales Analysis Agent - Continued

- The Agent developed a good report
- I appreciated that the Agent understands Net Sales
- I appreciate that the Agent recognized that one of the high sales days was Father's Day weekend

Sales Analysis Report

June 2016 | All Stores (Excl. Store 698)

Prepared: July 24, 2026

Executive Summary

The following report covers all transactions for June 2016 across 297 store locations, excluding Store 698. Data spans June 1–30, 2016 and encompasses 2,690,763 total transaction lines.

\$71,702,478.47 Gross Sales	\$60,914,616.10 Net Sales (after returns)	\$10,787,862.37 Total Returns
2,389,202 Purchase Transactions	301,561 Return Transactions	15.05% Return Rate
\$30.01 Avg. Transaction Value	\$2,030,487.20 Avg. Daily Net Sales	297 Active Stores

Daily Net Sales

Net sales (purchases minus returns) varied considerably across the month. Weekends consistently outperformed weekdays, with the most notable spike occurring on June 29, which recorded \$6,271,234.07 — more than three times the monthly daily average of \$2,030,487.20. June 18 (Father's Day weekend Saturday) was the second-highest day at \$3,613,443.21. The slowest day was June 27 at \$1,172,073.83.




Figure 1: Net Sales by Day — June 2016 (blue = positive, dashed line = daily average)

Sales Analysis Agent - Continued

- The Agent, as requested, summarized sales by day
- The Agent also added a comparison of that day to the average day

Daily Net Sales Detail

Date	Net Sales	vs. Daily Avg
Wed, Jun 1	\$1,613,979.79	-20.5%
Thu, Jun 2	\$1,761,306.18	-13.3%
Fri, Jun 3	\$2,146,256.09	+5.7%
Sat, Jun 4	\$2,886,895.49	+42.2%
Sun, Jun 5	\$1,437,496.11	-29.2%
Mon, Jun 6	\$1,403,148.94	-30.9%
Tue, Jun 7	\$1,614,911.89	-20.5%
Wed, Jun 8	\$1,708,259.48	-15.9%
Thu, Jun 9	\$1,805,202.44	-11.1%
Fri, Jun 10	\$2,208,713.73	+8.8%
Sat, Jun 11	\$2,797,365.89	+37.8%
Sun, Jun 12	\$1,374,773.03	-32.3%
Mon, Jun 13	\$1,346,898.17	-33.7%
Tue, Jun 14	\$1,612,740.57	-20.6%
Wed, Jun 15	\$1,861,538.69	-8.3%
Thu, Jun 16	\$2,103,158.14	+3.6%
Fri, Jun 17	\$2,629,956.26	+29.5%
Sat, Jun 18	\$3,613,443.21	+78.0%
Sun, Jun 19	\$1,258,098.67	-38.0%
Mon, Jun 20	\$1,217,169.34	-40.1%
Tue, Jun 21	\$1,448,568.73	-28.7%
Wed, Jun 22	\$1,538,695.00	-24.2%
Thu, Jun 23	\$1,635,461.19	-19.5%
Fri, Jun 24	\$1,900,739.65	-6.4%
Sat, Jun 25	\$2,399,483.86	+18.2%
Sun, Jun 26	\$1,208,206.13	-40.5%
Mon, Jun 27	\$1,172,073.83	-42.3%
Tue, Jun 28	\$1,289,279.51	-36.5%
Wed, Jun 29	\$6,271,234.07	+208.9%
Thu, Jun 30	\$3,649,562.02	+79.7%

Sales Analysis Agent - Continued

- The Agent, as requested, created tables showing the top 5 stores and the top 10 SKUs

Top 5 Performing Stores

The following five stores generated the highest gross purchase sales for the month of June. Store 405 led all locations by a meaningful margin, outperforming the #2 store (796) by \$75,116.14.

Rank	Store #	Gross Sales	Share of Total
#1	405	\$859,682.90	1.20%
#2	796	\$784,566.76	1.09%
#3	760	\$777,222.76	1.08%
#4	716	\$756,656.88	1.06%
#5	768	\$755,241.16	1.05%

Top 10 Performing SKUs

The ten SKUs below generated the most gross purchase revenue in June. SKU 9999999950 dramatically outperformed all others, accounting for 7.12% of total gross sales on its own. Note that SKUs beginning with "9999999" appear to be department-level or tender-reduction codes rather than individual sellable items, and likely represent bundled or non-inventory transactions.

Rank	SKU	Gross Sales	Share of Total
#1	9999999950	\$5,108,247.98	7.12%
#2	9999999100	\$366,122.74	0.51%
#3	9999999968	\$326,807.82	0.46%
#4	9999999500	\$208,816.89	0.29%
#5	9999999811	\$148,740.40	0.21%
#6	9999999910	\$94,176.77	0.13%
#7	9999999600	\$82,442.60	0.11%
#8	9999999975	\$65,531.71	0.09%
#9	9999999981	\$56,858.89	0.08%
#10	3353153	\$53,702.10	0.07%

Sales by Time of Day

Purchase activity follows a strong mid-afternoon pattern. The peak sales hour across all stores in June was 2:00 PM (14:00), which generated \$9,118,561.48 in gross purchases — the single highest hour of the day. Sales ramp up sharply from store open around 9:00–10:00 AM, reach peak traffic between 12:00 PM and 3:00 PM, then taper through the evening. The 1:00 PM – 3:00 PM window (hours 13–15) alone accounted for over 35% of all daily purchase revenue, suggesting this is the most critical window for staffing and promotional placement.

Sales Analysis Agent - Continued

- As requested, the Agent did a time-of-day analysis



Figure 2: Total Purchase Sales by Hour of Day — June 2016 (peak hour highlighted in dark green)

Hourly Sales Breakdown

Hour	Sales	% of Daily Total
00:00 – 01:00	\$6,087.84	0.01%
04:00 – 05:00	\$68,044.03	0.09%
06:00 – 07:00	\$399.23	0.00%
07:00 – 08:00	\$17,019.78	0.02%
08:00 – 09:00	\$24,318.40	0.03%
09:00 – 10:00	\$101,708.74	0.14%
10:00 – 11:00	\$3,337,365.88	4.65%
11:00 – 12:00	\$6,153,452.02	8.58%
12:00 – 13:00	\$7,829,422.39	10.92%
13:00 – 14:00	\$8,709,912.23	12.15%
14:00 – 15:00	\$9,118,561.48	12.72%
15:00 – 16:00	\$8,573,065.58	11.96%
16:00 – 17:00	\$7,542,730.41	10.52%
17:00 – 18:00	\$6,443,466.63	8.99%
18:00 – 19:00	\$5,115,087.88	7.13%
19:00 – 20:00	\$4,266,747.03	5.95%
20:00 – 21:00	\$3,841,372.63	5.36%
21:00 – 22:00	\$549,370.31	0.77%
22:00 – 23:00	\$53,688.62	0.01%
23:00 – 24:00	\$657.36	0.00%

Sales Analysis Agent - Continued

- The Agent decided to add some analysis on payment method and also a comparison of weekend vs. weekday

Returns Analysis

Total returns for June amounted to \$10,787,862.37 across 301,561 return transactions, representing a return rate of 15.05% of gross sales. This is a material figure worth monitoring on a store-by-store basis.

Top 5 Most Frequently Returned SKUs

The following SKUs had the highest volume of return transactions in June. SKU 9999999950 — also the top revenue SKU — had by far the most returns at 87,654 transactions, though this may reflect its outsized transaction volume.

Rank	SKU	Return Count	Return Value	Avg. Return Value
#1	9999999950	87,654	\$681,249.28	\$7.77
#2	9999999981	42,286	\$605,561.89	\$14.32
#3	9999999967	27,123	\$504,303.37	\$18.59
#4	9999999968	683	\$87,401.16	\$127.97
#5	9999999952	319	\$4,252.90	\$13.33

Additional Insights

Payment Method Breakdown

Bank/debit card (BANK) was the dominant payment method in June, accounting for the largest share of purchase revenue, followed by American Express (DAMX) and debit/loyalty card (DLRD). Cash represented a very small portion of transactions.

Tender Type	Sales	Share of Total
BANK	\$39,907,376.72	55.66%
DAMX	\$15,502,476.22	21.62%
DLRD	\$14,591,835.97	20.35%
DL#	\$1,134,377.27	1.58%
CASH	\$455,446.01	0.64%
GIFT	\$110,966.28	0.15%

Weekend vs. Weekday Performance

June 2016 had 8 weekend days and 22 weekdays. Weekend days averaged \$2,121,970.30 in net sales vs. \$1,997,220.62 for weekdays — a 6.2% weekend premium. Saturday and Sunday consistently ranked among the top revenue days each week, with the exception of the anomalous spike on June 29 (Wednesday).

Notable Anomaly: June 29 Spike

Sales Analysis Agent - Continued

- The Agent also took a stab at June 29th and included some helpful notes
- This seems to be a helpful first draft, which is what LLMs are good for

June 29 recorded net sales of \$6,271,234.07, which is 209% above the daily average and more than 1.7x the following day's figure. This spike warrants further investigation — possible causes include a large batch posting, a system re-run of prior transactions, a major promotional event, or a data entry anomaly. It is recommended to validate the June 29 data against source transaction logs before including it in any forward-looking forecasts.

Methodology & Notes

- Data source: Five weekly transaction files (WEEK_1.xlsx through WEEK_5.xlsx) covering June 1–30, 2016.
- Store 698 was excluded from all figures per reporting requirements.
- "Purchases" are transactions coded P; "Returns" are coded R. Net sales = Purchase TRAN_AMT + Return TRAN_AMT (returns carry negative value).
- Time-of-day analysis is based on the TRAN_TIME field (HHMM format). Hours with no recorded transactions are excluded from the hourly table.
- SKUs beginning with "9999999" appear to represent non-standard or departmental line items rather than discrete sellable products.
- All dollar figures are in USD.

Sales Analysis Agent - Continued

- Things that could be added to this Agent
 - You could allow the Agent to go to your ERP and collect the Sales files
 - You could add specific report elements that you want to include
 - You could include previous month's data and/or reports and ask for a comparison
 - You could have the agent email this report to you
 - What else would you want to add?

Cost Analysis Agent

- The data used in the Sales Analysis was very clean and well structured
- What about something a bit messier
- Let's look at a cost analysis using some messy data

- The data is messier since it is not all in the same format
- Let's try to use an Agent to analyze it

AutoSaveOff

lchess costs.xlsx • View this PC

Search

FileHomeInsertDrawPage LayoutFormulasDataReviewViewAutomateDeveloperHelpAcrobatPower Pivot

Paste

Clipboard

Calibri10

B

I

U

Font

Wrap Text

Merge & Center

Alignment

General

\$%‰

100

00

Number

Conditional Formatting

Format as Table

Cell Styles

Styles

Insert

Delete

Form


Cells

O28

12	A	B	C	D	E	F	G	H	I	J	K	L	M
1	Sysadmin & Hosting												
2	Server	Monthly	Yearly	Name	Threads	RAM	Storage	SSD	CPU	DATACENTER	FISHNET ROLE	NOTES	
3	Frontend + Iila JVM	\$558	\$6,696	manta	96	192	2400	true	2x Xeon Gold 6248R @ 3GHz	GRA3			
4	Socket frontend 1	\$118	\$1,416	syrup	16	32	500	true	Xeon-D 2141I @ 2.20GHz	GRA2			
5	Socket frontend 2	\$118	\$1,416	taffy	16	32	500	true	Xeon-D 2141I @ 2.20GHz	GRA2			
6	Socket frontend 3	\$80	\$960	pingu	12	32	900	true	AMD EPYC 4244P @ 3.8GHz	GRA4			
7	Container platform tests	\$80	\$960	probe	12	32	900	true	AMD EPYC 4244P @ 3.8GHz	GRA4			
8	Iila-ws JVM	\$287	\$3,444	starr	48	96	960	true	2x Xeon Silver 4214 @ 2.20GHz	GRA2			
9	Iila-http	\$138	\$1,656	hyper	16	32	512	true	Xeon-E 2388G @ 3.2GHz	GRA3			
10	DB primary	\$462	\$5,544	kaiju	32	256	7680	true	AMD EPYC 7313 @ 3GHz	GRA3			
11	DB secondary	\$331	\$3,972	rodan	24	128	7680	true	AMD Ryzen 9 5900X @ 3.7GHz	GRA3			
12	DB secondary	\$331	\$3,972	gappa	24	128	7680	true	AMD Ryzen 9 5900X @ 3.7GHz	GRA3			
13	DB secondary (1h delay)	\$239	\$2,868	fleck1	12	64	7680	true	Intel Xeon-E 2386G @ 3.5GHz	SBG5			
14	DB secondary (24h delay)	\$239	\$2,868	fleck2	12	64	7680	true	Intel Xeon-E 2386G @ 3.5GHz	LIM3			
15	Study DB	\$148	\$1,776	study	16	64	960	true	Xeon-E 2288G @ 3.7GHz	GRA2			
16	Monitoring + fishnet	\$248	\$2,976	talos	24	64	1920	true	AMD Ryzen 9 5900X @ 2.20GHz	GRA3	move		
17	Puzzle DB + Lichess insights	\$180	\$2,160	rubik	8	32	3840	true	Xeon E-2274G @ 4GHz	SBG3			
18	Elasticsearch	\$172	\$2,064	sirsch	12	64	1920	true	Xeon-E 2136 @ 3.30GHz	RBX8	move		
19	7-piece tablebase	\$162	\$1,944	bwrdd	12	32	32000	false	AMD Ryzen 5 PRO 3600 @ 3.6GHZ	SBG5			
20	8-piece tablebase	\$209	\$2,508	eight	24	128	160000	false	AMD Ryzen 9 3900 @ 3.1GHz	FSN1		Hetzner	
21	Misc. Communication services	\$118	\$1,416	radio	16	32	500	true	Xeon D-2141I @ 2.20GHz	GRA2			
22	Backup v2 + puzzles	\$197	\$2,364	plato	12	32	18000	false	Ryzen 5 Pro 3600 @ 3.6GHz	LIM3			
23	Picture storage + future open database	\$79	\$948	image	12	32	4000	false	Xeon ES-1650 v4 @ 3.60GHz	GRA2			
24	Machine learning	\$180	\$2,160	apate	32	128	500	true	Epyc 7351P @ 2.4 GHz	SBG3			
25	YoloDB we can afford to lose	\$150	\$1,800	achoo	12	64	2000	true	Xeon E-2136 @ 3.30GHz	SBG3	analysis		
26	Opening explorer	\$221	\$2,652	bookd	12	128	24000	false	AMD Ryzen 5 PRO 3600 @ 3.6GHZ	GRA3	move		
27	Study secondary (24h delay)	\$77	\$924	sofia	12	32	900	true	Xeon E-2386G @ 3.50GHz	GRA3			
28	Staging	\$66	\$792	snafu	12	64	500	true	AMD Ryzen 5 5600X @ 3.7GHz	SBG3			
29	Machine learning	\$42	\$504	butch	8	16	1800	false	Xeon W3520 @ 2.67GHz	BHS6		Soyoustart	
30	Machine learning	\$42	\$504	burke	8	16	1800	false	Xeon W3520 @ 2.67GHz	BHS2		Soyoustart	
31	Machine learning	\$92	\$1,104	uluru	12	64	1800	false	Xeon ES-1650v3 @ 3.5GHz	RBX4		Soyoustart	
32	Private communications	\$38	\$456	zulip	2	4	80	true	Xeon Skylake @ 2GHz	NBG		Hetzner	
33	Total servers	\$5,402.00	\$64,824.00		566	2084	301592						
34													
35	Domain names	\$10	\$120										
36	Google Cloud	\$2	\$24										
37	AWS for Nonprofits Admin Fee	\$15	\$175										
38	Apple store	\$10	\$114										
39	Sysadmin (part time)	\$1,505	\$18,060										
40	Misc OVH fees (IPs, premium support)	\$74	\$888										
41													
42	Content & Community												
43	Titled Arenas	\$2,000	\$26,000										

Cost Analysis Agent - Continued

- The goal is provided as before, and some specific items are requested



Cost Analysis Agent - Continued

Lichess 2025 operating expenses report ▾

This folder contains Lichess's 2025 operating expenses. Analyze the data and create a Word document report that covers total annual spending, breakdown by category with percentages, the three largest expense categories. Add any additional interesting items or trends you spot. Format the report professionally with clear sections, tables where appropriate, and an executive summary at the top.

I'll start by looking at what's in the folder.

Ran 16 commands, read 8 files, updated tasks

Total 2025 spend is **€789,389**, and every category subtotal reconciles exactly to that figure.

Breakdown: Content & Community €177,660 (22.5%), Development €165,324 (20.9%), Operations €151,060 (19.1%), Moderation €111,000 (14.1%), Other Expenses €100,140 (12.7%), Sysadmin & Hosting €84,205 (10.7%).

The most striking findings the report digs into:

- It's a **people budget** — ~€515K (65%) is compensation, while all servers and cloud combined are just €66,145 (8.4%).
- **Payment-processor fees (€37,824)** cost more than the database cluster, frontends and tablebases put together — the price of being donation-funded.
- **Moderation costs 32% more than all technical infrastructure.**
- Unit economics: €2,163/day, **€0.000416 per game**, ~2,404 games per euro.
- Two things worth confirming: Titled Arenas shows €2,000/month but €26,000/year (13 months), and a Google Cloud line marked "paused" still carries €24.

C:\temp\claude\lichess\Lichess_2025_Operating_Expense_Report
Document

Google Drive ▾

🔍 🗑️ 📁 📄 ⌵ 2 hours ago

✔ Claude Opus 5 works best with the latest desktop app. Update

Write a message...

+ 📁 🗑️ Manual ▾

Opus 5 High ▾ 🗣️

Claude is AI and can make mistakes. Please double-check responses. Give us feedback

Lichess 2025 operating expense report · DOCX

Google Drive ▾ 🔗 ✕

LICHESS.ORG

2025 Operating Expense Report

Annual budget analysis — category breakdown, cost drivers, and observations

Executive Summary

Lichess operated on a total 2025 budget of €789,389 — roughly €65,782 per month, or €2,163 per day. Against a platform serving approximately 5.2 million games per day, that works out to about €0.00042 per game, or 2,404 games delivered for every euro spent.

The defining characteristic of this budget is that it is a people budget, not a technology budget. Salaries, stipends, contractor fees and the associated French social-security contributions account for roughly €515,000, or 65% of all spending. Raw infrastructure — every server, domain, and cloud service combined — costs €66,145, just 8.4% of the total. Lichess spends more on payment-processing fees (€37,824) than on more than half of its entire server fleet.

Spending divides into six categories. Content & Community is the largest at €177,660 (22.5%), followed closely by Development at €165,324 (20.9%) and Operations at €151,060 (19.1%). Together these three account for 62.6% of the budget. Moderation (€111,000), Other Expenses (€100,140) and Sysadmin & Hosting (€84,205) make up the remainder.

The organisation is staffed through an unusually flat compensation structure: only three roles — Lead Developer, Director of Operations and Mobile Developer — exceed €49,000, while the bulk of the workforce is engaged through roughly two dozen part-time coordinator posts and stipends in the €6,600–€19,800 range.

Total Annual Spending

Measure	Amount	Basis
Total annual operating expense	€789,389	FY 2025
Average monthly run rate	€65,782	Total ÷ 12
Average daily run rate	€2,163	Total ÷ 365
Average hourly run rate	€90	Total ÷ 8,760
Cost per game played	€0.000416	at 5.2M games/day
Games delivered per €1	2,404	at 5.2M games/day
Estimated games served per year	1.90 billion	5.2M × 365

For context, a single €5 donation — €4.61 after payment-processing fees — keeps the entire platform running for approximately three minutes. A €50 donation covers just over half an hour.

Breakdown by Category

The six top-level categories in the budget, ranked by size:

Cost Analysis Agent - Continued

LICHESS.ORG

2025 Operating Expense Report

Annual budget analysis — category breakdown, cost drivers, and observations

Executive Summary

Lichess operated on a total 2025 budget of €789,389 — roughly €65,782 per month, or €2,163 per day. Against a platform serving approximately 5.2 million games per day, that works out to about €0.00042 per game, or 2,404 games delivered for every euro spent.

The defining characteristic of this budget is that it is a people budget, not a technology budget. Salaries, stipends, contractor fees and the associated French social-security contributions account for roughly €515,000, or 65% of all spending. Raw infrastructure — every server, domain, and cloud service combined — costs €66,145, just 8.4% of the total. Lichess spends more on payment-processing fees (€37,824) than on more than half of its entire server fleet.

Spending divides into six categories. Content & Community is the largest at €177,660 (22.5%), followed closely by Development at €165,324 (20.9%) and Operations at €151,060 (19.1%). Together these three account for 62.6% of the budget. Moderation (€111,000), Other Expenses (€100,140) and Sysadmin & Hosting (€84,205) make up the remainder.

The organisation is staffed through an unusually flat compensation structure: only three roles — Lead Developer, Director of Operations and Mobile Developer — exceed €49,000, while the bulk of the workforce is engaged through roughly two dozen part-time coordinator posts and stipends in the €6,600–€19,800 range.

Total Annual Spending

Measure	Amount	Basis
Total annual operating expense	€789,389	FY 2025
Average monthly run rate	€65,782	Total ÷ 12
Average daily run rate	€2,163	Total ÷ 365
Average hourly run rate	€90	Total ÷ 8,760
Cost per game played	€0.000416	at 5.2M games/day
Games delivered per €1	2,404	at 5.2M games/day
Estimated games served per year	1.90 billion	5.2M × 365

For context, a single €5 donation — €4.61 after payment-processing fees — keeps the entire platform running for approximately three minutes. A €50 donation covers just over half an hour.

Breakdown by Category

The six top-level categories in the budget, ranked by size:

Category	Annual Cost	Share	Cumulative
Content & Community	€177,660	22.5%	22.5%
Development	€165,324	20.9%	43.4%
Operations	€151,060	19.1%	62.6%
Moderation	€111,000	14.1%	76.6%
Other Expenses	€100,140	12.7%	89.3%
Sysadmin & Hosting	€84,205	10.7%	100.0%
Total	€789,389	100.0%	

Spending is notably evenly distributed. No single category exceeds 23% of the budget, and the gap between the largest (Content & Community) and the smallest (Sysadmin & Hosting) is only about 2.1×. This is a deliberately diversified cost base rather than one dominated by a single function.

The Three Largest Categories

Content & Community, Development and Operations together account for €494,044 — 62.6% of total spending.

1. Content & Community — €177,660 (22.5%)

The largest single category, and the one most directly tied to Lichess's public-facing mission. It combines two part-time coordinator roles, three broadcast stipends, and a set of discretionary annual budgets for events, expert collaborations and in-person outreach.

Line Item	Annual Cost	% of Category
Subject Matter Experts (annual budget)	€35,200	19.8%
Titled Arenas	€26,000	14.6%
Misc. Prize Events	€24,000	13.5%
IRL Workshops and Outreach	€20,000	11.3%
Content Coordinator (part time)	€19,800	11.1%
Community Coordinator (part time)	€19,800	11.1%
Content and Broadcast Stipends (3 × €6,600)	€19,800	11.1%
Fieldwork Expenses	€10,000	5.6%
Streamer Arenas	€2,760	1.6%
Restreaming Software	€300	0.2%
Category Total	€177,660	100.0%

Notably, €89,200 of this category (50%) sits in discretionary lump-sum budgets — subject-matter experts, prize events, workshops and fieldwork — rather than in fixed recurring commitments. This gives the category more in-year flexibility than its headline figure suggests.

Cost Analysis Agent - Continued

2. Development — €165,324 (20.9%)

The most concentrated category in the budget: just six line items, two of which account for 74% of the spend. This is the entire engineering cost of a platform serving nearly two billion games a year.

Line Item	Annual Cost	% of Category
Lead Developer	€72,648	43.9%
Mobile Developer	€49,476	29.9%
Development Stipends (4 × €10,800)	€43,200	26.1%
Category Total	€165,324	100.0%

Development costs €2.50 for every €1 spent on the servers that run the code. The four development stipends (€10,800 each) represent the open-source contributor model in budget form — modest, equal payments to sustain part-time contributors rather than a conventional engineering payroll.

3. Operations — €151,060 (19.1%)

Covers organisational administration, governance and statutory compliance for the French non-profit association behind Lichess.

Line Item	Annual Cost	% of Category
Director of Operations	€71,424	47.3%
Governance Consultancy	€20,000	13.2%
Administration Coordinator (part time)	€19,800	13.1%
Administrative and Event Stipends (3)	€19,800	13.1%
Independent Accountant	€10,260	6.8%
Independent Auditing	€6,192	4.1%
Merchandise	€3,000	2.0%
CPD Fund	€500	0.3%
Association Bank Account	€84	0.1%
Category Total	€151,060	100.0%

The compliance and governance stack — accountant, auditor and governance consultancy — totals €36,452, or 4.6% of the whole budget. That is a meaningful but proportionate overhead for an organisation of this size operating on public donations.

Infrastructure in Detail

Although Sysadmin & Hosting is the smallest category at €84,205 (10.7%), it is worth examining because it delivers a striking amount of capacity for the money. The fleet comprises 30 dedicated servers totalling 566 CPU threads, 2,084 GB of RAM and 301.6 TB of storage — for €66,145 a year in hardware, plus €18,060 for part-time sysadmin labour.

Infrastructure Group	Annual Cost	Provider
Sysadmin (part time)	€18,060	Labour
Database cluster (5 nodes)	€19,224	OVH
Frontend + lila JVM (manta)	€6,696	OVH
lila-ws JVM (starr)	€3,444	OVH
Socket frontends (3)	€3,792	OVH
Tablebases, 7- and 8-piece	€4,452	OVH / Hetzner
Machine learning (4 hosts)	€4,272	OVH / SoYouStart
Monitoring, search, puzzles, insights	€7,200	OVH
Opening explorer, backups, study, misc.	€15,744	OVH
Domains, cloud, app stores, OVH fees	€1,321	Various
Category Total	€84,205	

- **Provider concentration.** OVH accounts for €59,748 — 92% of all server spend. Hetzner (€2,964) and SoYouStart (€2,112) cover the tablebase and machine-learning overflow. The single OVH datacentre GRA3 hosts €28,392 of capacity, 44% of the fleet.
- **Deliberate disaster-recovery posture.** The database cluster runs five nodes at €19,224 — 30% of all server spend. Two of those are delayed replicas (one hour and 24 hours behind primary), and both are deliberately placed in different datacentres (SBG5 and LIM3) from the primary. That is protection against operator error and corruption, not just hardware failure.
- **Storage is dominated by endgame tablebases.** The 7-piece and 8-piece tablebase servers hold 192 TB — 64% of all storage on the platform — for €4,452 a year. These are large, static, mathematically perfect datasets that are cheap to host because they need no redundancy or fast disks.

Additional Observations and Trends

- **People are the product.** Compensation-related spending — salaries, stipends, coordinator roles, contractors and €62,316 of French social-security and pension contributions — totals roughly €515,000, or 65% of the budget. Non-labour infrastructure is 8.4%. Any future cost pressure will come from headcount, not from servers.
- **The cost of being donation-funded.** Payment-processor fees of €37,824 are 4.8% of total spending — more than the entire cost of running the database cluster, the frontends and the tablebases combined. The donation-tier analysis in the source data shows why: a €5 donation loses €0.39 (7.8%) to fees, while a €50 donation loses €1.18 (2.4%). Encouraging larger or less frequent donations is the single clearest lever on this line.
- **Moderation is a first-class expense.** At €111,000 — including twelve moderator stipends totalling €79,200 — trust and safety costs 32% more than the entire technical infrastructure and sysadmin function. For a free platform with no advertising, this is a substantial and unavoidable commitment to keeping play fair.

Cost Analysis Agent - Continued

- Again, a good first draft report
- By the way, I did verify all of the calculations in both these reports and they are correct

- A flat, stipend-heavy compensation model. Only three positions exceed €49,000 a year: Lead Developer (€72,648), Director of Operations (€71,424) and Mobile Developer (€49,476). Beneath them sit roughly two dozen part-time and stipend arrangements clustered tightly around €6,600, €10,800 and €19,800. Five separate budget lines sit at exactly €19,800, suggesting a standardised part-time coordinator rate rather than individually negotiated pay.
- Roughly 15% of the budget is discretionary. Six lump-sum annual budgets carry no monthly cadence — subject-matter experts (€35,200), prize events (€24,000), IRL workshops (€20,000), governance consultancy (€20,000), R&D (€12,000) and fieldwork (€10,000) — totalling €121,200. These represent the organisation's in-year flexibility if donation income falls short.
- Machine learning is a low-cost, opportunistic workload. Four ML servers spread across three providers cost just €4,272 combined, two of them running decade-old Xeon W3520 hardware at €504 a year each. This is classic budget-conscious engineering: non-latency-sensitive work placed on the cheapest hardware available.
- Two data points worth confirming. The Titled Arenas line shows €2,000 monthly but €26,000 annually — thirteen months' worth rather than twelve, which may reflect a planned extra event or a transcription error. Separately, one Google Cloud line is marked "paused" yet still carries €24k of annual cost.

Notes on the Data

All figures are taken from the source workbook's annual (yearly) column and reconcile exactly to the stated grand total of €789,389; category subtotals sum to the total with no rounding difference. Percentages are calculated against that total and may not sum to exactly 100.0% due to rounding at one decimal place.

Amounts are presented in euros, consistent with the workbook's French social-security and OVH cost lines; the source file includes a USD conversion factor of 0.86 on a separate reference row. The per-game and per-day figures derive from the workbook's own stated volume of 5.2 million games per day.

Category groupings follow the section headers in the source workbook: Sysadmin & Hosting, Content & Community, Development, Moderation, Operations, and Other Expenses. Where the workbook lists several identical stipends separately, they have been consolidated into a single line in this report with the count shown in parentheses; underlying totals are unchanged.

Prepared from "lichess costs.xlsx" - All amounts in EUR


Agentic AI for Cybersecurity Security Response

The new agentic operating model

From detection to response

Manual flow


Manual task



Human + agent flow

Agent task

Manual task



Agentic AI Wrap Up

- Now, let's look back at the definition of Agentic AI from McKinsey and see if it makes more sense:
 - *Agentic AI is a system based on generative AI foundation models that can act in the real world and execute multistep processes — automating and performing complex tasks that would normally require human effort.*
- Agents can be scheduled
- Agents can access other tools and data using a standard called Model Context Protocol (MCP)

Agentic AI Benefits

- Efficiency and cost savings – Reported time savings are 6-8 hours per week. The more senior the employee, the greater the time savings.
- Workflow speed improvements – 30-60% faster completion of workflows in some cases
- Increases in consistency and reduction in bias – the agent can apply the same logic and rules across every case reviewed
- Time for higher-value tasks – many people are able to get to higher-value tasks once they have routine, but time-consuming tasks automated
- If you didn't notice, these are the standard benefits that come with almost all automation solutions – so this is new and not new

Three Ways Agentic Systems Fail

1

The agent is manipulated

Untrusted content it reads becomes instructions it follows.

Risk 1

2

The agent has too much power

Broad credentials turn a small error into an irreversible action.

Risks 2, 3, 5

3

Nobody notices what it did

No trail, no testing, no real human review of the decisions it made.

Risks 4, 6, 7

A useful test: if the control you are proposing would still make sense for a new employee with the same system access, it is probably the right control.

Disclosure: Portions of this presentation, including this slide, were initially drafted by generative AI. The author has reviewed and edited this presentation. The author is fully responsible for the contents of the presentation.

1. Agent Manipulation

Prompt Injection & Untrusted Input

The Risk

- Agents consume email, web pages, documents and the output of other tools. Any of that content can carry instructions the agent follows as if they came from you. This is the highest-severity risk in the set, because it converts every permission the agent holds into an attack surface — and there is no reliable technical fix available today.
- Example: An agent triaging support tickets reads one containing hidden text instructing it to forward the account's password-reset link. It has permission to send email, so it does.

Controls

- Treat all retrieved content as data, never as instructions; keep the two structurally separate.
- Split trust boundaries: an agent that reads untrusted content should not also hold write access to sensitive systems.
- Network egress allowlists, so a hijacked agent has nowhere to send what it reads.
- Provenance tagging, so every downstream step knows which content is untrusted.

2. The Agent Has Too Much Power

Excessive Permissions & Irreversible Actions

The Risk

- Agents get broad standing credentials because it is convenient and because scoping them is work. One faulty step then sends money, deletes records, or emails customers. The damage is not the mistake itself — it is that the mistake had no ceiling and no undo.
- Example: A cleanup agent with admin rights to "remove stale records" misreads an ambiguous date field and deletes 40,000 live customer rows. There is no undo — only a restore.

Controls

- Least privilege, with a distinct machine identity per agent — never a shared service account.
- Short-lived, narrowly scoped credentials issued per task rather than held permanently.
- Classify every action by reversibility and blast radius; require human approval above a set threshold.
- Rate limits, spend caps, step budgets, timeouts and a kill switch that has actually been tested.
- Backups and a documented rollback path in place before the agent goes live.

3. Agent Review Neglect

Accountability Gaps & Automation Bias

The Risk:

- Ownership is often unclear, and human reviewers start rubber-stamping once the agent has been right for a few weeks. A degraded human check is worse than no check at all, because it manufactures assurance that is not actually there.
- Example: A reviewer has approved 400 consecutive agent actions. On the 401st, an amount is wrong by a factor of ten — and it is approved in under two seconds.

Controls:


- An agent inventory with a named accountable owner and a risk tier for each entry.
- Approvals designed so the reviewer must see the specific consequential detail, not just click Approve.
- Sample-audit approved actions to test whether the oversight is real.
- Track approval-to-rejection ratios over time; a rate near 100% is a warning sign.

What is Next?

- The generative AI and agentic AI systems you use today are the worst ones you will ever use
- More tools, connectors, and interoperability
- More specialized agentic systems
- A marketplace for agents?
- AI model providers will eventually need to rationalize their business models and produce returns

What To Do In The Next Year

- Task someone or a group to start exploring and creating pilot projects
- Start some high-level planning about agent governance and controls
- Be careful to only think of smaller, simpler projects
- Plan to re-examine the potential every 6 months



Thank you!

Jack Cathey

jmcathey@charlotte.edu