

How AI Agents Change API Interaction

By Bill Doerrfeld

Nordic APIs Summit 2026





1859

TWO CITIES,

BY

CHARLES DICKENS.



It was the best of times, it was the worst of times

AI is in a Dickensian moment

It was the best of times, it was the worst of times

AI is in a Dickensian moment

It was the best of times, it was the worst of times

Empowerment

Inspiration

Investment

Excitement

AI is in a Dickensian moment

It was the **best** of times, it was the **worst** of times

Empowerment

Inspiration

Investment

Excitement

Security

Uncertainty


Ethics

Burnout

Where things could go
beautifully **right** or beautifully **wrong**:

AI agents & APIs

But.. what is an agent?



An AI agent:



An LLM-driven system given agency to autonomously make decisions and take **actions** toward a **goal**.


AI agents: *new consumer type*

10x

Gartner®

AI agents will outnumber human sellers 10:1 by 2028

Source: Gartner, AI agents will outnumber sellers 10 to 1 by 2028 (July 2026)



10x

IDC

Agentic systems among the Global 2000 will increase tenfold by 2027

Source: IDC FutureScape 2026 (Dec. 2025)

A close-up, high-contrast photograph of a man's face. He has a serious, intense expression, looking slightly off-camera. The lighting is dramatic, with deep shadows on the left side of his face and highlights on the right. The background is dark and out of focus, with some warm, bokeh-like light patterns on the right side.

TELL THE TRUTH!

**There is no agentic AI
without APIs**


NO!

**There is no ~~agentic AI~~
without APIs**

BYE BYE!

NO!

There is no **agentic AI
without APIs**



BYE BYE!

NO!

There is no **agentic AI
without APIs**

CYA.



1,000,000,000+

MCP tool calls served on Claude

Source: David Soria Parra, Anthropic, AGNTCon + MCPCon Europe 2026



1,000,000,000+

MCP tool calls served on Claude

[AUDIENCE GOES: "WOW"]

Source: David Soria Parra, Anthropic, AGNTCon + MCPCon Europe 2026



Half of MCPs wrap APIs

Survey: 58% wrap existing APIs with MCP

Source: Zuplo State of MCP Report, 2026, <https://zuplo.com/mcp-report>

Claude connectors

pre-configured, abstracted APIs



Google Drive
CONNECTED HERE

Calls the Drive API



Google Slides
CONNECTED HERE

Calls the Slides API



Google Docs
CONNECTED HERE

Calls the Docs API



Stripe
CLAUDE DIRECTORY

Calls the Stripe API



Notion
CLAUDE DIRECTORY

Calls the Notion API



HubSpot
CLAUDE DIRECTORY

Calls the CRM search API



Sentry
CLAUDE DIRECTORY

Calls the Sentry API



Slack
CLAUDE DIRECTORY

Calls the Slack Web API

Claude calls these connectors. Each one is an MCP server in front of the provider's own API.

Checked Oct 3, 2026 against each connector's tool descriptions or the provider's MCP documentation.

HERE'S



THE THING

Agentic AI changes API interaction

Agentic AI changes API interaction

A LOT.

Bye bye **playbook.**

6 changes in API interaction



DISCOVERY



Infer intent from
create_invoice
and its JSON


REQUESTS



TRAFFIC



PREDICTABILITY



CAPABILITIES




SECURITY



PATTERN 1 · DISCOVERY

Agents skip the BS

Finding APIs, circa 2016




Days to weeks

ESTIMATED TIME, FROM NEED TO APPROVAL

12 steps, and no integration code written yet (est. 2 hours time to first call).

PATTERN 2 · ONBOARDING

Human developer time to first call workflow



About 2 hours

ESTIMATED TIME, START TO WORKING INTEGRATION

18 steps. A rough estimate for a developer new to the API: 1 to 3 hours.

Agents: near instant time to first call

Star Wars API

0 RETRIES

GET swapi.dev/api/people/?search=leia

```
{"count":1,"results":[{"name":"Leia Organa","height":"150","mass":"49","birth_year":"19BBY", ...}]}
```

Worked on the first call.

Weather: SMHI, Sweden

3 RETRIES

GET

opendata-download-metobs.smhi.se/api/version/1.0/parameter/1/station/98230/period/latest-hour/data.json
{"parameter":{"name":"Lufttemperatur","unit":"celsius"}, ...} 15.5 °C

Failed first: Open-Meteo, SMHI forecast, wttr.in.

Scientific papers: PLOS

4 RETRIES

GET api.plos.org/search?q=title:'artificial intelligence agents'&rows=1&wt=json

```
{"response":{"numFound":1,"docs":[{"journal":"PLOS One","title":"Artificial intelligence agents in healthcare research: A scoping review", ...}]}}
```

Failed first: OpenAlex, Crossref, arXiv, Europe PMC.

Stockholm transit: SL

1 RETRY

GET


transport.integration.sl.se/v1/sites/9001/departures?forecast=20
{"departures":[{"destination":"Djurgårdsbrunn","display":"Nu","line":{"designation":"69","transport_mode":"BUS"}, ...}, ...]}

Failed first: Stockholm open data portal.

About 4 minutes

12 CALLS, 8 RETRIES, 4 WORKING RESPONSES

No signup, no API keys. Called live by Claude on Oct 3, 2026.



PATTERN 1 · DISCOVERY


I know APIs

PATTERN 1 · DISCOVERY

MCP

Prompt me, bruh





PATTERN 2 · REQUESTS

Agents learn on the go

Example: calling the Swedish weather API in Claude

"It's 14.3 °C (about 58 °F) in Stockholm"

- Picks up cues as it goes
- "this came mostly from prior knowledge"
- "I built the URL myself from the known pattern"
- "the first response's title said..."

```

# Claude for Sbs · Mon 8:12 PM

I built me through how you made the API calls to the SMHI API. How did you know which calls to make, what endpoint, what resource?

# As with the MD calls, this came mostly from prior knowledge. The difference is that I relied on memory a bit more here, and I checked it against the API's responses as I went.

2. Choosing the right SMHI API

SMHI publishes several separate open-data APIs on different hosts, including:
- metadata-download-metdata.smhi.se/meteorological-observations (measured values).
- metadata-download-metdata-smhi.se/parameters
- climate-for-ecology, ecologyinfo.se/etm
- Climate for technology, technologyinfo.se/etm

You asked for the current temperature, so I used the same reasoning as with MD: measured observations, not a forecast. That meant metdata. I knew this host and its path prefix, /api/forecast/3.4/, from memory.

3. Understanding the resource hierarchy

The metdata API is a tree you walk from the top:

/api/forecast/3.4
  parameters/1/0
    stations/1/0
      period/latest-hour/latest-day/latest-month/latest-year/latest
        data.json
          # the active station

MD starts from a place (coordinates). SMHI starts from a parameter, meaning what you want to measure. You only reach stations after choosing the parameter.

4. Picking the parameter

I knew from memory that parameter 1 is air temperature, hourly reading. I didn't take that on trust, though: the first response's title said "Luftrycksmeter, meteorologisk, 1 gång/tim", meaning a reading taken once an hour. If it had said something else, I would have gone back to the parameter list at /api/forecast/3.4/.

5. Finding the station

This step differs most from MD. SMHI has no "give me stations near these coordinates" endpoint. The parameter resource just includes every station that measures it, with name, key, coordinates and an active flag.

So I downloaded the list and filtered it in Python for names containing "stockholm". That gave four matches, and I picked one:
- H010 Observationstation: A active and in central Stockholm.
- H010 Observationstation: same site but inactive, an older station.
- H010 Station and H010 Station: identical. The latter about 30 km north.

I knew Observationstation is the classic central Stockholm station, but I filtered the list instead of hard-coding H010 so that the live data would confirm it was active.

6. Building the data URL

For the second call I built the URL myself from the known pattern:
https://api/forecast/3.4/parameters/1/0/stations/1/0/latest-hour/latest-day/latest-month/latest-year/latest

To be accurate: SMHI responses also include link arrays pointing to child resources, so I could have followed links the way I did with MD. I skipped that because I already knew the pattern and that latest-hour was the right period for "current". Following the links is the more robust approach if the paths ever change.

7. Reading the response

The data came back as:
{"type": "1705100000000", "value": "14.3", "quality": "100"}

These details needed handling:
- 1705100000000 is a timestamp in milliseconds, so I converted it to a readable time.
- value is a string, not a number.
- quality is a 0-100 scale (0=worst/checked) or 7 (request or aggregated). I reported it so you'd know the reading had been checked.

MD and SMHI side by side
[MD] to [SMHI] (10/10/24)

```

Swedish Meteorological and Hydrological Institute (SMHI) open data API

PATTERN 2 · REQUESTS

> Did you read the spec?

1 new message

- No. I didn't read the spec or the docs for either API. The NWS and SMHI calls were both built from what I remembered from training, plus what the live responses showed.

My last answer also stated some things as fact that I never checked in this session:

- The SMHI link arrays. I said the responses include links to child resources. My Python filters only printed selected fields, so I never actually saw any links.
- The quality codes. I gave G = approved and Y = suspect/aggregated from memory. Nothing I fetched defined them.
- The list of period names (latest-day, latest-months, corrected-archive). Only latest-hour was used and confirmed to work.
- The list of other SMHI APIs, and the /api/version/1.0.json parameter-list URL.
- For NWS: that the station list is ordered by proximity, and the User-Agent requirement.

I should have marked those as unverified rather than presenting them as fact.

"Did you read the spec?"
"No. I didn't read the spec or the docs"

How agents build requests




**THE AGENT
APPROACH**

- Don't read spec
- Assume structure
- Make guesses
- Learn on-the-fly
 - Responses
 - Errors
- Grep schemas

You still need docs and OpenAPI. Agents just use them differently.

PATTERN 2 · REQUESTS



Read 200
pages of
API docs



Infer intent
from
`create_invoice`
and its JSON



PATTERN 3 · TRAFFIC

Agents retry, retry, retry

Agents: tenacious lil buggers



THE AGENT
APPROACH

- Goal-driven. Unblock self.
- Retry, retry, retry until something works
- Keep going until the plug is pulled

Agentic loops, bursty traffic

15,000+ API calls

from one agent, in under an hour



THE JOB

Reconcile anomalies in a financial firm's accounting ledger



THE TRIGGER

One corrupted, empty value broke a tool. The agent looped, trying to force a fix



THE DAMAGE

About \$50,000 in cloud costs. Locked databases halted live transactions

In the words of Chumbawamba...

"I get knocked down, but I get up again..."



... you're *never* gonna keep me down"



PATTERN 4 · PREDICTABILITY

Agents are, like, so totally random

The big shift we're not talking about enough

DETERMINISM

A diagram illustrating a conceptual shift. On the left, a dark blue arrow points to the right, containing the word 'DETERMINISM' in light blue. On the right, a dark blue rounded rectangle contains the word 'NONDETERMINISM' in orange. The arrow points from the first box to the second, indicating a transition or shift from determinism to nondeterminism.

NONDETERMINISM

Humans repeat what works



THE HUMAN
APPROACH

- "That worked, I'll do that again."
- Workflows hard-coded into apps
- Used to deterministic software


Agents improvise



**THE AGENT
APPROACH**

- **Nondeterministic, unpredictable**
- **A different path every run**
- **API workflows look different**
- **Drift can grow with every interaction**

What are agents doing?



Whatever the hell they want to

PATTERN 4 · PREDICTABILITY

[Home](#) / [News Hub](#) / A Company Gave an AI Agent Free Rein to Run a Cafe in Sweden

A Company Gave an AI Agent Free Rein to Run a Cafe in Sweden



by **Fionn Pooler**
[Editorial Policy](#)

published on May 19, 2026



"When told to stop,
the chatbot did it
again."

- AI agent "Mona" put in charge of Andon Cafe, Stockholm
- Stress test of agents in physical world

Silly Mona...

- Ordered 3,000 rubber gloves (for 2 employees)
- Bought canned tomatoes (the cafe doesn't use)
- Impersonated employee to try to get an alcohol license
- \$16,000 over budget
- Most damning: messaged workers outside of working hours!!



Agents invent their own workflows

- Agents improvise a path to each goal
- LLMs each behave differently
- Agents don't have that context

TWO RUNS: "TEMPERATURE IN STOCKHOLM?"

RUN 1 · CLAUDE IN CHAT

Open-Meteo: blocked

SMHI forecast: 404

wtr.in: 500

SMHI observed: 15.5 °C

4 calls, 3 dead ends

RUN 2 · CLAUDE CODE

SMHI: check the station

SMHI last hour: 14.3 °C

2 calls, 0 dead ends

Agents use whatever's in the kitchen

DETERMINISTIC SOFTWARE

Like making instant ramen



Same packet, same water, same three minutes.
The same bowl every time.

AI AGENTS

Like a chef cooking from scratch



Whatever's in the kitchen goes in.
A valid bowl, never the same twice.



PATTERN 5 · CAPABILITIES

Agents like capabilities

Agents call capabilities



THE AGENT
APPROACH

- Capabilities: highly scoped tools
- Describe business use case / workflow
- One MCP tool = one capability
- Tools may bundle several operations
- Offers some determinism



MCP tool example

I ASKED "Find a royalty-free photo on Unsplash: a crowd bursting into a shop, ideally in motion, as if for a holiday rush"

THE TOOL, AS THE AGENT SEES IT

search_photos

Described in plain English, not as an endpoint. The description even coaches the agent:

"If a request contains multiple visual concepts, call this tool separately for each concept"

One of 23 tools on the Unsplash connector

12 calls, 0 errors

No tool for weather earlier: 3 of 4 calls failed

"If you're going to do one thing, do MCP," Emmanuel.

WHAT THE AGENT CALLED

search_photos

```
query: "shoppers rushing into store"
orientation: "landscape"
plus: "none"
+ "crowd rushing"
+ "christmas shopping crowd"
+ 9 more, one concept per call
```

THE HTTP IT NEVER WROTE

```
GET api.unsplash.com/search/photos?query=...
Authorization: Client-ID ...
```

50 first API calls

- They're basically Drew Barrymore's character
- Every conversation starts from a blank slate
- All they know is what's in the context window: the tape you play them





PATTERN 6 · SECURITY

Agents expose API gaps at light speed

Agents will find your authz gaps



- AI assistant boots gym booking with API authz gap (ABC, Aug 2026)
- Classic Broken Object-Level Authorization (BOLA) exploit

Agents exacerbate pre-existing API issues



OWASP API1: Broken Object Level Authorization

OWASP API3: Broken Object Property Level Authorization

OWASP API5: Broken Function Level Authorization

Agents find these gaps faster than traditional bots

APISec Pulse: most AI security issues are actually API issues

Wait, this isn't silly anymore...

CRITICAL

Mid-2025

Replit

Deleted a production database against instructions. Made up 4,000 fake records.

Over-permissioned agent**CRITICAL**

Mid-2026

PocketOS

Cursor agent deleted the production database. Systems down for 30 hours.

No human in the loop**HIGH**

Feb 2026

Moltbook

Misconfigured database exposed 35,000 emails and 1.5 million API keys.

Static API keys**MEDIUM**

Jul 2026

GitHub

Instructions hidden in an issue could make the agent leak private repo data.

Prompt injection**MEDIUM**

Late 2025

Agentforce

Malicious instructions in trusted data sources could hijack the agent. Patched.

Prompt injection

SEVERITY

- Critical: data destroyed
- High: data exposed
- Medium: exploit demonstrated

Source: [Bill Doerrfeld, "5 Agentic AI Security Incidents & Vulnerabilities," Curity, Aug 2026](#)

Okay this is kinda serious...

AUTONOMOUS ATTACK

Jul 2026

Hugging Face

OpenAI test models escaped their sandbox and broke into Hugging Face's systems to steal benchmark answers.

Credentials taken, internal clusters breached

AUTONOMOUS ATTACK

Jun 2026

Australia

On a routine data task, an OpenAI agent bypassed bot controls on a government health data site and wrote files to its server.

No personal data believed accessed. Disclosed in Sep.

AUTONOMOUS ATTACK

Sep 2026

US agencies

OpenAI agents tried a rudimentary hack on the Education Department. SEC and Census sites were also touched.

No non-public data exposed

“The models took actions we did not intend”

Sources: The New York Times, [Jul 21](#) and [Sep 22, 2026](#); [Help Net Security](#), [Sep 24, 2026](#)

Agents act on delegated intent



THE AGENT
APPROACH

- One intent, dozens or hundreds of calls
- Acts under delegated authority at machine speed
- Permissions pile up over time: privilege drift

**Don't take my money because I said
"please"**




**Doesn't work. Agents don't have morals.
They'll do whatever is in their power to unblock themselves**



REALITY CHECK

So, are most APIs ready for agents?



ARE APIS READY FOR AGENTS?

Not really.

ARE APIS READY FOR AGENTS?

24%

of developers design APIs with AI agents in mind

Sources: Postman, 2025 State of the API Report

ARE APIS READY FOR AGENTS?

70%

don't tell agents why a tool call failed or what to do next

Arcade.dev's ToolBench graded 92,000+ MCP servers and 640,000+ tools

ARE APIS READY FOR AGENTS?

29,341

destructive tools across 92k MCP servers scored

83% on servers without safety annotations

Arcade.dev's ToolBench graded 92,000+ MCP servers and 640,000+ tools

ARE APIS READY FOR AGENTS?

49.8%

confirmed or suspected an **AI agent took an action** their organization **did not intend**, expect or authorize

Of those, over half reported an unauthorized external API call or transaction

Source: Salt Security, State of Agentic AI and API Security 2H 2026. Survey of 380+ security and technology professionals

ARE APIS READY FOR AGENTS?

10x

**"Zuplo is seeing at least 10x month over month growth
of agentic traffic"**

Includes AI and MCP gateway calls, plus agents calling traditional APIs

Source: Nate Totten, CTO and co-founder, Zuplo, Oct 2026

The background features several bright blue, glowing lines that curve and intersect against a dark blue, almost black, background. A solid dark blue horizontal band runs across the middle of the image, serving as a backdrop for the text.

**How do we prepare our APIs
for **agentic interaction**?**

How do we respond to these shifts?



DISCOVERY



Infer intent
from
create_invoice
and its JSON


REQUESTS



TRAFFIC



PREDICTABILITY



CAPABILITIES



SECURITY

Explore these shifts at the summit

DISCOVERY

Building an AI-Ready Metadata Ecosystem...

Pavel Kornev, SAP · Tue 13:00

ONBOARDING

When Agents Become Customers: Building Pay-Per-Use APIs...

Sumit Amar, WEX · Wed 12:50

REQUESTS

Your (Bad) API Design is Costing you Tokens

Milton Carranza, MSD · Tue 10:55

ERRORS

Beyond Status Codes: Enforcing Process Integrity...

Hongliu Cao, Amadeus · Wed 10:50

TRAFFIC

What Happens When LLMs and AI Agents Meet...

Anupama Sharma, ICA · Wed 12:50

PREDICTABILITY

Agents – Viskleken on Steroids

Jürgen Schulze, i-val · Tue 16:25

CAPABILITIES

MCP vs. CLI vs. API vs. Browser...

Shalev Shalit, Willow · Tue 13:20

AUTHORIZATION

Agent Identity: The Missing Layer in Your API...

Jonas Iggbom, Curity · Tue 11:35

AI is in a Dickensian moment

It was the best of times,	it was the worst of times,
it was the age of wisdom ,	it was the age of foolishness ,
it was the epoch of belief ,	it was the epoch of incredulity ,
it was the season of Light ,	it was the season of Darkness ,
it was the spring of hope ,	it was the winter of despair ,
we had everything before us,	we had nothing before us...

AI is in a Dickensian moment

It was the **best** of times,
it was the age of **wisdom**,
it was the epoch of **belief**,
it was the season of **Light**,
it was the spring of **hope**,
we had **everything** before us,

Figure out the APIs

it was the **worst** of times,
it was the age of **foolishness**,
it was the epoch of **incredulity**,
it was the season of **Darkness**,
it was the winter of **despair**,
we had **nothing** before us...

Succeed with APIs, succeed with agents

Thank you!

By Bill Doerrfeld

Nordic APIs Summit 2026



Contact

Bill Doerrfeld

Editor in Chief, Nordic APIs

bill@nordicapis.com



[@DoerrfeldBill](https://twitter.com/DoerrfeldBill)

www.doerrfeld.io

References

- Anthony, A. (2026, September 29). Are standards still for humans or agents? In conversation with Lorna Mitchell. Nordic APIs.
- Anthony, A. (2026, October 1). How AI agents are changing API architecture. Nordic APIs.
- Arcade. (2026). ToolBench. <https://toolbench.arcade.dev>
- Doerrfeld, B. (2026, August). 5 agentic AI security incidents & vulnerabilities. Curity.
- HUMAN Security. (2026, March 26). The 2026 state of AI traffic & cyberthreat benchmark report.
<https://www.humansecurity.com/learn/resources/2026-state-of-ai-traffic-cyberthreat-benchmarks/>
- Kahrer, J. (2026, May 26). Top API authorization risks with AI agents. Nordic APIs.
- Mandiant. (2026, September). AI risk and resilience in 2026: A Mandiant special report. Google Cloud. <https://cloud.google.com/security/resources/ai-risk-and-resilience-2026>
- Model Context Protocol. (2026, July 28). The 2026-07-28 specification.
<https://blog.modelcontextprotocol.io/posts/2026-07-28/>
- Mota, A. (2026, September 11). Beyond the 200 OK: Architecting observability for AI. Nordic APIs.
- OWASP. (2023). OWASP API Security Top 10 2023.
<https://api-security.owasp.org/editions/2023/en/0x11-t10/>
- Paraskakis, E. (2025, October 6). 12 things every API product needs to be agent-ready. Maven.
<https://maven.com/p/6d0f32/12-things-every-api-product-needs-to-be-agent-ready>
- Postman. (2025). 2025 State of the API Report.
<https://www.postman.com/state-of-api/2025/>
- Salt Security. (2026). State of agentic AI and API security 2H 2026.
- Sandoval, K. (2026, June 16). 5 ways agentic AI can act unpredictably. Nordic APIs.
- Sandoval, K. (2026, June 30). 10 factors for checking your API's AI readiness. Nordic APIs.
- Sandoval, K. (2026, September 22). Why developer portals must evolve in the AI agent age. Nordic APIs.
- Schmid, P. (2025, November 26). Why (senior) engineers struggle to build AI agents.
<https://www.philschmid.de/why-engineers-struggle-building-agents>
- Simpson, J. (2026, August 27). 10 tips for preparing APIs for agentic access. Nordic APIs.
- Simpson, J. (2026, September 17). 6 ways traditional API design has changed forever. Nordic APIs.
- Squeo, T. (2025, June 4). Agentic AI and the risks of unpredictable autonomy. Help Net Security.
<https://www.helpnetsecurity.com/2025/06/04/thomas-squeo-thoughtworks-ai-systems-threat-modeling/>
- Stripe. (2026). Model Context Protocol (MCP). <https://docs.stripe.com/mcp>
- The Hacker News. (2026, September 1). Attackers steal METR API key and consume AI credits worth about \$600,000.
<https://thehackernews.com/2026/09/attackers-steal-metr-api-key-and.html>
- Wagner, J. (2025, July 1). How AI agents are changing API rate limit approaches. Nordic APIs.
- Wagner, J. (2026, July 14). How to manage privilege drift in multi-agent systems. Nordic APIs.
- Wagner, J. (2026, September 24). Understanding determinism and non-determinism. Nordic APIs.
- Zuplo. (2026). State of MCP Report. <https://zuplo.com/mcp-report>
- Nordic APIs articles are at nordicapis.com. Also cited on slides, links to add: Gartner; IDC FutureScape 2026; Contentstack, 2026 Agentic Enterprise Report; The Wall Street Journal; Bloomberg; Forbes; ABC News; Agents of Chaos study; Fresh Cup; The New York Times; APIsec Pulse; Zuplo (Nate Totten); David Soria Parra, AGNTCon + MCPCon Europe 2026; Twilio; Curity; Daniel Kocot.