


Here comes the normies

How DevOps can handle the flood of AI slop-apps

Coté, October 5th, 2026 – DevOpsDays Prague

Load increase at GitHub over 2 years



“Accenture is trying to figure out how to stop non-technical workers from blowing through companies’ AI token budget on trivial tasks like **converting PDFs to presentation slides.**”

404 Media, June 2026

BBVA

4,000+

GPTs in regular use, out of 20,000+
created

US Federal Government

1,818

Active or planned AI use cases, +70% y/y

AIRFRANCE KLM

200,000

hours of manual work saved by ~170 bots

SIXT SE

1,000

citizen-built apps in three months.

Table 1: Occupations and Tasks With Highest and Lowest genAI Adoption Rates



Top 10 occupations		Top 10 tasks	
Occupation	Rate (%)	Task	Rate (%)
Computer and information research scientists	87.3	Read documents to gather technical information.	61.3
Information security analysts	85.4	Prepare research reports.	60.7
Network and computer systems administrators	82.4	Design integrated computer systems.	59.5
Computer programmers	81.2	Analyze data to identify trends or relationships among variables.	57.5
Public relations specialists	80.4	Update knowledge about emerging industry or technology trends.	57.1
Personal financial advisors	79.7	Conduct quantitative failure analyses of operational data.	55.8
Chief executives	79.7	Prepare data for analysis.	55.4
Software developers	78.8	Develop marketing plans or strategies.	55.0
Market research analysts and marketing specialists	78.7	Collaborate with others to determine design specifications or details.	54.8
Computer and information systems managers	78.2	Design websites or web applications.	54.6
Bottom 10 occupations		Bottom 10 tasks	
Occupation	Rate (%)	Task	Rate (%)
Stockers and order fillers	14.8	Arrange delivery of goods or services.	0.0
Janitors and building cleaners	13.8	Assist practitioners to perform medical procedures.	0.0
Other installation, maintenance, and repair workers	13.1	Collect biological specimens from patients.	0.0
Maids and housekeeping cleaners	11.3	Drive trucks or other vehicles to or at work sites.	0.0
Fast food and counter workers	11.3	Inspect electrical or electronic systems for defects.	0.0
Maintenance and repair workers, general	11.2	Maintain medical equipment or instruments.	0.0
Painters and paperhangers	10.4	Measure the physical or physiological attributes of patients.	0.0
Licensed practical and licensed vocational nurses	10.4	Prepare patient treatment areas for use.	0.0
Receptionists and information clerks	7.6	Present food or beverage information or menus to customers.	0.0
Animal caretakers	5.3	Record service or repair activities.	0.0

Source: Bick, Blandin, Deming & Schumacher, [“What Work Does Generative AI Do?”](#), Federal Reserve Bank of St. Louis Working Paper 2026-017A, August 2026, Table 1, page 16; n = 13,920 employed US respondents.



Côté

<https://www.cote.io/> | cote@broadcom.com



Tanzu
Catsup

RedMonk




451 Research
S&P Global
Market Intelligence

Pivotal

vmware

Tanzu
by Broadcom



Normie

read 600+ RSS feeds overnight

write the morning briefing

mirror the blog to the NAS

reindex 200,000 documents

prune old transcripts

sync the bank accounts

publish a blog post

tune Diane's own rules and memory

drain the summaries queue

post-produce a podcast episode

build or fix one of my own skills

prep a conference talk

file the paperwork

clean up files and free disk

summarize an article or paper

summarize a video or a talk

fix the overnight job that broke

track CFPs and submit talks

cut shorts and clips

write up a report

build a little web page

make an image or an icon

post a photo gallery

reconcile the finances

write the social posts

assemble the newsletter

work out what I'm paying for

sweep the task list

transcribe a recording

debug my own Mac

sweep Slack for links

file an expense claim

triage Gmail

write YouTube titles and descriptions

narrate something in my own voice

build a chart

tidy up the bookmarks

run the website analytics

What we built

- **feed, classify, summarize, publish**
- **assemble a recurring document**
- **search my own stuff**
- **an inventory I kept in a spreadsheet**
- **query & change a system of record**
- **jobs that run overnight**


What that is at work

- **the competitive-intel digest**
- **the weekly status report**
- **enterprise RAG over the shared drive**
- **the CMDB, also a spreadsheet**
- **ServiceNow, Workday, SAP, via MCP**
- **the batch job**

The one person “enterprise” needs:

- Authentication, authorization: credentials, keys, permissions.
- Governance: allowed, needs a human, never.
- Data policy: A list of what must never leave the “enterprise.”
- Audit: a list of what counts as leaving, what did leave.
- Cost: tokens, used, and left.


WHAT I HAVE PERSONALLY SPENT ON AI, BY MONTH



Sources: My own receipts and card ledger, Dec 2022 to Sep 2026; y-axis removed on purpose.

Harness building eats the fun

One person's AI-run D&D campaign, by quarter. Share of conversations spent building the machinery: 5% at the start, 53% by the end.



Note: 4,245 conversations across ChatGPT, Claude.ai, Claude Code and Gemini, of which 1,522 are about tabletop roleplaying. Classified by a local model, 87% accurate on a 30-case calibration set - the shape is reliable, individual counts less so. 2026 Q3 is omitted: it had not finished.

Source: librarian full-text index of Côté's own ChatGPT, Claude.ai, Claude Code and Gemini archives. Retrieved 2026-08-26.

Chart generated by AI on 2026-08-27. Likely not checked by a human, so it may have errors.



Enterprise requirements for AI

Build

Coding & compiling.
Registry & CLI install.
Hidden SDLC.
Review/feedback.
Guardrails injection.

Observe

Performance.
Debugging, etc.
Audit records.
Evals and sampling.

Run

Runtime environment.
Sandboxed execution.
Service binding.
The SDLC, hidden.

Connect

Service binding.
Data access.
Network access.

Security

Registry and CLI control.
CVEs and patching.
Credential rotation.
Auth & permissions
MCP gateway.


Spend

Model access, public & private
Token tracking.
Rate limits and quotas.


“Instead of developers ‘shifting left,’
they need to **shift down** and push more
workloads down onto the platforms
they’re already using.”

Richard Seroter, Google Cloud, June 2023

Agents in your CaaS



Guardrails at build time with a proper platform



Sources: [“Inside the Coding Agent-Powered SDLC,”](#) Adib Saikali & Frank Denneman, Explore 2026. [“When everyone can build: The platform behind 1000 citizen-built app,”](#) Boyan Dimitrov, CTO of SIXT, PlatformCon 2026, published 22 June 2026

🏠 < >

Codex ▾

🔍 🔔

📝 New chat

💬

🔗 Pull requests

🕒 Scheduled

⚙️ Plugins

⋮ Explore

Pinned

📁 spring-petclinic

Create Postgres database ⌘1

Projects >

Recents >

AS asaikali

⬇️

📁 Create Postgres database ...

📶 Share ⚙️ 🗨️ 📄

what command did you use to deploy the application?

The primary deployment command was:

</> Bash

cf push spring-petclinic-2026 \
-p target/spring-petclinic-4.0.0-SNAPSHOT.jar \
-m 1G -k 1G --no-start

➦ 📄

I then bound and started it:

</> Bash

cf bind-service spring-petclinic-2026 spring-petclinic-db
cf set-env spring-petclinic-2026 SPRING_PROFILES_ACTIVE postgres
cf start spring-petclinic-2026

⬇️

📄 👍 💬 ➦ 12:08 AM

Do anything

+

🗨️ Approve for me

⚡ 5.6 Sol High ▾

🎤

⬆️

Sources: ["Inside the Coding Agent-Powered SDLC,"](#) Adib Saikali & Frank Denneman, Explore 2026.

18

Codex

New chat

Pull requests

Scheduled

Plugins

Explore

Pinned

spring-petclinic

Create Postgres database

Projects >

Recents >

asaikali

Create Postgres database ... Share

I uploaded the executable JAR with `cf push`; the Java buildpack automatically supplied Java 17, configured the runtime, and produced the containerized droplet—no Dockerfile was needed.

What decisions did you have to make to deploy the application?

- Target the existing Tanzu org/space: `adib / vm-explore`.
- Use the Java buildpack with the tested executable JAR instead of building a Docker image.
- Name the app and route `spring-petclinic-2026`.
- Allocate 1 GB memory and 1 GB disk.
- Bind `spring-petclinic-db` and activate the `postgres` Spring profile.
- Disable repeat schema initialization after credential rotation because the database was already populated.
- Suppress Hibernate connection-detail logging to prevent database credentials appearing in logs.

Do anything


Approve for me

5.6 Sol High


Enterprise function

This Normie's Skills

Brand and editorial standards	the-style-guide, ncsa-mosaic, diane-ui, chart-style, teleprompter-script
Security and DLP	the-security, publish-guard
Configuration management, asset inventory	cmdb, udr7, dirs
Records management and retention	archive, soup-search, librarian
Model sourcing and cost control	local-llm, the-codex, the-gemini
Content supply chain	enclave, hugo-git-publish, photo-post, social-poster, buffer, newsletter, youtube-upload, summarize
Knowledge management	librarian, pinboard, triage-inbox, slack-link-sweep, apple-notes-import
Media production	transcribe, tts, the-video-maker, the-image-maker, the-screenshooter, vtt-server, video-share
Job scheduling	scheduler, the-couch
Work management	omnifocus, diane-calendar, flighty
Change management, for the harness itself	the-skill-builder, the-close-out
Line of business, not enterprise	solo-dm, dice-roller, draw-cards, dutch-anki



Deep dive: MCP server governance pattern



“MCP support will provide crucial governance and security for MCP.

— Broadcom



Do not treat MCP servers as casual plugins; treat them as privileged integration services.

Sources: Broadcom VCF AI-native announcement; ServiceNow MCP server governance docs; MCP specification




Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

21

CVEs reported by year

CVE records per calendar year, 2014 to 2026. 2026 is still running.



2026 has nine months on the board, not twelve - and has already beaten all of 2025 by half again.

Source: CVE.ICU, cve.icu, data as of 2026-09-28.


If you thought programmers were expensive, wait until everybody's eating tokens.

FIGURE 5

Deployment Location for the Development and Deployment of AI Models

Where does your organization primarily develop and deploy AI models?

(Percentage of respondents)




Note: Public cloud includes public cloud provider, managed services provider, and special-purpose cloud. On premises includes private cloud, traditional infrastructure, and colocation provider. n = 411; Source: IDC's *AI Infrastructure Survey*, July 2024.

For an accessible version of the data in this figure, see [Figure 5 Supplemental Data](#) in Appendix 1.

Workload placement over the years (Flexera)

■ Public Cloud ■ Private Cloud ■ Other



"Other" is "Non-cloud" in 2018 and 2019, "Additional workloads in public cloud in 12 months" in remaining years.

**“Treat every agent like it’s a human.
The identity controls, the access to
information, the ability to take actions,
the audit trail.”**

Stephen Elliott, IDC



<https://talks.cote.io/normieai/>



michael.cote@broadcom.com

Slides & stuff

