


THROUGHPUT

THROUGHPUT • MONTHLY • JULY 2026 • ISSUE M-1 • THE IT JOBS MONTH



Coverage: no new postings captured on 8 July days (3, 10, 11, 12, 16, 18, 28, 30); see p. 53.

Jobs to highlight

We scored 30 IT roles on eight measures of opportunity. Software engineer (general), ML & AI engineer and QA & manual testing lead July 2026. Here is why, and who hires for them.

BEST BET FOR FRESHERS

Software engineer (general): 14% of its postings ask 0–2 years

RISING

Software architect: +24% in share of IT hiring

COOLING: APPROACH WITH CARE

Frontend developer and Embedded & firmware lose share

IN THIS ISSUE

Contents

JULY 2026 · ISSUE M-1

3	Jobs to highlight Eight measures, one score, the leaders.
4	Why these roles Demand, skills, cities and who is hiring.
8	Best bets for you Freshers, mid, senior, switchers, remote.
9	Best bets by city Six metros, and smaller cities rising.
10	The map, cooling, the score Demand vs momentum; how the score works.
11	Your next move Starting roles and their bridges.
14	The month in numbers Slightly slower than June, and a patchy crawl.
15	Skills Data platforms rise; embedded and IoT fade.
16	Skills that come in packs Combinations, family by family.
17	Skills radar Adopt, trial, watch, fading.
18	The minimum stack What half the postings assume.
19	Skill-gap checklists Tick-box skill lists for the top 10 roles.
21	CV keyword guide The terms postings use, role by role.
22	30/60/90-day roadmaps Three-month learning paths for six roles.
24	Who's hiring IT The top 15, and the month's movers.
25	Who hires for each role Top employers per highlighted role.
26	What they might be building Speculation: domains by cohort.
27	Employer signals (speculation) Speculation: what employers may be building.
30	Where the IT jobs are Hyderabad gains a point.
31	Where each role lives Where each role over-indexes.
32	Roles and new roles Architects and enterprise apps gain share.
33	The league table Ranks, month by month.
34	Seniority and experience Senior-heavy, with little room for freshers.
35	The experience ladder Years asked, role by role.
36	Fresher's corner First jobs: roles, employers, skills, cities.
38	Remote vs office On-site still rules.
39	Hardest to fill Re-posts, and when to apply.
40	Hard problems worth reading about Editorial explainers, with reading lists.

43	Org watch Six cohorts compared, from services to banks.
44	Cohort org charts Six archetypes.
47	Org watch: company profiles Six big IT hirers, one inferred org chart each.
49	Glossary Plain-English meanings of terms used here.
51	References Every company named, with IT posting counts.
52	Data and sources Where the data comes from; rights and limits.
53	Methodology How we counted, and the crawl gaps.
54	The taxonomy How each posting was classified.
56	Numbers of the month Ten figures to take away.

From the editor

This issue answers one question for people in IT: which jobs should you go for? We scored the 30 IT roles with enough postings on demand, momentum, staying power, breadth of employers, share still open at the cut-off, entry level, flexibility and skill overlap, from 20,496 new IT postings first captured in July 2026.

The first pages give the ranking and a card for each leading role, with the employers hiring for it; then best bets by experience, city and work mode, and the roles cooling off. The rest of the issue backs it with skills, employers, cities and hiring evidence. There is no salary data, so pay is not scored. Boxes marked interpretation are our opinion. New this month: what employers might be building, read from titles and skills and labelled as speculation, and six hard problems worth reading about. Our crawl captured no new postings on 8 days this month; every comparison is per day or a share.

Also new in this issue: a skills radar, skill-gap checklists to print, CV keyword guides, 30/60/90-day learning roadmaps, a fresher's corner and a glossary. Advice in them is labelled as editorial. This is our first monthly edition, so there is no previous issue to compare with: the since-last-edition scorecard starts with the August issue.

— THE THROUGHPUT DATA DESK

For live data insights
CareerLens: careerlens.live · Track. Learn. Stay Ahead.

COVER STORY · JOBS TO HIGHLIGHT · JULY 2026

The jobs to go for

Which IT jobs should you go for? We scored all 30 IT roles with enough postings on eight measures of opportunity. These are the month's best bets, with the evidence.

The Highlight score asks eight questions of every role: how many new postings a day (demand), whether its share of IT hiring is rising (momentum), whether it stays near the top month after month (persistence), how many employers post it (breadth), what share is still marked open at the cut-off, how much of it asks three years or fewer

(accessibility), how often it is hybrid or remote (flexibility), and how many of its core skills it shares with neighbouring roles (skill-ladder fit).


Each answer is turned into a percentile rank among the 30 roles, then weighted: demand 20%, momentum, breadth, accessibility and skill-ladder fit 15% each, persistence 10%, still open at the cut-off and flexibility 5%

each. The formula and a sensitivity test are on p. 10. There is no salary data in our crawl, so the score says nothing about pay.

Software engineer (general) leads with 82 of 100, ahead of ML & AI engineer (81) and QA & manual testing (71). 10 of the top 10 stay in the top 10 under every one of 16 weight changes of plus or minus 25%.

Demand Momentum Persistence Breadth of employers Openness Accessibility Flexibility Skill-ladder fit

HIGHLIGHT SCORE, 0-100 (EACH SEGMENT = WEIGHT x PERCENTILE RANK OF ONE INPUT)



New IT postings first captured 1 Jul-31 Jul 2026 (20,496); roles with at least 150 postings (30 of 34). Percentile ranks among those roles; weights on p. 10. Bold: the top 10.

WHAT THIS MEANS FOR YOU · INTERPRETATION

Read the ranking as a shortlist, not a verdict: a role near the top has many employers, steady demand and room to enter. Use the cards that follow to check the skills, cities and employers for the ones that fit you.

JOBS TO HIGHLIGHT · ROLES 1-3 OF 12

Why these roles, and who hires for them

#1

Software engineer (general)

Software engineering · score 82/100

Demand: 103.6 new postings a day (15.7% of IT).

Momentum: -4% (vs June 2026 -4%, vs trailing -3%).

Employers: 672 a month; top five hold 10%.

Entry: 31% ask 0-3 years; 126 intern/fresher titles.

Flexibility: 19% hybrid or remote (of known).

Persistence: top 10 in 3 of 3 months.

Why: Strong on demand (100th percentile), breadth of employers (100th percentile), skill-ladder fit (100th percentile).

Skills to have: REST APIs, Agile, CI/CD, Python, Java, SQL, AWS.

Minimum stack (≥50%): no single skill reaches half of postings; the list above is the core.

Cities: Bengaluru 42%, Hyderabad 14%, Pune 10% of its postings.

Hiring most: Wipro 3%, Microsoft 2%, Capgemini 2%, Oracle 2%.

WHO'S HIRING FOR IT · JULY 2026

Wipro 82 postings · 2.6%

Microsoft 73 postings · 2.3%

Capgemini 58 postings · 1.8%

Oracle 51 postings · 1.6%

Dexter's Tech 50 postings · 1.6%

Share of the role's 3,212 new IT postings. 672 employers a month; the top five hold 10%.

#2

ML & AI engineer

Data & AI · score 81/100

Demand: 60.8 new postings a day (9.2% of IT).

Momentum: +4% (vs June 2026 +2%, vs trailing +6%).

Employers: 586 a month; top five hold 9%.

Entry: 28% ask 0-3 years; 44 intern/fresher titles.

Flexibility: 24% hybrid or remote (of known).

Persistence: top 10 in 3 of 3 months.

Why: Strong on demand (97th percentile), breadth of employers (97th percentile), flexibility (90th percentile).

Skills to have: Python, LLMs, REST APIs, Generative AI, Azure, Machine Learning, AWS.

Minimum stack (≥50%): Python 63%.

Cities: Bengaluru 35%, Hyderabad 15%, Delhi NCR 12% of its postings.

Hiring most: Teradata 2%, Alignerr † 2%, TalentGigs † 2%, Zorba Consulting India † 2%.

WHO'S HIRING FOR IT · JULY 2026

Teradata 39 postings · 2.1%

Alignerr † 34 postings · 1.8%

TalentGigs † 33 postings · 1.8%

Zorba Consulting India † 29 postings · 1.5%

Smartworkingsolutions † 29 postings · 1.5%

Share of the role's 1,885 new IT postings. 586 employers a month; the top five hold 9%.

#3

QA & manual testing

QA & testing · score 71/100

Demand: 24.7 new postings a day (3.7% of IT).

Momentum: -1% (vs June 2026 +1%, vs trailing -4%).

Employers: 332 a month; top five hold 11%.

Entry: 32% ask 0-3 years; 9 intern/fresher titles.

Flexibility: 19% hybrid or remote (of known).

Persistence: top 10 in 3 of 3 months.

Why: Strong on accessibility (86th percentile), persistence (84th percentile), breadth of employers (83th percentile). Weaker on openness.

Skills to have: REST APIs, Agile, Regression, Test Automation, CI/CD, Jira, Selenium.

Minimum stack (≥50%): no single skill reaches half of postings; the list above is the core.

Cities: Bengaluru 28%, Hyderabad 15%, Pune 13% of its postings.

Hiring most: Birlasoft 3%, VOIS 3%, Rockwell Automation 2%, TalentGigs † 2%.

WHO'S HIRING FOR IT · JULY 2026

Birlasoft 24 postings · 3.1%

VOIS 20 postings · 2.6%

Rockwell Automation 18 postings · 2.3%

TalentGigs † 12 postings · 1.6%

Agileengine 10 postings · 1.3%

Share of the role's 767 new IT postings. 332 employers a month; the top five hold 11%.

JOBS TO HIGHLIGHT · ROLES 4–6 OF 12

Why these roles

#4

Backend developer

Software engineering · score 71/100

Demand: 42.6 new postings a day (6.4% of IT).**Momentum:** -4% (vs June 2026 -1%, vs trailing -7%).**Employers:** 447 a month; top five hold 9%.**Entry:** 18% ask 0–3 years; 1 intern/fresher titles.**Flexibility:** 27% hybrid or remote (of known).**Persistence:** top 10 in 3 of 3 months.**Why:** Strong on flexibility (95th percentile), demand (93th percentile), breadth of employers (93th percentile).**Skills to have:** REST APIs, Java, CI/CD, SQL, Microservices, AWS, Agile.**Minimum stack (≥50%):** REST APIs 65%.**Cities:** Bengaluru 36%, Delhi NCR 12%, Hyderabad 11% of its postings.**Hiring most:** TalentGigs † 2%, Turingcom 2%, GlobalLogic 2%, Infosys 2%.

WHO'S HIRING FOR IT · JULY 2026

TalentGigs † 28 postings · 2.1%**Turingcom** 23 postings · 1.7%**GlobalLogic** 23 postings · 1.7%**Infosys** 22 postings · 1.7%**Iris Software** 21 postings · 1.6%

Share of the role's 1,320 new IT postings. 447 employers a month; the top five hold 9%.

#5

Data & BI analyst

Data & AI · score 65/100

Demand: 33.4 new postings a day (5.1% of IT).**Momentum:** -4% (vs June 2026 -3%, vs trailing -5%).**Employers:** 380 a month; top five hold 12%.**Entry:** 34% ask 0–3 years; 83 intern/fresher titles.**Flexibility:** 24% hybrid or remote (of known).**Persistence:** top 10 in 3 of 3 months.**Why:** Strong on accessibility (90th percentile), demand (86th percentile), breadth of employers (86th percentile). Weaker on skill-ladder fit.**Skills to have:** SQL, Dashboards, Power BI, Python, Excel, Data Analytics, Visualization.**Minimum stack (≥50%):** SQL 64%, Dashboards 52%.**Cities:** Bengaluru 29%, Hyderabad 14%, Delhi NCR 14% of its postings.**Hiring most:** Argo Intern 3%, Wake Up Whistle 3%, Dexter's Tech 2%, Synchrony Financial 2%.

WHO'S HIRING FOR IT · JULY 2026

Argo Intern 30 postings · 2.9%**Wake Up Whistle** 28 postings · 2.7%**Dexter's Tech** 26 postings · 2.5%**Synchrony Financial** 24 postings · 2.3%**HSBC** 18 postings · 1.7%

Share of the role's 1,036 new IT postings. 380 employers a month; the top five hold 12%.

#6

Data engineer

Data & AI · score 64/100

Demand: 40.6 new postings a day (6.1% of IT).**Momentum:** +12% (vs June 2026 +12%, vs trailing +13%).**Employers:** 413 a month; top five hold 14%.**Entry:** 16% ask 0–3 years; 2 intern/fresher titles.**Flexibility:** 17% hybrid or remote (of known).**Persistence:** top 10 in 3 of 3 months.**Why:** Strong on demand (90th percentile), breadth of employers (90th percentile), persistence (84th percentile). Weaker on openness and accessibility.**Skills to have:** SQL, Python, ETL, Data Pipelines, Data Quality, AWS, Azure.**Minimum stack (≥50%):** SQL 77%, Python 68%, ETL 64%, Data Pipelines 63%, Data Quality 53%.**Cities:** Bengaluru 33%, Hyderabad 19%, Pune 11% of its postings.**Hiring most:** Zorba Consulting India † 6%, Innova ESI † 3%, Cisco 2%, Tata Consultancy Services 2%.

WHO'S HIRING FOR IT · JULY 2026

Zorba Consulting India † 71 postings · 5.6%**Innova ESI** † 43 postings · 3.4%**Cisco** 24 postings · 1.9%**Tata Consultancy Services** 21 postings · 1.7%**Husky Technologies** 21 postings · 1.7%
Share of the role's 1,258 new IT postings. 413 employers a month; the top five hold 14%.

JOBS TO HIGHLIGHT · ROLES 7–9 OF 12

Why these roles

#7

Infrastructure, network & IT

Cloud, DevOps & SRE · score 63/100

Demand: 27.8 new postings a day (4.2% of IT).**Momentum:** +1% (vs June 2026 +1%, vs trailing +0%).**Employers:** 320 a month; top five hold 11%.**Entry:** 32% ask 0–3 years; 1 intern/fresher titles.**Flexibility:** 17% hybrid or remote (of known).**Persistence:** top 10 in 3 of 3 months.**Why:** Strong on persistence (84th percentile), accessibility (83th percentile), breadth of employers (78th percentile). Weaker on skill-ladder fit and flexibility.**Skills to have:** Python, Azure, Linux, Windows, AWS, TCP, ITIL.**Minimum stack (≥50%):** no single skill reaches half of postings; the list above is the core.**Cities:** Bengaluru 36%, Hyderabad 12%, Delhi NCR 12% of its postings.**Hiring most:** Tata Communication 3%, NTT Data 2%, Capgemini 2%, Amazon 2%.

WHO'S HIRING FOR IT · JULY 2026

Tata Communication 28 postings · 3.3%**NTT Data** 20 postings · 2.3%**Capgemini** 17 postings · 2.0%**Amazon** 16 postings · 1.9%**Hewlett Packard Enterprise** 16 postings · 1.9%

Share of the role's 861 new IT postings. 320 employers a month; the top five hold 11%.

#8

Full-stack developer

Software engineering · score 62/100

Demand: 24.7 new postings a day (3.7% of IT).**Momentum:** -9% (vs June 2026 -10%, vs trailing -8%).**Employers:** 320 a month; top five hold 12%.**Entry:** 19% ask 0–3 years; 23 intern/fresher titles.**Flexibility:** 27% hybrid or remote (of known).**Persistence:** top 10 in 3 of 3 months.**Why:** Strong on flexibility (95th percentile), skill-ladder fit (88th percentile), persistence (84th percentile). Weaker on momentum.**Skills to have:** REST APIs, React, CI/CD, JavaScript, Agile, Microservices, AWS.**Minimum stack (≥50%):** REST APIs 81%, React 64%, CI/CD 59%.**Cities:** Bengaluru 31%, Hyderabad 13%, Pune 13% of its postings.**Hiring most:** Bnyglobal 3%, Husky Technologies 3%, Smartworkingsolutions † 3%, Iris Software 2%.

WHO'S HIRING FOR IT · JULY 2026

Bnyglobal 25 postings · 3.3%**Husky Technologies** 20 postings · 2.6%**Smartworkingsolutions** † 20 postings · 2.6%**Iris Software** 16 postings · 2.1%**Programmers Io** 15 postings · 2.0%

Share of the role's 767 new IT postings. 320 employers a month; the top five hold 12%.

#9

Oracle, ERP & other platforms

Enterprise apps · score 61/100

Demand: 31.1 new postings a day (4.7% of IT).**Momentum:** +20% (vs June 2026 +19%, vs trailing +20%).**Employers:** 299 a month; top five hold 16%.**Entry:** 16% ask 0–3 years; 2 intern/fresher titles.**Flexibility:** 18% hybrid or remote (of known).**Persistence:** top 10 in 3 of 3 months.**Why:** Strong on momentum (97th percentile), persistence (84th percentile), demand (83th percentile). Weaker on skill-ladder fit and openness.**Skills to have:** REST APIs, Oracle, ServiceNow, Agile, SOAP, Workday, SQL.**Minimum stack (≥50%):** no single skill reaches half of postings; the list above is the core.**Cities:** Bengaluru 32%, Hyderabad 17%, Pune 10% of its postings.**Hiring most:** Oracle 5%, Teamware Solutions † 3%, SourceBae † 3%, Tata Consultancy Services 3%.

WHO'S HIRING FOR IT · JULY 2026

Oracle 46 postings · 4.8%**Teamware Solutions** † 30 postings · 3.1%**SourceBae** † 29 postings · 3.0%**Tata Consultancy Services** 25 postings · 2.6%**Nagarro** 24 postings · 2.5%

Share of the role's 963 new IT postings. 299 employers a month; the top five hold 16%.

JOBS TO HIGHLIGHT · ROLES 10–12 OF 12

Why these roles

#10

Security engineering (general)

Security · score 60/100

Demand: 22.3 new postings a day (3.4% of IT).**Momentum:** +4% (vs June 2026 +3%, vs trailing +5%).**Employers:** 309 a month; top five hold 14%.**Entry:** 30% ask 0–3 years; 11 intern/fresher titles.**Flexibility:** 17% hybrid or remote (of known).**Persistence:** top 10 in 0 of 3 months.**Why:** Strong on openness (81th percentile), accessibility (76th percentile), breadth of employers (72th percentile).**Skills to have:** Cybersecurity, Python, Cloud Security, Incident Response, SIEM, Azure, AWS.**Minimum stack (≥50%):** no single skill reaches half of postings; the list above is the core.**Cities:** Bengaluru 34%, Delhi NCR 15%, Hyderabad 13% of its postings.**Hiring most:** Wells Fargo 4%, NTT Data 3%, Verizon 2%, FIS 2%.

WHO'S HIRING FOR IT · JULY 2026

Wells Fargo 31 postings · 4.5%**NTT Data** 21 postings · 3.0%**Verizon** 17 postings · 2.5%**FIS** 16 postings · 2.3%**Knowdis** 14 postings · 2.0%

Share of the role's 309 new IT postings. 309 employers a month; the top five hold 14%.

#11

AI agent & GenAI engineer

Data & AI · score 56/100

Demand: 17.7 new postings a day (2.7% of IT).**Momentum:** +13% (vs June 2026 +6%, vs trailing +21%).**Employers:** 216 a month; top five hold 12%.**Entry:** 23% ask 0–3 years; 1 intern/fresher titles.**Flexibility:** 28% hybrid or remote (of known).**Persistence:** top 10 in 0 of 3 months.**Why:** Strong on flexibility (100th percentile), momentum (86th percentile). Weaker on openness.**Skills to have:** LLMs, Generative AI, Python, REST APIs, RAG, AWS, LangChain.**Minimum stack (≥50%):** LLMs 74%, Generative AI 73%, Python 72%, REST APIs 68%, RAG 57%.**Cities:** Bengaluru 36%, Delhi NCR 16%, Hyderabad 12% of its postings.**Hiring most:** Knowdis 3%, Iris Software 2%, PwC India 2%, Innova ESI † 2%.

WHO'S HIRING FOR IT · JULY 2026

Knowdis 14 postings · 2.6%**Iris Software** 13 postings · 2.4%**PwC India** 13 postings · 2.4%**Innova ESI †** 12 postings · 2.2%**Orioninnovation** 12 postings · 2.2%

Share of the role's 549 new IT postings. 216 employers a month; the top five hold 12%.

#12

Data scientist

Data & AI · score 54/100

Demand: 17.9 new postings a day (2.7% of IT).**Momentum:** -5% (vs June 2026 -5%, vs trailing -6%).**Employers:** 255 a month; top five hold 12%.**Entry:** 39% ask 0–3 years; 10 intern/fresher titles.**Flexibility:** 17% hybrid or remote (of known).**Persistence:** top 10 in 0 of 3 months.**Why:** Strong on accessibility (97th percentile).**Skills to have:** Python, Data Science, Machine Learning, Statistics, SQL, LLMs, Algorithms.**Minimum stack (≥50%):** Python 84%, Data Science 72%, Machine Learning 69%, Statistics 58%, SQL 55%.**Cities:** Bengaluru 40%, Hyderabad 13%, Delhi NCR 11% of its postings.**Hiring most:** Zorba Consulting India † 4%, Faceprep 3%, Amgen 2%, D Square Consulting Services Pvt Ltd 2%.

WHO'S HIRING FOR IT · JULY 2026

Zorba Consulting India † 20 postings · 3.6%**Faceprep** 18 postings · 3.2%**Amgen** 10 postings · 1.8%**D Square Consulting Services Pvt Ltd** 10 postings · 1.8%**Tata Consultancy Services** 8 postings · 1.4%

Share of the role's 556 new IT postings. 255 employers a month; the top five hold 12%.

JOB TO HIGHLIGHT · BEST BETS BY SEGMENT

Best bets for you

The same evidence, re-weighted for where you are: years of experience, a career switch, or a need for hybrid or remote work.

Freshers (0–2 years)

#	Role	Est. 0–2y	Ask 0–2y	Mom.
1	Software eng.	14.2/day	14%	-4%
2	Data/BI analyst	5.6/day	17%	-4%
3	ML/AI engineer	5.7/day	9%	+4%
4	Security eng.	4.2/day	19%	+4%
5	Data scientist	3.5/day	20%	-5%

Score: 35% postings/day asking 0-2 years (est. = per day x share asking 0-2), 20% share asking 0-2 years, 15% fresher/intern titles per week, 15% momentum, 15% employers.

Mid-career (3–7 years)

#	Role	Est. 3–7y	Ask 3–7y	Mom.
1	Data engineer	29.1/day	72%	+12%
2	Oracle/ERP	20.1/day	64%	+20%
3	Software eng.	66.0/day	64%	-4%
4	ML/AI engineer	36.0/day	59%	+4%
5	Backend	29.0/day	68%	-4%

Score: 40% postings/day asking 3-7 years, 20% share asking 3-7, 20% momentum, 20% employers.

Senior and lead (8+ years)

#	Role	Est. 8+y	Ask 8+y	Mom.
1	ML/AI engineer	19.2/day	32%	+4%
2	SAP	13.1/day	42%	+16%
3	Software architect	9.1/day	82%	+24%
4	Oracle/ERP	9.4/day	30%	+20%
5	Data engineer	9.7/day	24%	+12%

Score: 40% postings/day asking 8+ years, 20% share asking 8+, 20% momentum, 20% employers.

Career switchers (bridge skills)

#	Role	Shared core	Per day	Mom.
1	Software eng.	7.7/10	103.6	-4%
2	ML/AI engineer	6.0/10	60.8	+4%
3	Backend	6.7/10	42.6	-4%
4	Data engineer	5.0/10	40.6	+12%
5	Software architect	5.7/10	11.2	+24%

Score: 40% skill-ladder fit, 30% postings/day, 30% momentum.

Remote and hybrid seekers

#	Role	Hyb. +rem.	Remote	Mom.
1	GenAI/agents	28%	11%	+13%
2	ML/AI engineer	24%	11%	+4%
3	Backend	27%	12%	-4%
4	Data/BI analyst	24%	7%	-4%
5	Full-stack	27%	15%	-9%

Score: 50% hybrid+remote share of known work modes, 30% hybrid/remote postings per day, 20% momentum.

WHAT THIS MEANS FOR YOU · INTERPRETATION

Interpretation. Freshers: the best bets combine many postings asking 0–2 years with many employers. Switchers: favour roles whose core skills overlap with yours. Remote seekers: flexibility is rarer than it looks; read the work-mode line of each posting.

JOBS TO HIGHLIGHT · BY CITY

Best bets, city by city

Each city's best bets: the role's postings in the city, whether it is gaining share there, and its national score.

Bengaluru

243 new IT postings a day · 36.8% of India · -10% a day vs before

Role	Per day	City share	Δ share	vs India
Backend	15.5	6.4%	+6%	1.0×
Data engineer	13.4	5.5%	+12%	0.9×
ML/AI engineer	21.1	8.7%	+1%	0.9×
Software eng.	43.8	18.0%	-1%	1.1×
Oracle/ERP	10.1	4.2%	+11%	0.9×

Hyderabad

94 new IT postings a day · 14.2% of India · -1% a day vs before

Role	Per day	City share	Δ share	vs India
Data engineer	7.7	8.3%	+15%	1.3×
QA & testing	3.7	3.9%	+21%	1.0×
Oracle/ERP	5.2	5.6%	+18%	1.2×
Software eng.	15.0	16.0%	-8%	1.0×
ML/AI engineer	9.0	9.6%	-6%	1.0×

Pune

64 new IT postings a day · 9.7% of India · -11% a day vs before

Role	Per day	City share	Δ share	vs India
Software eng.	10.2	15.9%	+14%	1.0×
Data engineer	4.4	6.9%	+19%	1.1×
QA & testing	3.3	5.1%	+27%	1.4×
Data/BI analyst	3.2	5.0%	+29%	1.0×
ML/AI engineer	4.4	6.9%	-12%	0.8×

Chennai

53 new IT postings a day · 8.0% of India · -14% a day vs before

Role	Per day	City share	Δ share	vs India
Backend	4.5	8.5%	+66%	1.3×
ML/AI engineer	4.3	8.1%	+25%	0.9×
Software eng.	8.2	15.5%	+9%	1.0×
Infra & network	2.6	4.9%	+38%	1.2×
Oracle/ERP	2.4	4.6%	+31%	1.0×

Delhi NCR

71 new IT postings a day · 10.8% of India · +1% a day vs before

Role	Per day	City share	Δ share	vs India
ML/AI engineer	7.6	10.6%	+4%	1.2×
Security eng.	3.4	4.8%	+51%	1.4×
Infra & network	3.2	4.5%	+17%	1.1×
Backend	5.3	7.4%	-11%	1.1×
Oracle/ERP	3.1	4.3%	+28%	0.9×

Mumbai region

33 new IT postings a day · 5.0% of India · -7% a day vs before

Role	Per day	City share	Δ share	vs India
Software eng.	3.3	9.9%	+15%	0.6×
Data/BI analyst	2.8	8.6%	+23%	1.7×
Backend	2.7	8.3%	+7%	1.3×
ML/AI engineer	3.0	9.0%	-10%	1.0×
Full-stack	1.5	4.7%	+66%	1.3×

RIISING FROM A SMALL BASE

Emerging cities (outside the six metros, growth in new IT postings a day vs before, from a small base): Kochi +48% (2.9 a day); Jaipur +42% (1.7 a day); Ahmedabad +9% (3.4 a day); Thiruvananthapuram +7% (2.8 a day); Coimbatore -4% (2.8 a day).

City score: 40% the role's postings a day in the city, 30% change in its share of the city's IT postings, 30% national Highlight score. vs India = location quotient (city share ÷ national share). NCR = Gurugram, Noida, Delhi and neighbours.


WHAT THIS MEANS FOR YOU · INTERPRETATION

Interpretation. Bengaluru has the most of almost everything, so look at what each city over-indexes on: a smaller city where your role is concentrated can mean less competition for the same work.

JOBS TO HIGHLIGHT · THE MAP

Big, growing, cooling

Demand against momentum for every scored role. Big and growing is the sweet spot; big but cooling still hires, with more competition per posting.



COOLING: APPROACH WITH CARE · INTERPRETATION

Role	Per day	Momentum	Rank
Frontend developer	10.8	-21%	26
Embedded & firmware	7.0	-19%	30
Cloud engineer & architect	20.3	-15%	14
Salesforce & CRM	12.3	-13%	25
Test automation & SDET	10.4	-9%	22

Interpretation. These roles lost share of IT hiring against both the previous period and the trailing average. They still hire, often in volume, but a shrinking share usually means more applicants per opening. Approach with a sharper profile. Shown in pink on the map. Momentum = the geometric mean of the change in share vs the previous period and vs the trailing average.

HOW THE HIGHLIGHT SCORE WORKS · METHOD

Score = $\sum \text{weight} \times \text{percentile rank (0-100)}$ among the 30 IT roles with at least 150 new postings.

Demand 20%: new postings per day.

Momentum 15%: change in share of IT postings vs the previous period and the trailing average.

Persistence 10%: share of months (from May) in the top 10 roles.

Breadth of employers 15%: employers posting the role per month.

Openness 5%: share still marked open at the cut-off (the crawl's open flag).

Accessibility 15%: share asking 0-3 years (job-intel).

Flexibility 5%: hybrid + remote share of known work modes (job-intel).

Skill-ladder fit 15%: core skills (top 10) shared with its three nearest roles.

Sensitivity: each weight $\times 0.75$ and $\times 1.25$ (16 runs); 10 of the top 10 stay in the top 10 in every run (smallest overlap 10/10). In 2,000 random draws of all weights within $\pm 25\%$, ranks of the top 10 move by at most 5 places.

Limits: no salary data exists in the crawl, so pay is not scored. The open flag is updated in batches (near 100% for postings under ten weeks old), so openness barely separates roles; it describes the cut-off, not today. Job-intel inputs cover the analysed share of postings only.

YOUR NEXT MOVE · INTERPRETATION FROM SKILLS DATA

Your next move

For common starting roles: the highlighted roles closest in skills, the bridge skills to add, and the evidence.

IF YOU ARE A...

QA & manual testing

Now: 24.7 new postings a day, momentum -1%, Highlight rank 3.

→ **Software engineer (general)** (#1, skill overlap 59%)

Learn: Microservices (18% vs 0%), React (17% vs 0%), AWS (28% vs 12%).

Already shared: REST APIs, Agile, CI/CD. Evidence: 103.6/day, momentum -4%, 31% ask 0–3 y.

→ **Backend developer** (#4, skill overlap 52%)

Learn: Microservices (41% vs 0%), Spring (32% vs 0%), AWS (40% vs 12%).

Already shared: REST APIs, CI/CD, SQL. Evidence: 42.6/day, momentum -4%, 18% ask 0–3 y.

→ **ML & AI engineer** (#2, skill overlap 35%)

Learn: LLMs (50% vs 0%), Python (63% vs 25%), Machine Learning (37% vs 0%). Already shared: Python, REST APIs, CI/CD. Evidence: 60.8/day, momentum +4%, 28% ask 0–3 y.

IF YOU ARE A...

Infrastructure, network & IT

Now: 27.8 new postings a day, momentum +1%, Highlight rank 7.

→ **Software engineer (general)** (#1, skill overlap 48%)

Learn: REST APIs (45% vs 10%), Java (31% vs 4%), Agile (36% vs 10%).

Already shared: Python, AWS, Azure. Evidence: 103.6/day, momentum -4%, 31% ask 0–3 y.

→ **ML & AI engineer** (#2, skill overlap 40%)

Learn: LLMs (50% vs 0%), Generative AI (39% vs 0%), Machine Learning (37% vs 0%). Already shared: Python, Azure, AWS. Evidence: 60.8/day, momentum +4%, 28% ask 0–3 y.

→ **Backend developer** (#4, skill overlap 44%)

Learn: REST APIs (65% vs 10%), Java (48% vs 4%), Microservices (41% vs 0%). Already shared: AWS, Python. Evidence: 42.6/day, momentum -4%, 18% ask 0–3 y.

IF YOU ARE A...

Data & BI analyst

Now: 33.4 new postings a day, momentum -4%, Highlight rank 5.

→ **Software engineer (general)** (#1, skill overlap 30%)

Learn: REST APIs (45% vs 9%), CI/CD (32% vs 0%), Java (31% vs 0%).

Already shared: Python, SQL. Evidence: 103.6/day, momentum -4%, 31% ask 0–3 y.

→ **ML & AI engineer** (#2, skill overlap 29%)

Learn: LLMs (50% vs 0%), REST APIs (43% vs 9%), RAG (33% vs 0%).

Already shared: Python. Evidence: 60.8/day, momentum +4%, 28% ask 0–3 y.

→ **Data engineer** (#6, skill overlap 50%)

Learn: Data Pipelines (63% vs 11%), ETL (64% vs 13%), Spark (40% vs 0%).

Already shared: SQL, Python, Data Quality. Evidence: 40.6/day, momentum +12%, 16% ask 0–3 y.

IF YOU ARE A...

Software engineer (general)

Now: 103.6 new postings a day, momentum -4%, Highlight rank 1.

→ **Backend developer** (#4, skill overlap 91%)

Learn: Microservices (41% vs 18%), Spring (32% vs 10%), REST APIs (65% vs 45%). Already shared: REST APIs, Java, CI/CD. Evidence: 42.6/day, momentum -4%, 18% ask 0–3 y.

→ **ML & AI engineer** (#2, skill overlap 59%)

Learn: LLMs (50% vs 7%), Machine Learning (37% vs 0%), Generative AI (39% vs 4%). Already shared: Python, REST APIs, Azure. Evidence: 60.8/day, momentum +4%, 28% ask 0–3 y.

→ **QA & manual testing** (#3, skill overlap 59%)

Learn: Regression (43% vs 0%), Test Automation (35% vs 0%), Selenium (28% vs 0%). Already shared: REST APIs, Agile, CI/CD. Evidence: 24.7/day, momentum -1%, 32% ask 0–3 y.

IF YOU ARE A...

Backend developer

Now: 42.6 new postings a day, momentum -4%, Highlight rank 4.

→ **Software engineer (general)** (#1, skill overlap 91%)

Learn: TypeScript (10% vs 0%), Software Design (10% vs 0%), Web

Applications (10% vs 0%). Already shared: REST APIs, Agile, CI/CD.

Evidence: 103.6/day, momentum -4%, 31% ask 0–3 y.

→ **ML & AI engineer** (#2, skill overlap 56%)

Learn: LLMs (50% vs 9%), Generative AI (39% vs 0%), Machine Learning (37% vs 0%). Already shared: Python, REST APIs, AWS. Evidence: 60.8/day, momentum +4%, 28% ask 0–3 y.

→ **QA & manual testing** (#3, skill overlap 52%)

Learn: Regression (43% vs 0%), Test Automation (35% vs 0%), Jira (29% vs 0%). Already shared: REST APIs, Agile, CI/CD. Evidence: 24.7/day, momentum -1%, 32% ask 0–3 y.

IF YOU ARE A...

Frontend developer

Now: 10.8 new postings a day, momentum -21%, Highlight rank 26.

→ **Software engineer (general)** (#1, skill overlap 70%)

Learn: Python (31% vs 6%), Observability (14% vs 0%), Distributed Systems (14% vs 0%). Already shared: REST APIs, Agile, CI/CD. Evidence: 103.6/day, momentum -4%, 31% ask 0–3 y.

→ **Backend developer** (#4, skill overlap 58%)

Learn: Java (48% vs 17%), Python (35% vs 6%), Spring (32% vs 6%).

Already shared: REST APIs, CI/CD, Agile. Evidence: 42.6/day, momentum -4%, 18% ask 0–3 y.

→ **QA & manual testing** (#3, skill overlap 41%)

Learn: Regression (43% vs 0%), Test Automation (35% vs 0%), Jira (29% vs 0%). Already shared: REST APIs, Agile, CI/CD. Evidence: 24.7/day, momentum -1%, 32% ask 0–3 y.

Targets: the top-15 highlighted roles ranked by 50% skill overlap (cosine similarity of skill-share vectors × 100) and 50% Highlight score. Learn: the largest gaps in skill share (target vs start). A model of skills in postings, not of hiring decisions.

YOUR NEXT MOVE · INTERPRETATION FROM SKILLS DATA

Your next move, continued

For common starting roles: the highlighted roles closest in skills, the bridge skills to add, and the evidence.

IF YOU ARE A...

SAP

Now: 31.0 new postings a day, momentum +16%, Highlight rank 15.

→ **Software engineer (general)** (#1, skill overlap 16%)

Learn: CI/CD (32% vs 0%), Python (31% vs 0%), REST APIs (45% vs 18%).

Already shared: REST APIs, Agile. Evidence: 103.6/day, momentum -4%, 31% ask 0–3 y.

→ **QA & manual testing** (#3, skill overlap 16%)

Learn: Regression (43% vs 2%), Test Automation (35% vs 0%), CI/CD (32% vs 0%). Already shared: REST APIs, Agile. Evidence: 24.7/day, momentum -1%, 32% ask 0–3 y.

→ **Backend developer** (#4, skill overlap 15%)

Learn: REST APIs (65% vs 18%), CI/CD (46% vs 0%), Java (48% vs 4%).

Already shared: REST APIs, Agile. Evidence: 42.6/day, momentum -4%, 18% ask 0–3 y.

IF YOU ARE A...

Oracle, ERP & other platforms

Now: 31.1 new postings a day, momentum +20%, Highlight rank 9.

→ **Software engineer (general)** (#1, skill overlap 45%)

Learn: Python (31% vs 0%), Java (31% vs 0%), CI/CD (32% vs 10%). Already shared: REST APIs, Agile, SQL. Evidence: 103.6/day, momentum -4%, 31% ask 0–3 y.

→ **QA & manual testing** (#3, skill overlap 41%)

Learn: Regression (43% vs 3%), Test Automation (35% vs 0%), Selenium (28% vs 0%). Already shared: REST APIs, Agile, SQL. Evidence: 24.7/day, momentum -1%, 32% ask 0–3 y.

→ **Backend developer** (#4, skill overlap 40%)

Learn: Java (48% vs 0%), Microservices (41% vs 0%), CI/CD (46% vs 10%).

Already shared: REST APIs, SQL, Agile. Evidence: 42.6/day, momentum -4%, 18% ask 0–3 y.

IF YOU ARE A...

Test automation & SDET

Now: 10.4 new postings a day, momentum -9%, Highlight rank 22.

→ **Software engineer (general)** (#1, skill overlap 66%)

Learn: React (17% vs 0%), Distributed Systems (14% vs 0%), C++ (13% vs 0%). Already shared: REST APIs, Agile, CI/CD. Evidence: 103.6/day, momentum -4%, 31% ask 0–3 y.

→ **QA & manual testing** (#3, skill overlap 94%)

Learn: UAT (12% vs 0%), Defect Tracking (11% vs 0%), End-to-end Testing (10% vs 0%). Already shared: REST APIs, Agile, Regression. Evidence: 24.7/day, momentum -1%, 32% ask 0–3 y.

→ **Backend developer** (#4, skill overlap 61%)

Learn: Spring (32% vs 0%), Microservices (41% vs 10%), Spring Boot (26% vs 0%). Already shared: REST APIs, Java, CI/CD. Evidence: 42.6/day, momentum -4%, 18% ask 0–3 y.

IF YOU ARE A...

Embedded & firmware

Now: 7.0 new postings a day, momentum -19%, Highlight rank 30.

→ **Software engineer (general)** (#1, skill overlap 24%)

Learn: REST APIs (45% vs 6%), Java (31% vs 0%), SQL (30% vs 0%). Already shared: Python. Evidence: 103.6/day, momentum -4%, 31% ask 0–3 y.

→ **Backend developer** (#4, skill overlap 19%)

Learn: REST APIs (65% vs 6%), Java (48% vs 0%), SQL (43% vs 0%). Already shared: Python. Evidence: 42.6/day, momentum -4%, 18% ask 0–3 y.

→ **Infrastructure, network & IT** (#7, skill overlap 26%)

Learn: Windows (22% vs 0%), Azure (27% vs 5%), AWS (20% vs 4%).

Already shared: Python, Linux. Evidence: 27.8/day, momentum +1%, 32% ask 0–3 y.

Targets: the top-15 highlighted roles ranked by 50% skill overlap (cosine similarity of skill-share vectors × 100) and 50% Highlight score. Learn: the largest gaps in skill share (target vs start). A model of skills in postings, not of hiring decisions.



IT POSTINGS PER DAY, JULY VS JUNE

-3.9%

815 IT postings a day in July, against 848 in June

...in a month when our crawl missed 8 days

THE HEADLINE

The month in numbers

IT hiring eased slightly in July. The month is harder to read than most: on 8 days our crawl captured nothing new.

815

IT postings a day in July (-3.9% on June)

8 days


with no new postings captured (3, 10, 11, 12, 16, 18, 28, 30 July)

19%

were reposts or caught late: first seen before July (June: 15%)

41%


of all postings in our crawl were IT roles (June: 42%)



Bars: IT postings per day by posted date, July. Ticks: the same day of June. Line: IT postings first captured by our crawl that day; x marks a day with no new captures at all.

July's set holds 25,259 IT postings: 815 a day against 848 in June (-3.9%). Non-IT postings moved -0.7% a day. We count a posting by its posted date or the day we first captured it, so jobs posted on the 8 gap days still count when the crawl caught up. The by-date bars stay full on most gap days; the capture line drops to zero. The gaps do move one figure: postings first captured per day fell -8.4%, and the repost share rose to 18.9%. Some of those "reposts" are

July jobs we simply caught late. The busiest posting day was 31 July, with 1,336.



New = first captured in the month. Reposted = posted in the month but first seen earlier. With 8 gap days, July's new count is understated; read it with care.


WHAT THIS MEANS FOR YOU • INTERPRETATION

Demand held roughly level going into the second half of the year: a -3.9% dip in a month with a patchy crawl is within the noise. Interpretation: do not read July as a slowdown; August (next edition) is the cleaner test.

SKILLS

Skills rising and falling

Data warehousing, pipelines and observability gained most in July; embedded and IoT skills lost ground.



Change in each skill's share of IT postings, July vs June, relative to the typical skill: June postings listed fewer skills each, so every raw share rose (+54% for the median skill).

Data warehouses and pipelines: the month's biggest gainers.

June's postings listed fewer skills each, so the raw shares jump: cloud is named in 46.8% of July's IT postings that list skills, against 30.2% in June. We therefore rank skills by their change relative to the typical skill. On that measure data platforms lead: data warehouse +93%, Azure Data Factory +90%, data products +64%. Observability tools followed: CloudWatch +64%, Datadog +61%. Serverless rose +78%. Fading: enterprise software -71%, embedded systems -67%, IoT -54%.

Top IT skills	Ju1	Jun
Cloud	46.8%	30.2%
Security	39.7%	25.8%
REST APIs	37.4%	22.5%
Python	36.0%	25.4%
CI/CD	28.9%	18.6%
SQL	28.3%	18.7%
AWS	27.4%	19.2%
Agile	27.1%	17.1%
Azure	27.0%	18.0%
Analytics	23.4%	16.7%
DevOps	21.1%	15.3%
Java	18.4%	14.5%

Share of IT postings naming the skill. Soft skills left out.

WHAT THIS MEANS FOR YOU · INTERPRETATION

If you work in data, July rewarded the warehouse-and-pipeline stack (Data Factory, data products) and the tools that watch it run. Embedded and IoT skills cooled, in step with fewer embedded roles (p. 32). Treat one month as a hint, not a trend: compare with August's edition.

SKILLS ECONOMICS · COMBINATIONS

Skills that come in packs

Skills travel in packs. These are the combinations IT employers name together most often, family by family.

Software


6,387 postings; share of them naming both skills



Common trio: CI/CD + Git + REST APIs (16%).

Data & AI

5,600 postings; share of them naming both skills



Common trio: AWS + Azure + Python (19%).

Cloud & DevOps


2,290 postings; share of them naming both skills



Common trio: CI/CD + DevOps + Kubernetes (23%).

QA & testing


1,089 postings; share of them naming both skills



Common trio: CI/CD + Regression + REST APIs (25%).

Enterprise apps

2,304 postings; share of them naming both skills



Common trio: Agile + REST APIs + SOAP (8%).

Security

1,198 postings; share of them naming both skills



Common trio: AWS + Azure + GCP (12%).

Chip & hardware

847 postings; share of them naming both skills



Common trio: Perl + Python + SoC (10%).

Product & design

781 postings; share of them naming both skills




Common trio: Agile + REST APIs + Scrum (6%).

SKILLS · RADAR · INTERPRETATION

The skills radar

Which skills to adopt, trial, watch or let fade, from 20,294 IT postings in July 2026: share now against momentum.



Adopt	Share	Mom.	Trial	Share	Mom.	Watch	Share	Mom.	Fading	Share	Mom.
REST APIs	37.5%	+7%	Bedrock	2.1%	+39%	JWT	1.6%	+58%	IoT	1.7%	-40%
Python	35.5%	-5%	Integration testing	3.3%	+32%	ETL/ELT pipelines	1.3%	+57%	Software architecture	1.7%	-39%
CI/CD	28.9%	+1%	Containers	2.4%	+31%	Microsoft Fabric	1.0%	+55%	Computer vision	1.3%	-32%
SQL	28.6%	+3%	Relational databases	3.3%	+28%	AWS SageMaker	1.4%	+51%	Business intelligence	2.8%	-31%
Azure	27.0%	-2%	API integration	3.3%	+28%	SOAP APIs	1.3%	+50%	Frontend development	1.1%	-29%
AWS	27.0%	-5%	Data validation	2.8%	+26%	Azure OpenAI	1.7%	+49%	Cloud native	1.3%	-28%
Agile	26.6%	+3%	Query optimization	2.2%	+26%	Synapse	1.2%	+49%	System architecture	1.8%	-27%

Thresholds. Share = postings naming the skill ÷ IT postings listing any skill, July. Momentum = July against 16–30 June (provisional: earlier skill lists were shorter; shares normalised); the change in share, divided by the typical skill's change (+25%) so longer or shorter skill lists do not move every skill. Adopt: share ≥ 8% and momentum above +10%. Trial: share 2–8% and momentum ≥ +10%. Watch: 0.4–2%, momentum ≥ +20%. Fading: share ≥ 1% and momentum ≤ -15%. At least 40 postings now; up to 7 per ring (Adopt by share, the others by momentum). Soft skills and catch-all tags left out.

Provisional: June listed fewer skills a posting (20.7 against 24.7).

HOW TO READ IT · INTERPRETATION

Adopt is the core almost every IT team lists: keep it current. **Trial** skills are mid-sized and growing: a good next skill to add. **Watch** skills are small but rising fast: worth a weekend, not a career bet. **Fading** skills are losing share, not disappearing: keep them if you have them, but do not lead with them.


SKILLS ECONOMICS · THE MINIMUM STACK

The minimum stack

The minimum stack: skills named in at least half of a role’s postings that list skills. And the skills that separate senior postings from junior ones.


Role	Minimum stack (in ≥50% of postings)	Next most named
Software eng.	— (none reaches half)	REST APIs, Agile, CI/CD, Python, Java, S
ML/AI engineer	Python 63%	LLMs, REST APIs, Generative AI, Azure, M
QA & testing	— (none reaches half)	REST APIs, Agile, Regression, Test Autom
Backend	REST APIs 65%	Java, CI/CD, SQL, Microservices, AWS, Ag
Data/BI analyst	SQL 64%, Dashboards 52%	Power BI, Python, Excel, Data Analytics,
Data engineer	SQL 77%, Python 68%, ETL 64%, Data Pipelines 63%	AWS, Azure, Spark
Infra & network	— (none reaches half)	Python, Azure, Linux, Windows, AWS, TCP,
Full-stack	REST APIs 81%, React 64%, CI/CD 59%	JavaScript, Agile, Microservices, AWS, S
Oracle/ERP	— (none reaches half)	REST APIs, Oracle, ServiceNow, Agile, SO
Security eng.	— (none reaches half)	Cybersecurity, Python, Cloud Security, I
GenAI/agents	LLMs 74%, Generative AI 73%, Python 72%, REST APIs 68%	LangChain, Azure
Data scientist	Python 84%, Data Science 72%, Machine Learning 69%, Statis	LLMs, Algorithms, Generative AI
DevOps	DevOps 92%, CI/CD 82%, Kubernetes 73%, Python 64%	
Cloud engineer	— (none reaches half)	AWS, Azure, CI/CD, REST APIs, DevOps, Py

SKILLS THAT SIGNAL SENIORITY



Share of levelled IT postings naming the skill: junior = intern and junior/associate (909 postings), senior = senior and staff/lead/principal individual contributors (8,863); level ladder of the org pages. Ranked by the senior/junior ratio.

SKILLS COMMON IN JUNIOR POSTINGS



WHAT THIS MEANS FOR YOU · INTERPRETATION

Interpretation. Cover the minimum stack before you apply; it is what most postings for the role assume. The senior list shows what employers add as you climb: architecture, governance and platform work rather than more languages.

SKILLS · CHECKLIST · JULY

Skill-gap checklist

For each role to highlight: the skills most postings ask for, with shares. Tick what you can prove, then count.

Software engineer (general)

#1 to highlight · 3,176 postings with skills
Score yourself: __ / 8 nice-to-have
 Also common (15–25%): DevOps 23%, JavaScript 21%, Microservices 18%

MINIMUM STACK · named in ≥50% of postings

No single skill is named in half of this role's postings: the asks are spread out.

NICE TO HAVE · 25–50%

- | | | |
|--|------------------------------------|------------------------------------|
| ☐ REST APIs 45% | ☐ Agile 36% | ☐ CI/CD 32% |
| ☐ Python 31% | ☐ Java 31% | ☐ SQL 30% |
| ☐ AWS 28% | ☐ Azure 26% | |

ML & AI engineer

#2 to highlight · 1,868 postings with skills
Score yourself: __ / 1 minimum · __ / 9 nice-to-have
 Also common (15–25%): GCP 25%, SQL 24%, MLOps 23%

MINIMUM STACK · named in ≥50% of postings

- ☐ Python 63%

NICE TO HAVE · 25–50%

- | | | |
|------------------------------------|---|--|
| ☐ LLMs 50% | ☐ REST APIs 43% | ☐ Generative AI 39% |
| ☐ Azure 39% | ☐ Machine Learning 37% | ☐ AWS 36% |
| ☐ RAG 33% | ☐ CI/CD 31% | ☐ Orchestration 27% |

QA & manual testing

#3 to highlight · 759 postings with skills
Score yourself: __ / 9 nice-to-have
 Also common (15–25%): Python 25%, Java 23%, DevOps 22%

MINIMUM STACK · named in ≥50% of postings

No single skill is named in half of this role's postings: the asks are spread out.

NICE TO HAVE · 25–50%

- | | | |
|--|--|---|
| ☐ REST APIs 46% | ☐ Agile 44% | ☐ Regression 43% |
| ☐ Test Automation 35% | ☐ CI/CD 32% | ☐ Jira 29% |
| ☐ Selenium 28% | ☐ API testing 28% | ☐ SQL 28% |

Backend developer

#4 to highlight · 1,295 postings with skills
Score yourself: __ / 1 minimum · __ / 9 nice-to-have
 Also common (15–25%): PostgreSQL 22%, Kafka 22%, GCP 20%

MINIMUM STACK · named in ≥50% of postings

- ☐ REST APIs 65%

NICE TO HAVE · 25–50%

- | | | |
|--|-------------------------------------|-------------------------------------|
| ☐ Java 48% | ☐ CI/CD 46% | ☐ SQL 43% |
| ☐ Microservices 41% | ☐ AWS 40% | ☐ Agile 38% |
| ☐ Python 35% | ☐ Docker 32% | ☐ DevOps 30% |

Data & BI analyst

#5 to highlight · 1,027 postings with skills
Score yourself: __ / 2 minimum · __ / 3 nice-to-have
 Also common (15–25%): Data Quality 24%, Statistics 21%, Agile 16%

MINIMUM STACK · named in ≥50% of postings

- ☐ SQL 64%

NICE TO HAVE · 25–50%

- | | | |
|---------------------------------------|---|--|
| ☐ Power BI 43% | ☐ Dashboards 52% | |
| ☐ Python 39% | ☐ Tableau 32% | |

HOW TO SCORE · INTERPRETATION

Tick a box only if you could show the skill: a project, a repository, work you can talk through in an interview. Minimum boxes are what at least half of this role's postings name; ticking all of them means you meet the most common ask, not every employer's.

Nice-to-haves sort shortlists. If you tick fewer than half, pick two of them and use the 30/60/90 roadmaps on ???. Shares are of postings for the role in July 2026 that list any skill.

SKILLS • CHECKLIST • JULY

Skill-gap checklist: part 2

Roles 6–10. The same rules: minimum stack first, then the skills that sort a shortlist.

Data engineer #6 to highlight • 1,248 postings with skills Score yourself: __ / 5 minimum • __ / 9 nice-to-have Also common (15–25%): DevOps 25%, REST APIs 25%, Git 21%	MINIMUM STACK • named in ≥50% of postings ☐ SQL 77% ☐ Data Pipelines 63% NICE TO HAVE • 25–50% ☐ AWS 45% ☐ CI/CD 40% ☐ Data Modeling 32% ☐ Python 68% ☐ Data Quality 53% ☐ Azure 45% ☐ ELT 38% ☐ PySpark 32% ☐ ETL 64% ☐ Spark 40% ☐ Databricks 37% ☐ Data Governance 30%
---	--

HOW WE BUILT THE LISTS

Ten roles to highlight (ranked on ??). For each, the share of its postings in July 2026 whose skills list names each technical skill. Soft skills (communication, problem solving) and vague tags (cloud, security, AI tools) are left out; spelling variants are merged (GitHub with Git, LLM with LLMs). The same skills by family are on the minimum-stack page, ??.

No salary data exists in our crawl; nothing here is a statement about pay.

SKILLS · CV KEYWORDS · JULY

CV keyword guide

Which terms to put on a CV for each role to highlight: the skills tags and title words postings use most.

Software engineer (general) · skills tags and titles (3,212 postings; job text on 171)

Skills tags: REST APIs 45% · Agile 36% · CI/CD 32% · Python 31% · Java 31% · SQL 30% · AWS 28% · Azure 26% · DevOps 23% · JavaScript 21% · Microservices 18% · Git 18% · Kubernetes 17% · React 17%

Title words: “application” 5% · “programmer” 3% · “application engineer” 3% · “software developer” 2%

ML & AI engineer · skills tags and titles (1,886 postings; job text on 92)

Skills tags: Python 63% · LLMs 50% · REST APIs 43% · Generative AI 39% · Azure 39% · Machine Learning 37% · AWS 36% · RAG 33% · CI/CD 31% · Orchestration 27% · GCP 25% · SQL 24% · MLOps 23% · Prompt Engineering 22%

Title words: “ai” 72% · “ai engineer” 14% · “ml” 8% · “machine learning” 8% · “ai/ml” 7% · “artificial intelligence” 4%

QA & manual testing · skills tags and titles (767 postings; job text on 39)

Skills tags: REST APIs 46% · Agile 44% · Regression 43% · Test Automation 35% · CI/CD 32% · Jira 29% · Selenium 28% · API testing 28% · SQL 28% · Python 25% · Java 23% · DevOps 22% · Postman 22% · Software Testing 22%

Title words: “qa” 33% · “quality assurance” 24% · “test engineer” 17% · “testing” 16% · “qa engineer” 15% · “automation” 9%

Backend developer · skills tags and titles (1,320 postings; job text on 61)

Skills tags: REST APIs 65% · Java 48% · CI/CD 46% · SQL 43% · Microservices 41% · AWS 40% · Agile 38% · Python 35% · Docker 32% · DevOps 30% · Kubernetes 30% · Git 30% · Azure 28% · Spring Boot 26%

Title words: “java” 35% · “python” 22% · “backend” 21% · “java developer” 9% · “.net” 8% · “backend developer” 7%

Data & BI analyst · skills tags and titles (1,036 postings; job text on 40)

Skills tags: SQL 64% · Dashboards 52% · Power BI 43% · Python 39% · Tableau 32% · Data Quality 24% · Statistics 21% · Agile 16% · Data Governance 14% · Data Modeling 14% · ETL 13% · Azure 13% · KPIs 12% · Data Accuracy 11%

Title words: “data analyst” 32% · “bi” 15% · “power bi” 9% · “business intelligence” 6% · “sql” 5% · “reporting” 5%

Data engineer · skills tags and titles (1,258 postings; job text on 57)

Skills tags: SQL 77% · Python 68% · ETL 64% · Data Pipelines 63% · Data Quality 53% · AWS 45% · Azure 45% · Spark 40% · CI/CD 40% · ELT 38% · Databricks 37% · Data Modeling 32% · PySpark 32% · Data Governance 30%

Title words: “databricks” 9% · “snowflake” 6% · “etl” 5% · “aws” 5% · “azure data” 5% · “gcp data” 4%

Infrastructure, network & IT · skills tags and titles (861 postings; job text on 41)

Skills tags: Networking 45% · Python 29% · Azure 27% · Linux 26% · Windows 22% · AWS 20% · ITIL 16% · DNS 16% · Root Cause Analysis 15% · High Availability 15% · BGP 15% · TCP/IP 15% · Ansible 14% · CI/CD 13%

Title words: “network” 34% · “infrastructure” 16% · “network engineer” 13% · “linux” 7% · “networking” 5% · “storage” 5%

Full-stack developer · skills tags and titles (767 postings; job text on 44)

Skills tags: REST APIs 81% · React 64% · CI/CD 59% · JavaScript 50% · Agile 46% · Microservices 45% · AWS 44% · SQL 44% · DevOps 44% · Java 43% · TypeScript 42% · Docker 41% · Azure 39% · Git 38%

Title words: “full stack” 66% · “fullstack” 22% · “java” 16% · “.net” 11% · “react” 8% · “fullstack developer” 7%

Oracle, ERP & other platforms · skills tags and titles (963 postings; job text on 41)

Skills tags: REST APIs 43% · Oracle 42% · ServiceNow 28% · Agile 28% · SOAP 25% · Workday 20% · SQL 19% · HCM 19% · ITSM 19% · JavaScript 18% · ERP 18% · UAT 16% · Business Rules 15% · Solution Design 15%

Title words: “oracle” 36% · “servicenow” 26% · “workday” 18% ·

“functional” 10% · “servicenow developer” 8% · “erp” 7%

Security engineering (general) · skills tags and titles (690 postings; job text on 43)

Skills tags: Cybersecurity 37% · Python 27% · Cloud Security 25% · Incident Response 23% · SIEM 21% · Azure 21% · AWS 21% · CISSP 20% · Networking 20% · REST APIs 18% · Linux 15% · SOC 15% · GCP 15% · Risk Management 14%

Title words: “security” 70% · “security engineer” 19% · “cybersecurity” 10% · “security analyst” 4% · “management” 3% · “security consultant” 3%

USE THESE HONESTLY · EDITORIAL

Use only the words that describe work you have actually done, and say where you did it. Screening software may match words; the interviewer will ask about each one. A keyword you cannot defend costs more than it gains.

Mirror the posting’s exact spelling (CI/CD, PySpark), put the strongest terms in your summary and latest role, and do not paste lists of keywords in white text or hidden fields.

Job text exists for only 629 of these roles’ postings in July 2026 (under 200 for every role), so this page uses skills tags (share of postings listing skills) and job-title words (share of the role’s postings) instead of description text.

SKILLS · ROADMAPS · EDITORIAL

30/60/90-day roadmaps

Six roles to highlight, each split into three months: foundations first, then the core stack, then what senior postings ask for.

Editorial: a suggested path, not a guarantee.

Software engineer (general)

Median minimum experience asked: 5 years · 31% of postings ask 0–3 years · CI/CD + REST APIs are named together in 29% of Software engineering postings: learn them side by side.

DAYS 1–30 · FOUNDATION

Learn: REST APIs 45%, Agile 36%

Do: Build a small REST service with tests and push it to a Git repository with a CI check.

Free: MDN Web Docs: An overview of HTTP; Microsoft REST API Guidelines; The Scrum Guide (free)

DAYS 31–60 · CORE

Learn: CI/CD 32%, Python 31%, Java 31%, SQL 30%

Do: Package it in a container, deploy it to a free cloud tier, add logging and a dashboard.

Free: GitHub Actions documentation; GitLab CI/CD documentation; The Python Tutorial (official docs)

DAYS 61–90 · DIFFERENTIATORS

Learn: Cloud-native 13% (×7.8 senior), Observability 14% (×6.3 senior), Distributed Systems 14% (×6.1 senior)

Do: Split one part into a second service; write a one-page design note on trade-offs and failure modes.

Free: CNCF Cloud Native Landscape and glossary; OpenTelemetry documentation; Google SRE book (free online)

ML & AI engineer

Median minimum experience asked: 5 years · 28% of postings ask 0–3 years · Python + SQL are named together in 35% of Data & AI postings: learn them side by side.

DAYS 1–30 · FOUNDATION

Learn: Python 63%, LLMs 50%, REST APIs 43%

Do: Train and evaluate a classic model in Python on an open dataset; keep notebooks and code in Git.

Free: The Python Tutorial (official docs); Automate the Boring Stuff with Python (free online book); Hugging Face LLM Course (free)

DAYS 31–60 · CORE

Learn: Generative AI 39%, Azure 39%, Machine Learning 37%, AWS 36%

Do: Build a retrieval-augmented question-answering demo over a document set, with an evaluation set.

Free: Hugging Face LLM Course (free); DeepLearning.AI short courses (free); Microsoft Learn: Azure Fundamentals (AZ-900) learning path

DAYS 61–90 · DIFFERENTIATORS

Learn: Observability 19% (×6.3 senior), DevOps 18% (×4.5 senior), Kubernetes 19% (×4.5 senior)

Do: Wrap it as a service with monitoring, cost and latency numbers, and a written note on failure cases.

Free: OpenTelemetry documentation; Google SRE book (free online); DORA State of DevOps reports (free)

QA & manual testing

Median minimum experience asked: 5 years · 32% of postings ask 0–3 years · CI/CD + REST APIs are named together in 36% of QA & testing postings: learn them side by side.

DAYS 1–30 · FOUNDATION

Learn: REST APIs 46%, Agile 44%, Regression 43%

Do: Write test cases for a public web app from its requirements; log defects clearly.

Free: MDN Web Docs: An overview of HTTP; Microsoft REST API Guidelines; The Scrum Guide (free)

DAYS 31–60 · CORE

Learn: Test Automation 35%, Jira 29%, Selenium 28%, API testing 28%

Do: Automate a smoke suite for the same app and run it in a CI pipeline.

Free: Playwright documentation; Selenium documentation; Atlassian Jira guides

DAYS 61–90 · DIFFERENTIATORS

Learn: DevOps 22% (×4.5 senior), CI/CD 32% (×4.0 senior), AWS 12% (×3.0 senior)

Do: Add API tests and a simple performance check; report coverage and risk in one page.

Free: DORA State of DevOps reports (free); Google SRE book (free online); GitHub Actions documentation

How the order is set: Foundation = the role's most-named skills (in 35% or more of its postings in July 2026, or else its top two). Core = the next most-named skills (15% or more). Differentiators = skills the role's postings name less often (under 35%) but that senior postings across IT ask for far more than junior ones (×1.3 or more: share in senior individual-contributor postings ÷ share in junior postings). Percentages are shares of the role's postings listing skills. "Learn together" pairs come from skill co-occurrence in the

role's family. Projects and resources are the editors' suggestions. Resources are named, not linked: all are free to read or audit at the time of writing; search for them by name. Ninety days assumes steady part-time study; go faster or slower as you need. Nothing here promises a job or an interview.

SKILLS · ROADMAPS · EDITORIAL

30/60/90-day roadmaps: part 2

Three more roles, built the same way from this period's postings.

Editorial: a suggested path, not a guarantee.

Backend developer

Median minimum experience asked: 5 years · 18% of postings ask 0–3 years · CI/CD + REST APIs are named together in 29% of Software engineering postings: learn them side by side.

DAYS 1–30 · FOUNDATION

Learn: REST APIs 65%, Java 48%, CI/CD 46%

Do: Write a REST API with a relational database, input validation and unit tests.

Free: MDN Web Docs: An overview of HTTP; Microsoft REST API Guidelines; dev.java Learn (official Java tutorials)

DAYS 31–60 · CORE

Learn: SQL 43%, Microservices 41%, AWS 40%, Agile 38%

Do: Add a message queue between two services, containerise both and run them together.

Free: PostgreSQL documentation: the Tutorial chapter; SQLBolt interactive lessons; microservices.io pattern catalogue (Chris Richardson)

DAYS 61–90 · DIFFERENTIATORS

Learn: Kafka 22% (×6.1 senior), Distributed Systems 16% (×6.1 senior), Spring Boot 26% (×4.8 senior)

Do: Load-test the API, add caching and tracing, and write up what changed and why.

Free: Apache Kafka Quickstart (official docs); MIT 6.824 / 6.5840 Distributed Systems lectures (free); The System Design Primer (free, open source)

Data & BI analyst

Median minimum experience asked: 5 years · 34% of postings ask 0–3 years · Python + SQL are named together in 35% of Data & AI postings: learn them side by side.

DAYS 1–30 · FOUNDATION

Learn: SQL 64%, Dashboards 52%, Power BI 43%

Do: Answer five business questions on an open dataset with SQL; keep the queries tidy and commented.

Free: PostgreSQL documentation: the Tutorial chapter; SQLBolt interactive lessons; Fundamentals of Data Visualization (Claus Wilke, free online)

DAYS 31–60 · CORE

Learn: Python 39%, Tableau 32%, Data Quality 24%, Statistics 21%

Do: Build a dashboard with a clear headline metric, filters and a one-page written summary.

Free: The Python Tutorial (official docs); Automate the Boring Stuff with Python (free online book); Tableau free training videos

DAYS 61–90 · DIFFERENTIATORS

Learn: Snowflake 11% (×4.3 senior), Databricks 9% (×3.6 senior), ETL 13% (×3.1 senior)

Do: Add a simple statistical test or forecast and explain its uncertainty in plain words.

Free: Snowflake documentation: getting-started tutorials; Databricks documentation: tutorials; Data Engineering Zoomcamp (DataTalks.Club, free)

Infrastructure, network & IT

Median minimum experience asked: 5 years · 32% of postings ask 0–3 years · CI/CD + Python are named together in 26% of Cloud, DevOps & SRE postings: learn them side by side.

DAYS 1–30 · FOUNDATION

Learn: Networking 45%, Python 29%

Do: Build a home lab: a Linux server, DNS, DHCP and a firewall; document the network.

Free: Cisco Networking Academy free courses; The Python Tutorial (official docs); Automate the Boring Stuff with Python (free online book)

DAYS 31–60 · CORE

Learn: Azure 27%, Linux 26%, Windows 22%, AWS 20%

Do: Automate the setup with scripts or Ansible; add monitoring and alerting.

Free: Microsoft Learn: Azure Fundamentals (AZ-900) learning path; The Linux Command Line (William Shotts, free book); Linux Journey (free)

DAYS 61–90 · DIFFERENTIATORS

Learn: DevOps 11% (×4.5 senior), Kubernetes 11% (×4.5 senior), CI/CD 13% (×4.0 senior)

Do: Move one service to a cloud account with identity controls and a backup/restore test.

Free: DORA State of DevOps reports (free); Google SRE book (free online); Kubernetes Basics tutorial (official docs)

How the order is set: Foundation = the role's most-named skills (in 35% or more of its postings in July 2026, or else its top two). Core = the next most-named skills (15% or more). Differentiators = skills the role's postings name less often (under 35%) but that senior postings across IT ask for far more than junior ones (×1.3 or more: share in senior individual-contributor postings ÷ share in junior postings). Percentages are shares of the role's postings listing skills. "Learn together" pairs come from skill co-occurrence in the

role's family. Projects and resources are the editors' suggestions. Resources are named, not linked: all are free to read or audit at the time of writing; search for them by name. Ninety days assumes steady part-time study; go faster or slower as you need. Nothing here promises a job or an interview.

EMPLOYERS

Who's hiring IT


No employer dominates IT hiring: the ten biggest posters hold 7.3% of IT postings, and 3 of the top 15 are staffing firms.

#	Company	Jul/day	Jun/day	Change
1	Zorba Consulting India †	7.1	5.7	+24%
2	Oracle	6.9	7.7	-10%
3	Innova ESI †	6.1	5.8	+5%
4	Zensar	6.0	4.6	+31%
5	Google	5.8	5.5	+5%
6	Husky Technologies	5.7	11.9	-52%
7	Capgemini	5.5	5.5	-0%
8	EPAM Systems	5.5	5.9	-7%
9	YASH Technologies	5.4	5.4	-1%
10	Tata Consultancy Services	5.3	5.3	+0%
11	TalentGigs †	5.1	6.9	-26%
12	GlobalLogic	5.1	4.6	+12%
13	UST Global	4.9	5.6	-12%
14	IBM	4.7	5.5	-14%
15	Mastercard	4.5	4.4	+3%

A staffing firm, Zorba Consulting India, posted the most IT roles in July: 7.1 a day. Among direct employers Oracle led with 6.9, down -10%. Zensar rose 31%. Among employers with at least 40 IT postings, Persistent Systems moved fastest (1.1 to 3.7 a day), then Salesforce and S&P Global. Husky Technologies halved its pace (-52%).

Per day = IT postings in the month ÷ days. † staffing firm, recruiter or job platform (posts its clients' roles). Names tidied from the job platform's.

BIGGEST MOVERS: IT POSTINGS PER WEEK (≥40 THIS MONTH)



WHAT THIS MEANS FOR YOU · INTERPRETATION

Product companies and banks were adding IT postings fastest (Salesforce +95%, S&P Global +91% a day). A staffing-firm posting is a real job at an unnamed client: ask who the employer is before you apply. Every company's IT posting count is in the References, p. 51.

COMPANY × ROLE · WHO HIRES EACH HIGHLIGHTED ROLE

Who hires for each role

For each highlighted role, the employers posting it most, and whether a few firms dominate.

Software eng.

top 5 = 10% · top 10 = 16% · ~188 effective employers

Employer	Share
Wipro	2.6%
Microsoft	2.3%
Capgemini	1.8%
Oracle	1.6%
Dexter's Tech	1.6%
Nike	1.5%

ML/AI engineer

top 5 = 9% · top 10 = 14% · ~219 effective employers

Employer	Share
Teradata	2.1%
Alignerr †	1.8%
TalentGigs †	1.8%
Zorba Consulting India †	1.5%
Smartworkingsolutions †	1.5%
IBM	1.2%

QA & testing

top 5 = 11% · top 10 = 16% · ~154 effective employers

Employer	Share
Birlasoft	3.1%
VOIS	2.6%
Rockwell Automation	2.3%
TalentGigs †	1.6%
Agileengine	1.3%
Zensar	1.2%

Backend

top 5 = 9% · top 10 = 16% · ~160 effective employers

Employer	Share
TalentGigs †	2.1%
Turingcom	1.7%
GlobalLogic	1.7%
Infosys	1.7%
Iris Software	1.6%
Agoda	1.6%

Data/BI analyst

top 5 = 12% · top 10 = 19% · ~144 effective employers

Employer	Share
Argo Intern	2.9%
Wake Up Whistle	2.7%
Dexter's Tech	2.5%
Synchrony Financial	2.3%
HSBC	1.7%
Citi	1.6%

Data engineer

top 5 = 14% · top 10 = 21% · ~112 effective employers

Employer	Share
Zorba Consulting India †	5.6%
Innova ESI †	3.4%
Cisco	1.9%
Tata Consultancy Servi	1.7%
Husky Technologies	1.7%
IBM	1.6%

Infra & network

top 5 = 11% · top 10 = 18% · ~134 effective employers

Employer	Share
Tata Communication	3.3%
NTT Data	2.3%
Capgemini	2.0%
Amazon	1.9%
Hewlett Packard Enterp	1.9%
Infosys	1.7%

Full-stack

top 5 = 12% · top 10 = 20% · ~128 effective employers

Employer	Share
Bnyglobal	3.3%
Husky Technologies	2.6%
Smartworkingsolutions †	2.6%
Iris Software	2.1%
Programmers Io	2.0%
Innova ESI †	2.0%

Oracle/ERP

top 5 = 16% · top 10 = 25% · ~92 effective employers

Employer	Share
Oracle	4.8%
Teamware Solutions †	3.1%
SourceBae †	3.0%
Tata Consultancy Servi	2.6%
Nagarro	2.5%
Huronconsulting	2.4%

Security eng.

top 5 = 14% · top 10 = 20% · ~119 effective employers

Employer	Share
Wells Fargo	4.5%
NTT Data	3.0%
Verizon	2.5%
FIS	2.3%
Knowdis	2.0%
BT	1.4%

GenAI/agents

top 5 = 12% · top 10 = 21% · ~105 effective employers

Employer	Share
Knowdis	2.6%
Iris Software	2.4%
PwC India	2.4%
Innova ESI †	2.2%
Orioninnovation	2.2%
Infosys	2.0%

Data scientist

top 5 = 12% · top 10 = 18% · ~125 effective employers

Employer	Share
Zorba Consulting India †	3.6%
Faceprep	3.2%
Amgen	1.8%
D Square Consulting Se	1.8%
Tata Consultancy Servi	1.4%
Mastercard	1.3%

Share of the role's new IT postings in the period. Effective employers = $1 \div (\text{sum of squared employer shares})$ a month: how many equal-sized employers would give the same concentration. † staffing firm or job platform.

SPECULATION · INFERRED FROM PUBLIC JOB POSTINGS · NOT STATEMENTS BY OR ABOUT THE COMPANIES’ PLANS

WHAT THEY MIGHT BE BUILDING · JULY 2026

Where the work is going

What might IT employers be building? We read the titles and skills of 20,497 new IT postings for domain clues, cohort by cohort and employer by employer. Job text covers only 5% of this month’s postings, so the evidence is lighter than in our Q4 issue. Our reading of public clues, not anyone’s stated plans.

CELL = % OF THE COHORT'S POSTINGS MENTIONING THE TERM; SHADE = HOW MUCH MORE THAN ALL IT POSTINGS (LIFT)

	IT services	Product & internet	Semiconductors	Telecom	BFSI tech	GCCs (non-tech)
RTL / SoC design	0.6	0.4	32.5	3.7	0.3	0.0
Embedded firmware / RTOS / drivers	2.5	1.5	29.6	5.4	0.1	5.0
Data quality / governance / lineage	12.5	7.8	3.1	8.4	17.3	16.1
GPU / inference optimisation	1.1	1.7	14.5	1.6	1.0	1.0
AI accelerators / GPU architecture	0.3	0.4	13.4	1.2	0.4	0.4
Capital markets / trading	1.3	0.3	0.2	0.2	11.9	0.1
LLMs / large language models	10.6	17.5	7.1	13.0	9.0	12.0
SAP modules (FICO, MM, SD, ABAP, BTP)	11.7	2.9	4.1	4.7	0.5	10.3
Event streaming (Kafka, Flink, real-time)	8.7	12.5	2.0	7.2	8.8	5.4
DFT / test (scan, ATPG, BIST)	0.4	0.2	10.8	0.9	0.0	0.1
Generative AI	13.7	15.2	6.8	9.7	13.8	14.8
Identity and access (IAM, SSO, PAM)	7.6	10.7	1.9	8.1	5.5	5.7
Responsible AI / AI governance	3.7	4.8	11.8	3.9	4.3	6.1
AI agents / agentic systems	9.3	15.2	9.3	8.8	8.5	12.1
Networking / SDN / NFV	1.0	1.4	1.1	8.6	1.0	0.5
Snowflake / cloud data warehouse	9.2	7.5	1.7	6.3	8.7	9.5

under 0.77x all IT

0.77-1.3x

1.3-2x

2x or more

IT services
3,676 IT postings · base 3,676
Postings suggest a tilt towards SAP systems, SAP S/4HANA programmes.
Signals: SAP modules 429 (11.7%, 1.5×); SAP S/4HANA 170 (4.6%, 1.7×).

Product & internet
1,889 IT postings · base 1,889
Postings suggest a tilt towards applications built on large language models, AI agents that automate multi-step work.
Signals: LLMs / large language models 331 (17.5%, 1.4×); AI agents / agentic systems 288 (15.2%, 1.4×); Event streaming 236 (12.5%, 1.5×).

Semiconductors
636 IT postings · base 636
Postings suggest a tilt towards chip (RTL and SoC) design, embedded firmware and device software.
Signals: RTL / SoC design 207 (32.5%, 20.2×); Embedded firmware / RTOS / drivers 188 (29.6%, 8.8×); AI accelerators / GPU architecture 85 (13.4%, 15.3×).

Telecom
569 IT postings · base 569
Postings suggest a tilt towards networking software, cloud security.
Signals: Networking / SDN / NFV 49 (8.6%, 7.1×); Cloud security / posture 42 (7.4%, 2.1×); SRE / incident management / SLOs 47 (8.3%, 1.5×).

BFSI tech
1,395 IT postings · base 1,395
Postings suggest a tilt towards capital-markets and trading technology, data quality, governance and lineage.
Signals: Capital markets / trading 166 (11.9%, 6.1×); Data quality / governance / lineage 242 (17.3%, 1.5×); Regulatory reporting 67 (4.8%, 6.1×).

GCCs (non-tech)
1,518 IT postings · base 1,518
Postings suggest a tilt towards data quality, governance and lineage, lakehouse data platforms (Databricks, Delta, Iceberg).
Signals: Data quality / governance / lineage 245 (16.1%, 1.4×); Lakehouse / Databricks / Delta / Iceberg 144 (9.5%, 1.6×); SAP modules 156 (10.3%, 1.3×).

HOW WE READ THE CLUES · METHOD
We ran a curated lexicon of 75 domain terms (data/lexicon_domains.json, v2026-10-22.1) over the titles and skills of 20,497 new IT postings (job text exists for only 5%, so it is not used for the cards), first captured in the period. A posting counts once per term. Lift = an employer's or cohort's share ÷ the share across all IT postings. A signal needs ≥5 postings from ≥3 distinct job titles (a vacancy re-posted in many cities counts once), ≥3% of the employer's postings and lift ≥1.5. Staffing firms and unclassified employers are left out. Readings are ours; regexes and every count are in the data appendix.

SPECULATION · INFERRED FROM PUBLIC JOB POSTINGS · NOT STATEMENTS BY OR ABOUT THE COMPANIES' PLANS

WHAT THEY MIGHT BE BUILDING · EMPLOYERS 1–5 OF 15

Signals, what they may mean, and the skills you would need

Zensar

IT services · 154 IT postings
Hiring most: Data engineer 12%,
Software engineer (general) 10%,
Backend developer 9%.

SIGNALS

Snowflake / cloud data warehouse 22 of 154
(14%), 10 titles, 1.8×
Event streaming (Kafka, Flink, real-time) 22 of
154 (14%), 14 titles, 1.8×
Mainframe / legacy modernisation 5 of 154 (3%),
5 titles, 4.1×
FinOps / cloud cost 7 of 154 (4%), 5 titles, 1.7×
Stack: Test automation frameworks 12%

WHAT IT MAY MEAN · SPECULATION

Postings suggest work on cloud data warehousing,
event-streaming and real-time data systems and
mainframe and legacy-system work.
Skills you'd need: ci/cd pipelines, database,
root cause analysis, oracle, analytics, scrum,
microservices.

Oracle

Product & internet · 150 IT postings
Hiring most: Software engineer
(general) 34%, Oracle, ERP & other
platforms 31%, ML & AI engineer 12%.
Re-posted, not counted: Cloud migration (7
postings, 2 titles).

SIGNALS

SRE / incident management / SLOs 32 of 150
(21%), 17 titles, 3.9×
Generative AI 31 of 150 (21%), 7 titles, 1.7×
Prompt engineering 23 of 150 (15%), 4 titles, 2.9×
Healthcare / clinical / life sciences 16 of 150
(11%), 6 titles, 5.8×
Stack: Observability 32%, Microservices 31%,
Infrastructure as code 19%

WHAT IT MAY MEAN · SPECULATION

The postings point to site reliability and incident
management, generative-AI features and prompt
design for AI features.
Skills you'd need: oracle, cloud solutions, data
infrastructure, oci, oracle cloud, cloud
infrastructure, oracle cloud infrastructure.

Tata Consultancy Services

IT services · 144 IT postings
Hiring most: Oracle, ERP & other
platforms 17%, Data engineer 15%,
SAP 11%.

SIGNALS

Lakehouse / Databricks / Delta / Iceberg 18 of
144 (12%), 15 titles, 2.0×
SAP modules (FICO, MM, SD, ABAP, BTP) 17 of
144 (12%), 14 titles, 1.5×
Mainframe / legacy modernisation 7 of 144 (5%),
5 titles, 6.2×
ServiceNow / ITSM 14 of 144 (10%), 14 titles, 1.6×

WHAT IT MAY MEAN · SPECULATION

It appears to be investing in lakehouse data platforms
(Databricks, Delta, Iceberg), SAP systems and
mainframe and legacy-system work.
Skills you'd need: sap, databricks, analytics,
gcp, etl, llms, data quality.

Google

Product & internet · 137 IT postings
Hiring most: Cloud engineer &
architect 22%, ML & AI engineer 15%,
Software engineer (general) 10%.

SIGNALS

NLP / speech 31 of 137 (23%), 18 titles, 4.9×
MLOps / LLMops 21 of 137 (15%), 18 titles, 2.0×
RTL / SoC design 6 of 137 (4%), 6 titles, 2.7×
Networking / SDN / NFV 5 of 137 (4%), 5 titles,
3.0×

WHAT IT MAY MEAN · SPECULATION

Likely areas of work: language and speech technology,
platforms that deploy and monitor ML models and chip
(RTL and SoC) design.
Skills you'd need: natural language processing,
google cloud, vertex ai, artificial intelligence,
programming languages, system design,
distributed computing.

IBM

IT services · 130 IT postings
Hiring most: ML & AI engineer 17%,
Data engineer 15%, Software engineer
(general) 12%.
Re-posted, not counted: Healthcare /
clinical / life sciences (8 postings, 2
titles).

SIGNALS

Embedded firmware / RTOS / drivers 23 of 130
(18%), 17 titles, 5.3×
Event streaming (Kafka, Flink, real-time) 20 of
130 (15%), 9 titles, 1.9×
FinOps / cloud cost 13 of 130 (10%), 3 titles, 3.8×
Mainframe / legacy modernisation 8 of 130 (6%),
7 titles, 7.8×
Stack: Infrastructure as code 17%

WHAT IT MAY MEAN · SPECULATION

Postings suggest work on embedded firmware and
device software, event-streaming and real-time data
systems and cloud cost control.
Skills you'd need: cloud platforms, cloud,
firmware, kubernetes, orchestration,
observability, kafka.

SPECULATION · INFERRED FROM PUBLIC JOB POSTINGS · NOT STATEMENTS BY OR ABOUT THE COMPANIES' PLANS

WHAT THEY MIGHT BE BUILDING · EMPLOYERS 6-10 OF 15

Signals, what they may mean, and the skills you would need

YASH Technologies

IT services · 126 IT postings
Hiring most: SAP 40%, Software engineer (general) 10%, Backend developer 8%.

SIGNALS

SAP modules (FICO, MM, SD, ABAP, BTP) 59 of 126 (47%), 40 titles, 6.1×
SAP S/4HANA (migration / implementation) 18 of 126 (14%), 14 titles, 5.3×
Salesforce / CRM 13 of 126 (10%), 9 titles, 1.5×
Stack: Performance 6%, GRC 6%

WHAT IT MAY MEAN · SPECULATION

The postings point to SAP systems, SAP S/4HANA programmes and Salesforce and CRM platforms.
Skills you'd need: sap, abap, s/4hana, agile, sap s/4hana, sap abap, gcp.

Infosys

IT services · 126 IT postings
Hiring most: Backend developer 17%, Infrastructure, network & IT 12%, ML & AI engineer 10%.

SIGNALS

ServiceNow / ITSM 14 of 126 (11%), 12 titles, 1.9×
Forecasting / optimisation models 8 of 126 (6%), 8 titles, 4.4×
Computer vision 6 of 126 (5%), 5 titles, 3.2×
FinOps / cloud cost 5 of 126 (4%), 4 titles, 1.5×

WHAT IT MAY MEAN · SPECULATION

It appears to be investing in ServiceNow and IT service management, forecasting and optimisation models and computer-vision products.
Skills you'd need: digital transformation, knowledge management, react, servicenow, python, generative ai, quality assurance.

Microsoft

Product & internet · 118 IT postings
Hiring most: Software engineer (general) 62%, ML & AI engineer 10%, Product manager 6%.

SIGNALS

Snowflake / cloud data warehouse 33 of 118 (28%), 18 titles, 3.5×
Responsible AI / AI governance 17 of 118 (14%), 12 titles, 3.5×
GPU / inference optimisation 12 of 118 (10%), 10 titles, 5.9×
AI agents / agentic systems 20 of 118 (17%), 12 titles, 1.5×
Stack: Microservices 44%, ETL 31%, Observability 30%

WHAT IT MAY MEAN · SPECULATION

Likely areas of work: cloud data warehousing, AI governance and model-risk controls and GPU-based model serving and performance work.
Skills you'd need: c++, synapse, azure synapse, azure synapse analytics, synapse analytics, distributed systems, analytics.

Salesforce

Product & internet · 117 IT postings
Hiring most: Salesforce & CRM 36%, Software engineer (general) 12%, Data & BI analyst 11%.
Re-posted, not counted: Healthcare / clinical / life sciences (5 postings, 2 titles).
Standard text (not a signal): Salesforce / CRM 100%, AI agents / agentic systems 76%.

SIGNALS

Generative AI 26 of 117 (22%), 13 titles, 1.9×
LLMs / large language models 26 of 117 (22%), 15 titles, 1.8×
Data quality / governance / lineage 21 of 117 (18%), 14 titles, 1.5×
Prompt engineering 15 of 117 (13%), 11 titles, 2.4×
Stack: ETL 24%, GRC 8%

WHAT IT MAY MEAN · SPECULATION

Postings suggest work on generative-AI features, applications built on large language models and data quality, governance and lineage.
Skills you'd need: salesforce, agentforce, crm, ai agents, llms, artificial intelligence, generative ai.

Iris Software

IT services · 116 IT postings
Hiring most: Backend developer 18%, Test automation & SDET 17%, Full-stack developer 14%.
Standard text (not a signal): MLOps / LLM0ps 100%.

SIGNALS

Event streaming (Kafka, Flink, real-time) 36 of 116 (31%), 19 titles, 3.8×
Generative AI 22 of 116 (19%), 13 titles, 1.6×
Vector search / embeddings 12 of 116 (10%), 6 titles, 1.9×
Fine-tuning / model training 8 of 116 (7%), 5 titles, 2.6×
Stack: Microservices 37%, Test automation frameworks 22%, Performance 5%

WHAT IT MAY MEAN · SPECULATION

The postings point to event-streaming and real-time data systems, generative-AI features and semantic search over company data.
Skills you'd need: mlops, analytics, data analytics, kafka, database, microservices, ci/cd pipelines.

Counts: postings whose title or skills name the term, of the employer's IT postings; × = lift vs all IT postings. Employers ordered by IT postings first captured in the period (staffing firms left out). Interpretation is ours, from public postings only.

SPECULATION · INFERRED FROM PUBLIC JOB POSTINGS · NOT STATEMENTS BY OR ABOUT THE COMPANIES' PLANS

WHAT THEY MIGHT BE BUILDING · EMPLOYERS 11-15 OF 15

Signals, what they may mean, and the skills you would need

Persistent Systems

IT services · 116 IT postings
Hiring most: Backend developer 15%, Full-stack developer 10%, AI agent & GenAI engineer 8%.

SIGNALS

Generative AI 22 of 116 (19%), 19 titles, 1.6×
Identity and access (IAM, SSO, PAM) 16 of 116 (14%), 16 titles, 1.8×
RAG (retrieval-augmented generation) 15 of 116 (13%), 15 titles, 1.8×
Vector search / embeddings 12 of 116 (10%), 12 titles, 1.9×
Stack: Microservices 28%, Test automation frameworks 14%, GRC 7%

WHAT IT MAY MEAN · SPECULATION

It appears to be investing in generative-AI features, identity and access management and assistants that answer from internal documents (RAG).
Skills you'd need: generative ai, ci/cd pipelines, microservices, cloud-native, scrum, cloud platforms, backend services.

GlobalLogic

IT services · 115 IT postings
Hiring most: Backend developer 20%, ML & AI engineer 16%, Data engineer 13%.
Re-posted, not counted: Forecasting / optimisation models (5 postings, 1 title).

SIGNALS

Event streaming (Kafka, Flink, real-time) 29 of 115 (25%), 22 titles, 3.1×
IoT / edge / industrial automation 12 of 115 (10%), 10 titles, 4.6×
Snowflake / cloud data warehouse 18 of 115 (16%), 10 titles, 2.0×
Vector search / embeddings 14 of 115 (12%), 9 titles, 2.3×
Stack: Kubernetes 45%, Microservices 28%, ETL 24%

WHAT IT MAY MEAN · SPECULATION

Likely areas of work: event-streaming and real-time data systems, IoT and industrial-automation systems and cloud data warehousing.
Skills you'd need: kafka, docker, kubernetes, gcp, analytics, microservices, orchestration.

Deloitte

IT services · 113 IT postings
Hiring most: Backend developer 12%, SAP 11%, ML & AI engineer 10%.

SIGNALS

Security operations (SOC, SIEM, threat detection) 13 of 113 (12%), 8 titles, 2.5×
SAP modules (FICO, MM, SD, ABAP, BTP) 15 of 113 (13%), 13 titles, 1.7×
Vector search / embeddings 12 of 113 (11%), 8 titles, 2.0×
Data privacy / protection 9 of 113 (8%), 4 titles, 2.3×
Stack: Performance 4%

WHAT IT MAY MEAN · SPECULATION

Postings suggest work on security operations and threat detection, SAP systems and semantic search over company data.
Skills you'd need: analytics, machine learning, big data, data management, visualization, sap, microservices.

CGI

IT services · 109 IT postings
Hiring most: Data engineer 15%, Backend developer 10%, Software engineer (general) 7%.

SIGNALS

Lakehouse / Databricks / Delta / Iceberg 12 of 109 (11%), 8 titles, 1.8×
Fine-tuning / model training 6 of 109 (6%), 4 titles, 2.1×
Low-code / RPA / automation 5 of 109 (5%), 3 titles, 2.3×
Stack: ETL 27%, Test automation frameworks 12%, Performance 5%

WHAT IT MAY MEAN · SPECULATION

The postings point to lakehouse data platforms (Databricks, Delta, Iceberg), adapting models to its own data and low-code and process-automation platforms.
Skills you'd need: systems integration, database, scrum, ci/cd pipelines, oracle, ai tools, jira.

EPAM Systems

IT services · 109 IT postings
Hiring most: ML & AI engineer 19%, Software engineer (general) 16%, SAP 11%.
Re-posted, not counted: Credit risk / underwriting (5 postings, 2 titles).

SIGNALS

Generative AI 42 of 109 (38%), 19 titles, 3.3×
RAG (retrieval-augmented generation) 34 of 109 (31%), 15 titles, 4.3×
Vector search / embeddings 29 of 109 (27%), 12 titles, 4.9×
Platform engineering / internal developer platforms 24 of 109 (22%), 8 titles, 6.0×
Stack: Microservices 28%, Observability 21%, Infrastructure as code 16%

WHAT IT MAY MEAN · SPECULATION


It appears to be investing in generative-AI features, assistants that answer from internal documents (RAG) and semantic search over company data.
Skills you'd need: generative ai, platform engineering, rag, sap, mlops, gcp, microservices.

PLACES


Where the IT jobs are

Two in five IT postings named Bengaluru; Hyderabad gained the most share.

SHARE OF IT POSTINGS BY STATE Jul Jun



SHARE OF IT POSTINGS BY CITY Jul Jun



Bars: share of IT postings in July; ticks: June. Postings naming no state or city stay in the base.

62%

of July's IT postings named Bengaluru, Hyderabad or Pune.

+1.1

points of share gained by Hyderabad, the biggest move of any city.

Karnataka carried 38.8% of IT postings, Maharashtra 16.6% and Telangana 14.6%. Pune and Chennai each lost -0.4 points. 7.2% of IT postings named no state.

Bengaluru held 38.1% of IT postings, Hyderabad 14.3%, Pune 9.9%.

WHAT THIS MEANS FOR YOU • INTERPRETATION

Three cities hold three IT postings in five. If you are weighing Hyderabad against Pune, July tilted towards Hyderabad. Outside the big three, Chennai (7.8%) and Gurugram (5.3%) are the next markets.

CITY × ROLE · WHERE EACH ROLE CONCENTRATES

Where each role lives

Where each highlighted role concentrates, and which smaller cities are growing.

CELL = ROLE'S SHARE OF THE CITY'S NEW IT POSTINGS; SHADE = HOW CONCENTRATED VS INDIA (LOCATION QUOTIENT)

	Bengaluru	Hyderabad	Delhi NCR	Pune	Chennai	Mumbai	Kolkata	Ahmedabad
Software engineer (general)	18.0	16.0	10.5	15.9	15.5	9.9	8.1	10.5
ML & AI engineer	8.7	9.6	10.6	6.9	8.1	9.0	10.8	15.2
QA & manual testing	2.9	3.9	3.6	5.1	3.8	4.7	3.4	5.7
Backend developer	6.4	4.9	7.4	6.4	8.5	8.3	6.8	1.0
Data & BI analyst	3.9	5.1	6.7	5.0	3.7	8.6	6.8	4.8
Data engineer	5.5	8.3	5.0	6.9	7.7	5.8	12.8	1.0
Infrastructure, network & IT	4.1	3.5	4.5	4.7	4.9	6.3	4.1	6.7
Full-stack developer	3.1	3.3	3.3	4.9	4.0	4.7	4.7	1.0
Oracle, ERP & other platforms	4.2	5.6	4.3	4.9	4.6	3.3	4.7	3.8
Security engineering (general)	3.1	3.1	4.8	3.4	3.2	5.2	3.4	4.8
AI agent & GenAI engineer	2.7	2.2	4.0	1.6	2.6	2.3	4.7	0.0
Data scientist	3.0	2.4	2.8	2.0	3.2	3.6	4.7	1.9
DevOps & release	2.1	1.7	2.4	2.1	1.8	2.1	0.0	2.9
Cloud engineer & architect	2.8	3.6	3.1	2.8	4.3	2.4	2.7	2.9
SAP	4.5	5.0	4.0	6.8	3.5	4.6	4.1	1.9

under 0.8x India
0.8-1.2x
1.2-1.6x
1.6x or more

MOST OVER-REPRESENTED (VS INDIA)

Software eng.: Kochi 1.2×, Bengaluru 1.1×; ML/AI engineer: Ahmedabad 1.7×, Delhi NCR 1.2×; QA & testing: Kochi 3.0×, Pune 1.4×; Backend: Chennai 1.3×, Mumbai region 1.3×; Data/Bi analyst: Mumbai region 1.7×, Delhi NCR 1.3×; Data engineer: Kolkata 2.1×, Hyderabad 1.3×; Infra & network: Mumbai region 1.5×, Chennai 1.2×; Full-stack: Coimbatore 3.0×, Pune 1.3×.


SMALLER CITIES: NEW IT POSTINGS A DAY

City	Per day	Before	Change
Kochi	2.9	2.0	+48%
Jaipur	1.7	1.2	+42%
Ahmedabad	3.4	3.1	+9%
Thiruvananthapuram	2.8	2.6	+7%
Coimbatore	2.8	3.0	-4%
Kolkata	4.8	5.2	-8%

ROLES

Roles and new roles


Architecture, data and enterprise-platform roles gained share in July; embedded and frontend roles lost it.



Fastest-moving roles	Jul share	Change
Software architect	1.7%	+22%
Data architecture & governance	1.5%	+16%
Oracle, ERP & other platforms	4.5%	+13%
SAP	4.5%	+13%
Embedded & firmware	1.1%	-22%
Frontend developer	1.5%	-21%
Salesforce & CRM	1.7%	-18%

Change in share of IT postings vs June 2026; roles with ≥150 postings.

Software engineering is still the biggest family (31.5%). Data & AI gained 1.1 points to 27.0%, enterprise apps 0.6 to 10.7%. By role, software architects gained most (+22%), then data architecture (+16%) and SAP (+13%). Embedded (-22%) and frontend developers (-21%) lost most. AI agent engineer postings ran at 48.5 a week, against 36.2 over April–June.



New or surging titles	Posts	Before
Developer I3	68	62
Front-end development intern	24	21
Control software designer- develop	23	7
Pega developer	17	11
Abap developer	17	14
Developer I4	16	12

Normalised IT titles with at least 15 postings in July and at most three, or a third of the per-day rate, in April–June (a window that starts before the crawl steadied).

WHAT THIS MEANS FOR YOU · INTERPRETATION

Design and platform skills are being rewarded over hands-on screen-building; architects and SAP specialists gained, frontend developers lost. A senior developer moving towards architecture is moving with the market.

TRENDS · MONTH BY MONTH

The league table

Month by month, the rank of each highlighted role among the 34 IT roles, and its share of new IT postings.

	MAY	JUN	JUL	SHARE OF NEW IT POSTINGS	RANK Δ
Software engineer (general)	1	1	1	15.7%	0
ML & AI engineer	2	2	2	9.2%	0
QA & manual testing	8	10	9	3.7%	-1
Backend developer	3	3	3	6.4%	0
Data & BI analyst	5	5	5	5.1%	0
Data engineer	4	4	4	6.1%	0
Infrastructure, network & IT	6	6	8	4.2%	-2
Full-stack developer	9	7	10	3.7%	-1
Oracle, ERP & other platforms	10	9	6	4.7%	+4
Security engineering (general)	12	12	11	3.4%	+1
AI agent & GenAI engineer	18	15	15	2.7%	+3
Data scientist	14	14	14	2.7%	0
DevOps & release	17	16	16	2.0%	+1
Cloud engineer & architect	11	11	13	3.1%	-2
SAP	7	8	7	4.7%	0

Climbers since May: Oracle/ERP (#10 → #6), GenAI/agents (#18 → #15), Security eng. (#12 → #11). Slippers: Cloud engineer (#11 → #13), Infra & network (#6 → #8), Full-stack (#9 → #10).

Rank among the 34 IT roles by new IT postings first captured in the month (teal = top 10). May is the first full month after the crawl steadied (24 April); February–April were too thin to rank. July runs to 31 Jul.


WHAT THIS MEANS FOR YOU · INTERPRETATION

Interpretation. Persistence matters: a role that has stayed in the top ten every month is a safer bet than a one-month jump. A climber with momentum is worth a closer look; a slipper is still hiring, but with more competition.

SENIORITY

Seniority and experience

Most IT postings with a known level are senior or mid-level; entry-level demand stays thin.



Shares of IT postings with a known value, from our job-intel model. July vs June.

Senior roles were 49.2% of IT postings with a known level (June: 46.4%), mid-level 37.0%, lead 9.0% and junior 4.3% (from 8.0%). Experience asked: 0–1 years 3.4%, 2–3 years 20.3%, 4–5 years 30.8%, 6–9 years 29.8%, 10 or more 15.7%.

WHAT THIS MEANS FOR YOU • INTERPRETATION

Entry-level IT is thin: about one posting in thirty that states a minimum asks for 0–1 years. For freshers, internships are the visible door (a front-end development intern title was among July’s surging titles). For mid-career engineers, 4–9 years is where most demand sits.

Job-intel read 55% of July’s IT postings and 80% of June’s; level and work-mode shares between the two months partly reflect the model’s coverage, so treat month-on-month shifts as indicative.




THE EXPERIENCE LADDER

How many years each role asks

How many years each highlighted role asks for, and how steep its ladder is.

MINIMUM YEARS ASKED (JOB-INTEL): BOX = MIDDLE HALF, BAR = MEDIAN, WHISKERS = 10TH-90TH PCT



STEEPEST LADDERS: MOST 8+ ASKS PER 0-2 ASK

Role	Median	Ask 0-2 y	Ask 3-7 y	Ask 8+ y	Known
SAP	7 y	3%	55%	42%	335
Cloud engineer	6 y	7%	49%	44%	243
Oracle/ERP	6 y	5%	64%	30%	391
Data engineer	5 y	4%	72%	24%	443
Full-stack	5 y	7%	68%	26%	303
ML/AI engineer	5 y	9%	59%	32%	659
Backend	5 y	8%	68%	24%	503
GenAI/agents	5 y	9%	66%	25%	205

EASIEST FIRST RUNGS

Most postings asking 0-2 years: Data scientist 20%; Security eng. 19%; Data/BI analyst 17%; Infra & network 16%.

Years: the minimum asked, read by our job-intel model from the posting; shares of postings where it is known.

WHAT THIS MEANS FOR YOU • INTERPRETATION

Interpretation. A steep ladder means most openings want experienced people: enter such roles from an adjacent one rather than directly. Roles with many 0-2 year asks are the realistic first rungs.

FRESHER'S CORNER · 0-2 YEARS, INTERNS, TRAINEES


Where the first IT jobs are

1,168 of 20,497 new IT postings in July 2026 were open to freshers by our test: 5.7%. Here is who posts them, for which roles and where.

How we count. An entry-level posting is an IT posting first collected by our crawl in July 2026 whose title says intern, trainee, apprentice, fresher, graduate engineer, campus, entry level or GET, or whose job-intel analysis gives a minimum of 0-2 years. The minimum is known for 38% of postings, so

this is an undercount.

ENTRY-LEVEL POSTINGS BY ROLE




Bar: entry-level postings. Label: postings · share of the role's postings with a stated minimum that ask 0-2 years.

Employer	Listings	Postings	0-2 yrs	Titles
Wells Fargo	18	31	31	0
Qualcomm	14	15	15	0
PwC India	12	13	12	3
Argo Intern	10	35	0	35
American Express	10	15	15	0
Google	10	11	10	1
Future Interns	9	12	0	12
EY (Ernst & Young)	9	11	11	0
Dexter's Tech	8	107	0	107
Wake Up Whistle	8	45	0	45
Microsoft	8	9	5	4
Jobgether †	7	7	7	0

Ranked by listings: distinct title-city combinations in each fortnight, so a vacancy re-posted within a fortnight counts once. Postings: all entry-level postings. 0-2 yrs and Titles: the two tests (a posting can pass both). † recruiter or staffing firm. Exact counts.

ENTRY-LEVEL POSTINGS BY CITY



In July, 398 postings carried an intern, trainee or fresher title and 778 asked for 0-2 years (10.1% of postings with a stated minimum). Software engineer (general), Data & BI analyst and ML & AI engineer hold the most entry-level postings. Roles that lean junior (the largest share of postings with a stated minimum asking 0-2 years; at least 100 such postings): Data scientist 20%; Security engineering (general) 19%; Data & BI analyst 17%; Infrastructure, network & IT 16%. Across all IT the share is 10.1%.

READING THE CITY CHART


Shares are of all entry-level postings; 30% name no city. Bengaluru carries 28%. In brackets: the city's share of all new IT postings (same basis); a higher entry share means the city leans junior. Delhi NCR = Gurugram, Noida, Delhi and neighbours; Mumbai region = Mumbai, Navi Mumbai, Thane.

FRESHER'S CORNER · SKILLS, BEST BETS, ADVICE

What first jobs ask for

The skills entry-level postings list, the roles that score best for freshers, and what we would do with that.

Editorial: a suggested path, not a guarantee.



Share of entry-level postings that list skills (1,131) naming each skill; thin bar: share of all new IT postings that list skills (20,294). Soft skills and catch-all tags left out.

ASKED MORE OFTEN OF FRESHERS

Skill	Entry	All IT	Ratio
HTML	14.5%	4.2%	3.5×
CSS	14.6%	4.2%	3.5×
Excel	15.5%	6.2%	2.5×
Node.js	12.2%	4.9%	2.5×
Statistics	9.7%	4.6%	2.1×
Data visualization	5.6%	2.7%	2.0×

Ratio = entry-level share ÷ all-IT share; skills in at least 30 entry-level postings and 1.4× or more.

Degrees. Job text was collected for too few July 2026 postings to count degrees reliably (see the method page), so this edition does not report them.

FRESHER BEST BETS · FROM JOBS TO HIGHLIGHT

Role	Score	A day	0-2 yrs	Titles/wk	Momentum
Software engineer (general)	87	103.6	14%	28.5	-4%
Data & BI analyst	84	33.4	17%	18.7	-4%
ML & AI engineer	84	60.8	9%	9.9	+4%
Security engineering (general)	83	22.3	19%	2.5	+4%
Data scientist	74	17.9	20%	2.3	-5%

Score (0-100): the fresher segment score of Jobs to highlight. A day: new postings a day. 0-2 yrs: share of postings with a stated minimum. Titles/wk: fresher or intern titles a week. Momentum: change in share of IT postings. Full method ??.

EDITORIAL · INTERPRETATION: WHAT WE WOULD DO

Aim where the volume is. Software engineer (general), Data & BI analyst and ML & AI engineer carry the most entry-level postings; the best-bets table adds momentum and breadth of employers.

Get the basics visible. The everyday stack shows up in entry-level postings too (Python 34%, SQL 27%, Git 22%): put each on your CV with a project that uses it.

Show the junior-heavy tools. HTML, CSS, Excel, Node.js are asked of freshers more often than of IT postings overall: they are the cheapest way to look ready.

Check every internship. Some internship titles are re-posted many times, which is why we count listings. Before you commit, check the duration, the stipend and what you will work on; a genuine employer does not ask you to pay to join.

Use the 3-year rule loosely. 90% of postings with a stated minimum ask for 3 years or more. Many teams still consider strong freshers, so a short portfolio (two projects, a repository, one write-up) can stand in for a missing year.


WORK MODE

Remote vs office

Four in five IT postings with a known mode are on-site; hybrid is the main alternative.

WORK MODE (JOB-INTEL, KNOWN VALUES)

Jun Jul




Work mode among IT postings with a known mode, from job-intel. Illustration: Open Doodles (CC0).

Postings rarely state a work mode in a field (2.8% do), so we use our job-intel model's reading of the text, available for 47% of July's IT postings. Of those, 81.2% are on-site, 12.8% hybrid and 6.0% remote. June's figures (4.3% hybrid, 10.4% remote) rest on a smaller, different sample.

Data & AI roles were the most flexible (21% hybrid or remote); chip design and hardware the least (10%).

HYBRID + REMOTE SHARE BY IT FAMILY, JULY 2026 (KNOWN MODES)

Hybrid Remote



WHAT THIS MEANS FOR YOU • INTERPRETATION


Plan for the office: only about one IT posting in sixteen is fully remote. Hybrid, about one in eight, is the realistic ask, most often in data, software and QA roles.

HIRING DYNAMICS · HARDEST TO FILL · WHEN TO APPLY

Hardest to fill

Which roles employers struggle to fill (by how often they re-post them), and when IT postings appear.

HARD-TO-FILL PROXY, % OF THE ROLE'S POSTINGS · RANKED BY THE COMBINED INDEX




EMPLOYERS RE-LISTING THE MOST


Employer	Postings	Re-listed
Dexter's Tech	107	96%
Husky Technologies	163	95%
Harman India	75	73%
Nike	75	73%
Wipro	101	72%
TalentGigs †	158	67%
NetApp	72	61%
MSD	69	58%

Re-listed = postings beyond the first for the same employer, normalised title and city in a month. All IT roles: 21% re-listed; 22% of postings dated in the period were re-dated vacancies first seen earlier. † staffing firm or job platform. A proxy: a re-post can also mean several openings.

POSTINGS BY WEEKDAY, INDEX (100 = AVERAGE DAY)




BY WEEK OF THE MONTH, INDEX



When to apply: Postings cluster on Wed, Thu, Tue (index 121–129); Sun is the quietest (index 51). By week of the month, 1st-7th is busiest. Calendar: 89 posted days from 1 May.

Posting lifetimes, a caveat: we tried to measure how long postings stay open (a time-to-fill proxy). The crawl's open flag stays near 10% for postings up to ten weeks old, then drops for older ones: it is updated in batches, not when a vacancy closes. So we cannot yet say how fast roles fill; the re-post measures above are the better guide.

SHARE STILL MARKED OPEN AT THE CUT-OFF, BY WEEKS SINCE FIRST CAPTURE



HARD PROBLEMS WORTH READING ABOUT

Hard problems worth reading about

Six problems behind this month's postings: what each is, why it is hard, who works on it and where to start reading.

HARD PROBLEM 1 OF 6 · EDITORIAL EXPLAINER

Retrieval quality and grounding in RAG

Retrieval-augmented generation answers from your own documents. Most failures start in the retrieval, not the model.

Retrieval-augmented generation (RAG) finds passages in a document store that match a question and gives them to a language model to write the answer. It lets a model use private, current information without retraining, which is why it sits behind many enterprise assistants.

Why it is hard. Documents must be split into chunks, and the split matters: too small and the context is lost, too large and the relevant sentence is diluted. Embedding search finds meaning but can miss exact strings such as product codes, error numbers or names, so good systems combine it with keyword search and re-rank the results. Questions are vague or

multi-part. Sources contradict each other or are out of date. Users must only see documents they are allowed to see, so access control has to follow the data into the index. Even with the right passage retrieved, a model may overlook it (long contexts are read unevenly) or add claims the passage does not support.

So retrieval and generation are measured separately: did the right passage come back in the top few (recall), and is every claim in the answer supported by a cited passage (faithfulness)? Both need labelled examples, which are slow to build and easy to under-invest in.

THE EVIDENCE

Of 20,497 IT postings (titles and skills) in July 2026, RAG appears in 1,481 (7.2%), Vector search / embeddings appears in 1,106 (5.4%), LLMs / large language models appears in 2,505 (12.2%). Among the largest IT hirers: EPAM Systems 34 of 109, PwC India 16 of 82, Ericsson 11 of 57. Most mentions are in ML & AI engineer (43% of mentions) and AI agent & GenAI engineer (21% of mentions) postings. In July's titles and skills, RAG is named about as often in IT services as in product companies.

WHO WORKS ON IT

ML & AI engineer (p. 4); AI agent & GenAI engineer (p. 7); Software engineer (general) (p. 4); Data engineer (p. 5).

SKILLS

Python, embeddings and vector search, BM25 / keyword search, re-ranking, chunking strategies, evaluation metrics (recall@k, faithfulness).

WHERE TO START READING

- Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks" (NeurIPS 2020)
- Karpukhin et al., "Dense Passage Retrieval for Open-Domain Question Answering" (EMNLP 2020)
- Liu et al., "Lost in the Middle: How Language Models Use Long Contexts" (TACL 2024)
- Thakur et al., "BEIR: A Heterogeneous Benchmark for Zero-shot Evaluation of Information Retrieval Models" (NeurIPS 2021)
- Es et al., "RAGAS: Automated Evaluation of Retrieval Augmented Generation" (2023)

HARD PROBLEM 2 OF 6 · EDITORIAL EXPLAINER

Data platform migrations and data quality

Moving a warehouse to the cloud is easy to start and hard to finish: the numbers must still match on the day you switch.

Many companies are moving analytics from on-premise warehouses and Hadoop clusters to cloud warehouses or lakehouses built on open table formats such as Delta Lake and Apache Iceberg. The promise is cheaper storage, elastic compute and one copy of the data for reporting and machine learning.

Why it is hard. A mature warehouse carries thousands of pipelines, stored procedures and reports, many with business logic nobody documented. SQL dialects differ in small ways (null handling, date functions, decimal rounding, time zones) that change results quietly. Every migrated table must reconcile with the old one, not just by row count but by totals and edge cases, usually

by running both systems in parallel for weeks. Downstream consumers (dashboards, finance close, regulatory reports, ML features) have to cut over in the right order. Access rules and personal-data controls must move too, and the cost model changes from fixed hardware to pay-per-query.

Data quality does not end at cutover. Schemas drift, sources arrive late or twice, and a broken upstream field can poison every report. Teams add automated checks on freshness, volume and values, record lineage so they know what a change will break, and agree data contracts with the teams that produce the data.

THE EVIDENCE

Of 20,497 IT postings (titles and skills) in July 2026, Lakehouse / Databricks / Delta / Iceberg appears in 1,252 (6.1%), Snowflake / cloud data warehouse appears in 1,621 (7.9%), Data quality / governance / lineage appears in 2,398 (11.7%). Among the largest IT hirers: Microsoft 33 of 118, PwC India 17 of 82, Cisco 21 of 105. Most mentions are in Data engineer (42% of mentions) and ML & AI engineer (12% of mentions) postings. Lakehouse tools are named most, relative to the IT average, by the global capability centres.

WHO WORKS ON IT

Data engineer (p. 5); ML & AI engineer (p. 4); Software engineer (general) (p. 4); Data architecture & governance (p. 54).

SKILLS

SQL, Spark / PySpark, Databricks or Snowflake, dbt / Airflow, data quality testing, data modelling.

WHERE TO START READING

- Armbrust et al., "Delta Lake: High-Performance ACID Table Storage over Cloud Object Stores" (VLDB 2020)
- Armbrust et al., "Lakehouse: A New Generation of Open Platforms that Unify Data Warehousing and Advanced Analytics" (CIDR 2021)
- Breck et al., "Data Validation for Machine Learning" (MLSys 2019)
- Sculley et al., "Hidden Technical Debt in Machine Learning Systems" (NeurIPS 2015)
- Apache Iceberg documentation (iceberg.apache.org)

HARD PROBLEMS WORTH READING ABOUT · 3-4 OF 6

HARD PROBLEM 3 OF 6 · EDITORIAL EXPLAINER

Observability of distributed systems

When one user request touches forty services, 'it is slow' is a research problem.

In a microservice architecture a single request can pass through dozens of services, queues and databases, owned by different teams. Failures are partial: one dependency slows down, retries pile up, and the symptom appears somewhere else. Observability is the ability to ask new questions of a running system from its outputs: metrics, logs and traces (and increasingly profiles), correlated by a shared request identity.

Why it is hard. The data is enormous and costly. Adding detailed labels (user, endpoint, region) to metrics multiplies the number of time series, and storing every trace is unaffordable, so systems sample, and sampling can drop the one rare failing

request you needed. Context has to propagate across languages, thread pools and asynchronous queues, or traces break apart. The average hides the pain: users feel the slowest one percent of requests, and in a fan-out system rare slow components dominate that tail. Choosing service-level objectives that reflect what users experience, and alerting only when they are at risk, is a judgement call; too many alerts and on-call engineers stop reading them. Open standards such as OpenTelemetry have made instrumentation portable; the remaining work is deciding what to measure, what to keep and who acts on it.

WHERE TO START READING

- Sigelman et al., "Dapper, a Large-Scale Distributed Systems Tracing Infrastructure" (Google Technical Report, 2010)
- Dean and Barroso, "The Tail at Scale" (Communications of the ACM, 2013)
- Beyer et al., "Site Reliability Engineering" (O'Reilly, free to read at sre.google)
- OpenTelemetry documentation (opentelemetry.io)

THE EVIDENCE

Of 20,497 IT postings (titles and skills) in July 2026, Observability / monitoring appears in 2,882 (14.1%), SRE / incident management / SLOs appears in 1,110 (5.4%), Microservices / distributed systems appears in 3,526 (17.2%). Among the largest IT hirers: Oracle 48 of 150, Microsoft 35 of 118, Cisco 26 of 105. Most mentions are in Software engineer (general) (19% of mentions) and ML & AI engineer (13% of mentions) postings. Observability and monitoring tools are among the most-named terms in IT titles and skills.

WHO WORKS ON IT

Software engineer (general) (p. 4); ML & AI engineer (p. 4); DevOps & release (p. 54); Site reliability (SRE) (p. 54).

SKILLS

OpenTelemetry, Prometheus / Grafana, distributed tracing, SLOs and error budgets, Kubernetes, incident response.

HARD PROBLEM 4 OF 6 · EDITORIAL EXPLAINER

Modernising ERP and legacy cores

Moving decades of business logic to S/4HANA or off the mainframe is a migration of how a company works, not just of its software.

Enterprise resource planning (ERP) systems run finance, procurement, manufacturing and HR. SAP has set 2027 as the end of mainstream maintenance for its older ECC suite (with paid extended options), so many companies are moving to S/4HANA. Others are moving core systems off mainframes written in COBOL. Both are among the largest IT programmes a company runs.

Why it is hard. Old systems carry decades of custom code and local workarounds that encode how the business actually operates, often undocumented. S/4HANA changes the data model (finance and controlling merge into one journal, for example), so custom reports and interfaces must be rewritten or retired. Teams must choose between a

fresh implementation, a technical conversion of the existing system, or a selective mix, each with different costs and risks. Master data has to be cleansed before it moves. Dozens of integrations must keep working, cut-over windows are short (a plant cannot stop for a week), and end-to-end business processes, not just programs, must be tested. People who know both the old code and the business are scarce. Current practice keeps the ERP core close to standard and builds extensions outside it, and replaces legacy systems piece by piece rather than in one big bang.

THE EVIDENCE

Of 20,497 IT postings (titles and skills) in July 2026, SAP S/4HANA appears in 551 (2.7%), SAP modules appears in 1,577 (7.7%), Mainframe / legacy modernisation appears in 161 (0.8%). Among the largest IT hirers: YASH Technologies 59 of 126, EPAM Systems 23 of 109, EY (Ernst & Young) 11 of 64. Most mentions are in SAP (81% of mentions) and Data architecture & governance (3% of mentions) postings. S/4HANA is named most, relative to the IT average, by IT-services firms.

WHO WORKS ON IT

SAP (p. 54); Data architecture & governance (p. 54); Software engineer (general) (p. 4); Oracle, ERP & other platforms (p. 6).

SKILLS

SAP S/4HANA modules (FI/CO, MM, SD), ABAP and SAP BTP, data migration and cleansing, integration (APIs, middleware), COBOL / mainframe for legacy cores, business-process testing.

WHERE TO START READING

- SAP, "Simplification List for SAP S/4HANA" (SAP Help Portal)
- SAP, "Conversion Guide for SAP S/4HANA" (SAP Help Portal)
- SAP Activate methodology (SAP's implementation roadmap)
- Martin Fowler, "StranglerFigApplication" (martinfowler.com)

HARD PROBLEMS WORTH READING ABOUT • 5-6 OF 6

HARD PROBLEM 5 OF 6 • EDITORIAL EXPLAINER

GPU and inference cost

Training gets the headlines; serving a model to millions of requests is where the bill keeps arriving.

Every token a language model generates needs a pass through its weights, and that pass is limited more by memory bandwidth than by arithmetic. GPUs that are good at this are scarce and expensive, so the cost of serving a model (inference) often outgrows the cost of building it.

Why it is hard. Four goals pull against each other: low latency (how fast the first word appears), high throughput (requests served per GPU), answer quality and cost. Batching many requests together raises throughput but can slow individual answers. The memory that holds each conversation's context (the key-value cache) grows with prompt length and batch size, and running out of it forces smaller batches. Traffic is

bursty, so capacity sized for peaks sits idle at night. The main levers each carry a risk: quantisation (8- or 4-bit weights) cuts memory but can degrade accuracy in ways only an evaluation suite will show; distillation into a smaller model needs data and testing; speculative decoding and smarter attention kernels need specialist engineering.

In practice teams route easy requests to small models, cache repeated prompts, measure cost per successful task rather than per token, and keep a regression suite so that a cheaper configuration does not quietly become a worse one.

WHERE TO START READING

- Kwon et al., "Efficient Memory Management for Large Language Model Serving with PagedAttention" (SOSP 2023)
- Dao et al., "FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness" (NeurIPS 2022)
- Leviathan et al., "Fast Inference from Transformers via Speculative Decoding" (ICML 2023)
- Dettmers et al., "LLM.int8(): 8-bit Matrix Multiplication for Transformers at Scale" (NeurIPS 2022)
- vLLM project documentation (docs.vllm.ai)

THE EVIDENCE

Of 20,497 IT postings (titles and skills) in July 2026, GPU / inference optimisation appears in 354 (1.7%), MLOps / LLMops appears in 1,557 (7.6%), Fine-tuning / model training appears in 537 (2.6%). Among the largest IT hirers: EPAM Systems 21 of 109, PwC India 14 of 82, Google 21 of 137. Most mentions are in ML & AI engineer (40% of mentions) and Software engineer (general) (12% of mentions) postings. GPU skills concentrate in semiconductor employers, far above the IT average.

WHO WORKS ON IT

ML & AI engineer (p. 4); Software engineer (general) (p. 4); AI agent & GenAI engineer (p. 7); Platform engineer (p. 54).

SKILLS

Python, CUDA basics, model serving (vLLM, Triton), quantisation, Kubernetes, profiling and benchmarking.

HARD PROBLEM 6 OF 6 • EDITORIAL EXPLAINER

Zero trust security and identity

The office network is no longer a safe zone. Every request has to prove who and what it is, every time.

The traditional model trusted anything inside the corporate network and guarded the edge with firewalls and a VPN. With cloud services, remote work and partners connecting from everywhere, that edge has gone. Zero trust, as NIST describes it, grants no trust because of network location: every access request is authenticated and authorised for a specific resource, using the identity of the user or service, the health of the device and the context, with the least privilege needed, and is re-evaluated as conditions change.

Why it is hard. You cannot protect what you have not mapped, and most organisations lack a full inventory of users, service accounts, devices, applications and data

flows. Older applications cannot use modern single sign-on or multi-factor authentication. Identity becomes the new perimeter, so identity systems, privileged accounts and machine-to-machine credentials become prime targets. Fine-grained policies multiply and conflict. Every new check adds friction for users, and the change has to happen gradually without breaking the business.

Programmes usually start with strong identity (MFA, single sign-on, privileged access), then device posture, then segmenting networks and applications, with monitoring that can spot misuse of valid credentials.

WHERE TO START READING

- NIST Special Publication 800-207, "Zero Trust Architecture" (2020)
- NIST Special Publication 800-63, "Digital Identity Guidelines"
- CISA, "Zero Trust Maturity Model" (version 2.0, 2023)
- Ward and Beyer, "BeyondCorp: A New Approach to Enterprise Security" (;login:, USENIX, 2014)
- OWASP Application Security Verification Standard (ASVS) (owasp.org)

THE EVIDENCE

Of 20,497 IT postings (titles and skills) in July 2026, Zero trust appears in 6 (0.0%), Identity and access appears in 1,575 (7.7%), Security operations appears in 934 (4.6%). Among the largest IT hirers: Persistent Systems 16 of 116, GlobalLogic 15 of 115, EPAM Systems 13 of 109. Most mentions are in Software architect (33% of mentions) and Cloud & application security (17% of mentions) postings. The phrase zero trust is almost never a listed skill; identity, security operations and cloud security are.

WHO WORKS ON IT

Software architect (p. 54); Cloud & application security (p. 54); Backend developer (p. 5); GRC, identity & privacy (p. 54).

SKILLS

identity and access management, SSO, SAML, OAuth/OIDC, privileged access management, network segmentation, cloud IAM, SIEM and detection.

ORG WATCH


Org watch: who is building what

What are the big IT employers building? Their July hiring, de-duplicated and sorted onto the quarterly issue’s level ladder, gives a sketch.

We kept one posting per employer, title and city first captured in July, leaving 6,232 IT postings at the 151 employers we measured. Each was sorted into one of 34 IT roles and onto an eight-rung ladder from intern to VP. July’s crawl gaps thin these counts a little.
IT services firms post few manager roles (4.9% of levelled postings) and about 16 individual-contributor roles per manager role; banks post the most managers (20.2%).

Five of the month’s six biggest IT hirers are services firms; the sixth is Oracle.

COHORT COMPARISON · AVERAGE LEVEL PYRAMID, JULY 2026 · SHARE OF LEVELLED POSTINGS, %



Cohort	Firms	Postings	Mgrs	IC/mgr*	Largest family
IT services	49	2,611	5%	16	Software 29%
Product & internet	27	991	14%	7	Software 36%
Chips & hardware	14	481	6%	16	Chip & hardware 50%
Telecom & networking	8	398	12%	9	Cloud & DevOps 28%
Banking & finance	24	829	20%	4	Data & AI 35%
Non-tech GCCs	21	692	16%	6	Data & AI 30%

*Median IC postings ÷ manager postings (manager, director, VP/head rungs): a hiring-mix proxy, not a headcount ratio. Employers with ≥20 de-duplicated IT postings this month; equal-weight averages.

HOW TO READ THE PROFILES

Top box: the employer and its de-duplicated IT postings. Teal-edged boxes: its largest IT families; under each, its largest roles with their share of IT postings and most common level (Jr, Mid, Sr, Lead, Mgr, Dir, VP). Pyramids: individual contributors teal, managers amber. “Building now” compares family shares with the previous three months.

WHAT THIS MEANS FOR YOU · INTERPRETATION

Enterprise applications are where several big hirers lean: Oracle (36% of its IT postings), TCS (34%) and Capgemini (22%). If you know SAP, Oracle or Salesforce, July’s biggest employers were asking for you. If you want to manage people, banks post proportionally the most manager roles.

Inferred from job postings, not real org charts or headcount: a posting is a vacancy, not a person. Span = IC postings per manager posting, a hiring-mix proxy. Data: our crawl of public IT job postings in India first captured 1 Jul-31 Jul 2026, de-duplicated.


ORG WATCH • COHORTS 1-2 OF 6

Cohort org charts: the archetype of each cohort


LEVEL STRIP IN ROLE BOXES: Intern Junior Mid Senior Staff/Lead Manager Director VP/Head

IT services & consulting


49 employers • 2,611 de-duplicated IT postings • equal-weight average



Level pyramid, % of levelled



Function mix, % of postings




Largest families: Software 29%, Data & AI 27%, Enterprise apps 18%. Biggest roles: backend developer 10%, software engineer (general) 10%, sap 9%.


Managers 4.9% of levelled postings; median 16 IC postings per manager posting. Staff/lead 29%, intern and junior 2%. Largest members: Tata Consultancy Services, Zensar, Persistent Systems, GlobalLogic, Capgemini.

Product, software & internet


27 employers • 991 de-duplicated IT postings • equal-weight average



Level pyramid, % of levelled



Function mix, % of postings



Largest families: Software 36%, Data & AI 25%, Cloud & DevOps 15%. Biggest roles: software engineer (general) 14%, ml & ai engineer 12%, backend developer 9%.

Managers 14.5% of levelled postings; median 7 IC postings per manager posting. Staff/lead 28%, intern and junior 6%. Largest members: Oracle, Salesforce, Google, Microsoft, Agoda.

Inferred from job postings, not real org charts or headcount: a posting is a vacancy, not a person. Span = IC postings per manager posting, a hiring-mix proxy. Data: our crawl of public IT job postings in India first captured 1 Jul-31 Jul 2026, de-duplicated.


ORG WATCH · COHORTS 3-4 OF 6

Cohort org charts: the archetype of each cohort

LEVEL STRIP IN ROLE BOXES: Intern Junior Mid Senior Staff/Lead Manager Director VP/Head


Semiconductors & hardware

14 employers · 481 de-duplicated IT postings · equal-weight average




Other IT families: 10% of postings

Level pyramid, % of levelled




Function mix, % of postings



Largest families: Chip & hardware 50%, Software 18%, Data & AI 14%. Biggest roles: rtl, asic & fpga design 18%, design verification 18%, software engineer (general) 13%. Managers 5.7% of levelled postings; median 16 IC postings per manager posting. Staff/lead 42%, intern and junior 2%. Largest members: Qualcomm, AMD, Nvidia, NXP Semiconductors.


Telecom & networking

8 employers · 398 de-duplicated IT postings · equal-weight average




Other IT families: 18% of postings

Level pyramid, % of levelled



Function mix, % of postings



Largest families: Cloud & DevOps 28%, Software 21%, Data & AI 19%. Biggest roles: infrastructure, network & it 18%, security engineering (general) 12%, software engineer (general) 11%. Managers 11.6% of levelled postings; median 9 IC postings per manager posting. Staff/lead 22%, intern and junior 6%. Largest members: VOIS, Hewlett Packard Enterprise, Cisco.

Inferred from job postings, not real org charts or headcount: a posting is a vacancy, not a person. Span = IC postings per manager posting, a hiring-mix proxy. Data: our crawl of public IT job postings in India first captured 1 Jul-31 Jul 2026, de-duplicated.


ORG WATCH • COHORTS 5-6 OF 6

Cohort org charts: the archetype of each cohort

LEVEL STRIP IN ROLE BOXES: Intern Junior Mid Senior Staff/Lead Manager Director VP/Head


Banking, finance & insurance tech

24 employers • 829 de-duplicated IT postings • equal-weight average




Other IT families: 14% of postings

Level pyramid, % of levelled



Function mix, % of postings




Largest families: Data & AI 35%, Software 26%, Cloud & DevOps 14%. Biggest roles: software engineer (general) 13%, data & bi analyst 11%, security engineering (general) 8%.

Managers 20.2% of levelled postings; median 4 IC postings per manager posting. Staff/lead 26%, intern and junior 7%. Largest members: Citi, S&P Global, Wells Fargo, Mastercard, NatWest Group.


Global capability centres (non-tech)

21 employers • 692 de-duplicated IT postings • equal-weight average




Other IT families: 23% of postings

Level pyramid, % of levelled



Function mix, % of postings



Largest families: Data & AI 30%, Software 29%, Enterprise apps 10%. Biggest roles: software engineer (general) 21%, ml & ai engineer 10%, data engineer 6%.

Managers 15.9% of levelled postings; median 6 IC postings per manager posting. Staff/lead 25%, intern and junior 5%. Largest members: Ford, Caterpillar, Thermo Fisher Scientific.

Inferred from job postings, not real org charts or headcount: a posting is a vacancy, not a person. Span = IC postings per manager posting, a hiring-mix proxy. Data: our crawl of public IT job postings in India first captured 1 Jul-31 Jul 2026, de-duplicated.


The month’s biggest IT hirers

Pyramids: individual contributors teal, managers amber; rungs top to bottom VP/head, director, manager, staff/lead, senior, mid, junior, intern. Role boxes: share of postings · most common level (Jr, Mid, Sr, Lead, Mgr, Dir, VP).


IT SERVICES & CONSULTING

Tata Consultancy Services

128 IT postings (de-dup.) · level known for 50%



Level pyramid




Function mix: Enterprise apps 34%, Data & AI 31%, Software 13%, Cloud & DevOps 10%.
Levels: Mid 56%, Senior 44%, Intern 0%; managers 0.0%.
Span (proxy): n/a (no manager postings) IC postings per manager posting; hiring mix, not headcount.
Building now: Data & AI +9 pts, Enterprise apps +3 pts (share of its postings vs the previous three months).
Top cities: Chennai 34%, Hyderabad 22%, Bengaluru 21% (of postings naming a city).


IT SERVICES & CONSULTING

Zensar

116 IT postings (de-dup.) · level known for 60%



Level pyramid




Function mix: Software 27%, Enterprise apps 22%, Cloud & DevOps 20%, Data & AI 17%.
Levels: Mid 46%, Senior 37%, Staff/Lead 17%; managers 0.0%.
Span (proxy): n/a (no manager postings) IC postings per manager posting; hiring mix, not headcount.
Building now: Enterprise apps +6 pts, QA & testing +4 pts (share of its postings vs the previous three months).
Top cities: Pune 55%, Bengaluru 24%, Hyderabad 16% (of postings naming a city).


PRODUCT, SOFTWARE & INTERNET

Oracle

110 IT postings (de-dup.) · level known for 76%



Level pyramid



Function mix: Enterprise apps 36%, Software 30%, Cloud & DevOps 16%, Data & AI 13%.
Levels: Staff/Lead 35%, Mid 24%, Senior 23%; managers 18.1%.
Span (proxy): 4.5 IC postings per manager posting; hiring mix, not headcount.
Building now: Enterprise apps +12 pts (share of its postings vs the previous three months).
Top cities: Bengaluru 42%, Hyderabad 15%, Noida 14% (of postings naming a city).

Inferred from job postings, not real org charts or headcount: a posting is a vacancy, not a person. Span = IC postings per manager posting, a hiring-mix proxy. Data: our crawl of public IT job postings in India first captured 1 Jul-31 Jul 2026, de-duplicated.

ORG WATCH · COMPANY PROFILES 4-6 OF 6



Company profiles

Pyramids: individual contributors teal, managers amber; rungs top to bottom VP/head, director, manager, staff/lead, senior, mid, junior, intern. Role boxes: share of postings · most common level (Jr, Mid, Sr, Lead, Mgr, Dir, VP).

IT SERVICES & CONSULTING

Persistent Systems

106 IT postings (de-dup.) · level known for 53%





Function mix: Software 34%, Data & AI 26%, Security 11%, Cloud & DevOps 10%.
Levels: Staff/Lead 48%, Senior 27%, Mid 20%; managers 3.6%.
Span (proxy): 27 IC postings per manager posting; hiring mix, not headcount.
Building now: Security +11 pts, Data & AI +6 pts (share of its postings vs the previous three months).
Top cities: Bengaluru 75%, Hyderabad 14%, Mumbai 7% (of postings naming a city).

IT SERVICES & CONSULTING

GlobalLogic

105 IT postings (de-dup.) · level known for 73%




Function mix: Data & AI 36%, Software 35%, Cloud & DevOps 12%, QA & testing 9%.
Levels: Staff/Lead 39%, Senior 32%, Mid 18%; managers 10.4%.
Span (proxy): 8.6 IC postings per manager posting; hiring mix, not headcount.
Building now: Data & AI +12 pts (share of its postings vs the previous three months).
Top cities: Bengaluru 33%, Noida 23%, Hyderabad 20% (of postings naming a city).

IT SERVICES & CONSULTING

Capgemini

104 IT postings (de-dup.) · level known for 59%



Function mix: Software 31%, Data & AI 22%, Enterprise apps 22%, QA & testing 14%.
Levels: Mid 41%, Senior 30%, Staff/Lead 21%; managers 4.9%.
Span (proxy): 19 IC postings per manager posting; hiring mix, not headcount.
Building now: Enterprise apps +12 pts, QA & testing +2 pts (share of its postings vs the previous three months).
Top cities: Bengaluru 46%, Hyderabad 20%, Pune 16% (of postings naming a city).

Inferred from job postings, not real org charts or headcount: a posting is a vacancy, not a person. Span = IC postings per manager posting, a hiring-mix proxy. Data: our crawl of public IT job postings in India first captured 1 Jul-31 Jul 2026, de-duplicated.

REFERENCE · GLOSSARY

Glossary

The technical terms used in this issue, in plain English.

Agentic AI / AI agents AI systems that plan steps and call tools or APIs on their own to finish a task.

Agile A way of running software work in short cycles, with frequent feedback and re-planning.

API Application programming interface: a defined way for one program to ask another for data or actions.

AWS Amazon Web Services, the largest public cloud: rented servers, storage, databases and more.

Azure Microsoft's public cloud platform.

Backend The server side of an application: business logic, data storage and APIs.

BI (business intelligence) Reporting and dashboards that turn company data into figures managers use.

CI/CD Continuous integration and delivery: code is built, tested and released automatically on every change.

Cohort A group of employers of the same kind, such as IT services firms or banks.

Data pipeline A chain of automated steps that moves data from sources, cleans it and loads it somewhere useful.

Data quality Whether data is complete, correct, consistent and on time; checked with automated rules.

Data warehouse A database built for analysis, holding cleaned history from many systems.

Databricks A cloud platform for data engineering and machine learning built around Apache Spark.

DevOps Joining development and operations so teams build, release and run software together, with automation.

Embedded systems Computers built into devices such as cars, meters and routers, often with tight limits.

ERP Enterprise resource planning: one system for finance, purchasing, stock and HR, such as SAP or Oracle.

ETL / ELT Extract, transform, load: copying data out of systems, reshaping it and storing it for analysis.

Fresher In Indian hiring, a candidate with no full-time work experience yet, usually a new graduate.

Full-stack Working on both the front end (what users see) and the back end (servers and data).

Generative AI AI that produces new text, images or code from a prompt.

Hybrid work A job that mixes office days with days working from home.

Java A long-established programming language, common in enterprise and banking back ends.

Job intel Our structured reading of each posting's text: years asked, work mode and similar fields.

Kubernetes Open-source software that runs and scales containers across many machines.

LLM Large language model: an AI model trained on vast text to read and write language.

Machine learning Software that learns patterns from data to predict or classify, rather than following hand-written rules.

Median The middle value: half the figures are above it and half below.

Microservices Building an application as many small services that talk over APIs, each deployable on its own.

Momentum In this issue: the change in a role's share of IT postings against earlier periods.

Percentile The value below which a given share of figures fall; the 90th percentile is above 90% of them.

Python A general-purpose programming language, the default for data, AI and automation work.

QA Quality assurance: checking software works as intended before and after release.

RAG Retrieval-augmented generation: an AI answers using documents it looks up first, not memory alone.

Remote work A job done entirely away from the employer's office.

Repost / re-listed The same vacancy published again; we count it so it does not inflate demand.

REST A common style for web APIs that uses plain HTTP requests such as GET and POST.

S/4HANA SAP's current-generation ERP suite.

SAP The German company whose ERP software runs finance and operations at many large firms.

Share of postings The percentage of postings in a group that name something; our main unit of measure.

SOC In security, a security operations centre that watches for and responds to attacks; in chips, a system on a chip.

Speculation (in this issue) Our labelled interpretation of what postings may suggest; not a statement by any company.

SQL The standard language for querying and changing data in relational databases.

SRE Site reliability engineering: running production systems with software, error budgets and automation.

Staffing firm A company that recruits for other employers; its postings may stand for several clients.

Test automation Using code to run tests automatically instead of by hand.

45 terms that appear in this issue, in plain English. Definitions are the editors' own, kept short; they are not vendor descriptions.

REFERENCE · WORKSHEET

Your 90-day plan

A page to write on: pick a role, list your gaps, and set three checkpoints. Print it or annotate the PDF.

1 · Target role. Pick one role from the roadmaps pages and write why it suits you.

2 · Skills you have. Tick them off the role checklist; copy the ones you can already show with work.

3 · Gaps to close. The checklist items you can't show yet, most-asked first.

Days 1–30 · Foundation. What you will learn, and the one small thing you will build.

Days 31–60 · Core. The core stack for the role, and the project that proves it.

Days 61–90 · Differentiators. One differentiator, and where you will share the result (CV, portfolio, write-up).

CV keywords to add. From the CV keyword guide: the terms your CV should carry once they are true.

REFERENCES

References

Every company named in the hiring and org pages, A–Z, with its IT postings in July 2026 and the pages where it appears. We do not list or link individual job postings: they change and close quickly. Employers named only in the Speculation pages carry their counts there.

C

Capgemini · 170 IT postings · p. 24, 48
CGI · 129 IT postings · p. 24

E

EPAM Systems · 169 IT postings · p. 24

G

GlobalLogic · 159 IT postings · p. 24, 48
Google · 179 IT postings · p. 24

H

Hewlett Packard Enterprise · 68 IT postings · p. 24
Husky Technologies · 176 IT postings · p. 24

I

IBM · 147 IT postings · p. 24
Infosys · 138 IT postings · p. 24
Innova ESI · 188 IT postings · p. 24

M

Mastercard · 140 IT postings · p. 24

O

Okta · 107 IT postings · p. 24
Oracle · 215 IT postings · p. 24, 47

P

Persistent Systems · 116 IT postings · p. 24, 48

S

S&P Global · 75 IT postings · p. 24
Salesforce · 125 IT postings · p. 24

T

TalentGigs · 159 IT postings · p. 24
Tata Consultancy Services · 164 IT postings · p. 24, 47

U

UST Global · 153 IT postings · p. 24

Y

YASH Technologies · 166 IT postings · p. 24

Z

Zensar · 185 IT postings · p. 24, 47
Zorba Consulting India · 221 IT postings · p. 24

REFERENCES

Data and sources

Where the data comes from

Our own crawl of public job postings in India, structured by our pipeline every few minutes into a collection of job postings with title, company, place, posting date, skills and job text, and analysed by a job-intel model (work mode, seniority, experience). This edition reads the analytics copy of that collection, refreshed from the production database every six hours. Data to 20 Oct 2026, 17:47 IST. All queries were read-only; each is saved with its result in the edition’s data appendix (data/monthly/2026-07/queries.md).

What counts

An IT posting is one whose title rule (first match in the taxonomy) lands in an IT role: every role in the families Chip design & hardware, Security, Data & AI, Enterprise apps, Cloud, DevOps & SRE, QA & testing, Software engineering, plus Product manager and Product & UX design. Project and programme managers, business analysts, technical

support and graphic design are not counted, and nor are generic titles that match no rule ("Manager", "Lead", "Engineer"), so the IT count is conservative. The period: 1 July 2026 to 31 July 2026 (31 days), compared with 1 June 2026 to 30 June 2026 (30 days). The set is every IT posting posted or first captured in the period. Method: p. 53.

Pictures and licences

Illustrations: Open Doodles by Pablo Stanley (CC0 1.0), recoloured to the edition’s palette. Charts and tables: THROUGHPUT data desk, drawn from the data. No company logos are used.

How complete the data is

Field present in IT postings	Jul	Jun
Skills	99.1%	72.4%
State	92.8%	91.3%
City	89.1%	87.9%
Seniority (field)	47.6%	50.5%
Work type (field)	2.8%	7.2%
Employment type	7.8%	10.6%
Job description text	4.3%	2.0%
Salary	0.0%	0.0%
Job-intel analysis (work mode, seniority)	54.9%	80.3%

Salary is absent from every posting, so this edition prints no pay figures. June’s postings list fewer skills each and have a different job-intel sample, so July compares shares (skills: relative to the typical skill), never counts.

Copyright & disclaimer

© 2026 Michmer. All rights reserved. No part of this publication may be reproduced without permission, except brief quotations with attribution. Illustrations and images are used under their own licences; see Credits (Pictures and licences, above). Charts and analysis © Michmer, based on our crawl of public job postings. Data. The figures come from Michmer’s crawl of publicly listed job postings between 1 Apr 2026 and 31 Jul 2026 (cut-off 20 Oct 2026, 17:47 IST). They cover only what the crawl captured; they are not a census of the job market and may contain duplicates, reposts, errors or gaps (see the methodology, p. 53). Interpretation. Text labelled

interpretation is editorial opinion. Nothing in this edition is career, financial, legal or investment advice. Speculation. Sections marked Speculation are our interpretation of public job postings, not statements by or about the companies’ actual plans. Companies. Names and trademarks belong to their owners. Their appearance reflects job postings in our data, not endorsement, partnership or affiliation. Michmer is not affiliated with any job platform or company named. No company logos are used. Links. This issue does not link to individual job postings; postings

Publisher: Michmer

Contact: the publisher

change and close quickly. Errors. We have worked to be accurate, but errors are possible. Report corrections to the contact named in this box.

METHOD

Methodology

The set. Every IT posting posted or first extracted (captured by our crawl) in the period, counted once. July 2026: 1 Jul–31 Jul, 31 days; June 2026: 30 days. Months differ in length, so every comparison is per day, per week or a share, never a raw total.

New and reposted. New = first captured in the period. Reposted = in the set but first captured earlier and re-dated by the employer into the period.

Trailing average. The average day of the three months before the month (1 Apr–30 Jun), shown only when all three months fall after the crawl became steady on 24 Apr 2026. Not shown this month: the window starts before 24 Apr.

IT scope and taxonomy. The quarterly issue's taxonomy: ordered keyword rules on the job title, first match wins; 34 IT roles in eight families (product manager and UX design form Product & design). Generic titles ("Manager", "Lead") match no rule and are left out, so counts are conservative.

Skills. Share of IT postings naming the skill. The base is IT postings that list any skill. Rising and falling compare shares, for skills in at least 100 IT postings in both periods, relative to the typical skill (this period the median skill's share moved +54%), so a month whose postings list more skills each does not lift every skill; soft skills are left out. New or surging titles: normalised titles with at least 15 postings and at most 3 before, or three times the per-day rate.

Employers. Who's hiring counts raw IT postings per day by employer (the employer's platform identifier). † marks staffing firms, recruiters and job platforms, which post their clients' roles.

Org watch. One posting per employer, normalised title and city in each calendar month of first capture (de-duplication). Levels: the quarterly issue's eight-rung ladder, from the title, then the job-intel band, then title keywords, then years asked. Profiles: the six biggest IT hirers by de-duplicated postings, staffing firms excluded; cohorts

are equal-weight averages of employers with at least 20 postings.

Work mode and seniority. From our job-intel model's reading of the posting text, available for 54.9% of IT postings this period (80.3% before). Shares are of known values.

Crawl gaps. Our crawl captured no new postings on 8 July days: 3, 10, 11, 12, 16, 18, 28, 30 July (June had 5). Postings dated on those days still count when captured later, by posted date, but they show up as late captures, so July's new-listing count is understated and its repost share overstated. The trailing three-month average is not shown: April–June starts before the crawl steadied on 24 Apr.

Accuracy. Title rules were right for most postings in the quarterly issue's hand check; the level rule scored 74 of 75 on a blind check. Large tables were assembled from chunked queries; every query records an error bound (0 for every figure printed here unless stated).

Level ladder	Typical experience
Intern/Trainee	student / trainee
Junior/Associate	0–2 years
Mid	2–5 years
Senior	5–8 years
Staff/Lead/Principal	8+ years, IC
Manager	first-line manager
Sr Manager/Director	manages managers
VP/Head/C-level	runs a function

IT family	Share
Software engineering	31.5%
Data & AI	27.0%
Cloud, DevOps & SRE	11.4%
Enterprise apps	10.7%
Security	6.1%
QA & testing	5.1%
Chip design & hardware	4.4%
Product & design	3.9%

METHOD

The taxonomy

How each IT posting was classified: 34 roles in eight families, with this month’s shares.

Family	IT role	Share
Chip & hardware	Design verification	1.1%
Chip & hardware	Physical design & timing	0.5%
Chip & hardware	RTL, ASIC & FPGA design	1.3%
Chip & hardware	Embedded & firmware	1.1%
Chip & hardware	Hardware & electronics	0.5%
Security	Security operations & testing	0.7%
Security	Cloud & application security	0.8%
Security	GRC, identity & privacy	1.1%
Security	Security engineering (general)	3.5%
Data & AI	AI agent & GenAI engineer	2.6%
Data & AI	ML & AI engineer	9.3%
Data & AI	Data scientist	2.8%
Data & AI	Data engineer	5.9%
Data & AI	Data architecture & governance	1.5%
Data & AI	Data & BI analyst	4.9%
Enterprise apps	SAP	4.5%
Enterprise apps	Salesforce & CRM	1.7%

Family	IT role	Share
Enterprise apps	Oracle, ERP & other platforms	4.5%
Cloud & DevOps	Site reliability (SRE)	1.1%
Cloud & DevOps	Platform engineer	0.9%
Cloud & DevOps	DevOps & release	2.0%
Cloud & DevOps	Cloud engineer & architect	3.1%
Cloud & DevOps	Infrastructure, network & IT	4.4%
QA & testing	Test automation & SDET	1.5%
QA & testing	QA & manual testing	3.6%
Software	Engineering manager	1.1%
Software	Full-stack developer	3.7%
Software	Mobile developer	0.9%
Software	Frontend developer	1.5%
Software	Backend developer	6.3%
Software	Software architect	1.7%
Software	Software engineer (general)	16.4%
Product & design	Product manager	3.2%
Product & design	Product & UX design	0.7%

How we categorised

Rules. Each posting’s title is matched against ordered keyword rules (52 roles in 14 families across the whole market); the first match wins. An IT posting is one whose first match is an IT role. All rules are in taxonomy.py in the data appendix. Clustering, tested and dropped. The quarterly issue also tried assigning untitled-role postings by skill signature; checked by hand it was right only about half the time, so it is not used. Spot check. In the quarterly issue’s hand check of random postings, most rule assignments were right; errors were titles sharing a word with another field (“Quality” in process roles read as QA). One checker, one pass: treat accuracy as indicative. IT scope. An IT posting is one whose title rule (first match in the taxonomy) lands in an IT role: every role in the families Chip design & hardware, Security, Data & AI, Enterprise apps, Cloud,

DevOps & SRE, QA & testing, Software engineering, plus Product manager and Product & UX design. Project and programme managers, business analysts, technical support and graphic design are not counted, and nor are generic titles that match no rule (“Manager”, “Lead”, “Engineer”), so the IT count is conservative. Overlaps are cosine similarities of roles’ skill-share vectors. Emerging roles are title clusters matched by the issue’s regexes. Large tables were assembled from chunked queries with recorded error bounds.



CareerLens

Track. Learn. Stay Ahead.

IT hiring keeps moving, month by month, as this edition shows. CareerLens is built for job candidates: track your progress, take assessments, and sign in with Google or a candidate code.

careerlens.live →

VISIT CAREERLENS.LIVE OR SCAN THE CODE



Scan: `careerlens.live`

Numbers of the month

25,259

IT postings posted or first captured in July

815

IT postings a day, against 848 in June

8

July days with no new postings captured: crawl gaps

19%

of July's IT postings were first seen before July (June: 15%)

+93%

data warehouse's gain in share, relative to the typical skill

+230%

Persistent Systems' IT postings a day, the biggest mover

38.1%

of IT postings named Bengaluru

48.5

AI agent engineer postings a week

12.8%

hybrid; remote 6.0%

3.4%

of IT postings stating a minimum ask for 0–1 years

Data to 20 Oct 2026, 17:47 IST. Source: ARMed Jobs connector (structured public job postings in India), IT roles only. July 2026 vs June 2026; comparisons per day, per week or as shares. New in this issue: Skills radar p. 17 • Skill-gap checklists p. 19 • CV keywords p. 21 • Learning roadmaps p. 22 • Fresher's corner p. 36 • Glossary p. 49.

NEXT MONTH: AUGUST, WITH A FULL CRAWL