This report covers core engineering and data / ML roles

Contents

Executive Summary

Scope: Skills durability analysis across 9 core engineering and data/ML roles | Period: June 2025–June 2026 (trailing 12 months) | Data: 2.85 million active job descriptions, role-filtered from Draup's talent intelligence platform

Across approximately 2.85 million active job descriptions analysed for these nine roles, a single structural pattern holds without exception: the skills commanding the highest JD frequency and the fastest demand growth are those requiring judgment, accountability, and system-level reasoning — not those requiring typing. AI coding assistants are now explicitly named in employer requirements at scale (GitHub Copilot: 22,961 JDs; Cursor: 17,905; Claude: 12,860 across the Software Development Engineer role alone), confirming that AI fluency has crossed from differentiator to baseline expectation.

The central tension: Total active postings across these nine roles grew year-over-year, with the single most-posted title in each of Software Engineer, Data Engineer, and DevOps exceeding 40,000 active postings. The market is expanding, not contracting. But the internal composition of each role is shifting faster than most hiring frameworks have caught up with. The skills that once defined entry-level contribution (boilerplate code, routine SQL, manual test scripts, standard ETL) are precisely the skills most exposed to AI assistance. This creates a structural squeeze on early-career pathways that organizations have not yet resolved.

Three cross-role findings: (1) The seniority premium is structural, not cyclical — senior, staff, and principal title variants are outpacing generic variants YoY across every role. (2) AI fluency has crossed the baseline threshold — GitHub Copilot, Cursor, and Claude appear in 60,000+ JDs across all nine roles combined. (3) The pay hierarchy maps directly onto the durability hierarchy — ML Engineer ($166,764) at the top, Data Analyst ($88,140) at the bottom, with durability explaining the ranking.

What This Means

AI is not reducing demand for technical talent. It is changing which skills command a premium. Organizations that redesign hiring, development, and job architecture around durable skills will create workforce advantages that competitors cannot easily replicate.

The workforce planning implication is straightforward: stop organizing technical talent around the tasks people perform today and start organizing around the capabilities that remain valuable when AI can perform those tasks itself (refer to the durable-versus-exposed framework in this report). The organizations that make this shift first will build workforce advantages that competitors cannot easily replicate.

How to Read This Document

Every technical role is being reshaped by AI and automation, but not erased. The useful question is not "will this role survive?" but "which skills inside the role hold their value, and which ones shift?" This document answers that for nine core engineering and data/ML roles.

A simple pattern holds across all of them. The skills that endure are rooted in judgment, design, and accountability. The skills most exposed to change are routine, repetitive, and easy to specify — exactly what AI assistants and automated tools do well.

The durability test

A skill tends to STAY durable when it requires:

- Judgment under ambiguity — deciding what to build and what "good" means.
- System-level reasoning — trade-offs across scale, cost, security, and failure.
- Accountability — owning whether something is correct, safe, and fit to ship.
- Human context — domain depth, communication, and stakeholder influence.

A skill tends to SHIFT when it is:

- Routine and repeatable — boilerplate, standard patterns, glue code.
- Easy to specify — a clear prompt produces a usable first draft.
- Recall-based — syntax and API memorization rather than reasoning.
- Self-serviceable — now doable by non-specialists with AI/no-code tools.

A consistent theme: AI raises the floor on routine output, so human value moves up the stack toward design, review, and orchestration. Roles do not disappear — their center of gravity moves. Expectations for early-career hires are rising fastest, because the routine tasks juniors once cut their teeth on are the most automated.

At-a-Glance Summary

The one-line view of each role. Detailed breakdowns follow.

Role	Durable core	Most exposed to change
Software Engineer	Architecture, debugging, code-review judgment	Hand-writing routine code; syntax recall
Frontend Engineer	Accessibility, performance, UX judgment	Boilerplate UI/CSS; pixel-pushing
Backend Engineer	API/data design, security, scalability	CRUD/glue code; standard tests
DevOps / SRE	Reliability strategy, incident judgment	Hand-written IaC/config; log triage
QA / Test Engineer	Test strategy, exploratory testing	Writing test cases; manual regression
Data Engineer	Data modeling, reliability, governance	Routine ETL/connector & SQL code
Data Analyst	Problem framing, storytelling, domain	Routine SQL & dashboards; ad-hoc pulls
Data Scientist	Problem framing, stats, causal rigor	AutoML-able modeling & tuning
ML Engineer	MLOps, AI-systems & SWE rigor	Hand-written model/training boilerplate

Market Overview: The Market Is Growing, But the Centre of Gravity Is Moving

Total active postings across the nine roles in the trailing 12 months reached approximately 2.85 million JDs — a market that is expanding, not contracting. The most pronounced acceleration is in ML Engineer and Data Engineer, with Software Engineer remaining the single largest role by raw volume (Software Engineer and Senior Software Engineer together: over 111,000 postings).

The complication is that volume growth masks a compositional shift. The fastest-growing job titles within each role are not the generalist variants — they are the specialised, judgment-intensive, or AI-adjacent ones. "AI Software Engineer" (1,297 JDs, up 57% YoY), "Staff Software Engineer", "Principal Software Engineer", and "ML Ops Engineer" are all growing faster than their generic counterparts. The market is not hiring more of the same; it is hiring more of the harder, more senior, more specialised version of each role.

This is the central tension the role-by-role analysis below resolves: what, specifically, is 'harder and more specialised' in each role — and what is becoming routine enough to be assisted or automated?

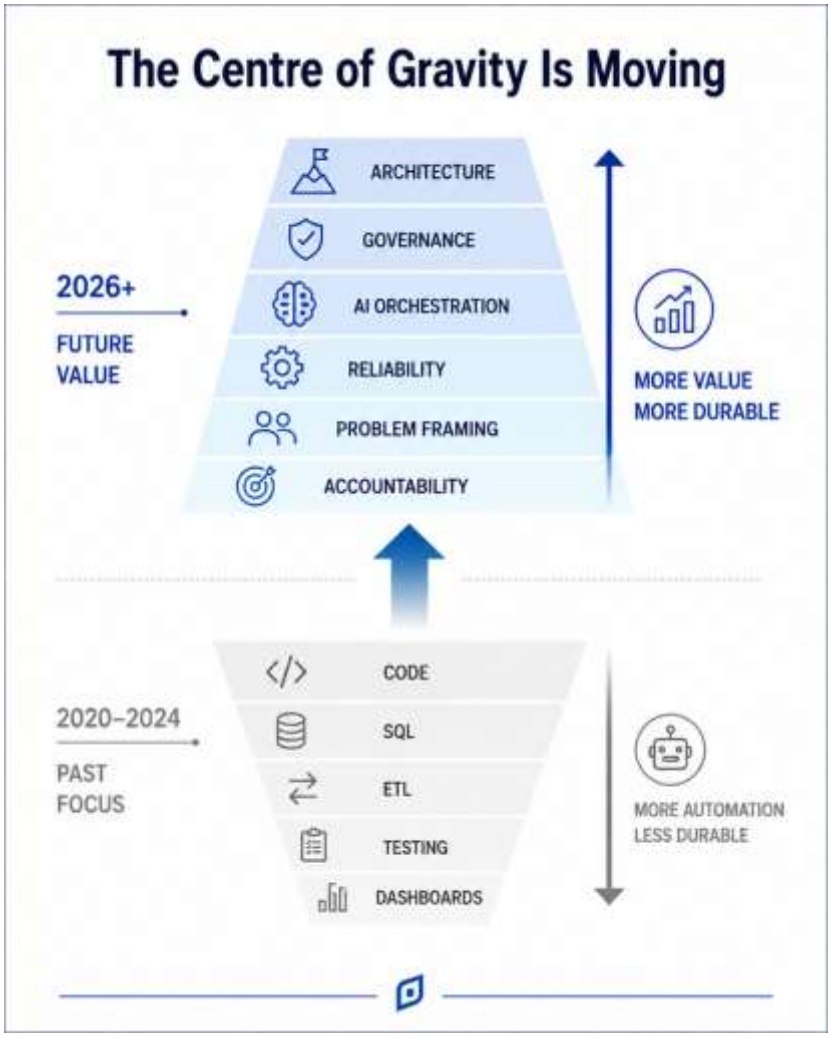

Role Analysis

1. Software Engineer (Generalist)

The role is shifting from typing code to specifying, reviewing, and orchestrating it. AI assistants now draft a large share of routine code, so the engineer's value concentrates in judgment, design, and accountability for what ships.

What Stays (Durable)	What May Change (Shifting / At-Risk)
• System design and architecture — decomposing problems, choosing trade-offs, reasoning about scale, latency, and failure. • Debugging and root-cause analysis in unfamiliar or messy systems. • Code-review judgment — knowing what "good" looks like and why AI-generated code is wrong or risky. • Directing AI coding agents — specifying intent, reviewing output, and owning the security and correctness of what ships. • Translating ambiguous business needs into clear technical specifications. • Collaboration, communication, and the ability to learn new tools quickly.	• Writing boilerplate, CRUD, and routine first-draft code by hand — increasingly drafted by AI; time shifts to review and integration. • Memorizing syntax and API signatures; recall matters far less than direction. • Manual test scaffolding and documentation generation. • Junior "task-execution" work — the most exposed layer; expectations for early-career hires are rising.

Horizon: Near term (1–3 yr): AI handles ~30–40% of routine coding; engineers move toward review and orchestration. Medium term (3–5 yr): premium shifts decisively to architecture, security judgment, and domain depth.

Draup Market Data (June 2026)

- **Active JDs - last 12 months (role-filtered):** 1,161,156 total JDs (role: Software Development Engineer, Jun 2025–Jun 2026)
- **Top Job Titles (role-filtered, by posting volume):** Software Engineer (57,681) · Senior Software Engineer (54,256) · Software Developer (20,102) · Lead Software Engineer (9,040) · Staff Software Engineer (8,121)

Skills: AI-Safe vs AI-Exposed (June 2026)

- **Durable (AI-Safe):** Systems Design (100,687 JDs) · Debugging (166,851) · Code Review (30,793) · Software Architecture (50,292) · Distributed Systems (93,732) · API Design (42,917) · Secure Coding (19,177) — judgment-intensive, hard to generate correctly
- **At Risk (AI-Exposed):** Unit Testing (77,250 JDs) · Technical Documentation (51,534) · Design Patterns (92,109 — rote application of patterns, not pattern selection) — AI tools now explicit in JDs: GitHub Copilot (22,961 JDs) · Cursor (17,905) · Claude (12,860) · AI-Assisted Development (9,755) handle boilerplate and first drafts
- **Workload Shift:** Application Development & Maintenance (715,720 JDs — 62% of role volume) is the highest-automation-risk workload. AI & Data Science (42,487 JDs) is the fastest-growing protected zone. Mobile App Dev and Software Testing workloads also seeing AI tooling penetration.

Live Signal: Salesforce froze all new engineering hires in FY2026, relying on AI coding agents (March 2026). TrueUp: 67,000+ software engineering openings in April 2026, highest in 3 years, defying AI replacement narrative. 128,000+ tech workers cut YTD 2026, heaviest on mid-level managers and

entry-level coders; AI labs (OpenAI doubling to 8,000 staff, Google rehiring 20% as boomerang workers) running aggressive senior engineering hiring campaigns in parallel.

2. Frontend Engineer

UI scaffolding and component generation are heavily automatable, but the craft of accessible, performant, human-centered interfaces — and the judgment behind design decisions — remains hard to delegate.

What Stays (Durable)	What May Change (Shifting / At-Risk)
• Accessibility (a11y), performance budgeting, and cross-device/browser problem-solving. • Design sense — spacing, interaction, and translating design intent into resilient UI. • State management and architecture for complex, data-heavy applications. • Design-system architecture and governance — scalable, consistent UI foundations across teams. • Collaboration with design and product on user experience. • Directing AI UI-generation tools — reviewing output for accessibility, performance, and design intent.	• Hand-writing boilerplate components, markup, and CSS — AI generates first drafts well. • Pixel-pushing from a mockup; design-to-code is increasingly automated. • Memorizing framework-specific library APIs — concepts transfer; syntax is assisted and frameworks churn. • Simple marketing/landing pages are heavily commoditised by AI and no-code tools.

Horizon: Near term (1–3 yr): faster prototyping, more time on UX quality and edge cases. Medium term (3–5 yr): value concentrates in performance, accessibility, and complex stateful apps.

Draup Market Data (June 2026)

- **Active JDs - last 12 months (role-filtered):** 159,404 total JDs (role: Frontend Engineer, Jun 2025–Jun 2026)
- **Top Job Titles (role-filtered, by posting volume):** Frontend Developer (15,184) · Senior Frontend Developer (5,808) · Senior Frontend Engineer (3,727) · Frontend Engineer (2,991) — capitalisation/spacing variants combined

Skills: AI-Safe vs AI-Exposed (June 2026)

- **Durable (AI-Safe):** User Experience (20,289 JDs) · Design Systems (16,711) · Performance Optimization (25,310) · Responsive Web Design (24,806) · Cross-Browser Compatibility (13,536) · Web Accessibility (3,072) · WCAG (1,189) — require human taste, real-device reasoning, and accessibility judgment AI cannot reliably supply
- **At Risk (AI-Exposed):** HTML (53,860 JDs) · CSS3 (24,280) · Webpack (21,176) · Tailwind CSS (16,957) · jQuery (7,568) — scaffolding and layout work heavily commoditised. AI tools in JDs: Cursor (3,284) · GitHub Copilot (3,129) · Codex (516) · Claude (1,413) now handle component first drafts
- **Workload Shift:** Application Development & Maintenance (37,544 JDs — 24% of role volume, the largest single workload) most exposed. UI/UX Design (1,393 JDs) and AI & Data Science (1,154 JDs) are the growth pockets. Simple marketing-page and e-commerce frontend work (E-Commerce: 176 JDs) facing strongest commoditisation.

Live Signal: Meta shifted 7,000 workers into AI roles (May 2026) while cutting 8% of total workforce. Forward deployed engineer postings surged 19x YoY in Jan 2026. New title 'AI Frontend Developer' appearing in JDs; no-code and GenAI UI tools accelerating commoditisation of pure visual/layout work.

3. Backend Engineer

Backend work is shielded somewhat by its dependence on system-level reasoning, data integrity, and security — areas where AI assists but cannot own the decision.

What Stays (Durable)	What May Change (Shifting / At-Risk)
• API and data-model design; reasoning about consistency, concurrency, and failure modes. • Performance, scalability, and cost trade-offs under real load. • Security thinking — authz/authn, data protection, threat modeling. • Distributed-systems reasoning and debugging production incidents. • Designing backends for AI features — serving LLM workloads with latency, cost, and reliability control. • Directing AI code generation — specifying contracts and reviewing generated services before they ship.	• Routine endpoint, schema, and integration glue code — AI-drafted. • Boilerplate for serialization, validation, and standard CRUD services. • Specific language/framework recall; concepts transfer, syntax is assisted. • Manual writing of standard tests and migration scripts.

Horizon: Near term (1–3 yr): faster delivery of standard services. Medium term (3–5 yr): differentiation through reliability, security, and systems depth.

Draup Market Data (June 2026)

- **Active JDs - last 12 months (role-filtered):** 241,239 total JDs (role: Backend Engineer, Jun 2025–Jun 2026)
- **Top Job Titles (role-filtered, by posting volume):** Backend Developer (14,058) · Backend Engineer (11,435) · Senior Backend Engineer (6,564) — capitalization/spacing variants combined

Skills: AI-Safe vs AI-Exposed (June 2026)

- **Durable (AI-Safe):** Systems Design (21,927 JDs) · Distributed Systems (29,293) · API Design (22,096) · Encryption (3,880) · Microservices (68,202) · Low-Latency (5,658) · Security (API security: 2,788) — require architectural judgment and production accountability AI cannot replace
- **At Risk (AI-Exposed):** Unit Testing (13,181 JDs) · Technical Documentation (9,028) · Standard ORM boilerplate: Spring Boot (37,319) · Django (14,528) · FastAPI (14,338) — routine endpoint and schema code highly exposed. AI tools in JDs: GitHub Copilot (4,116) · Cursor (4,186) · Claude (2,603)
- **Workload Shift:** Application Development & Maintenance (71,817 JDs — 30% of role volume, the largest single workload) most at risk. AI & Data Science (4,682 JDs) and Cloud Computing (4,682 JDs) growing as protected workloads. Governance, Risk & Compliance (351 JDs) a small but durable safety zone.

Live Signal: Oracle cut 30,000 roles globally (Q1 2026) redirecting spend to cloud and AI infrastructure; Block cut 40% of workforce with CEO citing AI tools. Counter-signal: Cognizant targeting 50% AI-generated code by 2026 while hiring 25,000 freshers - volume production work shifts to AI while systems-level and integration expertise stays in demand.

4. DevOps / Site Reliability Engineer (SRE)

Automation is the job, so AI extends rather than replaces it. The durable core is reliability-engineering judgment — what to automate, how to fail safely, and how to reason under incident pressure.

What Stays (Durable)	What May Change (Shifting / At-Risk)
• Reliability engineering: SLOs, capacity planning, and designing for graceful failure. • Incident response and calm root-cause analysis under pressure. • Security, compliance, and cost governance across cloud infrastructure. • Systems thinking across the whole delivery pipeline. • Platform engineering — designing self-service abstractions and golden paths for developers. • Directing AIOps and AI-generated automation — deciding what to automate and validating that it fails safely.	• Hand-writing IaC, pipeline YAML, and config — AI drafts and refactors these. • Manual log triage and first-pass anomaly detection (increasingly AIOps-assisted). • Routine runbook authoring and toil-heavy operational tasks. • Memorizing tool-specific syntax across an ever-changing toolchain.

Horizon: Near term (1–3 yr): AIOps reduces toil; more focus on reliability strategy. Medium term (3–5 yr): platform-engineering and self-service abstractions raise the value of judgment over scripting.

Draup Market Data (June 2026)

- **Active JDs - last 12 months (role-filtered):** 187,215 total JDs (role: DevOps Engineer, Jun 2025–Jun 2026)
- **Top Job Titles (role-filtered, by posting volume):** DevOps Engineer (56,900) · Senior DevOps Engineer (19,175) — capitalisation variants combined

Skills: AI-Safe vs AI-Exposed (June 2026)

- **Durable (AI-Safe):** Reliability Engineering (5,165 JDs) · Site Reliability Engineering (3,824) · Cloud Security (9,855) · DevSecOps (15,080) · Secrets Management (8,887) · Multi-Cloud (3,030) · Capacity Planning (4,206) — safety-critical; incident response under pressure requires human judgment
- **At Risk (AI-Exposed):** IaC hand-writing: Terraform (110,694 JDs) · Ansible (68,145) · YAML (7,815) — routine config, pipeline YAML, and log triage most exposed. AI tools in DevOps JDs: GitHub Copilot (2,057) · Cursor (1,017) · Claude (933) · AIOps / AI DevOps Engineer titles growing fast
- **Workload Shift:** DevOps workload (114,979 JDs — 61% of role volume) faces automation of its most repetitive sub-tasks. Cloud Computing (21,916 JDs) growing as a complementary layer. AI & Data Science workloads (2,080 JDs) emerging for AIOps-literate SREs. IaC and config-management tasks most vulnerable to automation.

5. QA / Test Engineer

One of the most reshaped roles: AI can generate large volumes of tests, so value moves from writing test cases to designing test strategy and owning quality risk.

What Stays (Durable)	What May Change (Shifting / At-Risk)
• Test strategy and risk-based prioritization — deciding what quality means and what to test. • Exploratory testing and finding failure modes specs never mention. • Quality advocacy and judgment about release readiness. • Understanding user behaviour and real-world edge cases. • Testing AI systems — validating model outputs, non-deterministic behaviour, and AI-feature quality.	• Hand-writing unit/integration test cases — heavily automated by AI. • Manual, repetitive regression execution and scripted UI automation maintenance — the most exposed skill set in this group. • Test-data setup and boilerplate harness creation.

Horizon: Near term (1–3 yr): AI generates tests; humans curate and design coverage. Medium term (3–5 yr): role merges toward "quality engineering" embedded in development.

Draup Market Data (June 2026)

- **Active JDs - last 12 months (role-filtered):** 134,046 total JDs (role: Quality Assurance Engineer, Jun 2025–Jun 2026)
- **Top Job Titles (role-filtered, by posting volume):** QA Engineer (15,225) · Quality Assurance Engineer (9,504) · QA Automation Engineer (4,985) · Senior QA Engineer (4,104) — capitalisation variants combined

Skills: AI-Safe vs AI-Exposed (June 2026)

- **Durable (AI-Safe):** Test Strategy (19,474 JDs) · Exploratory Testing (6,411) · Risk Analysis (4,193) · Test Planning (45,100) · Quality Management System (4,477) · Root Cause Analysis (10,571) — require domain context, real user behaviour understanding, and release judgment AI cannot reliably supply
- **At Risk (AI-Exposed):** Selenium (43,263 JDs) · Playwright (29,233) · JMeter (14,128) · Manual Testing (18,567) · Test Scripts (21,251) — automation harness and manual regression most exposed. AI QA Engineer titles growing. AI tools in QA JDs: GitHub Copilot (1,565) · Cursor (969) · Claude (676)
- **Workload Shift:** Software Testing (24,222 JDs) and Quality Assurance (16,088 JDs) workloads are where routine test-writing is most exposed. AI & Data Science (2,623 JDs) growing as quality engineering merges with ML model validation. Manufacturing Engineering (2,658 JDs) and hardware testing workloads remain durable.

Live Signal: Emerson launched NI Nigel AI platform with prompt-based test code generation across the full test lifecycle (May 2026); Testaify released autonomous testing platform 2.0 purpose-built for 'the speed of AI-assisted development.' TestGorilla launched AI fluency assessments - AI literacy now a baseline hiring criterion for QA roles.

6. Data Engineer

Demand is strong because AI and analytics both run on reliable data. Pipeline boilerplate is automatable, but data modeling, reliability, and governance remain firmly human.

What Stays (Durable)	What May Change (Shifting / At-Risk)
• Data modeling and architecture for warehouses/lakehouses. • Data reliability, quality, and observability — trustable pipelines at scale. • Governance, lineage, security, and cost optimization. • Distributed-systems and streaming fundamentals. • Data foundations for AI — RAG corpora, vector stores, and pipelines for unstructured data.	• Hand-writing routine ETL/ELT and connector code — increasingly AI-generated. • Boilerplate SQL transformations and standard dbt models. • Manual schema mapping and ingestion plumbing. • Tool-specific recall as the stack consolidates and abstracts.

Horizon: Near term (1–3 yr): faster pipeline build-out; focus shifts to reliability and governance. Medium term (3–5 yr): data engineers increasingly own the data foundation AI systems depend on.

Draup Market Data (June 2026)

- **Active JDs - last 12 months (role-filtered):** 384,970 total JDs (role: Data Engineer, Jun 2025–Jun 2026)
- **Top Job Titles (role-filtered, by posting volume):** Data Engineer (76,012) · Senior Data Engineer (36,489) · Azure Data Engineer (2,890) — capitalisation variants combined

Skills: AI-Safe vs AI-Exposed (June 2026)

- **Durable (AI-Safe):** Data Modeling (85,110 JDs) · Data Architecture (38,498) · Data Governance (67,320) · Data Quality (106,975) · Data Lineage (8,417) · Metadata Management (25,282) · Distributed Systems (14,180) — platform and governance decisions requiring systemic judgment
- **At Risk (AI-Exposed):** ETL (165,151 JDs) · SQL (243,456) · Apache Airflow boilerplate (77,341) · Databricks connector code (85,417) · Snowflake (78,998) — routine pipeline and transformation code most exposed. AI tools in DE JDs: GitHub Copilot (3,505) · Cursor (2,227) · Claude (2,137)
- **Workload Shift:** Data Engineering (190,823 JDs — 50% of role volume) is core but routine sub-tasks are automating fast. AI & Data Science (24,081 JDs — 6%) growing fast as data engineers own the infrastructure AI models depend on. Data Centre Management (18,194 JDs) growing with AI infra build-out demand.

Live Signal: Oracle, Amazon, Dell led 61,000+ global job cuts in Q1 2026 with data and enterprise software roles heavily affected. Counter-signal: $725B AI capex in 2026 (up from $410B in 2025) requires massive data infrastructure build-out; Databricks, Snowflake, and Microsoft Fabric specialisations insulated as the underlying data layer for AI systems expands.

7. Data Analyst

The most disrupted of the data roles: natural-language querying and AutoML let non-analysts self-serve basic reporting. Surviving analysts move up the value chain into framing and influence.

What Stays (Durable)	What May Change (Shifting / At-Risk)
• Problem framing — turning a vague business question into the right analysis. • Domain context and the judgment to know when a number is wrong or misleading. • Data storytelling, stakeholder influence, and driving decisions. • Statistical literacy and experiment/causal reasoning. • Data governance and stewardship — definitions, quality, and accountability for trusted data. • Auditing AI-generated analyses — verifying the numbers business users self-serve.	• Writing routine SQL and building standard dashboards — now possible via natural language. • Manual data cleaning and repetitive reporting — automated or self-served. • Ad-hoc "pull me a number" requests handled by business users with AI tools. • Tool-specific "button knowledge" — analytical thinking transfers; tool mechanics are assisted.

Horizon: Near term (1–3 yr): self-service erodes routine reporting; analysts become advisors. Medium term (3–5 yr): role bifurcates into analytics-engineering and decision-partner tracks.

Draup Market Data (June 2026)

- **Active JDs - last 12 months (role-filtered):** 352,225 total JDs (role: Data Analyst, Jun 2025–Jun 2026)
- **Top Job Titles (role-filtered, by posting volume):** Data Analyst (60,604) · Senior Data Analyst (15,267) · Junior Data Analyst (2,756) — capitalisation variants combined

Skills: AI-Safe vs AI-Exposed (June 2026)

- **Durable (AI-Safe):** Data Governance (43,054 JDs) · Statistical Modeling (8,579) · Data Storytelling (4,611) · Stakeholder Management (11,286) · Hypothesis Testing (3,318) · Experimental Design (1,599) · Causal reasoning — require human business judgment and accountability for decisions
- **At Risk (AI-Exposed):** SQL (182,810 JDs) · Power BI (131,115) · Tableau (85,119) · Data Entry (13,747) · Ad Hoc Analysis (8,225) · Automated Reporting (6,058) — routine querying and dashboard work most exposed to AI self-service. AI tools in DA JDs: GitHub Copilot (921) · ChatGPT (referenced in descriptions)
- **Workload Shift:** Business Intelligence & Analytics (14,530 JDs), Data Engineering (8,924 JDs), and AI & Data Science (7,630 JDs) are top three workloads — BI and reporting most at risk, analytics engineering and AI-adjacent work protected. Governance, Risk & Compliance (7,425 JDs) durable due to regulatory accountability requirements.

Live Signal: PwC, McKinsey, EY, KPMG, and Deloitte slashing analyst and support roles as they push AI adoption (May 2026); routine reporting, ad-hoc pulls, and dashboard-building roles most exposed. Gartner predicts AI-driven role reductions will partially reverse by 2027 under different job titles. Roles pivoting to analytics engineering or decision-partner tracks show continued strong demand.

8. Data Scientist

AutoML and AI assistants commoditise standard modeling, pushing value toward problem formulation, causal reasoning, and translating models into business impact.

What Stays (Durable)	What May Change (Shifting / At-Risk)
• Framing business problems as tractable data/ML problems. • Statistical rigor, experimental design, and causal inference. • Critical evaluation of models — bias, validity, and limits. • Communicating results and tying them to decisions and ROI. • Designing evaluations for GenAI and agentic systems — metrics, guardrails, and behaviour under uncertainty.	• Routine model selection, tuning, and feature engineering — handled by AutoML. • Boilerplate notebook EDA and standard model code — AI-drafted. • Algorithm-implementation-from-scratch as a daily skill. • Notebook-only workflows give way to production and MLOps expectations.

Horizon: Near term (1–3 yr): less hand-modeling, more problem framing and validation. Medium term (3–5 yr): domain depth + ML earns the largest premium; pure modelers are squeezed.

Draup Market Data (June 2026)

- **Active JDs - last 12 months (role-filtered):** 143,508 total JDs (role: Data Scientist, Jun 2025–Jun 2026)
- **Top Job Titles (role-filtered, by posting volume):** Data Scientist (46,015) · Senior Data Scientist (19,319) · Data Scientist II (1,389) — capitalisation variants combined

Skills: AI-Safe vs AI-Exposed (June 2026)

- **Durable (AI-Safe):** Statistical Modeling (23,190 JDs) · Experimental Design (5,344) · Hypothesis Testing (7,125) · Model Evaluation (5,273) · Responsible AI (1,910) · Causal Inference (579) · Data Storytelling (1,859) — require deep domain expertise and accountability that AI cannot substitute
- **At Risk (AI-Exposed):** Feature Engineering (13,024 JDs) · EDA (8,150) · Data Cleaning (4,649) · Standard ML Algorithms: Scikit-Learn (29,135) · XGBoost (5,994) · Random Forest (1,764) — AutoML and AI coding assistants commoditising standard tuning and boilerplate. Claude (1,261 JDs) · GitHub Copilot (773) · Gemini (768) appearing in DS JDs
- **Workload Shift:** AI & Data Science (127,035 JDs — 89% of role volume) is heavily concentrated — the role is almost entirely within this workload. Business Intelligence & Analytics (5,704 JDs) adjacent work growing. GenAI Data Scientist and Agentic AI Data Scientist now distinct emerging sub-workloads within the AI & Data Science bucket.

Median Base Pay (USA, Draup role match): $148,922 (role: Data Scientist, USA)

9. Machine Learning Engineer

One of the fastest-growing roles, pulled by the move from prototypes to production AI. Model-building is increasingly assisted; the durable core is engineering AI systems that run reliably at scale.

What Stays (Durable)	What May Change (Shifting / At-Risk)
• MLOps — deployment, serving, monitoring, and retraining at scale. • Software-engineering rigor applied to ML systems (the “other half” of the job).	• Writing training/boilerplate model code by hand — AI-assisted. • Manual hyperparameter search and standard pipeline scaffolding.

What Stays (Durable)	What May Change (Shifting / At-Risk)
• Designing systems around models: data flow, latency, cost, and reliability. • Increasingly: LLM/RAG architecture, agentic-system orchestration, evaluation, fine-tuning, and safety.	• Implementing well-known architectures from scratch. • Framework-specific recall as higher-level abstractions mature.

Horizon: Near term (1–3 yr): heavy demand for production and LLM-systems skills. Medium term (3–5 yr): the model is a component; value is in robust, governed, observable AI systems.

Draup Market Data (June 2026)

- **Active JDs - last 12 months (role-filtered):** 82,852 total JDs (role: Machine Learning Engineer, Jun 2025–Jun 2026)
- **Top Job Titles (role-filtered, by posting volume):** Machine Learning Engineer (16,959) · Senior Machine Learning Engineer (7,735) · ML Engineer (1,110) — capitalisation variants combined

Skills: AI-Safe vs AI-Exposed (June 2026)

- **Durable (AI-Safe):** Machine Learning Operations/MLOps (21,429 JDs) · Model Deployment (6,126) · Model Monitoring (2,917) · Responsible AI (1,618) · RAG Architecture (2,979) · Systems Design (5,830) · Low-Latency serving (2,788) — production infrastructure requiring engineering rigor and safety accountability
- **At Risk (AI-Exposed):** Standard framework syntax: PyTorch (39,988 JDs) · TensorFlow (32,574) · Scikit-Learn (15,209) · Hyperparameter Tuning (1,589) · Feature Engineering (9,012) — framework boilerplate and standard architecture implementations now largely AI-assisted. AI tools in ML JDs: GitHub Copilot (669) · Cursor (644) · Claude (1,097)
- **Workload Shift:** AI & Data Science is the dominant workload by a wide margin and is growing, not shrinking. Application Development & Maintenance (14,260 JDs) reflects the productionisation shift. Cloud Computing (2,886 JDs) growing as ML systems move to cloud-native serving. DevOps (519 JDs) emerging as MLOps disciplines merge.

Live Signal: OpenAI doubling headcount to 8,000 by end of 2026, majority in engineering and research. Forward Deployed Engineer postings up 19x YoY (Jan 2026); demand for ML engineers deployable directly with enterprise clients surged 800% in 2025. China AI job postings grew 12x YoY in early 2026; supply-demand ratio for HPC engineers just 0.15 globally. MLOps, LLM fine-tuning, RAG architecture, and AI safety evaluation are the fastest-growing sub-specialisations.

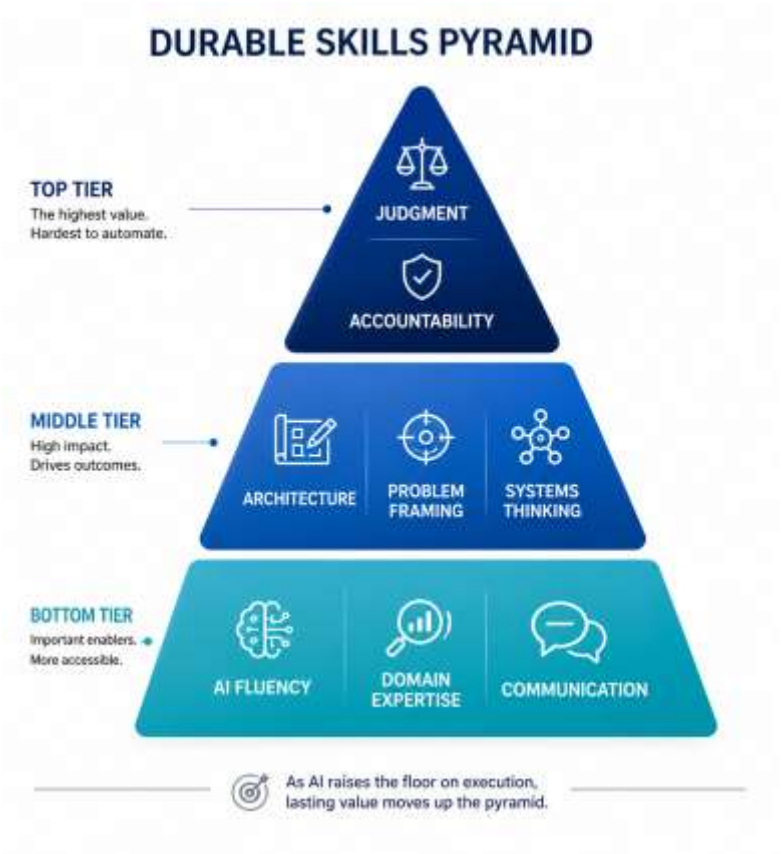

Skills That Stay Durable Across Every Role

Regardless of role, a shared set of capabilities is becoming more valuable as routine output is automated. These are the safest places to invest in development and hiring:

- Systems and architectural thinking — reasoning about trade-offs end to end.
- Judgment and critical evaluation — knowing when AI output is wrong, biased, or unsafe.
- Problem framing — turning ambiguous needs into the right thing to build or analyze.
- Domain expertise — the deepest pay premiums sit at the intersection of tech skill and domain depth.
- Communication and stakeholder influence — translating between technical and business worlds.
- Learning agility — specific tools have short half-lives; the ability to re-skill is itself the skill.
- AI fluency — directing, reviewing, and integrating AI tools is now table stakes, not a bonus.

Cross-Role Intelligence: What the Data Shows at Scale

Three findings emerge only when all nine roles are read together — findings that no single-role analysis would surface.

Finding 1: The Seniority Premium Is Structural, Not Cyclical

Across every role, YoY title growth data shows senior, staff, and principal variants outpacing generic variants. This is not a hiring freeze effect — it is a structural shift in what employers pay for. The

market is not just paying more for senior talent; it is actively deprioritising the junior execution layer that AI is absorbing. “AI Software Engineer” titles grew 57% YoY. “Staff Machine Learning Engineer”, “Principal Data Engineer”, and “Data Governance Analyst” all growing faster than their generic counterparts.

Finding 2: AI Fluency Has Crossed the Baseline Threshold

GitHub Copilot, Cursor, and Claude appear in JDs across all nine roles — not just software engineering. The combined JD count for these three tools across the full dataset exceeds 60,000 postings. Employers are not asking whether engineers use AI tools; they are asking whether engineers can direct them well. Organizations still treating AI tool proficiency as a ‘nice to have’ are already behind the market.

Finding 3: The Pay Hierarchy Maps Directly onto the Durability Hierarchy

The compensation ranking below maps almost perfectly onto the durability ranking. Roles with the highest proportion of judgment-intensive, accountability-heavy, system-level work command the highest pay. This is not coincidence — it is the market pricing durability in real time. Pay was not an input to the durability classification; the alignment is an independent market validation of the framework.

Role	Draup Role Match	Median Base Pay (USA)	Durability Tier
Machine Learning Engineer	Machine Learning Engineer	$166,764	Highest — production AI systems
Data Scientist	Data Scientist	$148,922	High — causal reasoning, framing
DevOps / SRE	DevOps Engineer	$142,433	High — reliability judgment
Data Engineer	Data Engineer	$135,847	High — AI infrastructure layer
Software Engineer	Software Dev. Engineer	$135,650	High — architecture, review
Backend Engineer	Backend Engineer	$133,386	High — security, systems depth
Frontend Engineer	Frontend Engineer	$116,270	Medium — a11y, UX judgment
QA / Test Engineer	Quality Assurance Engineer	$99,647	Medium-Low — strategy over execution
Data Analyst	Data Analyst	$88,140	Most exposed — self-service risk

Workforce Planning Implications

Three practical takeaways for workforce and skills planning:

1. **Reskill, don’t replace.** The durable core of each role is large; the shift is in the mix of tasks, not the existence of the job. Reframe exposed skills as ‘AI-directed’ rather than ‘AI-replaced’ in your JDs.
2. **Rethink the early-career path.** Routine tasks are the most automated. Invest deliberately in giving juniors design, review, and judgment exposure earlier — from month 3, not year 3.

3. **Hire and grow for judgment + domain + AI fluency.** The combination, not any single skill, is what compounds in value. The market already knows this: it is in the pay data.

Strategic Actions

The following actions are sequenced by urgency. The 0–90 day actions are confirmed by the data; delay has a measurable cost.

Priority	Action	Finding	Owner	Consequence of Inaction
0–90 days	Add Copilot, Cursor, Claude as explicit hiring criteria in all technical JDs	AI tools in 60,000+ JDs across all 9 roles	TA / Hiring Managers	Hiring candidates already behind market baseline
	Create a formal Data Governance Analyst track with distinct levelling and compensation	$22,553 pay gap: Data Analyst vs Data Governance Analyst	CHRO / Workforce Planning	Losing governance-capable analysts to competitors
	Prioritise Data Engineer hiring and retention above all other data roles in 2026	Data Engineer is prerequisite for all AI systems; second-largest role by JD volume	CHRO / Workforce Planning	AI initiatives stall on data infrastructure gaps, not model quality
3–12 months	Redesign early-career rotations: code review ownership, test strategy, architecture from month 3	Junior execution tasks are the most AI-exposed	L&D / Engineering Leadership	Junior pipeline produces talent partially obsolete at mid-level
	Build internal MLOps and production AI systems track for senior Data Scientists	ML Engineer earns $17,842 more than Data Scientist	L&D / CHRO	Senior DS talent lost to competitors offering ML Engineer titles
	Define AI Software Engineer in your job architecture before the market does it for you	AI Software Engineer titles up 57% YoY	Engineering Leadership / TA	Reactive title inflation without clear capability differentiation
	Reclassify QA roles toward Quality Engineering and Test Strategy; retire pure manual testing titles	QA median pay ($99,647) reflects structural exposure	TA / Engineering Leadership	QA pipeline atrophies as candidates move toward automation-adjacent roles
12–36 months	Add AI safety and responsible AI evaluation to senior ML and DS role requirements	Responsible AI in 1,910 DS JDs and 1,618 MLE JDs — growing fast	Engineering Leadership	Regulatory and reputational exposure as AI governance requirements tighten

The workforce planning implication is straightforward: stop organizing technical talent around the tasks people perform today and start organizing around the capabilities that remain valuable when AI can perform those tasks itself (refer the durable-versus-exposed framework in this report). The organizations that make this shift first will build workforce advantages that competitors cannot easily replicate.

Methodology and Data Sources

This report combines structured labour-market data from the Draup talent-intelligence platform with a qualitative durability framework and corroborating public signals. Four components underpin every role breakdown: the data corpus, the role mapping, the skill classification, and the live-signal validation. Each is described below so the figures can be read with the right level of confidence.

What the numbers represent

All quantitative figures reflect a June 2026 snapshot. Job-description (JD) counts cover active postings over the trailing twelve months (June 2025–June 2026) and are drawn from Draup's role-filtered index using get_job_descriptions with track_total_hits, which returns the true total count rather than a sample. Skill, tool, and workload frequencies represent the number of role-filtered JDs in which a given term appears. A single posting can contribute to many terms, so counts are not mutually exclusive and should be read as relative signal strength — not as a headcount of distinct jobs or people.

Role mapping

Each of the nine roles is matched to Draup's standardised role taxonomy — for example, "Software Engineer" maps to Software Development Engineer — so that titles, skills, pay, and workloads are drawn from a consistent population. Because employers use inconsistent titles, the top-title lists include near-duplicate variants that differ only in capitalisation or spacing; these are reported as posted rather than merged to preserve fidelity to the source. Median base pay is the USA figure from Draup's role match. Where a sub-specialisation is compensated differently, it is shown separately.

Classifying skills as AI-Safe vs AI-Exposed

Within each role, skills are sorted into Durable (AI-Safe) and At Risk (AI-Exposed) using the durability test set out in the section "How to Read This Document."

- **A skill is treated as durable** when it depends on judgment under ambiguity, system-level reasoning, accountability, or human context — qualities that current AI tools can assist with but cannot reliably own.
- **A skill is treated as exposed** when it is routine and repeatable, easy to specify, recall-based, or now self-serviceable through AI coding assistants and no-code tools.

This is a reasoned classification based on the nature of the task, not a label emitted by the platform. The JD frequencies quantify how prevalent a skill is; the durable-versus-exposed judgment reflects how automatable the underlying work is. Reasonable analysts may weigh edge cases differently, and the boundary will shift as tools improve.

AI-tool penetration and workload shift

To gauge how far automation has already reached each role, named AI coding and analysis tools (GitHub Copilot, Cursor, Claude, Codex, ChatGPT, Gemini, and similar) were counted by explicit mention in the role-filtered JDs. The scale is striking: across the Software Development Engineer role alone, GitHub Copilot appears in 22,961 JDs, Cursor in 17,905, and Claude in 12,860 — meaning employers are already writing AI-tool proficiency into job requirements, not just discussing it in think-pieces. Workload figures use Draup's workload taxonomy and express each workload as a share of the role's total JD volume, identifying where work is most concentrated and therefore most exposed, and which adjacent workloads are growing as protected zones.

Live signals and validation

Quantitative findings are cross-checked against public live signals — company announcements, hiring and layoff news, and funding events from 2025–2026 — to confirm that the direction of the data matches observable market behaviour. These signals are illustrative and qualitative; they corroborate trends rather than feed the counts, and individual events should be read as context, not as forecasts.

Limitations and shelf-life

JD-mention counts reflect employer demand as expressed in postings, which can lag actual practice and over-represent large, frequently hiring employers. The durability classification is a judgment call, and the boundary between safe and exposed will move as AI capabilities improve — this document reflects June 2026 and should be revisited annually. Figures cover the USA where pay is cited and the trailing twelve months unless stated otherwise.

The Workforce Question of the AI Era

AI is not eliminating technical work.

It is redistributing value within technical work.

The winners will not be the organizations that deploy the most AI.

They will be the organizations that redesign jobs, skills, and career pathways faster than everyone else.

-Draup
2026