Jobs Automated

What remains in a fully resourced multinational?

A conditional case study of 100,000 employees across twenty departments

Research cutoff: 2 October 2026. Future capability is evaluated as an operating present under stated assumptions. This is a research and organization-design case, not an observed company restructuring or capability-arrival forecast.

The answer

The reviewed F23 reference design reduces the fictional company's productive internal workforce from **100,000 to 15,819**, a **84.181% net reduction**. Its remaining organization comprises 14,309 core positions for residual work and coverage, plus 1,510 assumed system operations positions. It closes 84,559 original positions, retains 15,441 original workers and admits 378 outside operations hires. Original position closures and net internal reduction are different quantities.

These are calculated results of a specified future operating design. They show what follows if accepted automation, residual-work pooling, qualified coverage, revenue and cost assumptions hold. They do not establish that current AI has already removed these jobs, that these exact people can perform the remaining work, or that this is the smallest viable workforce. No actual workers are selected for release.

The remaining employee earns a seat through a necessary task, a valid decision right, a required response or compatible coverage. Seniority, degree, department, CXO title, operator title and a desire to maintain a junior pipeline confer no automatic retention. If a system can acceptably perform the work and the remaining duties fit elsewhere, the standalone position can disappear. If necessary work remains, the person must be qualified to deliver it with the new systems.

The company and the experiment

The baseline is a fictional global industrial and technology company with manufacturing, logistics, field service, software and ordinary corporate functions. It has exactly 100,000 employees. Board governance and unestimated supplier employment are outside that internal baseline. Eleven named group-level executives each own one baseline position in D01; department sponsor references add no people. Titles, grade distribution and reporting structure are illustrative rather than a claim about any real MNC.

The user-specified experiment assumes abundant compute and financial/technical resources, access to available enterprise tools, strong integration capability and continuing customer demand. These assumptions expand the feasible design space. They do not turn unavailable equipment, missing permissions, unaccepted system behavior or unknown worker competence into established capabilities.

The evidence base contains 69 primary-source records, 21 representative tool entries and 958 source/duty/capability claims. Licensed O*NET 31.0 statements anchor 244 illustrative role profiles and 1,952 selected task records, reusing 868 distinct occupational duties. All selected role scopes are reviewed for company relevance, grade authority and physical dependencies. Eight selected duties per role are substantial packets, but do not prove an exhaustive employer task inventory. See DATA_ATTRIBUTION.md, REVIEWED_ROLE_SCOPE.md and the detailed role annex.

What AI does to the work

The redesign combines reasoning/generation models, domain agents, deterministic calculations and controls, document processing, enterprise APIs, scoped identity, workflow orchestration and accepted physical equipment. A model does not need to perform arithmetic, a validated rule engine or a robot's motion control itself. Each component performs its suitable work and supplies evidence to the next stage. Tool availability, task fit and an integrated company deployment are distinct claims.

A typical chain ingests a grounded request, retrieves authorized records, generates or classifies an artifact, verifies policy and numerical constraints, executes only permitted actions, confirms the result and records an audit trail. An exception enters quarantine and bounded qualified recovery, then returns to acceptance. Success cannot bypass verification. Twelve proposed before/after chains and twenty department workbenches show this architecture without adding their overlapping savings to department totals. Explore GRAPH_OUTCOMES.html and WORKFLOW_GRAPHS.md.

Documented components include Oracle finance agents, Salesforce CRM/service agents, ServiceNow IT/HR workflows, Microsoft phishing triage, UiPath orchestration/document processing, professional legal tooling, Adobe creative workflows, industrial systems and foundation models. Their source-matched scope, alternatives and maturity appear in tool_registry.json. This representative stack is not an exhaustive list of every enterprise product or proof that the proposed connections run together.

From task automation to positions

For each task, baseline hours equal its role-group staffing times 1,600 net annual task hours times the disclosed task share. Automated hours equal baseline hours times eligibility times the scenario's autonomous labor fraction. Residual hours are baseline minus automated hours. Added review hours equal automated hours times the scenario review ratio. Residual/review work enters one primary compatible occupation/specialism pool. The model then rounds task and site/regional coverage demand and adds assumed system operations effort.

The eight task shares per role are 22%, 18%, 16%, 14%, 10%, 8%, 7% and 5%. These are synthetic weights, not an O*NET survey time measure or employer observation. Frequency survey categories are preserved separately and are not converted into company volumes. Every baseline hour has one owner. Twelve related chains are context, not extra savings.

A retired original position is compatible with some duties remaining: residual duties may consolidate in a smaller qualified pool. Conversely, a high task automation percentage does not establish that a whole occupation is eliminated. Every selected role retains some assumed residual work in F23, and no selected duty or whole role is marked as verified completely automated. FTE-equivalent work attached to a title is a workload measure, not a named person's retained job.

Three future operating states

State	Core positions	Assumed operations	Productive internal total	Net reduction	Steady profit, USD bn
F12	28,231	1,510	29,741	70.259%	3.271
F23	14,309	1,510	15,819	84.181%	3.141
U99	9,638	1,510	11,148	88.852%	1.502

F12 is the earlier conditional operating design; F23 is the primary unlocked future-as-present case; U99 is an extreme-capability stress design. Dates attached to these cases label hypothetical operating states, not predictions that models reach the specified capability on those dates. U99's lower staffing comes with more assumed automation spending, so the smallest headcount does not necessarily maximize profit.

F23 task class	Autonomous eligible baseline labor assumed
Routine digital	99%
Complex digital	96%
Physical repeatable	92%
Physical variable	60%
Decision relationship	70%

F23 adds review equal to 1.2% of modeled automated hours. Eligibility and capability fractions are aggressive assumptions. They are not benchmark accuracy, case-resolution rates, source-calibrated estimates or a statement that all tasks reach 99%-100% autonomy. Physical-variable and consequential/relationship work have distinct conditions.

Whole-company workforce map

Department	Original	Residual core	Assumed operations	Remaining productive	Net reduction
D01 Executive office, strategy and transformation	400	54	10	64	336
D02 Finance, accounting, tax and treasury	3,500	169	50	219	3,281
D03 Human resources and talent	2,000	160	35	195	1,805
D04 Legal, privacy and regulatory compliance	1,500	122	20	142	1,358
D05 Sales, account management and revenue operations	8,000	615	70	685	7,315
D06 Marketing, brand and digital commerce	2,500	174	30	204	2,296
D07 Customer support and customer success	10,000	576	80	656	9,344
D08 Product, research and physical engineering	10,000	1,118	100	1,218	8,782
D09 Software engineering and software QA	7,000	475	120	595	6,405
D10 IT, infrastructure and enterprise systems	5,000	595	150	745	4,255
D11 Cybersecurity	1,500	168	100	268	1,232
D12 Data, analytics and AI platforms	2,000	187	150	337	1,663
D13 Procurement and supplier management	2,500	210	35	245	2,255
D14 Supply chain, planning, warehousing and logistics	5,500	878	100	978	4,522
D15 Manufacturing and plant operations	26,000	4,393	250	4,643	21,357
D16 Quality, safety, environment and sustainability	3,500	2,470	50	2,520	980
D17 Field service and technical maintenance	6,000	1,250	90	1,340	4,660
D18 Facilities, administration and workplace operations	1,500	527	30	557	943
D19 Corporate communications and public affairs	600	102	10	112	488
D20 Internal audit, enterprise risk and corporate governance	1,000	66	30	96	904

The interactive original/remaining hierarchy displays every role, its level, selected duties, source and remaining pool. Baseline role counts sum to the department and company totals. Remaining pools replace titles rather than allocating fractional workload to named people. Sponsorship/dependency links add no extra headcount.

Department-by-department redesign

D01: Executive office, strategy and transformation

Reference staffing: 400 original; 54 residual core plus 10 assumed operations, leaving 64. 16 selected role groups span executive, head, manager, senior, junior, graduate. The sponsor is CEO, already counted in its owning executive group.

Scope and removal case: Company-level strategy, capital, delegated decisions and cross-functional capacity. Occupational management duties are adapted to a global group, not unrelated government office. Accepted research, forecasting, alternatives, reporting, coordination and bounded execution consolidate staff layers; even executive preparation and authorized decisions face automation tests.

Work that still requires a qualified owner: Explicit commitment ownership, contested strategic premises, risk acceptance, capital allocation and governance under required authority.

Proposed components: GPT-6.1 Sol; GPT-6 Astra. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Board evidence must reconcile to records; proposed actions must have an accountable owner. Principal failure: Unreliable strategic assumptions or an unowned consequential commitment. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Municipal appointments, public enforcement, automatic preservation of eleven CXO titles, or treating a specialist occupation as a complete executive duty inventory.

Occupational/capability anchors: S27: Chief Executives, 11-1011.00, S08: GPT-6.1 Sol model documentation, S09: GPT-6 Astra model documentation. Detailed source statements and all levels are in the role annex and hierarchy.

D02: Finance, accounting, tax and treasury

Reference staffing: 3,500 original; 169 residual core plus 50 assumed operations, leaving 219. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is CFO, already counted in its owning executive group.

Scope and removal case: Corporate accounting, payroll, tax, investment and risk duties selected in the prior reviewed finance pilot; physical cash custody stays explicit where selected. Accepted matching, ledger/payroll posting, reconciliation, forecast/report and payment workflows consolidate routine processing and preparation.

Work that still requires a qualified owner: Novel accounting/tax treatment, material control failure, contested records, authorized cash movement and independently reconciled totals.

Proposed components: Oracle Ledger Agent; Oracle Expenses Agent; Oracle Payables Agent; Oracle Payments Agent; UiPath agent/workflow orchestration; UiPath Document Understanding. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Every payment must match authorized records; totals reconcile; novel exceptions route to a named owner. Principal failure: Incorrect payment, incompatible records or an unsupported accounting treatment. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Unrelated bank lending/sales, government tax enforcement or treating regulatory reporting as proof of qualified signing authority.

Occupational/capability anchors: S28: Accountants and Auditors, 13-2011.00, S11: Overview of Payables Agent, release 26C, S21: Agents and workflows, S22: About Document Understanding. Detailed source statements and all levels are in the role annex and hierarchy.

D03: Human resources and talent

Reference staffing: 2,000 original; 160 residual core plus 35 assumed operations, leaving 195. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is CHRO, already counted in its owning executive group.

Scope and removal case: Company employment administration, justified staffing, compensation, employee relations and assessed learning; no automatic junior intake. Validated record, entitlement, scheduling, sourcing, bounded service and course-content workflows consolidate manual coordination and preparation.

Work that still requires a qualified owner: Contested employment facts, bargaining authority, appeals, competent vacancy coverage and independent learning outcomes.

Proposed components: Oracle Expenses Agent; Oracle HR agentic applications; ServiceNow HR agents. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Rights, records and approvals must remain traceable; contested cases need an explicit resolution route. Principal failure: Contested decisions or incorrect employee entitlements. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Public licensing decisions, medical decisions and an assumed compulsory hiring/training quota are outside these company duties.

Occupational/capability anchors: S31: Human Resources Specialists, 13-1071.00, S12: Fusion Agentic Applications for HR announcement, S17: Agentic workflows in HR Service Delivery. Detailed source statements and all levels are in the role annex and hierarchy.

D04: Legal, privacy and regulatory compliance

Reference staffing: 1,500 original; 122 residual core plus 20 assumed operations, leaving 142. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is General Counsel, already counted in its owning executive group.

Scope and removal case: Private-company commercial legal, privacy/compliance, evidence and document operations. Broad lawyer examples are restricted to applicable company matters. Grounded retrieval, verified clause/document generation, obligation tracking and records workflows remove manual searching, indexing and routine preparation when accepted.

Work that still requires a qualified owner: Novel or contested obligations, jurisdiction, professional standing, authorized commitments and independently checked factual/legal evidence.

Proposed components: Thomson Reuters CoCounsel; UiPath agent/workflow orchestration. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Citations and jurisdiction must be checked; commitments and representations require defined ownership. Principal failure: Wrong authority, jurisdiction mismatch or a missed material obligation. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Issuing government licenses, collecting licensing fees, testing public applicants, criminal prosecution, probate and drafting public enforcement programs. Regulatory interviews gather company permit/obligation facts; they do not exercise agency authority.

Occupational/capability anchors: S32: Lawyers, 23-1011.00, S19: CoCounsel Legal product scope, S21: Agents and workflows. Detailed source statements and all levels are in the role annex and hierarchy.

D05: Sales, account management and revenue operations

Reference staffing: 8,000 original; 615 residual core plus 70 assumed operations, leaving 685. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is Chief Commercial Officer, already counted in its owning executive group.

Scope and removal case: Industrial/technical products, services, revenue records, commercial pricing and delivery promises within the approved catalog and delegated terms. Accepted prospect research, bounded conversations, pricing/configuration, proposals, CRM entries and order processing consolidate manual sales support.

Work that still requires a qualified owner: Nonstandard customer needs, technical fit, relationship/commitment authority and verified delivery capacity.

Proposed components: Salesforce Agentforce. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Contracts, pricing and product suitability must meet the modeled commercial rules. Principal failure: A nonstandard customer need or an unauthorized deal promise. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Pharmacy sales, unrelated retail location expansion and assuming that a digital promise itself fulfills a physical order. Product demonstrations or installation coordination require actual delivery/engineering evidence.

Occupational/capability anchors: S33: Technical and Scientific Sales Representatives, 41-4011.00, S15: Agentforce platform overview. Detailed source statements and all levels are in the role annex and hierarchy.

D06: Marketing, brand and digital commerce

Reference staffing: 2,500 original; 174 residual core plus 30 assumed operations, leaving 204. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is CMO, already counted in its owning executive group.

Scope and removal case: Company product/brand strategy, factual copy, digital graphics, campaign execution and valid incremental measurement. Accepted content/asset generation, research synthesis, segmentation and campaign workflows consolidate production, administration and routine analysis.

Work that still requires a qualified owner: Claim authorization, experimental validity, original customer evidence, material rights and correction of failed commercial outcomes.

Proposed components: Salesforce Agentforce; Adobe GenStudio. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Claims trace to product evidence; results use valid measures; consequential releases have owners. Principal failure: Unsubstantiated claims, misleading content or poorly defined success metrics. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Novels, probate/estate work, unverified product or environmental claims and assuming campaign output causes incremental revenue. Trade-show work remains physical where selected.

Occupational/capability anchors: S34: Marketing Managers, 11-2021.00, S23: GenStudio for Performance Marketing guide, S15: Agentforce platform overview. Detailed source statements and all levels are in the role annex and hierarchy.

D07: Customer support and customer success

Reference staffing: 10,000 original; 576 residual core plus 80 assumed operations, leaving 656. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is Chief Customer Officer, already counted in its owning executive group.

Scope and removal case: Remote industrial-product/service support and customer-success work reviewed in the prior pilot; difficult cases, reopening and non-case work remain explicit. Accepted knowledge-backed conversation, troubleshooting, permitted account actions and case-record workflows consolidate queues and routine support.

Work that still requires a qualified owner: Novel defects, contested/high-impact cases, qualified escalation and verification that the customer's issue is actually resolved.

Proposed components: Salesforce Agentforce; Intercom Fin. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Resolution must be independently checked with repeat-contact and reopened-case metrics, not silence alone. Principal failure: Silent abandonment, reopened cases or unsafe/unapproved actions. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Insurance-policy coverage adjudication, remote completion of physical repair, and converting a vendor case-resolution percentage directly into labor reduction.

Occupational/capability anchors: S29: Customer Service Representatives, 43-4051.00, S15: Agentforce platform overview, S20: Intercom reporting metrics and attributes. Detailed source statements and all levels are in the role annex and hierarchy.

D08: Product, research and physical engineering

Reference staffing: 10,000 original; 1,118 residual core plus 100 assumed operations, leaving 1,218. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is CTO, already counted in its owning executive group.

Scope and removal case: Industrial product design, CAD/CAE, engineering analysis and actual prototype/laboratory measurement. Accepted generative design, calculation, simulation, drawing/BOM generation and automated test cells consolidate drafting, analysis and repetitive testing.

Work that still requires a qualified owner: Novel physical failure, representative test coverage, responsible release authority and qualified laboratory recovery.

Proposed components: GPT-6.1 Sol; GPT-6 Astra; Siemens industrial AI stack. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Required performance must be demonstrated on representative physical systems before release. Principal failure: A simulation-real-world gap, an invalid test or an unidentified failure mode. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Assuming a simulation validates a physical product, or treating a design/installation/test duty as digital-only merely because CAD or software is mentioned.

Occupational/capability anchors: S35: Mechanical Engineers, 17-2141.00, S25: Industrial AI technologies at CES 2026, S08: GPT-6.1 Sol model documentation. Detailed source statements and all levels are in the role annex and hierarchy.

D09: Software engineering and software QA

Reference staffing: 7,000 original; 475 residual core plus 120 assumed operations, leaving 595. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is CTO, already counted in its owning executive group.

Scope and removal case: Requirements, software architecture/code, independent tests, authorized release, technical documentation and observed production behavior. Accepted coding, test generation/execution, defect triage, release workflows and documentation consolidate developer/test/support preparation; no senior developer title is exempt.

Work that still requires a qualified owner: Ambiguous requirements, independent behavioral/security acceptance, novel production incidents and recovery authority.

Proposed components: GPT-6.1 Sol; GPT-6 Astra; Managed agent execution infrastructure. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Outcome tests, security checks and production recovery must meet the modeled service standard. Principal failure: Passing narrow tests while failing real user or production requirements. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Using the same unchallenged model as sole author/test oracle; assuming generated code or unit-test passes establish service outcomes; mixing software installation with physical hardware repair.

Occupational/capability anchors: S30: Software Developers, 15-1252.00, S08: GPT-6.1 Sol model documentation, S09: GPT-6 Astra model documentation, S04: Early-2025 developer productivity experiment, S05: Developer productivity experiment design update, S49: Scaling Managed Agents: Decoupling the brain from the hands. Detailed source statements and all levels are in the role annex and hierarchy.

D10: IT, infrastructure and enterprise systems

Reference staffing: 5,000 original; 595 residual core plus 150 assumed operations, leaving 745. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is CIO, already counted in its owning executive group.

Scope and removal case: Digital IT administration and enterprise integration plus explicitly physical endpoint/server/inventory work. Accepted remote troubleshooting, configuration, monitoring, databases, access workflows and inventory records consolidate manual administration and ticket processing.

Work that still requires a qualified owner: Qualified local hardware recovery, permission/identity exceptions, interdependent service incidents and verified restoration.

Proposed components: ServiceNow ITSM agents; UiPath agent/workflow orchestration; Managed agent execution infrastructure. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Identity, rollback, recovery and service coverage must survive the modeled incidents. Principal failure: A widespread outage, broken access boundary or unhandled hardware failure. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Calling physical cable installation, equipment inspection or defective component replacement a text-only support resolution. Remote help and on-site work are distinguished.

Occupational/capability anchors: S36: Computer User Support Specialists, 15-1232.00, S16: Agentic AI in IT Service Management, S21: Agents and workflows. Detailed source statements and all levels are in the role annex and hierarchy.

D11: Cybersecurity

Reference staffing: 1,500 original; 168 residual core plus 100 assumed operations, leaving 268. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is CISO, already counted in its owning executive group.

Scope and removal case: Authorized defense, security architecture, least privilege, bounded containment, independent attack simulation and verified recovery. Accepted monitoring, enrichment, bounded access changes, authorized attack tests and incident/report workflows consolidate alert triage and security administration.

Work that still requires a qualified owner: Novel/adversarial incidents, disproportionate containment, independent evidence and recoverable privileged actions.

Proposed components: Microsoft phishing triage agent. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Attack simulations, permissions and recovery need measurable acceptance gates. Principal failure: Correlated agent errors or an adversary exploiting trusted automation. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Unauthorized penetration testing, assuming attacker behavior stays stationary, equating alert closure with containment, or preserving AI/security operators by title alone.

Occupational/capability anchors: S37: Information Security Analysts, 15-1212.00, S18: Security Copilot Phishing Triage Agent, S26: Generative AI Risk Management Profile, AI 600-1. Detailed source statements and all levels are in the role annex and hierarchy.

D12: Data, analytics and AI platforms

Reference staffing: 2,000 original; 187 residual core plus 150 assumed operations, leaving 337. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is Chief Data and AI Officer, already counted in its owning executive group.

Scope and removal case: Dataset preparation, database integration, analytic/model development and held-out performance evaluation. The former broad MLOps title is narrowed to Model Performance Analyst because its selected source duties do not cover an entire production operations role. Accepted querying, transformations, database design/integration, feature/model experiments, report generation and repeatable evaluation consolidate routine analyst/engineering work.

Work that still requires a qualified owner: Metric/lineage integrity, distribution shift, experimental validity, high-impact interpretation and accountable platform recovery.

Proposed components: GPT-6.1 Sol; GPT-6 Astra; Managed agent execution infrastructure. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Schemas, held-out evaluation and calculation lineage must support the actual decision. Principal failure: Invalid definitions, leakage or misleading correlations. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Treating statistical association as causality, data cleaning as complete governance, model benchmarking as global operational proof, or inventing deployment/rollback duties absent from this source selection.

Occupational/capability anchors: S38: Data Scientists, 15-2051.00, S08: GPT-6.1 Sol model documentation, S26: Generative AI Risk Management Profile, AI 600-1, S49: Scaling Managed Agents: Decoupling the brain from the hands. Detailed source statements and all levels are in the role annex and hierarchy.

D13: Procurement and supplier management

Reference staffing: 2,500 original; 210 residual core plus 35 assumed operations, leaving 245. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is Chief Procurement Officer, already counted in its owning executive group.

Scope and removal case: Company sourcing, specification fit, tenders, approved purchasing, supplier performance and records. Accepted bid analysis, catalog/PO processing, order tracking, price comparison and bounded supplier communications consolidate clerical and analytic work.

Work that still requires a qualified owner: Supplier failure, contested claims, novel technical substitutions, physical supplier evidence and authorized commercial commitments.

Proposed components: Oracle Payables Agent; SAP supply-chain agentic capabilities; UiPath agent/workflow orchestration; UiPath Document Understanding. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Supplier facts, specifications and commitments must be independently reconcilable. Principal failure: Unavailable supply, misleading supplier evidence or hidden contract risk. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Assuming supplier labor vanished, a proposal equals delivered material, or a supplier plant visit is a digital record task.

Occupational/capability anchors: S39: Purchasing Agents, 13-1023.00, S13: Building the Autonomous Supply Chain, S14: The Future of the Enterprise Is Autonomous, S21: Agents and workflows. Detailed source statements and all levels are in the role annex and hierarchy.

D14: Supply chain, planning, warehousing and logistics

Reference staffing: 5,500 original; 878 residual core plus 100 assumed operations, leaving 978. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is COO, already counted in its owning executive group.

Scope and removal case: Industrial logistics, executable inventory/transport planning and warehouse stock/pick/pack handling. Accepted scheduling/routing, WMS/TMS records, automated handling and independent inventory/delivery verification consolidate planners, clerks and pickers.

Work that still requires a qualified owner: Irregular goods, route/site disruption, qualified physical recovery, imports/exports authority and verified actual delivery.

Proposed components: SAP supply-chain agentic capabilities; ABB/NVIDIA physical AI; Siemens industrial AI stack; Industrial robots and conventional controls. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Plans must be physically executable and meet the same delivery and safety criteria. Principal failure: A physical bottleneck, changed route conditions or an infeasible plan. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Retail cash-register/advertising duties, remote digital claims of physical goods movement, and assuming freight-provider workers count as eliminated internal labor.

Occupational/capability anchors: S40: Logisticians, 13-1081.00, S13: Building the Autonomous Supply Chain, S14: The Future of the Enterprise Is Autonomous, S24: ABB and NVIDIA physical AI announcement. Detailed source statements and all levels are in the role annex and hierarchy.

D15: Manufacturing and plant operations

Reference staffing: 26,000 original; 4,393 residual core plus 250 assumed operations, leaving 4,643. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is COO, already counted in its owning executive group.

Scope and removal case: Industrial production, electromechanical controls, actual assembly/handling/inspection and machinery maintenance. Accepted MES scheduling, robot/control execution, repetitive assembly/material movement and sensor-backed inspection consolidate operators and production support.

Work that still requires a qualified owner: Novel faults, actual throughput/yield, safe local recovery, qualified control changes and representative defect testing.

Proposed components: SAP supply-chain agentic capabilities; ABB/NVIDIA physical AI; Siemens industrial AI stack; Industrial robots and conventional controls. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Yield, uptime, quality, throughput and recovery must match the baseline requirements. Principal failure: Failure outside the designed cell or unsafe physical recovery. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Treating a robot program as completed physical production, ignoring cell limits/changeovers, or defaulting equipment maintenance to routine digital work.

Occupational/capability anchors: S41: Industrial Production Managers, 11-3051.00, S24: ABB and NVIDIA physical AI announcement, S25: Industrial AI technologies at CES 2026, S13: Building the Autonomous Supply Chain. Detailed source statements and all levels are in the role annex and hierarchy.

D16: Quality, safety, environment and sustainability

Reference staffing: 3,500 original; 2,470 residual core plus 50 assumed operations, leaving 2,520. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is COO, already counted in its owning executive group.

Scope and removal case: Independent quality/safety/environmental measurement and evidence, corrective action and supported sustainability reporting. Accepted inspection cells, measurement/data analysis, controlled reporting, records and corrective-action tracking consolidate routine quality/safety preparation.

Work that still requires a qualified owner: Severe/novel hazards or defects, independent representative measurement, stopping authority and verified remedial performance.

Proposed components: ABB/NVIDIA physical AI; Siemens industrial AI stack; Industrial robots and conventional controls. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Representative ground truth and severity-weighted detection must validate the system. Principal failure: A systematic missed defect or unverified physical condition. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Government prosecution or regulatory adjudication, clinical diagnosis, paper approval as proof of physical safety and environmental claims without denominators/boundaries.

Occupational/capability anchors: S42: Inspectors, Testers, Sorters, Samplers, and Weighers, 51-9061.00, S24: ABB and NVIDIA physical AI announcement, S25: Industrial AI technologies at CES 2026, S26: Generative AI Risk Management Profile, AI 600-1. Detailed source statements and all levels are in the role annex and hierarchy.

D17: Field service and technical maintenance

Reference staffing: 6,000 original; 1,250 residual core plus 90 assumed operations, leaving 1,340. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is COO, already counted in its owning executive group.

Scope and removal case: Remote diagnosis and dispatch are separate from site access, parts handling and actual industrial equipment repair and testing. Accepted diagnostics, parts/work-order coordination, routing and capable physical repair systems consolidate administration and repeatable service work.

Work that still requires a qualified owner: Inaccessible/modified assets, physical parts/repair, qualified safe access and independently tested restoration.

Proposed components: SAP supply-chain agentic capabilities; Salesforce Agentforce; Siemens industrial AI stack; Industrial robots and conventional controls. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Equipment function, safe access and verified repair quality must meet acceptance. Principal failure: A repair that cannot be reached or verified remotely or robotically. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Police/emergency dispatch, military-specific work, unsupported repair completion, and pretending a text-only diagnostic agent replaces qualified manipulation.

Occupational/capability anchors: S43: Industrial Machinery Mechanics, 49-9041.00, S15: Agentforce platform overview, S13: Building the Autonomous Supply Chain, S25: Industrial AI technologies at CES 2026. Detailed source statements and all levels are in the role annex and hierarchy.

D18: Facilities, administration and workplace operations

Reference staffing: 1,500 original; 527 residual core plus 30 assumed operations, leaving 557. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is COO, already counted in its owning executive group.

Scope and removal case: Company building services, facilities contracts, administrative records, actual maintenance/HVAC and cleaning. Accepted scheduling, records/requests, building controls and capable cleaning/maintenance systems consolidate routine workplace administration and physical service.

Work that still requires a qualified owner: Occupied-site constraints, unusual equipment faults, safe local recovery and verified building service/cleanliness.

Proposed components: ServiceNow ITSM agents; UiPath agent/workflow orchestration; Industrial robots and conventional controls. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Sites remain usable, safe and serviced under the stated coverage requirements. Principal failure: Unstructured site faults or a service outsourced without labor accounting. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Residential boiler/groundskeeping specialties outside selected sites, assumed clinical authority, and omitting facilities contractor labor from automation claims.

Occupational/capability anchors: S44: Facilities Managers, 11-3013.00, S21: Agents and workflows, S16: Agentic AI in IT Service Management. Detailed source statements and all levels are in the role annex and hierarchy.

D19: Corporate communications and public affairs

Reference staffing: 600 original; 102 residual core plus 10 assumed operations, leaving 112. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is CEO, already counted in its owning executive group.

Scope and removal case: Accurate company/public-affairs communications, consented media research, factual corporate writing and software/document localization. Accepted factual copy, media monitoring, publication workflows and terminology-grounded translation consolidate content/reporting and routine language work.

Work that still requires a qualified owner: Contested facts, reputational crises, material commitments, rights and independent meaning/claim validation.

Proposed components: GPT-6.1 Sol; GPT-6 Astra; Adobe GenStudio. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Facts and commitments remain traceable; contested releases follow accountable decisions. Principal failure: A false public claim or an unowned institutional commitment. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Fiction/novels, tourist or school interpreting, creating unsupported claims, and assuming fluent generated language guarantees factual or cultural fidelity.

Occupational/capability anchors: S45: Public Relations Specialists, 27-3031.00, S08: GPT-6.1 Sol model documentation, S23: GenStudio for Performance Marketing guide. Detailed source statements and all

levels are in the role annex and hierarchy.

D20: Internal audit, enterprise risk and corporate governance

Reference staffing: 1,000 original; 66 residual core plus 30 assumed operations, leaving 96. 12 selected role groups span head, manager, senior, junior, entry, graduate. The sponsor is Chief Risk Officer / audit committee, already counted in its owning executive group.

Scope and removal case: Company-independent assurance, control testing, private fraud evidence, enterprise risk analysis and audit records. Accepted population extraction, reconciled tests, anomaly detection, evidence indexing and report preparation consolidate manual sampling/records and repeated assurance analysis.

Work that still requires a qualified owner: Independent challenge, contested evidence, material risk acceptance, validated scope and verified remediation.

Proposed components: GPT-6.1 Sol; GPT-6 Astra; Thomson Reuters CoCounsel; UiPath Document Understanding. The model proposes these components; feature fit, permissions and combined execution remain untested.

Acceptance and failure: Independent evidence, conflict checks and accountable findings must survive adversarial review. Principal failure: Automation evaluating itself with shared blind spots. Qualified recovery is bounded and returns to verification before success.

Source-scope boundary: Arresting suspects, serving public-authority subpoenas, unrelated bank lending, government tax enforcement and circular assurance by the same unchecked operating agent.

Occupational/capability anchors: S46: Compliance Officers, 13-1041.00, S28: Accountants and Auditors, 13-2011.00, S26: Generative AI Risk Management Profile, AI 600-1, S19: CoCounsel Legal product scope. Detailed source statements and all levels are in the role annex and hierarchy.

Who remains, and why

The core comprises 120 proposed occupation/specialism pools. Tax and privacy specialisms remain separated; unrelated occupations are not pooled merely because they share a broad soft skill. The remaining work is residual production, exceptions, review and coverage under the scenario. The design adds twenty assumed operations pools. Every pool has owned tasks and declared domain/authority requirements.

The core contains 10,729 task-work seats under existing local rounding, with 3,580 additional coverage seats. Removing coverage leaves 12,239 positions including unchanged operations, but that counterfactual does not establish global viability. Doubling coverage produces 22,933 positions. Increasing review to 5% produces 18,503; reduced physical autonomy produces 21,618. These are assumption tests, not confidence intervals.

Occupation-specific presence floors need particular scrutiny. Three quality/environment pools increase beyond baseline under separate 120-site coverage assumptions. A physical duty does not prove that its occupation needs separate continuous coverage at every site. Assessed workers may share compatible duties; simultaneous posts, qualification, authority, language, site, leave, travel and failure demand may also require more people. The 1,510 operations seats are equally subject to automation and shared-platform challenge.

Twenty-one synthetic assignment cases test competence, authority, valid credentials, language, jurisdiction, physical location, availability, daily/weekly capacity, rest and simultaneous work. A qualified graduate can pass; an unqualified executive fails. These tests demonstrate declared assignment rules, not an actual international roster. The matcher checks one assignment and rest; it is not an exhaustive minimum-staffing optimizer. See GLOBAL_WORKFORCE_REPORT.md.

Retired work and necessary admissions

Original position IDs enter a terminal conditional-retirement state and have no reopen transition. New productive admissions require a necessary open pool seat, accepted conditions, qualifications and bounded residual/review work. A renamed role cannot authorize the old automated manual core. Training is admitted against credible future vacancies and does not count as productive capacity. Recovery is incident-specific, temporary and bounded, with restoration and re-verification.

The local admission prototype passes twenty golden cases and six stateful reservation checks for closed IDs, work budgets, last-seat capacity, idempotent requests, training limits and accumulated incident recovery. It uses declared synthetic evidence and single-process state. Actual HR, payroll, purchasing, authenticated work evidence and concurrent reservations are not integrated. Terminal IDs prove a control rule in this model, not that changed real-world work can never create a legitimate new need. See DURABLE_WORKFORCE_MODEL.md.

Talent, learning and company continuity

There is no mandatory junior pipeline. The company can recruit assessed future-ready entrants, retrain compatible original workers, buy visible qualified temporary capacity or train novices for a credible remaining need. Company-specific practice, conduct and culture need assessment rather than a fictional pass percentage treated as evidence. A trainee enters a productive seat only after the competence/culture gates, job acceptance and a matching vacancy.

The ready-hiring continuity design fills 2029 targets after 1,208 modeled productive exits, while using 693.3 qualified outside backfill FTE-years for hiring delays and renewal. It yields USD 3.031bn operating profit under fixed output/revenue. The optional novice route enrolls 158 and uses 20 mentors, giving 15,997 total internal employees and USD 3.019bn profit. Trainees and mentors sit outside the productive-core headline.

The seven five-year alternatives include AI-supported mentoring, zero-human-mentor and labor-scarcity hypotheses. Constant learning outcomes in an AI-mentoring sensitivity are assumptions. Zero expected attrition/growth leads to no cohort; zero competence success halts enrollment. A scarce-hiring scenario can appear profitable because revenue is fixed, yet fails the labor condition. It must not be promoted as successful workforce reduction. Renewal, mentor turnover, recruiting and outside backfill are counted. Later physical-asset cash replacement and real qualification/hiring supply remain open conditions.

Economics and transition

F23 assumes USD 20bn revenue, USD 10bn other operating costs, the disclosed remaining-worker costs, USD 2.5bn annual automation operating cost and USD 12bn automation capex depreciated over five years. Its steady profit before financing, tax and continuity adjustments is USD 3.14132bn. The transition includes upgraded retained salaries, qualification fees, exit payments, new-hire recruiting, outside implementation labor, capex, financing and a restricted replacement reserve.

The base needs USD 8.914bn additional funding beyond opening cash to maintain the assumed buffer. It has a full-year 2027 loss and twelve loss months. Cumulative modeled profit through 2030 trails the

unchanged-workforce comparator by USD 1.074bn. Higher steady profit does not establish transition profitability, cash payback or superior cumulative value.

Twelve financial variants expose equity/debt, lower exit cost, cost/revenue stress and operator automation. A lower exit-cost hypothesis meets modeled monthly profit, but actual obligations and funding access are not known. Tax, working capital, supplier quotes, quality/revenue loss, vendor labor and transition training time are unestimated. Restricted reserve is cash, not a second depreciation expense. See TRANSITION_CONTINUITY_REPORT.md and its tables.

What the evidence supports and challenges

Primary evidence supports useful task productivity and bounded substitution. It also distinguishes ordinary assistance from integrated redesign, paid-work spending from employee headcount, accepted case coverage from labor autonomy, and retrospective benchmarks from deployed outcomes. The critical challenge tests fifteen broad claims and records evidence that would change each conclusion. Source affiliations, inspected sections and limits remain visible; no vendor resolution metric calibrates this company's future labor rate.

A contracted-online-labor payments preprint provides bounded substitution evidence; a broad executive survey reports small ordinary-adoption effects. A real single-disease clinic trial demonstrates autonomous value, while a later diagnostic benchmark shows selective autonomy with substantial deferral. These results can coexist because they study different work and interventions. S63: Firm Data on AI, March 2026 revision, S64: Payrolls to Prompts: Firm-Level Evidence on the Substitution of Labor for AI, S65: On-premise medical AI agents for reliable clinical decision-making, S66: Autonomous AI increases real-world specialist clinic productivity.

The project does not select a reduction percentage to prove skeptics wrong. It challenges blanket claims of impossibility and applies the same scrutiny to blanket claims of complete enterprise autonomy. Evidence that raises remaining work or reveals a failed condition changes the result. The detailed design rose from 13,295 to 15,819 after the all-role scope review corrected physical dependencies and site demand; the earlier result is preserved as history.

Industries and the next operating stages

Digital services, financial services, healthcare and physical industry require different task inventories, staffing weights, permissions, asset fleets and outcome gates. The industry comparison describes these changes and a conditional 2027-2030 deployment sequence. It does not transfer the hybrid company's percentage to other industries or predict capability arrival. A separate numeric industry case requires its own source-grounded duties, work volumes, coverage and economics.

The staged design first validates bounded work, then integrates chains and qualified pools, sustains vacancies/renewal/learning, and re-tests changed work, systems and replacement costs. A company must repeatedly pass operating conditions. A five-year synthetic flow and unchanged customer-demand assumption do not prove perpetual viability.

Model choice for this project

GPT-6.1 Sol at extra-high reasoning is the current research/implementation base. Keep it for orchestration and consistent artifact generation. GPT-6 Astra is a documented alternative for demanding reasoning, but no project-specific head-to-head test establishes that switching improves this report. A useful comparison would blind-review the same source packets, difficult task classifications and numerical counterarguments before paying for a broader run. Model branding does not replace source inspection, deterministic checks or domain review. No model switch or paid API experiment was performed. S08: GPT-6.1 Sol model documentation, S09: GPT-6 Astra model documentation.

Success and the remaining acceptance conditions

A successful conditional graph requires accepted output, necessary human coverage and rights, positive relevant economics, funded transition and qualified recovery. A local green state shows these conditions assumed; a failed condition blocks the state. Real success requires traceable executed outputs and independent acceptance. The hierarchy and graph label consolidation candidates and assumed pools, while keeping verified complete role removal at zero.

To convert this research case into an operating claim, supply the actual inventory/volumes, deployed task/chain acceptance, exception/review time, qualified site/shift/authority rosters, authentic retirement/admission enforcement and measured total costs/output. These are open acceptance conditions. They do not prevent a complete transparent scenario report, but they prevent asserting observed automation, a proven global minimum or inevitable displacement.

Reproduction and reading map

The standard-library detailed pipeline is research, evidence registers, role model, workflow graphs, role report/source verifier, transition/continuity, workforce controls, graph outcomes, global assignment and evidence/industry challenge. The main report and hierarchy are presentation outputs. Rebuilding role research resets derived fields, so use the documented order. Inputs, outputs and runtime/export checks have separate evidence states; changed presentation/data hashes invalidate old browser proof.

- Main narrative: this report. Full occupational duties and company-role analysis: ROLE_LEVEL_REPORT.md.

- Finance/support details: PILOT_ANALYSIS.md. Role-scope review: REVIEWED_ROLE_SCOPE.md.

- Interactive hierarchy: JOBS_AUTOMATED.html. Accepted-success chains: GRAPH_OUTCOMES.html.

- Remaining skills/coverage: GLOBAL_WORKFORCE_REPORT.md. Necessary admission: DURABLE_WORKFORCE_MODEL.md.

- Economics/talent continuity: TRANSITION_CONTINUITY_REPORT.md. Evidence challenge: RESEARCH_CHALLENGE.md. Industry stages: INDUSTRY_OUTLOOK.md.

- Sources, claims, tools and editable assumptions: JSON/CSV registries linked in README.md. Requirement status and exact verification limits: COMPLETION_AUDIT.md.

Primary-source reference register

Cutoff 2 October 2026. Selected source sections were inspected; full methodology, independent replication and company deployment are distinct checks. Source limitations and inspection depth are in the accompanying ledger. O*NET statements are licensed under CC BY 4.0 and adapted; no US Department of Labor endorsement.

S01: ILO global exposure overview
International Labour Organization. 2025-05-20.
https://www.ilo.org/resource/news/one-four-jobs-risk-being-transformed-genai-new-ilo%E2%80%93nask-global-index-shows

S02: Canaries in the Coal Mine, revised publication
Stanford Digital Economy Lab. 2026-08-12. https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/

S03: Generative AI at Work, revised abstract
Brynjolfsson, Li and Raymond / arXiv. 2024-11-06. https://arxiv.org/abs/2304.11771v2

S04: Early-2025 developer productivity experiment
METR. 2025-07-10. https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/

S05: Developer productivity experiment design update
METR. 2026-02-24. https://metr.org/blog/2026-02-24-uplift-update/

S06: Economic Index report: Cadences
Anthropic. 2026-06-26. https://www.anthropic.com/research/economic-index-june-2026-report

S07: Scenarios for our Economic Future
Anthropic Institute. 2026-09. https://www.anthropic.com/institute/econ-scenarios

S08: GPT-6.1 Sol model documentation
OpenAI. Date not established. https://developers.openai.com/api/docs/models/gpt-6.1-sol

S09: GPT-6 Astra model documentation
OpenAI. Date not established. https://developers.openai.com/api/docs/models/gpt-6-astra

S10: OpenAI model selection overview
OpenAI. Date not established. https://developers.openai.com/api/docs/models

S11: Overview of Payables Agent, release 26C
Oracle. Date not established. https://docs.oracle.com/en/cloud/saas/financials/26c/fappp/overview-of-payables-agent.html

S12: Fusion Agentic Applications for HR announcement
Oracle. 2026-04-09. https://www.oracle.com/news/announcement/oracle-introduces-fusion-agentic-applications-for-hr-2026-04-09/

S13: Building the Autonomous Supply Chain
SAP. 2026-05. https://news.sap.com/2026/05/more-autonomous-supply-chain/

S14: The Future of the Enterprise Is Autonomous
SAP. 2026-05. https://news.sap.com/2026/05/future-enterprise-autonomous/

S15: Agentforce platform overview
Salesforce. Date not established. https://www.salesforce.com/agentforce/

S16: Agentic AI in IT Service Management
ServiceNow. 2026-09-10. https://www.servicenow.com/docs/r/it-service-management/now-assist-for-it-service-management-itsm/now-assist-itsm-ai-agents-use-cases.html

S17: Agentic workflows in HR Service Delivery
ServiceNow. 2026-09-10.
https://www.servicenow.com/docs/r/employee-service-management/now-assist-for-hrsd/now-assist-hrsd-ai-agents-use-cases.html

S18: Security Copilot Phishing Triage Agent
Microsoft. Date not established. https://learn.microsoft.com/en-us/defender-xdr/phishing-triage-agent

S19: CoCounsel Legal product scope
Thomson Reuters. Date not established. https://legal.thomsonreuters.com/en/products/cocounsel-legal

S20: Intercom reporting metrics and attributes
Intercom. Date not established. https://www.intercom.com/help/en/articles/7022438-reporting-metrics-attributes

S21: Agents and workflows
UiPath. Date not established. https://docs.uipath.com/agents/automation-cloud/latest/user-guide/agent-vs-workflows

S22: About Document Understanding
UiPath. Date not established.
https://docs.uipath.com/document-understanding/automation-cloud/latest/user-guide/about-document-understanding

S23: GenStudio for Performance Marketing guide
Adobe. 2026-03-09. https://experienceleague.adobe.com/en/docs/genstudio-for-performance-marketing/user-guide/home

S24: ABB and NVIDIA physical AI announcement
ABB Robotics. Date not established.
https://new.abb.com/news/detail/134030/prsrl-abb-robotics-partners-with-nvidia-to-deliver-industrial-grade-physical-ai-at-scale

S25: Industrial AI technologies at CES 2026
Siemens. 2026-01.
https://press.siemens.com/global/en/pressrelease/siemens-unveils-technologies-accelerate-industrial-ai-revolution-ces-2026

S26: Generative AI Risk Management Profile, AI 600-1
NIST. 2024. https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence

S27: Chief Executives, 11-1011.00
O*NET OnLine. 2026. https://www.onetonline.org/link/summary/11-1011.00

S28: Accountants and Auditors, 13-2011.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/13-2011.00

S29: Customer Service Representatives, 43-4051.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/43-4051.00

S30: Software Developers, 15-1252.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/15-1252.00

S31: Human Resources Specialists, 13-1071.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/13-1071.00

S32: Lawyers, 23-1011.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/23-1011.00

S33: Technical and Scientific Sales Representatives, 41-4011.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/41-4011.00

S34: Marketing Managers, 11-2021.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/11-2021.00

S35: Mechanical Engineers, 17-2141.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/17-2141.00

S36: Computer User Support Specialists, 15-1232.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/15-1232.00

S37: Information Security Analysts, 15-1212.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/15-1212.00

S38: Data Scientists, 15-2051.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/15-2051.00

S39: Purchasing Agents, 13-1023.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/13-1023.00

S40: Logisticians, 13-1081.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/13-1081.00

S41: Industrial Production Managers, 11-3051.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/11-3051.00

S42: Inspectors, Testers, Sorters, Samplers, and Weighers, 51-9061.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/51-9061.00

S43: Industrial Machinery Mechanics, 49-9041.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/49-9041.00

S44: Facilities Managers, 11-3013.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/11-3013.00

S45: Public Relations Specialists, 27-3031.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/27-3031.00

S46: Compliance Officers, 13-1041.00
O*NET OnLine. Date not established. https://www.onetonline.org/link/summary/13-1041.00

S47: Impact of modeling assumptions on time horizon results
METR. 2026-03-20. https://metr.org/notes/2026-03-20-impact-of-modelling-assumptions-on-time-horizon-results/

S48: Measuring AI agent autonomy in practice
Anthropic. Date not established. https://www.anthropic.com/news/measuring-agent-autonomy

S49: Scaling Managed Agents: Decoupling the brain from the hands
Anthropic. 2026-04-08. https://www.anthropic.com/engineering/managed-agents

S50: O*NET OnLine content license
U.S. Department of Labor / O*NET. Date not established. https://www.onetonline.org/help/license

S51: O*NET 31.0 occupational database
U.S. Department of Labor / O*NET. 2026-08. https://www.onetcenter.org/database.html

S52: Building Customer Support AI Agents at 100M-User Scale, v2
Gupta and coauthors / arXiv. 2026-06-13. https://arxiv.org/abs/2606.08867v2

S53: Four generally available Oracle ERP agents
Oracle Fusion Development. 2026-06-22. https://blogs.oracle.com/fusioninsider/agentic-ai-in-erp-four-agents-you-can-use-today

S54: Salesforce Agentforce sales implementation
Salesforce. 2025-11-06. https://www.salesforce.com/customer-stories/agentforce-for-sales-implementation/

S55: Fin resolutions and outcomes
Intercom. 2026-03-12. https://www.intercom.com/blog/from-resolutions-to-outcomes-evolving-how-fin-delivers-value/

S56: Intelligence Center X industrial deployments
Siemens. 2026-06-01. https://news.siemens.com/en-us/siemens-intelligence-center-x/

S57: Agentforce customer metrics
Salesforce. Date not established. https://www.salesforce.com/agentforce/metrics/

S58: FORCE-Bench enterprise finance benchmark
Pauli and coauthors. 2026-07-11. https://arxiv.org/abs/2607.19409

S59: EnterpriseBench strategic reasoning benchmark
Yang and coauthors. 2026-09-29. https://arxiv.org/abs/2609.37658

S60: READY or Not: Reliable Enterprise Agent Deployment
Research authors / arXiv. 2026-09-02. https://arxiv.org/abs/2609.02095

S61: World Robotics 2026 president report
International Federation of Robotics. 2026-09-24. https://ifr.org/ifr-press-releases/news/presidents-report-by-jane-heffner-3-2026

S62: Shifting Work Patterns with Generative AI
Dillon, Jaffe, Immorlica and Stanton / NBER. 2025-11. https://www.nber.org/papers/w33795

S63: Firm Data on AI, March 2026 revision
Yotzov and coauthors / NBER. 2026-03. https://www.nber.org/papers/w34836

S64: Payrolls to Prompts: Firm-Level Evidence on the Substitution of Labor for AI
Ryan Stevens / arXiv. 2026-01-28. https://arxiv.org/abs/2602.00139v1

S65: On-premise medical AI agents for reliable clinical decision-making
Zhang and coauthors / Nature Medicine. 2026-09-15. https://www.nature.com/articles/s41591-026-04609-x

S66: Autonomous AI increases real-world specialist clinic productivity
Abramoff and coauthors / npj Digital Medicine. 2023-10-04. https://www.nature.com/articles/s41746-023-00931-7

S67: Does Generative AI Narrow Education-Based Productivity Gaps?
Cruces and coauthors / NBER. 2026-05. https://www.nber.org/papers/w34851

S68: Generative AI and Sales Productivity: Field Experiments in Online Retail, v6
Fang, Yuan, Zhang, Donati and Sarvary / arXiv. 2026-06-29. https://arxiv.org/abs/2510.12049v6

S69: Artificial intelligence in UK financial services, 2024 survey
Bank of England and Financial Conduct Authority. 2024.
https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024