AGI Labor Transition & Enterprise Agentic Deployment Runbook

Version 4.0 — Unified Operational Manual

Author Nathan Lim
Contact contact@nathanlim.io
Date March 2026
Version 4.0
Scope Business model, cloud architecture, workforce pipeline, governance, cybersecurity, GTM, and execution phases

Not legal advice. This runbook summarizes strategic planning and includes governance and legal risk considerations. Employment law, privacy, and procurement rules vary by jurisdiction; verify with qualified counsel for any production deployment.

Table of Contents

I. Executive Summary & Key Metrics

The AGI Transition Framework is an operating system for governed AGI deployment and human workforce transition. It aligns two converging forces: enterprises deploying agentic AI that need safety, compliance, and operational coverage; and workers (displaced, underemployed, and new graduates) who need paid pathways into credible, measurable roles.

This runbook is the unified operational manual. It merges strategic thesis, cloud architecture, NVIDIA stack integration, cybersecurity threat modeling, governance frameworks, workforce pipeline design, and go-to-market sequencing into a single reference document designed for a CEO or founding team to use as both a decision framework and an execution guide.

Core Thesis

> *Do not sell 'AGI' as a buzzword. Sell reliable outcomes: throughput, quality, auditability, and controlled risk. Use AGI vendors as components, not as the core product. The durable moat is governance, workflow design, outcome data, and workforce pipeline design — not the base model.*

Market Snapshot (2025–2026)

88% Enterprises using AI in 1+ function (McKinsey Nov 2025)	**7%** Have fully scaled AI enterprise-wide (McKinsey Nov 2025)	**23%** Scaling agentic AI in 1+ function (McKinsey Nov 2025)	**$4.44M** Global avg cost of a data breach (IBM 2025)
$107B Q3 2025 cloud infra spend (Synergy Research)	**63%** Market held by AWS+Azure+GCP (Q3 2025)	**97%** AI-breached orgs lacking controls (IBM 2025)	**$670K** Added breach cost from shadow AI (IBM 2025)

Three-Year Strategic Sequence

Year	Cloud Anchor	Market Focus	Deployment Priority	Key Milestone
Year 1	AWS-first	US / SMB consulting	IT/SecOps, Customer Support	First paid pilots, SOC 2 initiated
Year 2	Azure expansion	US enterprise, single-tenant	Back-Office, Compliance Ops	Managed ops retainers, ISO 27001

Year	Cloud Anchor	Market Focus	Deployment Priority	Key Milestone
Year 3	GCP + multi-cloud	Global expansion	Recruiting Ops, cross-vertical	Platform subscription, FedRAMP path

What the Framework Builds (In Sequence)

- AGI Deployment Engine: vendor/tool selection, workflow tailoring, integrations, safe execution.
- Governance Engine: approvals, audit logs, escalation rules, evaluation harness, incident response.
- Talent Transition Engine: paid training, tiered roles, promotion path into higher-value work.

II. Market Intelligence: 2025–2026 Signals

Enterprise AI Adoption

McKinsey's November 2025 State of AI survey (1,993 respondents, 105 countries) confirms that AI adoption is broad but scaling remains the primary constraint. 88% of organizations report regular AI use in at least one business function, up from 78% in 2024. However, only 7% report full enterprise-wide scaling, and approximately one-third have begun scaling in any capacity.

Agentic AI is emerging but early. 23% report scaling an agentic system somewhere in the enterprise; 39% are experimenting. In any single business function, no more than 10% report scaling agents. The functions where agentic use is most commonly reported are IT and knowledge management — precisely the workflow families the AGI Transition Framework targets first in its FIFO sequence.

Among 25 organizational attributes tested, workflow redesign had the single largest effect on EBIT impact from AI. High performers (5%+ EBIT from AI, ~6% of respondents) are 2.8x more likely to report fundamental workflow redesign and far more likely to have human-in-the-loop validation (65% vs 23%).

Enterprise AI Adoption Trajectory (% Using AI in 1+ Function)

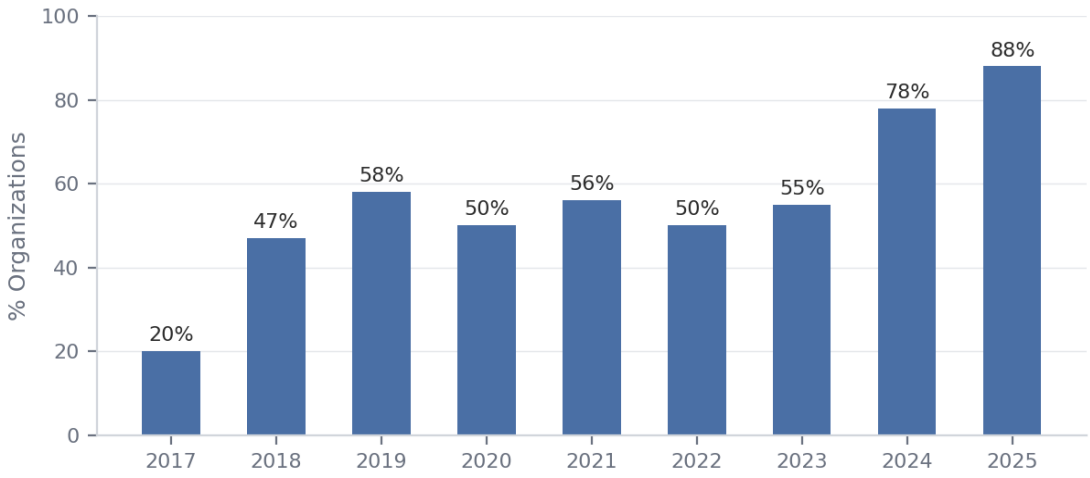

Source: McKinsey Global Survey on the State of AI, 2017–2025

Generative AI Adoption Wave

The Federal Reserve Bank of St. Louis documented generative AI adoption among US adults aged 18–64 reaching 54.6% by August 2025. Microsoft's AI Economy Institute reported broad global adoption in the second half of 2025 with uneven regional distribution — a leading indicator for how shadow AI emerges when governance cannot keep pace.

Cloud Infrastructure Market

The global cloud infrastructure market reached $107 billion in Q3 2025, up 28% year-over-year and crossing the $100B quarterly threshold for the first time (Synergy Research Group). Full-year 2025 cloud revenues are projected to exceed $400 billion. GenAI-specific cloud services grew 140–180% in Q2 2025.

Cloud Infrastructure Market Share — Q3 2025

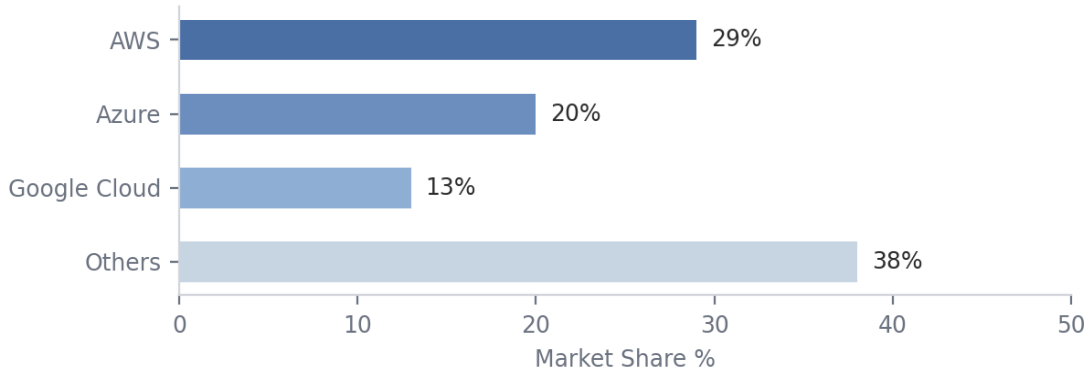

Source: Synergy Research Group, Q3 2025. Total market: $107B/quarter.

AWS leads at 29%, Azure at 20%, Google Cloud at 13% — collectively ~63% of the market. Procurement reality: most enterprises standardize on one hyperscaler, so portability must be engineered. This validates the AGI Transition Framework's control-plane / runtime-plane / data-plane separation architecture.

Cybersecurity Cost Landscape

IBM's 2025 Cost of a Data Breach report (600 organizations, 17 industries) found the global average dropped to $4.44M — down 9% from $4.88M — driven by faster AI-powered detection (mean 241 days to identify/contain, lowest in nine years). However, AI adoption is outpacing security governance:

- 97% of organizations with AI-related security incidents lacked proper AI access controls.
- 63% of breached organizations have no AI governance policies in place.
- Shadow AI was involved in 20% of breaches, adding $670K to average breach cost.
- US breach costs rose to $10.22M despite global average declining.
- Healthcare costliest sector for 14th consecutive year at $7.42M average.
- Extensive AI-powered security saved $1.9M per breach on average.

> *This validates the Framework's governance-first approach: the market gap is not AI capability — it is AI oversight. Boards will increasingly treat AI governance as part of security posture.*

III. System Architecture: What the Framework Is

The AGI Transition Framework aligns two converging forces: (a) enterprises deploying agentic AI/AGI needing safety, compliance, and operational coverage; (b) workers (laid-off, underemployed, new graduates) needing paid pathways and credible experience.

Definitions and Primitives

Term	Definition
Agentic Workflow	A multi-step process where an AI agent uses tools (APIs, databases, ticketing, browser) to produce or execute outcomes.
Human-in-the-Loop	Structured checkpoints where humans approve, correct, or escalate outputs.
Cell	The smallest scalable delivery unit (engineer + domain lead + QA/governance + associates).
Tier	A level of task complexity and authority. Movement is competency-based.
Lane	An intake pathway into the tier system (new grads, displaced workers, advancement).

Problems Solved

For Enterprises

- Vendor selection confusion — which model/agent stack fits the workflow.
- Integration burden — tooling, identity, permissions, data flow.
- Trust gap — auditability, approvals, rollback, monitoring.
- Human coverage for exceptions and policy interpretation.
- Change management — rollouts, training, measurement, accountability.

For Workers (Displaced + Underemployed)

- 'Training warehouse' failure: courses without jobs at the end.
- Experience translation failure: domain skills not expressed as employer-ready evidence.
- Time-to-income problem: long retraining cycles create financial instability.

> *Counter-design: Employment-first, paid placement, training while producing real deliverables, measurable progression.*

For Students / New Grads

- Alternative to scarce internships: paid Tier 1 AGI ops roles.

- Portfolio of audited work artifacts (rubric-scored outputs, QA logs).
- Ladder into implementation, governance, solutions, and leadership roles.

Compounding Moat

- Workflow templates + SOPs by function and vertical.
- Evaluation harnesses, scenario libraries, and regression suites.
- Governance controls (approval matrices, escalation rules, audit exports).
- Training curricula tied to real workflows and quality metrics.
- Outcome datasets: supervision mins/100 tasks, rework rates, audit readiness time.
- Cloud-portable control planes and security-constrained autonomy.

IV. Workflow FIFO Prioritization

All workflow families are targeted. Enterprise leadership sequences by three factors: measurability and baseline availability, integration and rollback feasibility, and legal/regulatory exposure.

Priority	Workflow Family	Example Use Cases	Why This Position	Risk Profile
1st	IT / SecOps	Ticket enrichment, KB ops, alert triage, incident response	Instrumented data, strong governance norms, reversible changes, clear KPIs	Low — highly measurable, rollback-safe
2nd	Customer Support	Triage, response drafting, QA, escalation routing	High volume, immediate ROI, natural human checkpoints	Low-Medium — customer-facing but controllable
3rd	Back-Office	Procurement, AP/AR, billing, claims processing	High ROI but more ERP coupling and financial correctness constraints	Medium — agent decides, automation executes, controls verify
4th	Compliance Ops	Evidence collection, audit prep, control testing	Dramatically stronger after audit logs and versioning are mature	Medium — requires mature logging infrastructure
5th	Recruiting Ops	Screening assistance, scheduling, interview packets	Touches employment decisions and discrimination risk	High — NYC LL144, CO SB24-205, EEOC scrutiny

Decision Framework

- Measurability: Does a quantitative baseline exist? Can before/after deltas be calculated within 4–8 weeks?
- Integration feasibility: Can the workflow be instrumented without deep ERP/HRIS coupling? Is rollback safe?
- Legal/regulatory exposure: Does it touch employment decisions, protected classes, financial obligations, or regulated data?
- Stakeholder readiness: Is there a budget owner who will sponsor the pilot and accept measured results?

Workflows scoring high on measurability and feasibility but low on regulatory exposure move first. Employment-adjacent workflows move last regardless of ROI potential.

V. Prototyping Lifecycle & Enterprise Gates

A workable lifecycle preserving startup speed while meeting enterprise expectations. Each gate enforces minimum requirements before the next phase begins.

Gate	Phase	Scope	Key Deliverables	Exit Criteria
Gate 0	Scope + Risk Tier	Define workflow boundaries, success metrics, data classification, tool inventory, autonomy level	Workflow spec, risk tier, autonomy limits (draft-only / read-only / propose / execute-low-risk)	Signed-off scope doc with stakeholder approval
Gate 1	Sandbox Prototype	No sensitive data, no writes. End-to-end orchestration with synthetic data	Scenario libraries, logging/eval harnesses, regression tests, orchestration validated	All test scenarios pass; eval harness produces repeatable scores
Gate 2	Production Pilot	Humans decide. Add SSO, RBAC, audit exports, approval queues, incident playbooks	Working workflow, approval gates, audit logs, dashboard, runbook, rollback plan	Baseline vs measured deltas documented; no unresolved P0/P1
Gate 3	Controlled Expansion	Expand autonomy with evidence only. Rate limits, circuit breakers, canary releases	Auto-rollback on drift, regression monitoring, expanded coverage	Sustained improvement 30+ days; governance review complete

Autonomy Progression Model

Level	Agent Authority	Human Role	Evidence to Advance
Draft-Only	Produces outputs, takes no action	Reviews and decides all actions	Scenario test pass rate > 95%
Read-Only Tools	Reads from APIs/DBs, cannot write	Approves any state changes	No data leakage in 100+ test runs
Propose Actions	Recommends actions with confidence scores	Approves or rejects each proposal	Acceptance rate > 80% over 2 weeks
Execute Low-Risk	Executes pre-approved action classes autonomously	Monitors; handles exceptions	Zero P0 over 30 days; rework < 5%
Execute Standard	Handles full workflow with circuit breakers	Spot-checks; governs; handles edge cases	Sustained KPIs; governance audit passed

VI. Cloud Reference Architectures

Portability: Control Plane vs Runtime Plane vs Data Plane

A robust multi-cloud strategy separates three architectural layers. This makes 'AWS-first now, multi-cloud later' feasible because the control plane stays consistent while runtime and data mapping changes by cloud.

Plane	Responsibility	Examples	Portability
Control	Policy engine, tenant config, approvals, audit exports, evaluation gates	Custom policy service, approval workflows, config mgmt	Cloud-agnostic; deploy as containers
Runtime	Agent orchestration, tool runners, inference endpoints, queues	EKS/AKS/GKE, Lambda/Functions, Bedrock/Azure AI/Vertex	Abstracted via container orchestration; inference via NIM
Data	RAG corpora, workflow state, artifacts, logs	PostgreSQL, vector stores, object storage, SIEM export	Standard DB protocols; object storage is portable

AWS Year-One Reference Stack

Layer	AWS Services	Purpose
Identity	IAM, IAM Identity Center, scoped service accounts	Fine-grained RBAC, SSO federation, least privilege
Network	VPC segmentation, PrivateLink, KMS (customer-managed)	Tenant isolation, encrypted transit, data boundaries
Compute	EKS, Lambda, Step Functions, ECS	Long-lived runtime, glue code, deterministic orchestration
AI/Inference	Bedrock, SageMaker, OpenSearch / Aurora pgvector	Managed models, custom tuning, vector store / RAG
Security	GuardDuty, Security Hub, CloudTrail, Config, Detective	Threat detection, compliance monitoring, SIEM export
Secrets	Secrets Manager, SSM Parameter Store, multi-account	Secrets rotation, parameter mgmt, disaster recovery
Observability	CloudWatch, X-Ray, OTel-based tracing	Metrics, logs, distributed tracing

Azure Year-Two Integration

Layer	Azure Services	Purpose
AI	Azure OpenAI, Azure AI model access	Microsoft-aligned environments

Layer	Azure Services	Purpose
Identity	Entra ID	Federation and approval workflows
Compute	AKS, Azure Container Apps	Runtime portability from EKS
Security	Key Vault, Defender for Cloud, Sentinel, Monitor	Security posture, SIEM, observability
Isolation	Confidential computing, isolated workloads	Regulated workloads, procurement needs

GCP Year-Three Expansion

Layer	GCP Services	Purpose
AI	Vertex AI, Agent Builder	Customer-aligned deployments, agent governance
Compute	GKE	Runtime portability from EKS/AKS
Security	Cloud IAM, Secret Manager, SCC, Chronicle	Governance, SIEM, threat detection
Isolation	CMEK, private service connectivity	Regulated / high-sensitivity deployments

VII. NVIDIA Enterprise Agentic Stack

NIM: Run-Anywhere Inference as Portability Lever

NVIDIA NIM microservices are prebuilt, optimized inference containers for deploying AI models on NVIDIA-accelerated infrastructure across cloud, data center, and edge. This maps directly onto the Framework's requirement to support customer VPC and future on-prem/air-gapped routes.

NeMo: Agent Lifecycle Governance

NVIDIA NeMo is a modular suite for the agent lifecycle: data processing, fine-tuning, evaluation, policy enforcement, and observability across cloud, on-prem, and hybrid.

NeMo Guardrails: Programmable Policy Enforcement

Acts as a programmable guardrail layer between applications and model endpoints. The implication: central policy enforcement with versioning, auditability, and consistent enforcement across teams.

Practical NVIDIA Enterprise Pathway

Stage	Focus	NVIDIA Components	When	Rationale
Prototype	Reduce infra drag	Managed model APIs	Mo 0–6	Speed > optimization; validate workflow value first
Expansion	Portability + cost	NIM-backed inference	Mo 6–12	GPU economics justify NIM where volume warrants
Platform	Lifecycle controls	NeMo, NeMo Guardrails, eval pipelines	Mo 9–18	Standardize eval, fine-tuning, policy enforcement
Advanced	Performance	Triton, TensorRT-LLM	Mo 12–24	Latency-critical or high-throughput optimization
Robotics	Physical AI	Omniverse, simulation	Mo 24+	Simulation-first; constrained pilot second; scale last

Fortune-Scale Automation Interop

> *The practical pattern: agents interpret, prioritize, summarize, or propose; deterministic systems execute; governance systems verify.*

Platform	Role	Integration Pattern
ServiceNow	IT, SecOps, approvals, workflow routing	Agent enriches tickets / proposes resolutions; ServiceNow executes
UiPath / Automation Anywhere	Deterministic execution	RPA handles structured execution; agents handle interpretation
SIEM / ERP / HRIS	System of record	Agents read; write-back requires approval gates + audit logging

VIII. Deployment Model Decision Matrix

Model	Best For	Strengths	Trade-Offs	Industries
SaaS Multi-Tenant	SMBs, tech-forward, fastest onboarding	Best unit economics, fastest time-to-value	Requires mature tenant isolation, audit partitioning	Tech, Media, Professional Services
Single-Tenant Cloud	Regulated / high-sensitivity data	Dedicated resources, easier compliance narrative	Higher cost, more ops overhead	Healthcare, Finance, Insurance, Legal
Customer VPC / VNet	Large enterprises, fastest close	Runtime inside customer's network boundary	Customer infrastructure dependency	Fortune 500, Gov contractors, Banking
On-Prem / Air-Gapped	Defense, sovereign, extreme sensitivity	Complete data isolation, sovereign control	GPU portability, highest delivery cost	Defense/IC, FedRAMP High, Critical Infra

IX. Governance, Compliance & Risk

Framework Anchors

Framework	Type	Applicability	Timeline
NIST AI RMF 1.0	AI risk management	All industries; US-anchored	Voluntary; increasingly expected
NIST GenAI Profile	Gen AI-specific	All industries using GenAI	Extension of AI RMF
NIST CSF 2.0	Cybersecurity	All industries	Governance as first-class function
ISO/IEC 42001	AI management systems	Global enterprises	Growing procurement requirement
SOC 2 Type II	Security trust signal	All SaaS/cloud; baseline	12–18 month timeline; table stakes
ISO 27001	InfoSec management	Global enterprises	Often required with SOC 2 for international
HIPAA	Healthcare privacy	Healthcare, health-adjacent	Required for PHI handling
PCI DSS	Payment security	Financial, retail, payments	Required for cardholder data
FedRAMP	US government cloud	Government, defense	Multi-year; confirm pipeline first
EU AI Act	AI regulation	Any serving EU market	Build traceability early vs retrofit

Compliance by Industry

Industry	Required / Expected	Recommended	Long-Horizon
Technology / SaaS	SOC 2 Type II	ISO 27001, NIST AI RMF	ISO 42001, EU AI Act readiness
Healthcare	HIPAA, SOC 2 Type II	ISO 27001, NIST CSF 2.0	FedRAMP (if gov health)
Financial Services	SOC 2 Type II, PCI DSS	ISO 27001, NIST AI RMF	FFIEC, EU AI Act
Government	FedRAMP (Moderate+)	NIST CSF 2.0, NIST AI RMF	FedRAMP High, IL4/IL5
Education	FERPA compliance	SOC 2 Type II	State-level AI regulations

Risk Register

Risk	Trigger	Mitigation
License mismatch	Reselling without rights	BYO-license first; explicit agreements

Risk	Trigger	Mitigation
Employment discrimination	AI-ranked actions without safeguards	Human review; rubrics; impact tests; audit logs
Misclassification	High control under 'contractor' labels	Clear employment model; counsel review
Security / blast radius	Overbroad permissions; input injection	Least privilege; sandboxing; monitoring
Training warehouse	Training without placement	Demand-side pilots first; train on paid work
Shadow AI gap	Unapproved AI tools	Access controls; governance policies; audits
Regulatory gap	Operating without certification	Phase certs by customer mix; SOC 2 first

X. Cybersecurity Threat Model for Agentic Systems

OWASP LLM Top 10 + MITRE ATLAS

Threat	Attack Vector	Impact	Control
Prompt Injection	Malicious input manipulates agent	Unauthorized actions, exfiltration	Input validation, NeMo Guardrails, allowlists
Tool Manipulation	Exploiting tool-use interfaces	Lateral movement, priv escalation	Least privilege, separate R/W creds, sandboxing
Info Disclosure	Agent surfaces confidential data	Data breach, compliance violation	Output filtering, DLP, audit logging
Excessive Agency	Agent exceeds intended scope	Financial loss, disruption	Rate/spend limits, kill switches, circuit breakers
Supply Chain	Compromised model or dependency	Backdoor access, poisoned outputs	Approved catalog, version pinning, signing
Insecure Connector	Vulnerable tool integration	RCE, data compromise	Security review, sandboxing, segmentation
Output Handling	Unsanitized output chains to other systems	Cross-system exploit	Output validation, injection prevention at boundaries

Excessive Agency Controls (First-Class Security)

Agency scope is a security boundary. Autonomy is constrained by policy, permissions, and auditability.

- Separate read vs write credentials for all tool use.
- Require approvals for high-risk: money movement, identity changes, account closure, employment decisions.
- Enforce spend limits, rate limits, and kill switches at orchestration layer.
- Action allowlists and explicit deny paths.
- Sandboxed tool runners. Log every tool call, policy check, approval, override, rollback.

Model Supply Chain Controls

- Approved model catalog with version pinning.
- Evaluation gates before production promotion.
- Secure artifact signing and provenance.

- Formal change control for prompts, tools, schemas, guardrails.
- Retention rules and customer-configurable logging boundaries.
- SIEM export and compliance API integration.

XI. Workforce Pipeline Design

The workforce pipeline is a product capability, not an HR side note. Scaling automation increases demand for oversight, exception handling, QA, governance, and incident response — even when direct task labor shrinks.

Task Tier System

Tier	Scope	Examples	Controls	Advancement Evidence
0	Sandbox only	Synthetic cases; docs; non-sensitive QA	No sensitive data; supervised	Scenario pass rate; training complete
1	Production review	Validation; triage; audit evidence	Approvals; audit logs	QA pass rate > 95%; supervision declining
2	Exception resolution	Complex exceptions; client summaries	Escalation thresholds	Escalation quality; resolution accuracy
3	Policy/governance	Rubric ownership; adversarial tests	Formal signoff; change control	Governance contributions; rubric improvements
4	Architecture/leadership	Program ownership; scale playbooks	Executive review; KPIs	Revenue impact; team performance

Multi-Lane Intake

Lane	Source	Entry Point	Design Principle
A	High school	Pre-apprenticeship; Tier 0 sandbox	Educational pathway only. No sensitive data, no hiring decisions, strict supervision
B	College/new grad	Core Tier 1 roles	Fastest way to solve entry-level gap. Builds audited work portfolio
C	Displaced workers	Reviewer/escalation (Tier 1–2)	Highest early leverage. Domain context creates immediate value
D	Advancement	Tier 2+ into governance, solutions, leadership	Competency-gated promotion with signed-off evidence

Core Measurements

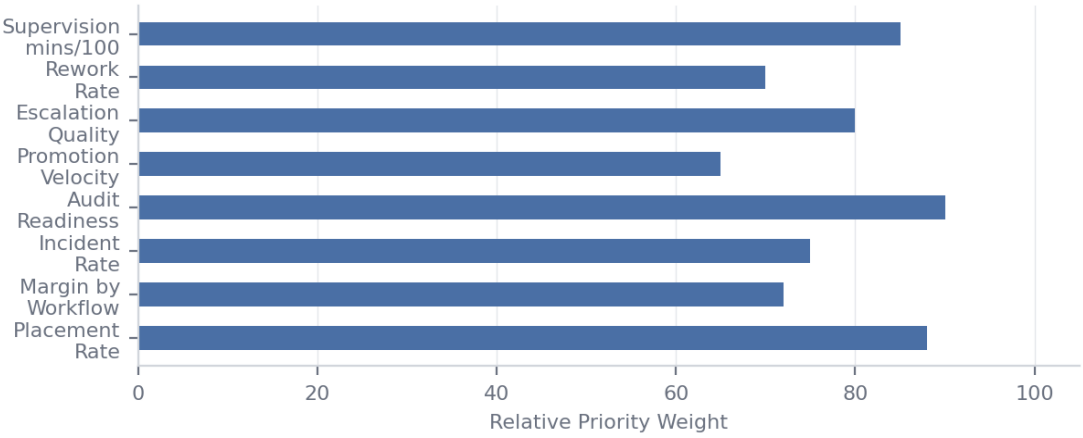
Supervision mins/100
Rework Rate
Escalation Quality
Promotion Velocity
Audit Readiness
Incident Rate
Margin by Workflow
Placement Rate
0
20
40
60
80
100
Relative Priority Weight

XII. Hiring OS: Human vs AGI Comparable Evaluation

The Hiring OS creates a comparable evaluation framework for human and AGI candidates applying to the same workflow roles. Not about replacing humans — about making comparison structured, auditable, and bias-aware.

Resume-Review Rubric (100 Points)

Category	Wt	Human Evidence	AGI Evidence
Role relevance	20	Prior roles; project similarity	Validated task coverage in same workflow
Outcomes	20	Metrics; impact stories	Cycle time, throughput, cost saved, SLA results
Reliability	15	References; consistency	Error rate, rework, incident history, uptime
Communication	10	Writing/interview quality	Escalation quality, audit trail completeness
Problem solving	15	Edge-case stories	Transfer performance; exception handling
Risk/compliance	10	Judgment signals	Policy compliance, approvals, containment
Team fit	5	Collaboration history	SOP adherence; coordination quality
Cost-effectiveness	5	Comp vs impact	TCO; supervision load

> *In real deployments, risk/trust often gets higher effective weight for AGI because scaling amplifies downside.*

AGI Resume Example: Astra-7 (Template)

Summary: High-autonomy generalist agent optimized for knowledge work, operations execution, and multi-step decision workflows. Demonstrated performance across support, research synthesis, software implementation, reporting, scheduling, procurement, and process automation.

Core capabilities: Reasoning/planning (task decomposition, dependency tracking), Tool use (APIs, databases, ticketing, browser), Communication (drafting, summarization, handoffs), Analysis (KPI tracking, anomaly triage), Execution (SOP-based operations, exception routing).

Known limits (must disclose): Degrades under ambiguous policy constraints unless clarified. Requires permissions architecture and action sandboxing. Benchmarks can misrepresent edge cases.

XIII. Company Operating Model: Cells + Tiers

Cell Structure (Delivery Unit)

Role	Primary Responsibility	Outcome Metric
Workflow Engineer	Orchestration, integrations, permissions, monitoring	Deployment success; incident rate
Domain Lead	Workflow mapping, edge cases, acceptance criteria	Throughput + quality vs baseline
QA / Governance	Rubrics, audits, escalation rules, change control	Rework; compliance; audit quality
Ops Associates (xN)	Validation, triage, exception handling, docs	Supervision mins/100; QA pass rate

XIV. Revenue Model & Unit Economics

Phase	Revenue	Why It Works	Early Moat Created
0	Advisory / audits	Fast sales; low build cost	Workflow knowledge + buyer language
1	Pilot projects	Proof of value; case studies	Playbooks + baseline metrics
2	Managed ops retainers	Recurring; sticky ops	Outcome data + trained workforce
3	Platform subscription	Scales beyond services	Templates + governance engine
4	Specialist consulting	Premium upsell	Deep vertical expertise

Vendor / Licensing Strategy

Pattern	Mechanism	Best For	Risk
BYO-License	Client procures; you sell platform + ops	Enterprise with existing vendor relationships	Lowest licensing risk
Embedded SaaS	Bundle vendor usage into pricing	Mid-market / SMB turnkey	Guard against cost blow-ups
Authorized Resale	Resale with vendor authorization	Distribution leverage	Requires partner agreements

Service Lines

- AI Workflow Readiness Audit (2 weeks): inventory, risk tiering, ROI estimate, 90-day roadmap.
- Agentic Ops Pilot (4–8 weeks): 1 workflow with approvals, audit logs, metrics, runbooks.
- Security/Governance Hardening: least privilege, secrets, monitoring, threat model, incident playbook.
- Managed Agent Ops (retainer): tuning, monitoring, exceptions, reporting, continuous eval.

XV. GTM Strategy & Customer Acquisition

US-only initial rollout, globalization goal modeled on Microsoft/Amazon expansion. Start niche consulting for SMBs as SaaS, expand to all industries. Healthcare, finance, education, and government serve as validation tracks for obstacle mapping.

Customer Profile: Who Buys First (Values-Based)

Rank	Buyer Profile	Why First	Decision Maker
1st	Growth-stage tech (50–500 employees)	Value speed; cloud-native; low procurement friction; scale without headcount	CTO / VP Engineering
2nd	Mid-market professional services	High labor cost/task; clear workflows; immediate ROI visibility	COO / Managing Partner
3rd	IT / managed service providers	Already sell ops; understand SLA delivery; can embed governed agent ops	CEO / VP Service Delivery
4th	Enterprise IT/SecOps (Fortune 500)	Largest budgets; clear ticket/alert pain; strong governance norms; measurable baselines	CISO / CIO
5th	Healthcare / finance (regulated)	Highest value/workflow but longest cycles; buy after proof; require compliance packaging	CISO / Compliance / CTO

Order of Operations

Bad: deck → mass recruiting → open source → travel → raise → find product.

Good: wedge → demo → pilot → case study → recruit selectively → raise from evidence.

XVI. Competitive Landscape

Category	Estimated Count	Nature
Directly similar thesis	Hundreds to low-thousands	Governed AGI deployment + workforce transition
Near-adjacent	Tens of thousands	AI consulting, RPA, staffing, AI SaaS
Budget-competition field	~15,000	Practical competitive awareness number

Advantage: one wedge workflow + measurable outcomes + repeatable delivery cells.

Open Source Strategy

- Open: evaluation schemas, governance checklists, demo harnesses, non-core templates.
- Private: customer playbooks, outcome data, tuning methods, compliance delivery.

XVII. Timeline & Execution Phases

Phase	Timeline	Key Actions	Exit Deliverables
Foundation	2 weeks	Pick wedge (vertical + workflow + buyer). 1-page thesis + pilot offer. Minimal demo: tool chain + approval gate + audit log + dashboard stub. Tier 1 role + competency gates.	Thesis, pilot offer, demo, role brief
Proof	0–90 days	20–30 discovery calls. Close 1 pilot (paid). Internal case study (baseline vs after).	1 paid pilot, case study, refined pitch
Delivery	3–9 months	1–3 workflows; standardize cell model. First cohort (new grads + displaced). Instrument supervision/rework.	Cells, measurement, first cohort
Systematize	9–18 months	Training tracks; scenario library; eval harness. Productize modules. Shift to retainers.	Curriculum, modules, recurring revenue
Scale	18–36 months	Expand workflow families, then verticals. Platform subscription. Enterprise readiness (SSO, DPA, SLAs).	Platform, enterprise-ready, multi-vertical

First Hires

- AI / workflow engineer (orchestration + integrations).
- Domain ops lead (maps reality and acceptance criteria).
- QA / governance lead (rubrics, audits, change control).
- Curriculum / training lead (workflows into competency ladders).

XVIII. Robotics Roadmap (Post-24 Month)

Not a year-one dependency. Operating principle mirrors digital agent deployment: simulation first, constrained pilot second, scaled autonomy last. Modeled on Tesla Optimus staging.

Phase	Timeline	Focus	Key Activities
Digital Only	Mo 0–24	Digital workflows exclusively	All resources on agentic software; build governance + data foundation
Simulation	Mo 18–24	Digital twin experimentation	NVIDIA Omniverse; repeatable environments; no physical deployment
Limited Pilots	Mo 24–36	Structured physical environments	Warehouse/logistics; human override; safety zoning; measurement-first
Expansion	36+ months	Expand where justified	Only after simulation fidelity, economics, and governance confirm viability

XIX. Hypothetical: AGI Autonomous Employment Decisions

Current Legal Baseline (2026)

Under current US law, employment decisions require human decision-makers. Title VII, ADA, ADEA, and state equivalents assign liability to employers. NYC LL144 and CO SB24-205 specifically regulate automated employment tools, requiring bias audits, notice, and human oversight.

> *Current position: Humans decide. AGI assists with information gathering and structured recommendations. All employment decisions require human judgment, approval, and accountability.*

Graduated Automation Scenario

Stage	AGI Role	Human Role	Legal Basis	Risk
Current (2026)	Info assembly, screening, scheduling	All decisions; review required	Title VII, ADA, NYC LL144, CO SB24-205	Low — fully compliant
Near-Term	Non-biased recommendations with confidence scores	Reviews, decides, retains override	EEOC guidance; bias audit required	Medium — ongoing testing needed
99% Ethical Parity	Equivalent or superior ethical consistency	Governance oversight; appeals; periodic audit	No current framework; requires legislation	High — regulatory/trust barriers
Full Automation	End-to-end employment lifecycle	Policy design; governance; appeal adjudication	Requires fundamental legal change	Very High — multi-decade horizon

XX. Operational Checklists & Templates

Workflow Intake Checklist

- Define workflow boundaries and success metrics (throughput, quality, cycle time).
- Map data sensitivity and risk tier (low / medium / high).
- List tools required; design least-privilege permissions.
- Define approval points and escalation thresholds.
- Define audit outputs (what must be logged and exportable).

Pilot Delivery Checklist

- Baseline measurement captured (before state).
- Agentic workflow built with safe tool calls and rollback.
- Human review queue operational with rubrics.
- Monitoring/alerts enabled; incident playbook drafted.
- Post-pilot report: metrics, failures, fixes, next scope recommendation.

Governance Checklist (Ship Every Time)

- Versioning (model, prompts, tools, rubrics).
- RBAC + secrets management; approvals for sensitive actions.
- Audit trail completeness (inputs, outputs, reviewer decisions).
- Regression tests and scenario replay suite.
- Data retention policy and PII handling rules documented.

One-Page Thesis Template

Problem: [who is suffering, what breaks]

Customer: [buyer persona + budget owner]

Solution: [what you deploy + what humans do + what controls exist]

Wedge Workflow: [one workflow]

Moat: [templates + governance + data + workforce pipeline]

Success Metrics: [cycle time, error rate, supervision mins/100 tasks]

First 90 Days: [demo → pilot → case study]

Pilot Offer Template

Scope: 1 workflow, 1 team, 1 environment (4–8 weeks).

Deliverables: working workflow, approval gates, audit logs, dashboard, runbook, rollback plan.

Metrics: baseline + target improvements, QA thresholds, escalation criteria.

Client responsibilities: access, SMEs, change approvals, security review window.

Pricing: fixed pilot fee + optional retainer.

Role Brief Template

Role: [title]

Mission: [outcome they own]

First 30/60/90 Days: [deliverables]

Interfaces: [who they work with]

Success Criteria: [metrics + expectations]

Non-Goals: [what they must not do]

XXI. References

1. McKinsey — The State of AI in 2025: Agents, Innovation, and Transformation (November 2025)
2. Federal Reserve Bank of St. Louis — State of Generative AI Adoption in 2025
3. Microsoft AI Economy Institute — Global AI Adoption Report, 2025
4. IBM — Cost of a Data Breach Report 2025
5. NIST — AI Risk Management Framework 1.0
6. NIST — Generative AI Profile (AI 600-1)
7. NIST — Cybersecurity Framework 2.0
8. ISO/IEC 42001 — AI Management Systems Standard
9. NYC Local Law 144 — Automated Employment Decision Tools
10. Colorado SB24-205 — High-Risk Artificial Intelligence Systems
11. EEOC — AI and Employment Discrimination Guidance (US)
12. AWS — Well-Architected Generative AI Lens
13. AWS — Agents for Amazon Bedrock; Amazon Bedrock Guardrails
14. Microsoft Azure — Azure AI Content Safety; Azure OpenAI Data/Privacy
15. Google Cloud — Vertex AI Agent Builder and Platform Guidance
16. NVIDIA — NIM Microservices; NeMo; NeMo Guardrails; Triton; TensorRT-LLM
17. OWASP — Top 10 for LLM Applications
18. MITRE — ATLAS / Adversarial Threat Landscape for AI Systems
19. World Economic Forum — Future of Jobs Report 2025
20. Synergy Research Group — Cloud Infrastructure Market Data (Q1–Q3 2025)
21. Canalys — Global Cloud Infrastructure Spending Reports (2025)
22. US Department of Labor — WARN Act; YouthRules; FLSA Internships
23. Apprenticeship.gov — Registered Apprenticeship Resources
24. OECD — Employment Outlook (Displacement and Earnings Effects)

End of Document

Nathan Lim — AGI Labor Transition and Enterprise Agentic Deployment Runbook v4