



AIGW Policy Observatory

Global AI Governance & Workforce Transformation

Global AI Governance and Workforce Transformation Policy Observatory

AIGW WORKFORCE TRANSITION DEEP DIVE

BNY: Building a Workforce That Builds and Manages AI

How one global bank moved from AI access to employee-built agents, digital employees and an AI-native early-career operating model

October 2026

A decorative graphic in the bottom right corner consisting of several concentric, curved lines in shades of blue and purple, creating a sense of motion and depth.

CONTENTS

Contents

Executive Summary	03
Key Findings	07
Case Context	08
Digital Employees	12
Human Authority	17
Outcome Metrics	24
Workforce Model	27
Evidence and Conclusion	29
Notes and References	39

EXECUTIVE SUMMARY

Executive Summary

BNY offers one of the clearest examples of an enterprise moving beyond the idea that artificial intelligence is a tool used by a specialist technology team.

Its model is broader.

Employees across technical and nontechnical functions are expected to use AI. Many are expected to build agents. Some teams now work alongside what BNY calls digital employees — multi-agent AI systems with distinct identities, credentials, supervisors and dedicated workflows. Early-career analysts are being introduced to AI as part of their formal development architecture rather than as an optional later-career skill.[1][2]

This makes the BNY case especially important for workforce transformation.

The central change is not simply:

“people use AI.”

It is:

“people increasingly design, supervise and work alongside AI systems that perform parts of the operating workflow.”

That is a different organizational model.

BNY's enterprise AI platform, Eliza, was designed as a shared environment for secure AI access, approved models, agent building, learning and governance. By October 2025, BNY reported that 99% of employees had been fully trained and onboarded onto Eliza, up from just over one-third in January of that year. The company also reported that the majority of its AI builders were coming from outside engineering teams.[1]

By the end of 2025, formal annual-report disclosures showed:

100% of employees had access to Eliza; BNY delivered 171,000 AI learning hours in 2025; nearly 50% of employees were building AI agents; 160 enterprise AI solutions were in production; 134 digital employees were live.[2].

These are unusually strong indicators of behavioral adoption.

They still require careful interpretation.

Access is not proficiency.

Training hours are not role transition.

Agent building is not equivalent to independent deployment authority.

And 134 digital employees do not represent 134 human jobs eliminated.

EXECUTIVE SUMMARY

The case becomes more concrete inside operations.

BNY reported that digital employees were embedded in payment teams, working alongside people to automate routine work. The company described these systems as having distinct personas, credentials and supervisors.[1] Later vendor-hosted reporting in September 2026 added more detailed operational evidence: digital employees were handling more than 10% of payment-repair issues globally; AI-assisted client settlement inquiries were being processed 80% faster; and AI-assisted onboarding was reported to be 20% faster for the relevant cases.[3]

Those figures should be treated as later customer-story evidence, not audited company financial metrics.

But they help reveal the operating model.

The human role is shifting from:

performing every task directly

toward combinations of:

defining use cases; building agents; supervising digital employees; handling exceptions; providing judgment.

redesigning workflows; maintaining controls.

OpenAI's December 2025 BNY case described this transition explicitly. BNY's Chief Data and AI Officer said that for some tasks, rather than handling the task first, the human operator becomes the trainer or nurturer of the digital employee.[4]

That is the core workforce mechanism in this case.

AI does not only change what employees do.

It changes the relationship between employees and the systems performing the work.

AIGW calls this transition the Builder–Supervisor Shift:

“employee as task performer → employee as AI user → employee as AI builder → employee as supervisor / exception handler → employee as workflow owner”

Not every employee moves through all five stages.

Not every role should.

But BNY's workforce architecture increasingly supports that direction.

The most important evidence is organizational, not anecdotal.

EXECUTIVE SUMMARY

BNY has built role-based AI learning pathways across engineering, operations, product management, client coverage and corporate functions. It has analyst bootcamps for early-career employees and AI leadership programmes for managers. Its current 24-month Analyst Program explicitly tells incoming analysts they will build AI skills through pathways such as AI Explorers, AI Builders and AI Pioneers, alongside job rotations, mentoring and progression toward associate roles.[1][5]

This means AI is entering the entry architecture of the firm.

A new analyst is not expected to first learn the organization in an AI-free environment and add AI later.

AI is part of how the firm now defines professional development.

The model also creates a governance challenge.

BNY says its digital employees have:

personas; credentials; supervisors; guardrails; telemetry.

approvals; continuous oversight.[1].

Its Responsible AI commitment adds enterprise-level expectations for:

approval processes; risk management; model review; human oversight; explainability.

lifecycle monitoring.[6].

However, the public evidence does not disclose a complete action-level decision-right matrix for specific payment workflows.

We do not know from the reviewed sources:

which payment-repair actions a digital employee can complete independently; which require human approval; who can override the system; who makes the final decision when the digital employee and human operator disagree; how exceptions are categorized by risk.

That distinction matters because supervision is not the same as final authority.

The case therefore connects directly to AIGW's wider governance work:

“agentic operating models require both workforce redesign and explicit authority design.”

BNY is also useful because its leaders frame AI as a capacity creator, not simply a headcount-reduction mechanism. CEO Robin Vince has publicly rejected the idea that AI adoption should be read automatically as a headcount story, while the annual report describes AI efficiency as creating space for reinvestment in products, platforms and people.[2][7]

EXECUTIVE SUMMARY

That framing should not be treated as proof that AI will have no future employment effect.

The public record does not provide a standardized AI-attributable headcount decomposition.

The correct conclusion is narrower:

“BNY is using AI to expand operational capacity, redesign workflows and create an AI-building workforce, while public evidence does not establish a direct AI-driven workforce reduction.”

This Deep Dive argues that BNY's most important contribution to the AI workforce debate is not its

number of agents.

It is the institutional sequence behind them:

“platform → access → learning → building → governed deployment → digital workforce”

The company did not begin with autonomous agents and then ask employees to adapt.

It built a platform, trained the workforce, created role-specific learning, enabled employees to build, and then scaled agentic deployment.

That sequence may be one of the most important transferable lessons in enterprise AI adoption.

KEY FINDINGS

Key Findings

1. 1. BNY treated AI as an enterprise workforce capability, not a specialist engineering tool: By late 2025, nearly all employees were trained and onboarded onto Eliza, and BNY said the majority of AI builders came from outside engineering.[1] By year-end, the bank reported that nearly half of employees were building agents.[2]

2. 2. The strongest transition is from AI user to AI builder: Employees across technical and nontechnical functions can create Eliza agents, turning domain expertise into automated workflows.[1] This is stronger workforce evidence than simple tool adoption.

3. 3. Digital employees are part of the operating model: BNY had 134 digital employees live at year-end 2025, described as multi-agentic AI systems operating autonomously alongside human colleagues.[2] BNY's earlier reporting says they have distinct personas, credentials and supervisors.[1]

4. 4. Payment operations provide a concrete production example: Digital employees are embedded in payment teams.[1] Later September 2026 vendor-hosted reporting says one payment-validation digital employee handled more than 10% of payment-repair issues globally.[3] This is process-scale evidence, not a human-job substitution count.

5. 5. Capability building is role-specific and embedded in work: BNY uses: role-based learning pathways; certifications; responsible-AI curriculum; problem-driven bootcamps; hands-on prototyping. hackathons; peer communities; AI Pioneers.[1][8]. The model aims to move learning from awareness toward deployment.

6. 6. Early-career development is being redesigned around AI: BNY's Analyst Program now explicitly includes AI learning pathways, rotations, mentoring and a transition toward associate roles. [5] AI is therefore becoming part of the first rung, not an optional overlay.

7. 7. Governance is embedded in the platform, but workflow-level authority is still incompletely public: BNY documents: guardrails; telemetry; approvals; supervisors; human oversight. approval and model-review structures.[1][6]. The reviewed evidence does not publish a complete payment-workflow approval / override / final-authority matrix.

8. 8. The case is behavioral before it is economically causal: The evidence strongly establishes: adoption; agent building; training; changed workflows; digital employees. It is less clean on: AI-specific ROI; AI-attributable headcount effects; promotion outcomes; long-term role displacement. That is why the case is strongest as an operating-model and capability-formation case.

CASE CONTEXT

BNY as a Workforce Case

1. Why This Case Matters: Enterprise AI strategies often fail at the point between access and use.

An organization can buy licenses for thousands of employees and still change very little.

The technology exists.

The workflow remains the same.

BNY's case is different because it tries to move employees through several stages:

access → fluency → application → agent building
→ workflow integration → digital-employee
supervision

The organizational challenge becomes less:

“How do we get people to try AI?”

and more:

“How do we turn domain experts into people who can redesign work with AI?”

This is an important workforce shift.

A bank has deep pools of expertise in:

payments; custody; settlement; compliance; risk.

client service; operations; contracts.

Engineers may understand the technology.

But operational employees understand where work is:

repetitive; risky; exception-heavy; dependent on context; constrained by regulation.

If those employees can build or shape agents, AI transformation can move closer to the work itself.

That is the strategic significance of BNY's “AI for everyone” model.

It is not only democratizing access to AI.

It is democratizing workflow redesign capability.

2. Institutional Context: A Regulated Global Operating System: BNY sits at the infrastructure layer of global finance.

The bank provides services across custody, settlement, payments, collateral, asset servicing, investment management and wealth. Its 2025 reporting emphasizes that the organization operates at large scale across global markets and positions AI as part of a wider transformation toward a platform operating model.[2]

CASE CONTEXT

This context matters for two reasons.

First, operational risk is high: Payment, settlement and custody workflows cannot be treated like casual productivity experiments.

Errors can affect:

money movement; client obligations; market operations; compliance; financial risk.

That makes governance a production requirement.

Second, knowledge is distributed: The expertise needed to redesign a payment workflow is not located only in an AI team.

It is distributed across:

operations; engineering; risk; compliance; legal.

product; client teams.

An enterprise platform can therefore create value if it lets those groups build within controlled boundaries.

That is the role Eliza plays in BNY's architecture.

3. Eliza: The Platform Before the Workforce Model: BNY's sequence matters.

The company did not begin by buying isolated AI products for each department.

It built a common enterprise platform.

Eliza provides employees with access to approved AI capabilities and models from multiple providers while applying the firm's data, risk, legal and compliance standards.[1][2]

BNY's own 2025 explanation emphasizes three reasons for the platform approach:

strategic value and differentiation; flexibility to integrate different models as technology changes; efficiency through reuse rather than duplicated point solutions.[1].

This infrastructure creates a common environment in which employees can:

use AI; access approved data; build agents; share practices; move solutions toward production.

operate inside common governance.

Why platform design is a workforce issue: A fragmented AI environment creates fragmented learning.

Different teams may:

learn different tools; apply different standards; build incompatible solutions; duplicate work; struggle to move pilots into production.

CASE CONTEXT

A common platform can make capability more portable inside the organization.

The employee learns not only one chatbot interface but a repeatable way to:

identify a business problem; work with enterprise data; select an approved model; build an agent; test the workflow.

deploy through controls.

That is a much more consequential workforce architecture than tool access alone.

Why platform design is a workforce issue — AIGW interpretation: The important unit is not the AI application.

It is the AI-enabled operating environment.

Once the environment becomes standard, employees can move from consumers of technology toward co-designers of workflows.

4. From Awareness to Building: BNY's Capability Architecture: BNY's learning system is unusually useful because it shows several stages of workforce enablement rather than one generic training statistic.

Its October 2025 framework included six elements.
[1]

AI access: A secure, unified enterprise platform gives employees a common place to use AI in everyday work.

Access is the base condition.

It is not the outcome.

E-learning: BNY describes role-based AI learning paths with certifications and a responsible-AI curriculum.[1]

The stated goal is both:

firmwide fluency; role-specific depth.

Bootcamps: Problem-driven bootcamps are intended to turn priority use cases into deployable minimum viable products while simultaneously upskilling cross-functional teams.[1]

This is important because the learning artifact is not only a test score.

It can be a working solution.

Hands-on learning: Employees can prototype, test and deploy agents with:

guardrails; telemetry; approvals.[1].

This connects learning to production discipline.

CASE CONTEXT

Gamified learning: Hackathons, promptathons, weekly demonstrations and leaderboards are used to create experimentation and discover useful ideas.[1]

Community: AI champions, peer learning circles and practitioners help spread adoption and common standards.[1]

Community – 2026 maturity: Explorer → Builder → Pioneer: A September 2026 Microsoft customer story describes a more explicit three-level workforce architecture:[8]

Explorer

AI fundamentals; confidence using the tools.

Builder

more advanced use; workflow application.

Pioneer

a selected group completing a 40-hour agentic-AI bootcamp; intended to help teams rewire workflows with AI.

Microsoft reported that Pioneer bootcamp cohorts had expanded from about 200 to about 800 participants and that BNY had delivered more than 400,000 hours of AI training to date.[8]

This later figure should not be merged with the 171,000 hours reported for calendar 2025.

They measure different periods.

Why the progression matters: Many enterprise programmes stop at:

“know how to use AI.”

BNY's architecture increasingly asks:

“can you build with AI?”

and then:

“can you help redesign the team's workflow with AI?”

That is a much more consequential capability target.

DIGITAL EMPLOYEES

From AI Access to Human-Supervised Digital Work

The most important workforce transition in the BNY case is the movement from task execution toward AI-building and AI-supervision.

The shift can be represented as five levels.

Evidence boundary: Levels 1–4 are supported by BNY's public descriptions of use, building and digital employees.

Evidence boundary: Level 5 is an AIGW analytical implication of the operating model, not a published BNY job-family classification.

Why this matters: A worker can remain employed while the economic content of the role changes substantially.

Why this matters: The person may perform fewer transactions directly but become responsible for a larger system of work.

Why this matters: That is role redesign even when the job title stays unchanged.

6. Digital Employees: What BNY Means by the Term: The phrase digital employee can easily create confusion.

It sounds like a one-for-one digital replacement for a human worker.

BNY's evidence supports a different interpretation.

Its digital employees are multi-agent AI solutions embedded into workflows.[2]

BNY says they have:

distinct personas; credentials; supervisors.[1].

OpenAI describes advanced Eliza agents as having identities, access controls and dedicated workflows.[4]

Independent reporting from the Wall Street Journal and Axios similarly describes digital employees with organizational identities, logins and human managers.[7][9]

What the number does mean: At year-end 2025, BNY reported 134 digital employees live.[2]

That establishes:

deployment beyond pilots; agentic systems embedded in production; a formal management concept for AI systems.

What the number does not mean: It does not mean:

“134 human jobs were eliminated.”

A digital employee can:

DIGITAL EMPLOYEES

handle only part of a workflow; operate at a different scale than a human employee; support many people; require human supervision; perform work that previously accumulated across multiple roles.

Human and digital headcount are therefore not directly comparable units.

What the number does not mean — AIGW rule: Do not convert digital-worker counts into human-job equivalents without workload and staffing evidence.

This is especially important as firms increasingly use anthropomorphic language for agentic systems.

7. Payment Validation: The Most Concrete Operating Example: Payments provide the clearest view of BNY's digital-workforce model in action.

Payment instructions sometimes cannot be processed straight through because information is missing or incomplete.

Traditionally, those cases require operational review and repair.

BNY says it created digital employees to help validate and repair such payments.[3][10]

2025 evidence: BNY's October 2025 account states that digital employees were already part of payment teams, working alongside human staff to accelerate processing.[1]

The company also said more than 100 digital employees had been deployed across routine work. [1]

2026 follow-on evidence: A Microsoft customer story published in September 2026 reports that a payment-validation digital employee was handling more than 10% of payment-repair issues globally. [3]

The same source reports that:

more than 10% of certain client inquiries were resolved or assisted by AI; those AI-assisted inquiries were processed 80% faster; 25% of new onboardings in the reported period were AI-assisted; those cases were 20% faster on average. [3].

Evidence boundary: These are vendor-hosted customer-story metrics based on BNY statements.

They are useful operational evidence.

They are not:

audited productivity metrics; a full payment-process evaluation; proof of staffing reduction; proof of equivalent risk or error rates.

The workforce transition: The payment example suggests a three-part operating model:

digital employee

DIGITAL EMPLOYEES

handles recurring validation / repair work.

human operator

manages exceptions, supervision and context.

workflow owner / control functions

define boundaries, risk and governance.

The public sources do not disclose the exact staffing distribution among these layers.

8. Human Authority: Supervisors Are Visible, Final Decision Rights Are Not Fully Public: BNY is a useful complement to the Deutsche Telekom case.

Unlike a generic agent deployment, BNY publicly describes several governance components around digital employees.

These include:

distinct credentials; human supervisors; guardrails; telemetry; approvals.

continuous oversight.[1][6].

BNY's Responsible AI commitment states that the organization operates:

governance and accountability structures; approval processes; risk management; model

review; lifecycle monitoring.

appropriate human oversight.[6].

That is meaningful evidence of governance architecture.

But it is still not a complete action-level authority map.

The unanswered questions: For payment repair, the reviewed sources do not specify:

which action the digital employee can complete without a human gate; which repair requires approval; whether a digital employee can release a repaired instruction autonomously; whether humans sample completed actions after execution; who can override or rollback an action.

who owns final accountability for an erroneous repair.

AIGW interpretation: BNY sits between two of our earlier governance archetypes.

AIGW interpretation — Kaiser Permanente: Human review and final clinical authority are explicit.

AIGW interpretation — Deutsche Telekom RAN Guardian: Autonomous execution is explicit, while use-case-level human approval / override rights are less public.

DIGITAL EMPLOYEES

AIGW interpretation — BNY: Organizational supervision, approvals and human oversight are explicit, but the detailed action-level rights for specific digital-employee workflows remain incomplete publicly.

This means the correct description is:

“governed agentic deployment with incomplete public workflow-level authority specification”

—not:

“humans approve every action.”

9. Capability Formation Is Becoming Part of the Operating Model: One of the strongest aspects of BNY's approach is that capability formation is not treated as a separate HR programme.

It is integrated into adoption.

BNY's learning architecture links:

use-case discovery; training; prototyping; deployment; peer learning.

governance.[1].

That changes the role of learning and development.

Traditional model: Training occurs before or outside work.

Success is measured through:

attendance; course completion; certification.

BNY model: Learning increasingly happens through:

identifying a business problem; building an agent; testing it; moving it toward deployment; sharing it.

supervising how it performs.

BNY model — AIGW concept: production-linked AI learning: AIGW defines production-linked AI learning as capability development in which the learning pathway is connected to real workflows and produces a usable artifact, agent or operating change.

The important metric becomes:

“What can the employee redesign after the training?”

not only:

“Did the employee finish the training?”

BNY's evidence moves meaningfully in this direction.

ROLE TRANSITION

5. The Builder–Supervisor Shift

FIGURE 1

role levels

From task performer to workflow owner

The most important workforce transition in the BNY case is the movement from task execution toward AI-building and AI-supervision.

Role level	Work responsibility
Level 1 — Task performer	The employee performs the work directly. Examples: review a payment exception; investigate an inquiry; conduct onboarding research; reconcile information. [1] [2] [4]
Level 2 — AI user	The employee uses AI for: research; preparing text; summarization; analysis; coding. [1] [2] [4]
Level 3 — AI builder	The employee creates or configures an agent around a specific work problem. BNY says employees from engineering and nonengineering functions are doing this. [1] [2] [1] [2]
Level 4 — Digital-employee supervisor / exception handler	The AI performs recurring work and the human increasingly: monitors; coaches; handles exceptions; validates performance; changes rules or context. OpenAI's BNY case describes human operators becoming trainers or nurturers for some digital employees. [4] [4]
Level 5 — Workflow owner	The human is no longer primarily responsible for performing individual instances of the task. The role becomes responsibility for: workflow outcomes; risk; controls; escalation; continuous improvement. how human and digital labor are allocated. [1] [2] [4]

Source: AIGW synthesis from BNY public reporting [\[1\]](#) [\[2\]](#) [\[4\]](#)

Note: The shift is represented here as five levels. Levels 1–4 reflect BNY public descriptions. Level 5 is an AIGW analytical implication, not a BNY job-family classification or a mandatory career sequence.

HUMAN AUTHORITY AND WORKFORCE TRANSITION

Human Authority and Workforce Transition

10. Early-Career Redesign: AI Enters the First Rung: BNY's case is especially important after PwC and SimplifyNext because it shows a third way the first rung can change.

PwC: AI compresses traditional junior work, creating an apprenticeship problem.

SimplifyNext: An AI-native apprenticeship places young workers directly into real AI-enabled delivery.

BNY: A traditional large-employer analyst programme is being redesigned so AI building becomes part of the mainstream early-career pathway.

BNY's current Analyst Program is a 24-month development pathway.[5]

It includes:

induction; rotations; mentoring; cohort learning; business exposure.

progression toward associate roles.

The programme now explicitly tells analysts they will build AI skills through pathways including:

AI Explorers; AI Builders; AI Pioneers.[5].

BNY's 2025 AI strategy also notes analyst bootcamps targeted specifically at early-career

employees.[1]

What has changed?: The new analyst is not only learning:

“how the bank works.”

The analyst is increasingly expected to ask:

“how should this work be redesigned with AI?”

That is an important change in professional identity.

What has changed? — From apprentice analyst to workflow improver: AIGW interprets the emerging early-career role as combining:

domain learning + AI fluency + agent building + workflow questioning + human judgment

This is a stronger model than generic AI literacy.

It also creates a developmental risk.

If analysts automate tasks before they understand the underlying process, they may improve throughput without building sufficient institutional judgment.

The answer is not to exclude them from AI.

It is to connect AI building to:

HUMAN AUTHORITY AND WORKFORCE TRANSITION

rotations; mentoring; domain learning; accountable review.

BNY's 24-month programme architecture provides the beginnings of that system.[5]

11. The New Role of the Manager: Digital employees create a new management problem.

Traditional managers manage people.

An agentic team may require managers to oversee both:

human colleagues; digital systems.

BNY's language of supervisors for digital employees makes this change unusually visible.[1]

The manager may need to answer new questions: Is the digital employee operating inside its approved scope?; Is performance degrading?; Are exceptions increasing?; Does the underlying policy need to change?; Should access be expanded or restricted?.

Who should handle a disputed case?; What should be automated next?; Which work should remain human?.

This is a new form of management work.

Management capability therefore changes: A manager may increasingly need:

AI fluency; workflow literacy; performance interpretation; risk judgment; exception design.

escalation discipline; change leadership.

BNY's AI leadership series for managers is consistent with this shift.[1]

Management capability therefore changes — AIGW concept: hybrid span of control: Traditional span of control asks:

“how many people report to this manager?”

Agentic organizations may need another measure:

“what combination of people, agents and automated workflows sits under this manager's accountability?”

BNY's digital-employee model makes that question operational rather than theoretical.

12. Operating Model Redesign: AI Sits Inside a Broader Transformation: BNY's AI transition is happening inside a wider redesign of how the company operates.

Its 2025 year-end investor materials report that more than 70% of employees were working in a new platforms operating model by year-end and that roughly 50% of employees transitioned into that model during 2025.[11]

HUMAN AUTHORITY AND WORKFORCE TRANSITION

The same materials also report approximately \$550 million of efficiency savings and approximately \$500 million of incremental investment in new client solutions, technology and people.[11]

These figures are highly relevant context.

They should not be coded as AI outcomes.

Why the boundary matters: The platform operating-model transformation includes:

organizational redesign; technology modernization; simplification; commercial changes; process improvements.

AI.

The public reporting does not provide a clean causal decomposition showing what share of the \$550 million came from Eliza or digital employees.

Why the boundary matters — Evidence-safe formulation: “BNY's agentic AI programme is being scaled inside a much broader platform operating-model transformation. The company reported substantial efficiency savings and reinvestment at enterprise level, but those financial outcomes should not be attributed solely to AI.”

This discipline is important because enterprise transformation produces many overlapping

effects.

The strongest BNY AI evidence is not the enterprise savings figure.

It is the direct evidence of:

workforce training; agent building; digital employees; production workflows; early-career redesign.

13. Outcomes: Separate Adoption, Capability, Process and Enterprise Metrics: The BNY evidence is unusually rich but only if different metric types are kept separate.

Level A — Access and adoption: These measures show reach and usage.

They do not prove mastery.

Level B — Learning and capability formation: These are stronger than simple license counts.

They still do not show what proportion of builders created production-grade agents.

Level C — Production deployment: The later numbers show continued scaling.

They should not be merged into a single time series unless definitions remain stable.

Level D — Workflow outcomes: Later Microsoft customer-story evidence reports:

HUMAN AUTHORITY AND WORKFORCE TRANSITION

more than 10% of payment-repair issues handled by a digital employee; client inquiries processed 80% faster when AI-assisted; relevant onboarding processes 20% faster on average.[3].

These are process outcomes.

They are not proof of headcount reduction or company-wide productivity.

Level E — Enterprise transformation metrics: The 2025 investor presentation reports:

>70% of employees in the platform operating model; ~\$550 million efficiency savings; ~\$500 million incremental investments.[11].

These are enterprise transformation measures.

They should remain outside the specific AI causal chain.

Level E — Enterprise transformation metrics — AIGW interpretation: A mature enterprise AI scorecard should distinguish:

Access → Adoption → Capability → Deployment → Workflow Outcome → Workforce Outcome → Enterprise Outcome

Most firms collapse these levels.

BNY provides enough evidence to show why they should remain separate.

14. Capacity Creation Is Not the Same as Headcount Reduction: BNY leaders repeatedly frame AI as a source of capacity.

CEO Robin Vince described AI as a “superpower” and pushed back publicly on the idea that AI adoption should be understood primarily through headcount reductions.[7]

The annual report uses similar logic: efficiency can mean saving capacity in one place so that investment can be made elsewhere.[2]

This is a consequential strategic choice.

Three possible uses of released capacity: When AI reduces effort, the organization can:

Three possible uses of released capacity — 1. Remove capacity: eliminate positions; reduce hiring; use attrition or non-backfill.

Three possible uses of released capacity — 2. Reallocate capacity: shift workers to different tasks; deepen client service; redeploy to growth areas.

Three possible uses of released capacity — 3. Absorb growth: handle more volume with the same or slower-growing workforce.

HUMAN AUTHORITY AND WORKFORCE TRANSITION

The BNY public evidence emphasizes the second and third logics.

It does not provide a complete AI-attributable workforce-flow table.

Why the distinction matters: A company can become dramatically more productive without announcing an “AI layoff.”

The labor-market effect may appear through:

slower future hiring; higher output per employee; different role mix; reduced contractor demand; internal movement.

None of those mechanisms should be inferred without evidence.

Why the distinction matters — Evidence-safe conclusion: “BNY presents AI primarily as capacity creation and reinvestment. Public evidence reviewed for this case does not establish a direct AI-driven workforce reduction, but the long-run staffing implications remain an open research question.”

15. AI Building as a New Form of Domain Expertise: One of BNY's most important claims is that the majority of AI builders are no longer coming from engineering.[1]

That suggests a different definition of an AI expert.

Traditional AI expert: Value comes from:

model development; data science; machine learning engineering; infrastructure.

Embedded AI expert: Value comes from combining:

deep knowledge of a business domain; understanding of process failure; ability to identify automatable work; enough AI capability to build or shape the solution.

Examples can include:

payments operations; legal; risk; finance; HR.

product; client coverage.

Embedded AI expert — AIGW concept: domain-to-agent translation: The critical capability is increasingly:

“turning domain knowledge into agent instructions, workflows, rules, data access and control boundaries.”

That is a translation skill.

It sits between:

business expertise; technology; governance.

The BNY case therefore reinforces a pattern already visible in LONGi and JD:

HUMAN AUTHORITY AND WORKFORCE TRANSITION

domain expertise does not disappear when the task is automated. It can become part of the system that performs the task.

16. Comparison with Earlier AIGW Cases: BNY becomes more analytically useful when placed next to the existing Signature Case portfolio.

DBS — formal operating-model transformation: DBS identifies AI-affected roles, defines future roles and links learning to redeployment.

BNY adds a different mechanism:

“employees themselves become AI builders inside the operating model.”

JD — customer-service worker becomes AI trainer: JD shows an individual transition from service execution into AI operations.

BNY scales a similar principle institutionally by enabling nontechnical employees to build and supervise agents.

LONGi — inspector becomes AI trainer: LONGi demonstrates tacit manufacturing expertise becoming model-training knowledge.

BNY shows domain expertise becoming agentic workflow design.

Deutsche Telekom — autonomous agents act on infrastructure: Deutsche Telekom gives a strong case of autonomous execution.

BNY adds a more visibly formal organizational-management layer around digital employees through credentials, supervisors, guardrails and approval structures.

SimplifyNext — AI-native entry pathway: SimplifyNext builds a new first rung around real AI project work.

BNY shows how a large incumbent institution can retrofit an existing analyst programme so AI building becomes part of mainstream professional formation.

SimplifyNext — AI-native entry pathway — Cross-case pattern: Across these cases, AI capability increasingly moves from:

“specialist technical function”

toward:

“distributed institutional capability embedded in domain teams.”

BNY is one of the clearest examples of that transition at enterprise scale.

17. AIGW Framework: The AI-Enabled Workforce Operating Model: The BNY case suggests a six-layer operating model.

Layer 6 — Redesign — Why this framework matters: Many organizations invest heavily in Layers 1 and 2.

HUMAN AUTHORITY AND WORKFORCE TRANSITION

They provide:

licenses; generic training.

But the largest workforce transformation occurs in Layers 3–6.

That is where roles, management and decision rights begin to change.

18. The Manager–Agent Contract: Once a digital employee has a supervisor, the organization needs something equivalent to a role charter.

AIGW calls this the Manager–Agent Contract.

It should define:

Accountability — BNY evidence: BNY publicly documents several pieces of this structure:

credentials; supervisors; guardrails; telemetry; approvals.[1].

The full contract is not public.

But the organizational direction is visible.

Accountability — Workforce implication: Managers may increasingly manage designed capability, not only people.

This changes what effective management means.

OUTCOME METRICS

AI Access and Capability Metrics

FIGURE 2

reported measures

Reach, learning and employee agent-building

Company and vendor-hosted snapshots retain their original periods

Measure	Reported value	Period / source or boundary
Employees trained/onboarded on Eliza	99%	October 2025 ^[1]
Employees with Eliza access	100%	year-end 2025 ^[2]
Employees trained on Eliza	100%	September 2026 vendor-hosted customer story ^[8]
Daily or “super” AI users	almost two-thirds	September 2026 vendor-hosted customer story ^[8]
AI learning hours	171,000	2025 annual-report period ^[2]
AI training hours	400,000+	cumulative / “to date” vendor-hosted 2026 reporting ^[8]
Employees building agents	nearly 50%	year-end 2025 ^[2]
AI Pioneers bootcamp	40 hours	2026 programme description ^[8]

Source: BNY annual reporting and attributed company/vendor reporting ^[1] ^[2] ^[8]

Note: Access, training, learning hours and agent-building are different measures. The 2025 and 2026 values are separate snapshots, not a comparable time series or proof of proficiency.

OUTCOME METRICS

Production, Workflow and Enterprise Measures

FIGURE 3 reported measures

Production scale across distinct reporting periods

Reported counts and workflow outcomes are not a single time series

Measure	Reported value	Period / source or boundary
Enterprise AI solutions in production	160	year-end 2025 ^[2]
Digital employees live	134	year-end 2025 ^[2]
AI solutions	300+	September 2026 vendor-hosted reporting ^{[3] [8]}
Digital employees	140+ reported in journalism by August 2026	journalistic update ^[12]

Source: BNY annual reporting, vendor-hosted reporting and journalism ^{[2] [8] [12]}

Note: Definitions and reporting periods differ. These production counts should not be merged into a single time series.

OUTCOME METRICS

Workflow and Enterprise Outcome Measures

FIGURE 4

reported measures

Reported process and enterprise measures

Distinct outcomes retain their source and causal boundary

Measure	Reported result and boundary
Payment-repair issues handled by a digital employee	More than 10%; Microsoft customer story, September 2026 ^[3]
AI-assisted client settlement inquiries	80% faster; Microsoft customer story, September 2026 ^[3]
Relevant AI-assisted onboarding processes	20% faster on average; Microsoft customer story, September 2026 ^[3]
Employees in the platform operating model	>70%; year-end 2025 investor presentation ^[11]
Enterprise efficiency savings	~\$550 million; 2025 investor presentation, broader transformation metric ^[11]
Incremental investments	~\$500 million; 2025 investor presentation ^[11]

Source: Microsoft customer story and BNY investor reporting ^[3] ^[11]

Note: Workflow results are company-reported process outcomes, not proof of headcount reduction or company-wide productivity. Enterprise savings and investment are broader transformation measures, not AI-only effects.

WORKFORCE MODEL

AIGW Framework: The AI-Enabled Workforce Operating Model

FIGURE 5

framework layers

Access to workflow redesign

The BNY case suggests a six-layer operating model

Layer	Capability and diagnostic question
Layer 1 — Access	Employees can securely use approved AI capabilities. Question: Can people use the technology without creating uncontrolled shadow AI? ^[1]
Layer 2 — Fluency	Employees understand core capabilities, limitations, responsible use, and where AI may fit their work. Question: Can people identify useful applications? ^[1]
Layer 3 — Build	Domain employees can create or configure agents. Question: Can business knowledge be translated into working AI systems? ^[1]
Layer 4 — Deploy	Promising agents can move into production through testing, telemetry, approvals, and governance. Question: Can experiments become reliable operating assets? ^[1]
Layer 5 — Supervise	Humans monitor digital employees, manage exceptions and improve performance. Question: Who is responsible for the agent after deployment? ^[1]
Layer 6 — Redesign	Teams restructure the workflow around the new human-agent division of labor. Question: Are we merely adding AI to old work, or changing how the work should be organized? ^[1]

Source: AIGW analytical framework [1] [2] [4] [5] [8] [15]

Note: This is an AIGW framework informed by public BNY evidence. It is not a BNY-published maturity model or a claim that every organization follows a fixed sequence.

AGENT GOVERNANCE

The Manager–Agent Contract

FIGURE 6 contract dimensions

A role charter for supervised digital work

AIGW governance instrument informed by the BNY case

Dimension	Question the contract must answer
Purpose	What outcome is the agent responsible for? [1] [6]
Scope	What tasks may it perform? [1] [6]
Access	Which data and systems may it use? [1] [6]
Autonomy	Which actions may it take without approval? [1] [6]
Escalation	When must the agent stop and hand work to a human? [1] [6]
Performance	How is quality measured? [1] [6]
Review	Who evaluates the agent and how often? [1] [6]
Change control	Who can alter prompts, rules, access or workflows? [1] [6]
Accountability	Which human role remains responsible for the result? [1] [6]

Source: AIGW framework informed by public BNY governance descriptions [\[1\]](#) [\[6\]](#)

Note: The contract is an AIGW analytical proposal. BNY publicly describes credentials, supervisors, guardrails, telemetry and approvals; its complete action-level authority matrix is not public.

LEADERSHIP AND EVIDENCE

Leadership Implications, Evidence and Conclusion

19. The Early-Career Paradox: Build AI Before You Fully Know the Work?: BNY's approach is promising, but it also raises a serious developmental question.

If new analysts are encouraged to identify opportunities for AI and build agents early, how do they know what should be automated?

This is the same apprenticeship problem raised in the PwC case, but in a different form.

The opportunity: AI can help analysts:

explore unfamiliar topics; test ideas; automate administrative work; learn faster; contribute earlier.

The risk: An analyst may automate a process without fully understanding:

why the process exists; where exceptions matter; what the regulatory logic is; what an experienced operator notices; which apparent inefficiency is actually a control.

The risk — The answer is not delayed AI exposure: BNY's rotational, mentoring and cohort structures are important because they can pair:

AI building with domain formation.

The strongest first-rung model should require both.

The risk — AIGW design rule: “Do not ask early-career workers to automate a process before they can explain the purpose, risk and exception structure of the process.”

That rule preserves the upside of early AI fluency without allowing output speed to substitute for institutional understanding.

20. What Makes the BNY Sequence Distinctive: An independent Forrester assessment in May 2026 characterized BNY's approach as effectively sequencing the transformation as:

“platform first → people enablement → agentic automation.[13]”

That sequence is analytically useful.

Alternative sequence: automation first: An organization may:

1. buy an agent; 2. automate a workflow; 3. later explain the change to employees; 4. later create training.

This can create:

resistance; weak ownership; poor domain integration; shadow workarounds.

BNY sequence: The visible BNY trajectory is closer to:

LEADERSHIP AND EVIDENCE

1. build common platform; 2. create governance; 3. expand employee access; 4. develop fluency; 5. let employees build.

6. move solutions into production; 7. scale digital employees.

That does not guarantee success.

But it changes who participates in the transformation.

Employees are not only recipients of automation.

They can become authors of it.

BNY sequence — AIGW concept: participatory automation: AIGW uses participatory automation to describe automation in which affected domain workers have an active role in identifying, building, testing or supervising the systems that change their work.

BNY provides unusually strong enterprise evidence for that model.

Track builder participation, not only AI adoption:
Ask:

How many employees use AI?; How many can redesign a workflow with it?; How many production solutions came from domain teams?.

Adoption and transformation are different metrics.

Separate capacity creation from cost reduction: If AI releases capacity, decide explicitly whether it will be:

removed; redeployed; invested in growth.

Do not let the outcome emerge accidentally.

Govern digital employees as organizational actors:
A production agent should have:

identity; owner; scope; permissions; performance standards.

escalation; accountability.

Require workforce evidence alongside AI ROI:
Boards should ask not only:

“What did the agent save?”

but:

“What changed in the human role?”

Build an AI capability ladder: A useful workforce architecture can distinguish:

Explorer → User → Builder → Pioneer →
Supervisor / Workflow Owner

Not every employee needs to reach the final level.

LEADERSHIP AND EVIDENCE

But roles should have explicit expectations.

Redesign role profiles: If employees increasingly supervise digital labor, job descriptions should include:

AI orchestration; exception management; verification; workflow redesign; risk awareness.

Track internal mobility: As routine work changes, measure whether employees move toward:

client work; analytical work; AI building; supervision; product roles.

Training without role movement is incomplete evidence.

Redesign manager development: Managers need to learn how to manage:

people; automation; agent performance; human-agent handoffs.

Build a platform, not a collection of experiments: Shared infrastructure can standardize:

model access; data access; controls; monitoring; reuse.

Put domain experts close to building: The person who understands the process should have a structured way to shape the agent.

Create a path from prototype to production: Hands-on learning should connect to:

testing; telemetry; approvals; deployment.

Otherwise hackathons become innovation theater.

Treat governance as a scaling mechanism: BNY's own approach makes this point clearly: guardrails are not only restrictions.

They are infrastructure that allows more people to build safely.[1][6]

Do not stop at AI literacy: Teach workers to:

identify a problem; map the workflow; build or configure an agent; test it; evaluate failure.

understand controls.

Connect learning to production: A bootcamp should ideally produce:

a prototype; a use case; a workflow redesign; measurable capability.

Pair early AI building with domain formation: For analysts and other new entrants:

rotations; mentoring; process education; reflection; supervised building.

should develop together.

LEADERSHIP AND EVIDENCE

Measure the next role: Ask whether AI learning changes:

assignments; responsibility; promotion; mobility; pay.

That evidence is currently much thinner than training evidence.

25. What the Evidence Establishes: The reviewed evidence establishes that:

BNY created a centralized enterprise AI platform, Eliza; Eliza provides employees access to multiple approved models and AI-building capabilities; By October 2025, BNY reported 99% of employees trained and onboarded onto Eliza; By year-end 2025, BNY reported 100% employee access to Eliza; BNY delivered 171,000 AI learning hours during 2025.

Nearly 50% of employees were reported to be building AI agents by year-end 2025; BNY had 160 enterprise AI solutions in production at year-end 2025; BNY had 134 digital employees live at year-end 2025; Digital employees are multi-agentic AI systems operating alongside human colleagues; BNY describes digital employees as having distinct personas, credentials and supervisors.

Digital employees were embedded in payment teams; BNY uses role-based AI learning across multiple functions; BNY uses problem-driven bootcamps, hands-on building, hackathons and peer learning; Early-career analyst bootcamps

were part of the AI capability architecture; BNY's current 24-month Analyst Program explicitly includes AI Explorers, Builders and Pioneers.

BNY's AI governance includes guardrails, telemetry, approvals and continuous oversight; BNY's Responsible AI commitment includes approval processes, model review and appropriate human oversight; Later September 2026 vendor-hosted reporting says BNY had trained 100% of employees on Eliza and delivered more than 400,000 AI training hours to date; The same later reporting says almost two-thirds of employees were daily or "super" users of AI; Later vendor-hosted reporting says BNY had more than 300 AI solutions.

Later vendor-hosted reporting says a digital employee handled more than 10% of payment-repair issues globally; Later vendor-hosted reporting says certain AI-assisted client inquiries were processed 80% faster and relevant AI-assisted onboarding cases 20% faster on average.

26. What the Evidence Does Not Establish: The reviewed evidence does not establish:

that 99% or 100% onboarding equals advanced AI proficiency; that every employee builds agents; that every employee building an agent has deployed one into production; the exact number of employees whose roles have materially changed because of AI; the number of human jobs replaced by digital employees.

LEADERSHIP AND EVIDENCE

that 134 digital employees equal 134 human FTEs; that BNY reduced headcount because of Eliza; that AI will have no future effect on BNY headcount; the share of company-wide efficiency savings directly caused by AI; that the platform operating-model transformation was caused by AI alone.

the error rate of payment digital employees; the full risk profile of automated payment repair; a complete action-level approval / veto / override matrix for payment digital employees; who holds final authority in every payment exception; the percentage of early-career analysts who become AI Builders or Pioneers.

whether AI training increases promotion probability; whether AI builders receive higher pay; whether AI capability improves long-term retention; whether employees who do not become strong AI users face disadvantages; whether agent-building capability is equally distributed across roles, regions or demographic groups.

an independently audited causal estimate of BNY's AI productivity ROI.

27. Research Gaps to Watch: 1. How many employees are active AI builders rather than occasional users?; 2. What qualifies someone as an AI builder?; 3. How many employee-built agents reach production?; 4. What is the prototype-to-production conversion rate?; 5. What are the most common source functions for successful agents?;

6. How many agents are reused by teams other than their original builder?; 7. How are AI builders recognized in performance management?; 8. Does agent building affect promotion or compensation?; 9. How many employees become AI Pioneers?; 10. How are Pioneers distributed across regions and functions?.

11. What does an AI Pioneer do after the 40-hour bootcamp?; 12. How much work time is allocated to agent building and learning?; 13. What share of digital employees have dedicated human supervisors?; 14. What is the manager-to-digital-employee ratio?; 15. How is a digital employee's performance reviewed?.

16. Can a human manager disable a digital employee immediately?; 17. Which digital-employee actions require pre-approval?; 18. Which payment-repair actions can be completed autonomously?; 19. How are exceptions routed?; 20. What is the error / rework rate for automated payment repair?.

21. What happens when an agent and human operator disagree?; 22. Who owns final accountability for an erroneous payment repair?; 23. What proportion of payment work has shifted from execution to supervision?; 24. Has human staffing in payment operations changed?; 25. Has replacement hiring changed?.

LEADERSHIP AND EVIDENCE

26. Have digital employees absorbed volume growth without equivalent hiring?; 27. How have analyst job descriptions changed over time?; 28. What percentage of analysts build agents during the 24-month programme?; 29. Does AI reduce or increase time spent learning core financial operations?; 30. How does the firm test that analysts understand the process they automate?.

31. What are the 12- and 24-month career outcomes of AI-intensive analyst cohorts?; 32. How portable are the AI-building skills outside BNY?; 33. Can independent research validate the 2026 workflow speed claims?; 34. How much of enterprise efficiency savings can be causally attributed to AI?; 35. How will the human role change as digital employees become more autonomous?.

28. Method and Source Quality: This Deep Dive separates six evidence types.

A. Formal company reporting: BNY's 2025 Annual Report and year-end investor materials provide the strongest evidence for:

employee access; learning hours; share of employees building agents; production AI solutions; number of digital employees.

enterprise operating-model context.[2][11].

These are first-party disclosures, but they have higher institutional weight than marketing case

material.

B. First-party AI implementation reporting: BNY's October 2025 Eliza article provides detailed evidence on:

training architecture; role-based pathways; analyst bootcamps; hands-on building; guardrails.

digital employees in payment teams; supervisors and credentials.[1].

It is the core behavioral source for this case.

C. Responsible-AI policy evidence: BNY's public Responsible AI commitment provides enterprise governance principles including:

accountability; approval processes; risk management; model review; appropriate human oversight.[6].

It establishes governance intent and structure.

It does not establish the exact control design of every digital employee.

D. Current early-career programme evidence: BNY's current Analyst Program provides direct evidence that AI pathways have entered the formal 24-month early-career architecture.[5]

This is stronger than inferring entry-path redesign from general training.

LEADERSHIP AND EVIDENCE

E. Vendor-hosted customer-story evidence:
Microsoft's September 2026 BNY customer stories provide later evidence on:

400,000+ training hours to date; 100% training; almost two-thirds daily / super users; 300+ AI solutions; payment-repair share.

faster inquiries; faster onboarding.[3][8].

These sources are useful follow-on evidence but are promotional vendor/customer stories rather than independent evaluations.

F. External analytical and journalistic evidence: Forrester, Axios, the Wall Street Journal and other reporting help corroborate the emergence of BNY's digital-workforce model and provide external context.[7][9][12][13]

These sources are supporting, not substitutes for primary evidence.

Training is not proficiency: Do not say:

"100% of BNY employees are AI experts."

Say:

"BNY reports enterprise-wide access/training and a large share of active builders, with deeper tiered pathways for more advanced capability."

Builder is not deployer: An employee who creates an agent is not necessarily authorized to deploy it into production.

Production requires additional controls.

Digital employee is not FTE: Do not convert digital-employee counts into job-loss estimates.

Workflow speed is not headcount reduction: An 80% faster inquiry process does not establish an 80% staffing reduction.

Enterprise efficiency is not AI ROI: BNY's broader efficiency savings reflect a larger operating-model transformation.

Do not attribute them solely to Eliza.

Supervisor is not automatically final authority: Public evidence of a human supervisor does not tell us the complete approval / override architecture.

Current 2026 metrics do not replace 2025 metrics silently: The annual-report figures and September 2026 vendor-hosted figures are snapshots from different periods and sometimes use different definitions.

30. AIGW Interpretation: The Workforce Is Becoming a System Designer: The deepest shift in the BNY case is not that AI performs more tasks.

LEADERSHIP AND EVIDENCE

It is that more employees are being asked to participate in designing how work is performed.

This changes the workforce model in three ways.

From execution to configuration: Employees increasingly specify how an agent should behave.

From individual productivity to workflow productivity: The target is no longer only:

“Can this person work faster?”

It becomes:

“Can this team redesign the system so the work itself flows differently?”

From fixed roles to adaptive roles: The employee may alternate among:

domain expert; builder; verifier; supervisor; exception handler.

client-facing professional.

That is a more fluid role architecture.

From fixed roles to adaptive roles — The organizational implication: AI capability becomes less like a software skill and more like a form of institutional agency.

Workers who can redesign workflows influence:

where labor is used; what gets automated; where judgment remains; how client service is delivered; how controls are embedded.

That makes AI fluency a workforce-governance issue, not only a productivity skill.

1. Sequence matters: Platform → people → agents is usually safer than agents → people adaptation.

2. Democratize building, not governance: Many employees can build.

Production authority should still be controlled.

3. Treat domain expertise as an AI asset: The best agent builder may be the person who understands the workflow deeply, not the person with the most advanced machine-learning degree.

4. Design a progression from user to builder: AI literacy alone does not redesign work.

5. Manage digital employees explicitly: Give every production agent:

identity; owner; scope; permissions; performance standards.

escalation; accountability.

LEADERSHIP AND EVIDENCE

6. Redesign manager capability: Managers need to supervise systems as well as people.

7. Bring AI into early-career learning deliberately: Young workers should learn to build with AI and understand the underlying business process.

8. Measure real capability: Training hours and course completion are leading indicators.

Production behavior is stronger evidence.

9. Separate capacity from employment outcome: Decide how released capacity will be used.

10. Publish better workforce-flow evidence: The next step in transparency is to show how AI changes:

role mix; hiring; internal mobility; staffing; progression.

Conclusion: BNY's AI transformation is not best understood as a story about a bank deploying hundreds of AI tools.

It is a story about a bank trying to change who is allowed to participate in automation.

Eliza gives employees a common enterprise AI environment.

Training moves from general awareness toward role-specific pathways and agentic bootcamps.

Employees outside engineering increasingly build agents.

Digital employees operate alongside human teams.

Early-career analysts enter a development system in which AI building is already part of professional formation.

Managers begin to supervise not only people, but digital capability.

The sequence can be summarized as:

“Access → Fluency → Build → Deploy → Supervise → Redesign”

That is the central workforce architecture of the case.

The evidence is strongest on behavior.

BNY has unusually strong public evidence of:

enterprise adoption; learning; agent building; digital employees; operational deployment.

early-career integration.

The evidence is weaker on the downstream employment outcome.

LEADERSHIP AND EVIDENCE

We still do not know with precision:

how many roles will become smaller; how hiring patterns will change; how many employees will shift from execution into supervision; whether AI building changes promotion or pay; which autonomous actions require human approval.

Those are not minor omissions.

They are the next frontier of enterprise AI workforce disclosure.

The BNY case nevertheless establishes something important.

A workforce does not become AI-ready simply because every employee has access to AI.

It becomes AI-ready when people can understand the technology, connect it to their domain, build responsibly, move useful systems into production, supervise what those systems do and redesign work around the new division of labor.

That is a much higher bar.

And it suggests a different definition of the AI-era employee:

“not merely a user of intelligent tools, but a participant in designing the intelligent operating system of the organization.”

NOTES & REFERENCES

Notes and References

1. BNY. *Unlocking Value with BNY's Enterprise AI Platform*. 20 October 2025. Company / first-party implementation reporting.
Supports: Eliza platform approach; 99% employee training/onboarding; role-based AI pathways; analyst bootcamps; hands-on agent building; guardrails, telemetry and approvals; 75% quarter-over-quarter increase in production AI solutions; digital employees in payment teams; 100+ digital employees with personas, credentials and supervisors.
Boundary: company-reported adoption and behavioral evidence; does not provide a standardized AI productivity or workforce-reduction estimate.
<https://www.bny.com/corporate/global/en/insights/unlocking-potential-enterprise-ai-platform-bny.html>
2. BNY. *Annual Report 2025*. February 2026. Formal company annual reporting.
Supports: 100% employee access to Eliza; 171,000 AI learning hours in 2025; nearly 50% of employees building agents; 160 enterprise AI solutions in production; 134 digital employees; AI as capacity creation and part of enterprise transformation.
Boundary: annual-report AI metrics should not be combined with later 2026 customer-story metrics without preserving the time difference.
<https://www.bny.com/corporate/global/en/investor-relations/annual-report-2025.html>
3. Microsoft. *Frontier Firm BNY resolves client inquiries 80% faster with Microsoft AI-powered Eliza*. 24 September 2026. Vendor-hosted customer story.
Supports: 300+ Eliza solutions; digital employee handling >10% of payment-repair issues globally; >10% of certain client inquiries AI-resolved/assisted and 80% faster; 25% of new onboardings AI-assisted and 20% faster on average.
Boundary: later vendor/customer evidence; not an independent audit and not a staffing metric.
<https://www.microsoft.com/en/customers/story/27322-bny-microsoft-365-copilot>
4. OpenAI. *BNY builds "AI for everyone, everywhere" with OpenAI*. 12 December 2025. Technology-provider customer case.
Supports: 20,000 employees actively building agents at the reported point; 125+ live use cases; digital employees with identities, access controls and workflows; BNY executive description of human operator shifting toward training/nurturing digital employees.
Boundary: provider-hosted customer case; figures reflect the publication date and should not silently replace annual-report metrics.
<https://openai.com/index/bny/>
5. BNY. *Analyst Program*. n.d.. Current company early-career programme evidence.
Supports: 24-month analyst pathway; rotations; mentoring; progression toward associate; explicit AI Explorers, AI Builders and AI Pioneers learning pathways.
Boundary: current programme architecture; does not provide completion, promotion or salary outcomes attributable to AI learning.
<https://www.bny.com/corporate/global/en/about-us/careers/students/analyst-program.html>

NOTES & REFERENCES

6. BNY. *Our Commitment to the Responsible, Ethical Use of Data and AI*. n.d.. Company governance commitment.

Supports: accountability structures, approval processes, model review, lifecycle monitoring, employee training, transparency, explainability and appropriate human oversight.

Boundary: enterprise governance commitment; does not establish the action-level authority matrix for every digital employee.

<https://www.bny.com/corporate/global/en/about-us/technology-innovation/artificial-intelligence/data-ai-ethical-commitment.html>

7. Axios. *Over 100 “digital employees” work at this Wall Street bank*. 17 October 2025. Attributable business journalism.

Supports: external reporting on digital employees with managers/logins; CEO Robin Vince's statement that BNY does not frame AI adoption primarily as a headcount-reduction strategy.

Boundary: executive framing does not prove AI will have no future staffing effect.

<https://www.axios.com/2025/10/17/ai-wall-street-digital-workers>

8. Microsoft. *To innovate faster, BNY upskills 100% of staff on Eliza*. 24 September 2026. Vendor-hosted workforce / customer case.

Supports: 400,000+ AI training hours to date; Explorer/Builder/Pioneer tiers; 40-hour Pioneer bootcamp; cohort expansion from ~200 to ~800; 100% employee training; almost two-thirds daily/super users; governance-by-design description; 4,000 AI project ideas and 300+ production use cases.

Boundary: current vendor-hosted customer-story evidence; cumulative training figure is not comparable to the 171,000 hours reported for 2025 alone.

<https://www.microsoft.com/en/customers/story/27320-bny-microsoft-foundry>

9. Wall Street Journal. *Digital Workers Have Arrived in Banking*. 2025. Independent journalism.

Supports: external reporting that BNY digital employees have organizational identities, logins and human managers and perform tasks including payment validation and code-related work.

<https://www.wsj.com/articles/digital-workers-have-arrived-in-banking-bf62be49>

10. BNY Institute. *Future of Finance: Key Developments in 2026*. n.d.. Company institutional / market report.

Supports: AI-enabled payment validation for payments that cannot be straight-through processed; settlement prediction; broader 24/7 financial-infrastructure context.

Boundary: descriptive institutional context; not a staffing outcome.

<https://www.bny.com/content/dam/bnymellon/documents/pdf/institute/future-of-finance.pdf>

NOTES & REFERENCES

11. BNY. *Fourth Quarter 2025 / Full-Year Strategic Update*. n.d.. Company investor reporting.
Supports: >70% of employees in new platforms operating model by year-end 2025; ~50% transitioned during 2025; ~\$550m efficiency savings; ~\$500m incremental investments; 99% employees on Eliza; 160 AI solutions; 134 digital employees.
Boundary: operating-model and financial metrics reflect broader enterprise transformation and should not be attributed solely to AI.
<https://www.bny.com/content/dam/bnymellon/documents/pdf/investor-relations/earnings/quarterly-update-presentation-4q-2025.pdf>

12. Axios. *BNY CEO tells incoming analysts to embrace AI*. 18 August 2026. Journalism / early-career evidence.
Supports: CEO message encouraging incoming analysts to experiment with AI and question existing processes; reports 140+ digital employees at the publication point.
Boundary: journalistic snapshot; does not measure analyst capability or employment outcomes.
<https://www.axios.com/2026/08/18/bny-ceo-vince-financial-analysts-ai>

13. Forrester. *BNY Built Its Digital Workforce Backward — And It's Working*. 12 May 2026. External analyst commentary based on BNY case research.
Supports: external interpretation that BNY sequenced platform foundation and workforce enablement before scaling 130+ digital employees; regulated-bank agentic operating-model context.
Boundary: analyst interpretation; financial performance cannot be causally attributed to the digital-workforce model from this source alone.
<https://www.forrester.com/blogs/bny-built-its-digital-workforce-backward-and-its-working/>

14. BNY and The University of Manchester. *Future of Work Alliance*. 21 May 2026. Company / university research partnership.
Supports: five-year £5 million initiative focused on responsible human-led digital transformation and new models for AI application in large data-intensive organizations.
Boundary: research commitment, not evidence of a completed workforce outcome.
<https://www.bny.com/corporate/global/en/about-us/newsroom/press-release/bny-and-the-university-of-manchester-launch-the-future-of-work-alliance.html>

15. BNY. *Artificial Intelligence*. n.d.. Company technology overview.
Supports: Eliza as a central enterprise platform for approved data, AI solutions, workflow enablement and governance.
<https://www.bny.com/corporate/global/en/about-us/technology-innovation/artificial-intelligence.html>