WHOLESALE BANKING · THE 2027–2031 DELIVERY PROGRAM

# One machine. A multi-year program.

A target operating model for the transformation — and the three instruments each turn of its loop produces: Foresight senses the market, Incubation pilots the portfolio, Transformation scales the results. Built on Citi's Q2 2026 public results, paced to Citi's own disclosed glidepath, across 2027–2031.

**Paced to what Citi can realistically deliver.** The efficiency and RoTCE glidepath is Citi's own, verified from Investor Day 2026; consent-order removal, which gates autonomy in regulated flows, sits with the regulators. Every ° figure is illustrative — Citi has published no AI-attributable savings schedule, so the dollar run-rate is kept deliberately below JPMorgan's ~$2B/yr leader benchmark: a discussion starter to react against, never a claim. Built on the public record only, and adversarially reviewed — each instrument closes with the questions it should be asked, several conceded outright.

CONTENTS

# Six sections, one loop.

The case, then the machine that runs it, then the three artefacts one quarterly cycle produces — Foresight, Incubation, Transformation. Figures throughout are shown at the **Q1 2027** cycle on the **base** pace with **supervised** autonomy; the scenario in force is stated at the top of each section that depends on it.



This edition is generated directly from the interactive portal, page for page — it is a rendering of that material, not a separate document. **Illustrative concept — independent analysis based on Citi's public disclosures. Not official Citi material; figures are directional, used to frame discussion, and are not Citi projections.**

THE MACHINE

RUN FROM AN OFFICE OF TWO

# The Delivery Mandate — a Target Operating Model

The mandate drawn as one loop: an AI-run machine (sense → pilot → scale) inside a thin human governance rim. A business management system, not a technology program.

- Run from an office of two — principal + one analyst; the businesses, finance and controls already sign here
- Every mechanic is live — run the sensing pass, move the selection weights, advance the quarter and watch it bank

**"As the delivery engine of the transformation,**

**we convert AI capacity into the outcomes the firm has committed to,**

**by running three disciplines that keep delivery FIT for purpose."**

**Foresight — sense**

**Incubation — pilot**

**Transformation — scale**

DOES NOT MOVE WITH THE QUARTER

## The standing case

The three reports below are readings. This is the argument they are readings *of* — where the firm stands on its own published numbers, the capability already proven elsewhere, and the questions the case should be asked before anyone funds it.

Where the firm stands · The franchise spine · Four currencies · Five levers · One flow, end to end · Capability evidence · Pressure-test · 16 questions

## THE OPERATING-LEVERAGE GLIDEPATH — **EFFICIENCY DOWN · ROTCE UP · 2027 → 2031**

planning now · Q3 2026 — delivery opens Q1 2027

AUTONOMY: Assisted | **Supervised** | Autonomous

PACE: Conservative | **Base** | Accelerated

**= Citi's disclosed glidepath** · sub-55% ≈ Q3 2029

**Q1 2027** · plan · stand up · efficiency **58.5%↓** · RoTCE **11.2%↑** · run from **a desk of ~6** · AI run-rate net-investing (J-curve) illustrative

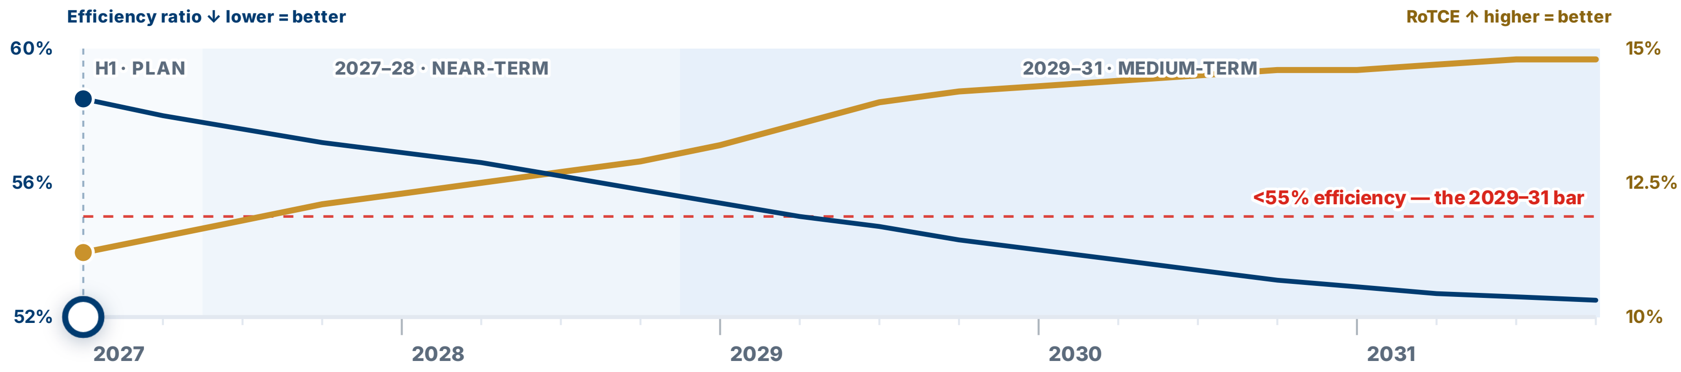


## WHAT THE LOOP PRODUCES IN **Q1 2027** · PLAN · STAND UP

**F · SENSE** — cadence set

### Foresight

**Radar v1**

the inaugural scan — regulated-flow autonomy still gated

Inside: the radar, and the shortlist it hands to Incubation.

**I · PILOT** — chartering

### Incubation

**planning**

portfolio v1 ranked · first charter shaped

Inside: eight opportunity families, and a live ROI model.

**T · SCALE** — net-investing

### Transformation

**58.5%°**

AI run-rate: net-investing (J-curve) · RoTCE target 11–13% (2027–28)

Inside: the quarter's close, the leverage path, the scorecard.

THE STANDING CASE · THE ARGUMENT BEHIND THE LOOP · NOT A CYCLE OUTPUT

# What holds whichever quarter you are in.

The three instruments — Foresight, Incubation, the leverage review — are quarterly reports, and a quarterly report should carry only what the quarter changes. Everything permanent was pulled out and put here. There is no quarter selector on this page — that absence is the point.

**4**
currencies transformation value can be paid in

**5**
business levers where that value lands

**15**
questions the case invites, answered

THE FRAME HAS CHANGED

## Three facts set the context.

A record quarter, public medium-term targets, and a productivity agenda that now runs out of the CFO organization. These are the fixed points every cycle is argued against — they do not move when the quarter does.

1 · THE TURN IS DELIVERED

2 · THE NEXT CHAPTER IS PUBLIC

3 · THE MACHINERY MOVED

### Q2 2026 was a record.

| | |
|---|---|
| Revenues | **$24.8B · +14%** |
| Expenses | **$14.2B · +5%** |
| Operating leverage | **>9 points** |
| RoTCE / efficiency | **13.0% · 57.4%** |
| Capital return | **$30B buyback · +12% dividend** |

The “can Citi fix itself” discount is closing. The conversation has moved on.

### Investor Day set the arithmetic.

| | |
|---|---|
| RoTCE 2027–28 | **~11–13%** |
| RoTCE medium term | **~14–15%** |
| Efficiency ratio | **<55%** |
| Incremental investment | **~$5B · 2026–28** |
| Transformation expense | **~50% roll-off near term** |

Every number is an operations statement: expenses growing materially slower than revenue, every year, through an investment cycle.

### Productivity now runs from the finance chair.

| | |
|---|---|
| Operator CFO | **since March 2026** |
| Processes mapped for automation | **100+ · weekly review** |
| Severance YTD | **>$800M** |
| Headcount | **219,000** |
| Employees on AI tools | **~9 in 10** |

A portfolio with P&L accountability and “returns discipline” as the stated first principle — no longer a collection of technology projects.

**The durability question.** Q2's operating leverage was revenue-led — a +14% quarter does the heavy lifting, and a strong Markets tape flatters any ratio. The targets assume mid-single-digit revenue; the year Markets normalizes, the ratio is carried by the expense line or not at all. Separating the revenue-led from the productivity-led share of operating leverage — and growing the second — is the difference between a good quarter and a structural sub-55%. That is a business-transformation problem; technology is the instrument.

What each quarter's cycle does about it is the subject of the three instruments.

## WHERE VALUE CONCENTRATES

# Three businesses, one operations spine.

The wholesale franchise is three revenue engines sharing a single processing backbone. The backbone is where the addressable expense sits — and therefore where the radar points, every cycle.

### Services — the engine

Treasury & Trade Solutions (cross-border / international payments, USD clearing) + Securities Services. The highest-return, most process-heavy franchise — $416B of cross-border value moves through it.

### Markets

Fixed Income + Equities (incl. Prime). High-volume, low-latency; exception- and reconciliation-heavy.

### Banking

Investment Banking (M&A, ECM, DCM) + Corporate Banking. Document- and approval-heavy relationship work.

The shared spine — onboarding and KYC, reference & client data, trade & payment processing, reconciliations, settlement, controls, and regulatory reporting.

**FY2025 REVENUE VS. OPERATING EXPENSE** PUBLIC

Revenue · Operating expense

**Services** 50.7% eff. — $21.3B rev · $10.8B opex

**Markets** 64.1% eff. — $22.0B rev · $14.1B opex

**Banking** 54.9% eff. — $8.2B rev · $4.5B opex

Bars are scaled to the largest revenue line; the segment efficiency ratios are simple arithmetic on the two disclosed figures. Markets carries the ratio drag, but the densest *convertible* cost sits in the cross-border exception flow inside Services — which is why that flow is the recurring lighthouse call.

FROM CAPACITY TO P&L

# The four currencies of transformation value.

The discipline that governs every cycle: AI creates capacity; capacity becomes value only through explicit decisions. “Savings” is just one of the four currencies that value can be paid in — and it is not the one most of today's capacity is being paid in.

¤1

### Expense — release

Released positions, avoided backfills, reduced vendor and BPO spend, location-mix shifts. The only currency that lands directly on the expense line — and the one severance has been buying. The transformation task: convert one-off severance into a structurally different shape.

¤2

### Operating leverage — redeploy

Capacity moved into growth work: volume absorbed without hiring, expenses held flat while revenue compounds. Invisible as “savings,” decisive for the ratio — and the currency most of the ~100K weekly developer-hours are currently paid in.

¤3

### Revenue speed — enable

Onboarding cycle time is time-to-first-revenue; payment certainty and exception latency are client experience — share of wallet in Services. Small percentages on a wholesale franchise running ~$15.3B a quarter (62% of firm revenue) outweigh large percentages on the cost base.

¤4

### Risk & remediation productivity — protect

Cheaper, faster, audit-ready evidence and higher validation throughput protect the planned ~50% roll-off of transformation expense — the same roll-off that part-funds the investment envelope. “A large body of work passing internal audit validation” is a trajectory to defend, not just report.

WHERE THE P&L MOVES

# Five business levers.

The opportunities in Incubation are the supply side — where capacity comes from, and which of them is chartered changes every cycle. These five levers are the demand side — where capacity lands, on whose line, measured how — and they do not change. Sizing is gross, overlapping, one lever inside the committed program. Click any card.

L1

SERVICES · MARKETS

### Unit cost-to-serve — payments & post-trade

Measured by: Cost per item by flow · STP rate · exception aging

+ details

L2

CROSS-LOB → REVENUE

### Onboarding & KYC as a revenue accelerator

Measured by: Time-to-revenue · WIP aging · cost per KYC file

+ details

L3

CONTROLS

### Remediation & controls productivity

Measured by: Cost per validated milestone · evidence cycle time · re-work rate

+ details

L4

THE BRIDGE

### Workforce shape — severance into structure

Measured by: Hours redeployed vs. released · backfill rate · structural vs. cyclical saves

+ details

L5

TECHNOLOGY

### Technology unit economics

Measured by: Change throughput per $ · decommission rate · run-cost per application

+ details

SEE IT IN ACTION

# What one wholesale flow looks like, end to end.

Before the portfolio arithmetic means anything, the mechanic has to be concrete. Two flows, from copilot to multi-step agent — grounded in Citi data and bank-grade governance, with a human sign-off gate on any regulated action. Pick a scenario and run the agent.

## FLOW 1 · CROSS-BORDER PAYMENT EXCEPTION

1. **Ingest**
   Read the failed MT103 + payment record
2. **Diagnose**
   Classify failure (sanctions screening hit)
3. **Context**
   Pull entity, beneficiary & sanctions context
4. **Resolve**
   Draft repair + rationale; clear false positive
5. **Act**
   Auto-release or route to human for sign-off

| BEFORE · MANUAL | AFTER · AGENTIC |
|---|---|
| ~45 min · ~$28/case | ~90 sec · ~$2/case |

**BASIS**

Shape drawn from HSBC × Google Cloud AML screening (see the capability evidence below); agentic repair lifts straight-through processing.

Resolved by the agent — against ~45 min · ~$28/case manually

## FLOW 2 · KYC ONBOARDING

1. **Ingest**
   Read incorporation docs, UBO forms, IDs
2. **Extract**
   Resolve entities, ownership & control chain
3. **Screen**
   Adverse-media & sanctions triage
4. **Draft**
   Generate the KYC narrative & risk rating

| BEFORE · MANUAL | AFTER · AGENTIC |
|---|---|
| ~4–6 weeks | ~2 days |

**BASIS**

Shape drawn from JPMorgan COIN document review (see below); KYC automation cuts onboarding 40–70%.

Resolved by the agent — against ~4–6 weeks manually

5 **QC**
Quality-check, evidence & file for approval

A demonstration of the mechanic, not a benchmarked result — the timings are illustrative and the basis line names the external evidence each shape is drawn from.

WHY THE SIGNALS ARE CREDIBLE

# Capability evidence — what has already been done elsewhere.

A claim that something is “ready now” has to be backed by a result achieved somewhere, not a vendor promise. These are the documented quantum-leaps the whole system weighs — used as bounds on what is possible, never as forecasts for Citi. Including the failure, which is a design input too.

**AMAZON Q DEVELOPER**

30,000+ apps modernized → ~4,500 developer-years and ~$260M/yr saved — bound on change-capacity per dollar.

**HSBC × GOOGLE CLOUD**

AML transaction monitoring: −60% false positives, 2–4× more financial crime caught across 1B+ transactions/month.

**JPMORGAN COIN**

360,000 lawyer-hours/yr of agreement review compressed to seconds — document-heavy work collapses.

**MORGAN STANLEY / GOLDMAN**

Firmwide assistants at near-total adoption; 30-min lookups → seconds over hundreds of thousands of documents.

**OPERARTIS RECONCILIATION**

Auto-match 79%→94% at a tier-1 IB, freeing a 20+ person recon desk; ~90% manual-time cuts.

**KLARNA (THE CAUTION)**

Work of ~700 agents automated, then partly walked back — why human-in-the-loop gates every judgment-dense step.

The modeled savings in Incubation sit deliberately below these bounds, and Klarna's over-automation of judgment work is precisely why human-in-the-loop tiers gate every regulated or judgment-dense step in this design.

PRESSURE-TEST

# The sixteen questions this case should be asked.

Same rule across the system: a serious review will — and should — attack this. The challenges it expects, answered plainly — several conceded outright, because the concessions are the point. Click any card.

**ON THE SIZING AND THE PORTFOLIO** — eight challenges to the opportunity case

FAIR — GROSS VS. NET

**"The headline potential exceeds Citi's entire $2–2.5B savings program — overreach?"**

+ answer

CONCEDED — DESIGNED IN

**"The 100,000 hours a week is capacity, not P&L — you're conflating the two."**

+ answer

BY DESIGN · ROM

**"Citi doesn't disclose operations cost by process — where do the pools come from?"**

+ answer

TRUE — READ AS LANDSCAPE

**"The opportunity rows double-count — KYC, screening, and controls overlap."**

+ answer

**TRUE — USED AS BOUNDS**

**"The benchmarks are survivor stories — Klarna walked back its agent claims; HSBC's result is classic ML, not agentic."**

+ answer

**AGREED — SEQUENCED FOR IT**

**"Consent orders on data and controls are open — autonomous agents in sanctions flows is reading the room badly."**

+ answer

**DELIBERATELY OUT OF SCOPE**

**"There's no cost-to-achieve and no J-curve — what does this cost?"**

+ answer

**DIFFERENT MODEL**

**"Cross-LOB KYC consolidation has a graveyard — the external utilities failed."**

+ answer

## ON THE P&L AND THE OWNERSHIP — eight challenges to the business case

**FAIR — AND THE POINT**

**"Operating leverage is already above 9 points — why does this need a program?"**

+ answer

**THE STRUCTURAL FLOOR**

**"Efficiency already beat the full-year guide — won't the targets take care of themselves?"**

+ answer

**DELIBERATELY**

**"The business case is the opportunity case re-labeled."**

+ answer

**CONCEDED — THE WHOLE GAME**

**"Capacity vs. P&L is still the weak joint."**

+ answer

EVENT VS. ANNUITY

**"Severance is doing the work — why layer transformation on top?"**

+ answer

OWNER & ARITHMETIC

**"Why should this sit with the CFO organization rather than the businesses?"**

+ answer

OWNER & COMPOSITION

**"These are CIO programmes. This is not CFO work."**

+ answer

CONCEDED — SECOND ORDER

**"The revenue-speed claims are soft."**

+ answer

WHAT THE CYCLE DOES WITH ALL THIS

# The standing case is the argument. The instruments are the work.

Nothing above moves when the quarter moves — which is exactly why it is not in the quarterly reports. Each instrument below carries only what its cycle changes: what the radar sees now, what is chartered now, what has banked now.

F · SENSE

**Foresight →**

The market & capability radar — dated signals, re-read every cycle.

I · PILOT

**Incubation →**

The opportunity portfolio — what is queued, chartered, piloting, in production.

T · SCALE

**Leverage review →**

What banked, what it did to the ratio, and what is still to come.

THE DELIVERY MACHINE · THE TARGET OPERATING MODEL · 2027–2031 DELIVERY PROGRAM

# A firm-scale mandate. An office of two.

This is a management system, not a technology program — the CFO organization's transformation run the way a portfolio is run. AI does the labor around the loop; two people steer it; and every signature that governs it already belongs to someone on the firm's payroll.

> **“As the delivery engine of the transformation, we convert AI capacity into the outcomes the firm has committed to — by running three disciplines that keep delivery FIT for purpose.”**
>
> Foresight — sense · Incubation — pilot · Transformation — scale. Re-weighted here to run on agents.

| CYCLE | CURRENT PRIORITY | NEXT TURN | LEGEND |
|---|---|---|---|
| ● Q1 2027 | Charter № 2027-01 | in Q2 2027 | sample |


HOW TO READ THIS

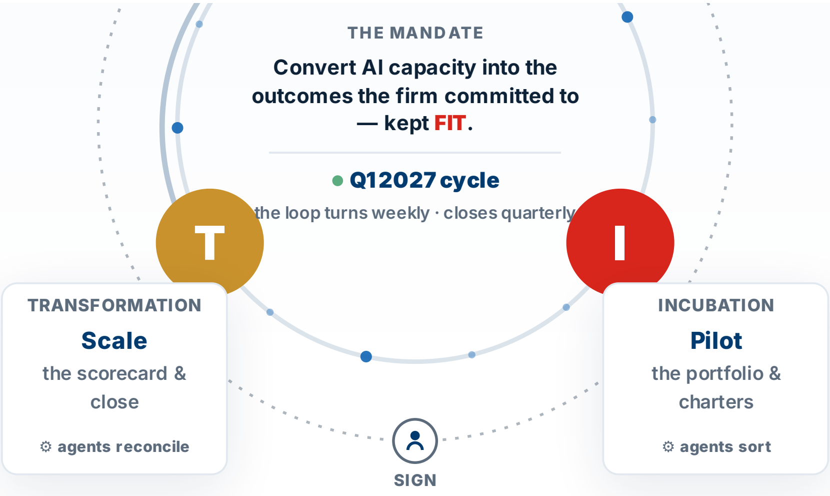


The inner ring is the machine — it runs the loop. The outer rim is the humans — **thin by design**: three gates, and the office adds only the first.

# One loop. The machine runs it; two people steer.

Scroll, and each discipline lights up on the ring and opens its live artifact here. Nothing on this page sits outside the loop: Foresight produces a market radar, Incubation a sorted portfolio and signed charters, Transformation a banked scorecard — and the reconciliation closes back to Foresight so the next turn is sharper than the last. Every artifact below is populated for the current Q1 2027 cycle; where only the inside view can supply a number, the cell is flagged sample.

**F** FORESIGHT · SENSE

## The market radar — drafted by agents, blessed in one meeting.

Sensing agents read earnings, the Street, regulators and market structure, and draft dated signals — each one landing on a business flow, a date and a decision. The office does not research; it approves. ⚙ agents draft ✋ weekly review approves

→ full Foresight instrument

7 approved signals · this cycle

| SIGNAL — LATEST READ | SO-WHAT — THE BUSINESS DECISION | FLOW · WHEN |
|---|---|---|
| ● Instant rails (Citi Payments Express & peers) are collapsing the tolerance for payment-exception latency | Convert the cross-border exception flow this cycle, before volume forces it | Cross-border · **Q1 2027** convert now |
| ● AI now drafts exception repairs to a standard humans will approve — proven on historical cases sample | Move from study to charter; a human approves every action | Cross-border · **Q1 2027** convert now |
| ● Consent-order & control expectations are converging on the same client-data foundations the business needs anyway | Sequence data-dependent flows behind foundations already funded — complementary, not competing | Client data · **Q1 2027** align now |
| ● Peers face productivity-disclosure questions from the Street, not tool counts — the question is coming to Citi | Have the productivity-led share of operating leverage measured before it is asked | Scorecard · **Q2 2027** prepare |
| ● Corporate treasurers increasingly price payment certainty; exception latency is appearing in RFP scoring | Fold exception-aging into cross-border client SLAs at the next renewal | TTS · **2028** watch |
| ● The ~$5B transformation envelope funds the first wave with no new budget ask | Draw lighthouse cost-to-achieve from the envelope, not a new line | Portfolio · **Q1 2027** fund |
| ● Severance is running ahead of structure; the backfill rate is the real test of whether it sticks | Track backfill as a first-class number from Q3 — the honesty check on released cost | Workforce · **Q1 2027** track |

Row marks: ● what the outside demands · ● what moves the business · ● what funds the engine.

I INCUBATION · PILOT

# The portfolio sorts itself — then one charter, signed.

A scoring agent ranks the firm's flows on three business tests. Move the weights — the tiers re-rank live; the office arbitrates disputes, not spreadsheets. Work-in-progress caps lighthouses at two: a small machine ships, it does not queue.

⚙ agents sort & assemble ✋ existing owners sign → full Incubation instrument

Cost concentration 1.0× Conversion readiness 1.0× Control payoff 1.0×

| BUSINESS FLOW | COST | READY | CONTROL | SCORE | TIER |
| --- | --- | --- | --- | --- | --- |
| **Cross-border payment exceptions** | ●●● | ●●● | ●●● | 9.0 | Lighthouse — chartered |
| **KYC periodic refresh** | ●●● | ●●○ | ●●● | 8.0 | Lighthouse — next |
| Sanctions screening false positives | ●●● | ●●○ | ●●● | 8.0 | Pipeline |
| Reconciliations & breaks — Markets | ●●● | ●●○ | ●●○ | 7.0 | Pipeline |
| Trade & working-capital processing | ●●● | ●●○ | ●●○ | 7.0 | Pipeline |
| Regulatory-reporting lineage | ●●● | ●○○ | ●●● | 7.0 | Watch · data-dependent |

| BUSINESS FLOW | COST | READY | CONTROL | SCORE | TIER |
|---|---|---|---|---|---|
| Control testing & evidence | ●●○ | ●●○ | ●●● | 7.0 | Pipeline |
| Client onboarding & narratives | ●●○ | ●●● | ●●○ | 7.0 | Pipeline |
| Corporate-actions processing | ●●○ | ●●○ | ●●○ | 6.0 | Watch |
| Servicing & investigations — middle office | ●●○ | ●●● | ●○○ | 6.0 | Watch |
| Credit-memo drafting — Banking | ●●○ | ●●● | ●○○ | 6.0 | Watch |
| Collateral management | ●●○ | ●○○ | ●●○ | 5.0 | Watch |

## CHARTER № 2027-01 — Cross-border payment exceptions

Assembled by the charter agent · every action human-approved

### WHERE THE HANDLE TIME SITS SAMPLE

| | |
|---|---|
| Intake & classification | 15% |
| Context — entity, sanctions, history | 35% |
| Repair drafting & rationale | 30% |
| Approval & release | 12% |
| Client communication | 8% |

Agents target the red 65%. Baseline ~45 min, ~$28/case sample.

### THE BRIDGE — DATED, WITH A NUMBER TO BEAT

**Q1 27** Assisted triage live in one corridor, ~20% of volume; measure handle time, override rate, aging

**Q2 27** Scale to ~60%; target cost per case −40% sample; first freed-capacity log

**Q3 27** Redeploy-or-release decided, finance co-signed; banked as a share of the leverage review's payments row

Business owner: **Services — cross-border ops**

Co-sign the bridge: **Finance**

| | |
|---|---|
| Clear the risk: **Controls** | Delivery: **the agents** — no new hire |

T

TRANSFORMATION · SCALE

## The machine closes the quarter itself.

Collector agents pull the metrics, assemble the scorecard and draft the reconciliation; finance co-signs what banks. Q2 2026 actuals are the honest baseline — advance the clock and watch the close fill itself.

⚙ agents assemble & reconcile ✋ finance co-signs

| ALTITUDE | METRIC | LATEST ACTUALS · Q2 2026 | STATUS |
|---|---|---|---|
| **Impact** the Street | Efficiency ratio → the <55% bar | 57.4% — record-revenue quarter | on track · cyclical |
| | Operating leverage — productivity-led share | >9 pts total; split not yet measured | baseline set |
| | RoTCE bridge contribution | 13.0% firmwide; engine's line at zero | first entry Q3 |
| **Results** operations | Cost per cross-border exception | ~$28 sample · target −40% by Q4 | baseline set |
| | STP rate by flow | baselining across corridors | baseline set |
| | Freed capacity — redeployed vs. released | capacity exists; routing unlogged | first log Q3 |

| ALTITUDE | METRIC | LATEST ACTUALS · Q2 2026 | STATUS |
|---|---|---|---|
| **Resources** the portfolio | Investment drawn vs. ~$5B envelope | on plan per Investor Day pacing | on track |
| | Model-risk approval cycle time | SLA to be agreed | SLA set Q3 |
| | Severance conversion — backfill rate | >$800M YTD; backfill is the test | tracking Q3 |

**The quarter-close reconciliation — latest-actuals shell (Q2 2026)**

| | |
|---|---|
| Capacity claimed this cycle | **disclosed capacity ≈ 2,500 FTE-equivalents** |
| Routed — redeployed (operating leverage) | destination log starts Q3 |
| Routed — released (expense) | finance co-sign starts Q3 |
| Banked vs. bridge | first entries land with № 2027-01 |
| Claims retired — asserted, never routed | the honesty line — starts Q3 |

Deliberately a shell: a machine proposed mid-year cannot show a close it never ran. Advance the quarter to see the mechanics fill it.

✦ THE HUMAN RIM

## Run from an office of two.

This is the answer to thin staffing — not a workaround for it. The machine is the team; the two humans steer; and everyone else who governs the loop already works at the firm. Low headcount is not the constraint the model fights — it is the model's whole premise.

## 2

### The office — what only these two do

**Principal** — holds the mandate, chairs the weekly review, arbitrates what matters

**Analyst** — curates what the agents surface, preps the one pack the review runs on

That is the entire standing team. ~30 min sample a week is the steering cost.

## +0

### New hires — everyone else already signs here

**The businesses** own their flows and sign their charters — already their job

**Finance** co-signs what banks — already the returns-discipline role

**Controls & Model Risk** clear the risk tier — already the mandate they hold

### The machine is the rest of the team.

That is how a firm-scale transformation mandate runs — at the highest autonomy the control environment allows — without standing up a program office. And the sunset test holds: when the unit-cost curve stays flat for four quarters without the office touching it, even the two step back.

THE DELIVERY LOOP · OUTPUT: FORESIGHT — THE MARKET & CAPABILITY RADAR · 2027–2031 DELIVERY PROGRAM

# Where the opportunity is — and when it's ready to convert.

The first discipline of the loop: sense. Agents read the market, the regulators, the Street and the capability curve, and hand up a short list of dated signals — each one saying which business flow is ripe, and when. Foresight decides what to look at; Incubation decides what to do about it.

**7**

dated signals this cycle

**Q127**

the window the radar opens — governance, gate first

**~5%** sample

of the engine's effort; sensing is small by design — it protects the other 95%

THE CORE ARTIFACT

## The radar — agents draft it, one meeting approves it.

Not a trends essay: a signals log. Every entry must land on a business flow, a date and a decision — or it is commentary, and the agent discards it before a human ever reads it.

**WHAT THIS CYCLE IS SENSING FOR — Q1 2027 · PLAN**

Stand up the sensing cadence while the consent-order data-governance work finishes. The first radar names the lighthouse; capability is proven, but autonomy in regulated flows waits on the regulators.

- Weekly sensing cadence established
- Lighthouse named — cross-border payment exceptions
- Regulated-flow autonomy gated on consent-order completion (regulators' timing)

⚙ sensing agents — draft ✋ weekly review — approve

7 approved signals · Q1 2027 cycle

| Signal — current scan | So-what — the business decision | Flow · when |
| --- | --- | --- |
| The cross-border exception pool is the densest convertible cost in the wholesale spine | Charter it as the lighthouse — its early saves fund the next wave | Cross-border · **Q3 2027** convert next |
| AI now drafts exception repairs to a standard humans will approve — proven on historical cases sample | Move from study to charter; a human approves every action | Cross-border · **Q3 2027** convert next |
| Consent-order & control expectations are converging on the same client-data foundations the business needs anyway | Sequence data-dependent flows behind foundations already funded — complementary, not competing | Client data · **Q2 2027** align now |
| Autonomy inside regulated flows stays gated until the 2020 OCC / Federal Reserve data-governance work closes | Design every wave-1 step human-approved; do not model autonomous savings the regulators have not unlocked | Governance · **Q1 2027** gate |
| The ~$5B 2026–28 transformation envelope funds the first wave with no new budget ask | Draw lighthouse cost-to-achieve from the envelope, not a new line | Portfolio · **Q2 2027** fund |
| Severance is running ahead of structure; the backfill rate is the real test of whether released cost stays out | Track backfill as a first-class number from cycle one — the honesty check on released cost | Workforce · **Q2 2027** track |

| Signal — current scan | So-what — the business decision | Flow · when |
|---|---|---|
| Peers face productivity-disclosure questions from the Street, not tool counts — the question is coming to Citi | Have the productivity-led share of operating leverage measured before it is asked | Scorecard · **Q4 2027** prepare |

Every "ready now" call is bounded by a result already achieved elsewhere — the capability evidence, and the cost pressure the radar is aimed at, sit in the standing case, because neither changes when the quarter does.

WHAT FORESIGHT HANDS UP

# A readiness-ranked shortlist for Incubation.

Sensing ends here. The radar collapses to one prioritized list, which Incubation sizes, ranks on the selection triangle, and sequences into the portfolio.

| | | |
|---|---|---|
| 1 | **Cross-border payment exceptions — densest convertible cost, capability proven** | charter · Q3 27 |
| 2 | **Client data foundations — align behind the work the consent orders already require** | align · Q2 27 |
| 3 | **Governance pattern for human-approved steps in regulated flows** | agree · Q2 27 |
| 4 | **KYC periodic refresh — golden-source coverage approaching pilot threshold** | prepare · Q4 27 |
| 5 | **Backfill & released-cost tracking — the honesty check** | stand up · Q3 27 |

THE DELIVERY LOOP · OUTPUT: INCUBATION · 2027–2031 DELIVERY PROGRAM

# The opportunity portfolio — sized, ranked, sequenced.

The second discipline of the loop: pilot. Foresight hands up what is ripe; Incubation sizes each flow, ranks it on the selection triangle and sequences it into a portfolio of chartered bets — grounded in Citi's disclosed figures and real benchmarks.

**$246–461M**

Tier 1 — gross annual run-rate at maturity, on pools the CFO organization owns

**~0.4–0.8%**

of the ~$56.0B FY2026 expense base guided at Q4'25 earnings · one lever inside the ~$2–2.5B program, not additive to it

**65.6%**

of the change weight on these eight cards is signed inside the CFO organization

▼ Scroll to explore

## THE PORTFOLIO RIGHT NOW

# Tier 1 at Q1 2027.

The eight programmes the CFO organization signs for, decomposing into their share of the ~24 convertible flows on the curve. Where each sits below is read off the Q1 2027 position at **base** pace, autonomy set to **supervised** — not asserted. The firm-wide tier is a separate ledger further down the page.

| 6<br>IN THE FUNNEL | 2<br>CHARTER SHAPING |
|---|---|

**PROGRAMMES CHARTERED**
**0 of 8**
none yet — plan phase

**PRODUCTION FLOWS LIVE**
**0**
of ~24 convertible · firm-wide

**TIER 1 RUN-RATE BANKED**
**net-investing**
illustrative ° · no Citi schedule

**PORTFOLIO STAGE**
**plan**
portfolio v1 · toward $1.6–3.4B gross across both tiers

Charter the first lighthouse and sequence the pipeline behind the data foundations. Nothing scales yet — this is the plan, not the payoff.

- Portfolio v1 ranked on the selection triangle
- Charter № 2027-01 shaped — cross-border exceptions
- Pipeline sequenced behind data-governance readiness

*The ROI model below sizes each programme at **maturity** — the ceiling, reached only once it is in production. The run-rate above is Tier 1's share of what the firm-wide curve says is banked by Q1 2027. The gap between them is the delivery job.*

## TIER 1 — THE CFO ORGANIZATION OWNS THE OUTCOME

# Eight programmes the CFO signs for.

One test picks these eight: **when this process breaks, who does the regulator, the auditor or the street call?** If that person reports to the CFO, it is here. T&BE is the delivery partner on every card — but the accountable business owner is Finance. Each estimate is built bottom-up from a disclosed cost pool × an addressable share × a conservative savings rate, anchored to a documented benchmark. Click any card to expand.

**READ THE BAR AS A LIST OF SIGNATURES**

Every card below distributes 100 points of **change weight** — how much of the work of getting it into production sits in each leg. These are judgements calibrated against what the card actually has to do to go live, published so they can be argued with. The order is the argument: the first three legs cannot move without a decision from someone who reports to the CFO.

**THREE LEGS THE CFO ORGANIZATION SIGNS FOR**

- **People** — **The function head** owns role redesign, the span-of-control change and the severance line.
- **Process** — **Finance** owns the policy, threshold or approval change that makes the work smaller before anything is automated.
- **Controls** — **The control owner** re-performs and re-attests. The change is not done until the control is re-rated.

**TWO LEGS T&BE BUILDS AND RUNS**

- **Data** — **T&BE with the CDO** own lineage, authoritative sources and the data-quality rules.
- **Technology** — **T&BE** builds it, hosts it and runs it under model risk.

Across the eight cards **65.6%** of the change weight is signed inside the CFO organization, **17.5%** is shared with the CDO and **16.9%** is T&BE's build. Fold the five legs back into People / Process / Technology and the boundary does not move — **65.6 / 34.4** either way — but the Technology column then reads **34.4%** when the Technology leg is **16.9%**, and **35.6%** of the work (Controls and Data) loses its owner. That is why the bar has five segments and not three.

#1

REG REPORTING CHARTER SHAPING

## Regulatory report production & attestation

FINANCE · CONTROLLERS — REGULATORY REPORTING

**$75–140M**

est. annual run-rate · charter 2027–28

¤4 **Risk & remediation** protect then ¤1 Expense

15%

**60%** signed inside the CFO organization · **25%** shared with the CDO · **15%** is T&BE's build

+ details

#2

CONTROLLERSHIP CHARTER SHAPING

## Close, account substantiation & journal entries

FINANCE · CONTROLLERS

**$40–76M**

est. annual run-rate · charter 2027–28

¤1 **Expense** release then ¤2 Operating leverage

15%

**75%** signed inside the CFO organization · **10%** shared with the CDO · **15%** is T&BE's build

+ details

#3

FP&A IN THE FUNNEL

## Forecasting, management reporting & the productivity ledger

FINANCE · PLANNING & ANALYSIS

**$37–70M**

est. annual run-rate · charter 2027–28

¤2 **Operating leverage** redeploy then ¤1 Expense

20%

**65%** signed inside the CFO organization · **15%** shared with the CDO · **20%** is T&BE's build

+ details

#4

CAPITAL IN THE FUNNEL

## RWA data quality & capital reporting

FINANCE · CAPITAL PLANNING & MANAGEMENT

**$26–50M**

est. annual run-rate · charter 2028–29

¤4 **Risk & remediation** protect then ¤2 Operating leverage

15%

**55%** signed inside the CFO organization · **30%** shared with the CDO · **15%** is T&BE's build

+ details

**#5**

TAX · IN THE FUNNEL

### Tax provision, transfer pricing & information reporting

FINANCE · TAX

## $21–38M

est. annual run-rate · charter 2028

¤1 **Expense** release then ¤4 Risk & remediation

20%

**70%** signed inside the CFO organization · **10%** shared with the CDO · **20%** is T&BE's build

+ details

**#6**

PRODUCT CONTROL · IN THE FUNNEL

### Product control & independent price verification

FINANCE · VALUATION CONTROL

## $23–42M

est. annual run-rate · charter 2028

¤4 **Risk & remediation** protect then ¤1 Expense

15%

**70%** signed inside the CFO organization · **15%** shared with the CDO · **15%** is T&BE's build

+ details

**#7**

TREASURY · IN THE FUNNEL

### Liquidity reporting & ALM analytics

FINANCE · TREASURY

## $17–32M

est. annual run-rate · charter 2028–29

¤4 **Risk & remediation** protect then ¤2 Operating leverage

15%

**60%** signed inside the CFO organization · **25%** shared with the CDO · **15%** is T&BE's build

+ details

**#8**

EXTERNAL REPORTING · IN THE FUNNEL

### SEC reporting, disclosure tie-out & IR

FINANCE · EXTERNAL REPORTING & INVESTOR RELATIONS

## $7–13M

est. annual run-rate · charter 2029

¤4 **Risk & remediation** protect then ¤2 Operating leverage

20%

**70%** signed inside the CFO organization · **10%** shared with the CDO · **20%** is T&BE's build

> People 10% — unchanged, deliberately. You cannot shrink the population that certifies the 10-K; this is the one card where headcount is not the point.

+ details

Each card carries the documented result its estimate is sized against. The full evidence base, and the franchise cost split those pools are cut from, sit in **the standing case** — they are the same in every cycle, so they are stated once.

TIER 2 — THE FIRM-WIDE ARITHMETIC

# Eight more the CFO funds and scores, but does not run.

Apply the same test and these eight land elsewhere: when payments screening breaks the regulator calls Services; when the SDLC slips it calls T&BE. They are still the CFO organization's business, because the capital and the productivity claim reconcile on one scorecard — but here Finance has exactly two jobs, so those are the only two things stated below. **No composition bar appears on this tier, deliberately:** publishing a People / Process / Technology split on an engineering programme inside an engineering estate does not make it read as business transformation. It publishes, in a chart, that somebody else owns it.

| PROGRAMME · THE FUNCTION THAT ANSWERS FOR IT | COST POOL ALLOCATED | PRODUCTIVITY ATTRIBUTED | PAID IN |
|---|---|---|---|
| **AI-augmented SDLC & legacy modernization**<br>Technology · cross-LOB | **$2.00B**<br>at a 22% savings rate | **$300–600M**<br>gross at maturity · charter 2027–28 | **¤2 Operating leverage**<br>redeploy · then ¤1 Expense |
| **Payments operations & screening**<br>Services · TTS | **$2.75B**<br>at a 13% savings rate | **$250–450M**<br>gross at maturity · charter 2027–28 | **¤3 Revenue speed**<br>enable · then ¤1 Expense |
| **Client lifecycle & KYC/AML**<br>Cross-LOB | **$1.25B**<br>at a 26% savings rate | **$200–450M**<br>gross at maturity · charter 2027–28 | **¤3 Revenue speed**<br>enable · then ¤1 Expense |
| **Post-trade processing & reconciliation**<br>Markets | **$4.50B**<br>at a 7% savings rate | **$200–400M**<br>gross at maturity · charter 2028 | **¤1 Expense**<br>release · then ¤4 Risk & remediation |

| | | | |
|---|---|---|---|
| **Custody & asset-servicing ops**<br>Services · Sec. Services | **$1.75B**<br>at a 11% savings rate | **$150–250M**<br>gross at maturity · charter 2028 | **¤1 Expense**<br>release · then ¤2 Operating leverage |
| **Controls, reg reporting & transformation**<br>Controls · cross-LOB | **$1.80B**<br>at a 15% savings rate | **$150–400M**<br>gross at maturity · charter 2028–29 | **¤4 Risk & remediation**<br>protect · then ¤1 Expense |
| Control evidence and data lineage for regulatory reporting and capital are Tier 1 №1 and №4. $101–190M of this band is counted there, not here. | | | |
| **IB knowledge work**<br>Banking | **$1.60B**<br>at a 12% savings rate | **$100–300M**<br>gross at maturity · charter 2028 | **¤3 Revenue speed**<br>enable · then ¤2 Operating leverage |
| **Client servicing & surveillance copilots**<br>Middle office | **$1.50B**<br>at a 12% savings rate | **$100–250M**<br>gross at maturity · charter 2029 | **¤2 Operating leverage**<br>redeploy · then ¤3 Revenue speed |
| **Tier 2 gross, across a $17.15B pool — allocated and attributed by the CFO organization, delivered by eight other functions.** | **$1.45–3.10B**<br>gross at maturity · midpoint $2.27B | **$1.6–3.4B**<br>both tiers combined, net of the $101–190M counted once in Tier 1 | |

Read the two tiers across, never down: they are separate landscapes, and one card in each perimeter overlaps the other. The boundary is written into the row it affects, and the combined figure below it is stated net of that overlap — counted once, in the tier that signs for it.

LIVE ROI MODEL

# Pressure-test every assumption.

Pick an opportunity and adjust the addressable cost pool and AI savings rate. The financial impact and the portfolio total recompute live. Defaults reflect the conservative ranges in the analysis.

Opportunity

#1 · Regulatory report production & attestation

Addressable cost pool: **$0.52B**

$0.05B $1.0B

AI savings rate: **20.0%**

0% 40%

This opportunity → **$104M / yr** / yr run-rate

Benchmark: Sized against the steady-state cost of producing these reports — the people who will still be closing FR Y-9C in 2030 — and deliberately not against transformation expense. That distinction is the whole card. Citi's 2026 Investor Day says transformation expense peaked at $3.3B in 2025, that about 50% comes out in the near term "with most of this reduction showing up in Corporate/Other", and that what remains in the businesses "will be reduced by further productivity saves over time from additional investments in Technology and AI". Citi has already promised investors this saving once. Anything here that lands inside the roll-off is not incremental — it accelerates a commitment already made — which is why the card is drawn outside it.

Tier 1 potential — capacity + expense (all 8 CFO-owned cards, at maturity)

**$345M**

≈ **0.62%** of the ~$56.0B FY2026 guided expense base

Gross potential at maturity, not a plan. It reaches the P&L only through a redeploy-or-release decision, and ~78% of gross survives to net by year three. The firm-wide tier below is a separate landscape — read across the boundary, never add the two. AI is one lever inside Citi's ~$2–2.5B program.

PRIORITIZATION

# Impact vs. feasibility.

Position is a standing judgement — impact and feasibility are properties of the flow, not of the calendar — so nothing here moves. Size is run-rate impact *at maturity*. Only the **colour** moves: each family is shaded by the rung it is standing on at the selected cycle. Hover any bubble for detail.

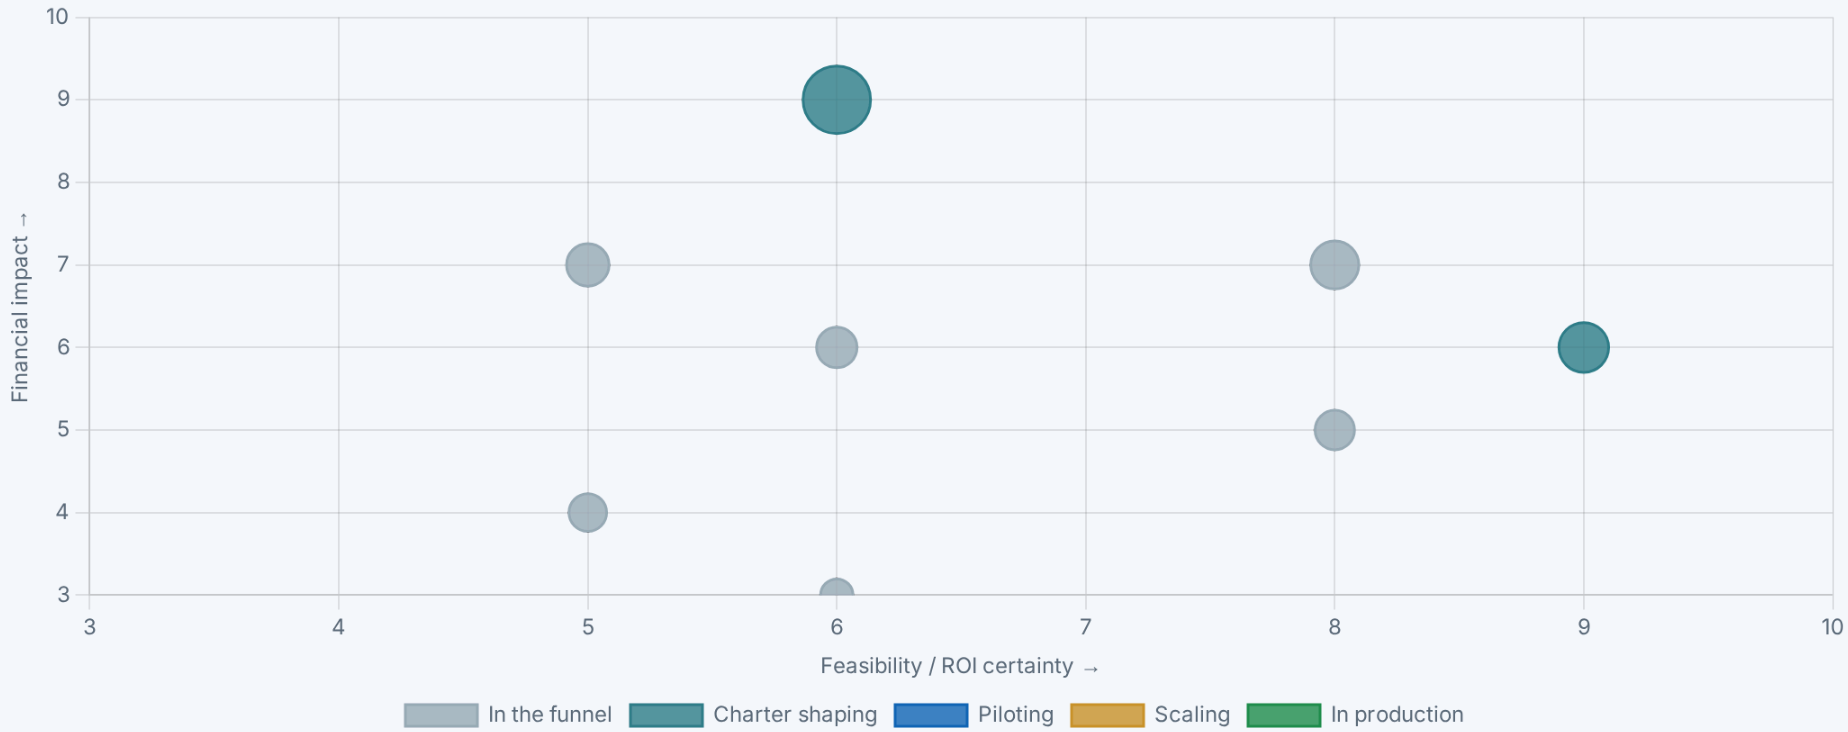


**Q1 2027** — 6 in the funnel · 2 charter shaping. The plot is the judgement; the colour is the delivery state.

What this portfolio is challenged on — and the end-to-end flow one opportunity runs through, step by step — do not change with the cycle. Both are in **the standing case**.

citi

THE DELIVERY LOOP · OUTPUT: TRANSFORMATION — THE OPERATING-LEVERAGE REVIEW · 2027–2031 DELIVERY PROGRAM

# From AI capacity to operating leverage.

The capacity map shows where agentic AI creates capacity. This review — the loop's banking instrument — asks the CFO-organization question: where does it bank — on which line, signed by whom, visible in which ratio?

**>9 pts**
positive operating leverage in Q2 2026 — mostly revenue-led

**<55%**
Investor Day medium-term efficiency-ratio bar

**~$1B**
value of each full efficiency-ratio point at today's run-rate

THE CLOSE · RENDERED FROM THE S-CURVE

## The Q1 2027 close.

Every figure below is the S-curve read at Q1 2027 under the selected scenario; the reconciliation is the plan phase. Move the strip above and the whole scorecard re-renders.

| EFFICIENCY RATIO | ROTCE | AI RUN-RATE | PRODUCTION FLOWS LIVE | OPERATING LEVERAGE | RUN FROM |
|---|---|---|---|---|---|

| | | | | | |
|---|---|---|---|---|---|
| **58.5%↓** toward the <55% bar (2029–31) | **11.2%↑** 11–13% (2027–28) | **net-investing** illustrative · no Citi schedule | **0** of ~24 convertible | **net-investing (J-curve)** productivity-led share | **a desk of ~6** supervised autonomy |

**The quarter-close reconciliation — Q1 2027**

Baselines set; scorecard stood up. AI is net-investing — the J-curve trough — so banked savings are near zero while the ~$5B 2026–28 envelope funds build-out.

- Scorecard live across three altitudes
- Baselines set — cost per flow, STP, aging
- Net-investing (J-curve): savings back-loaded by design

Why this is a business-transformation problem rather than a technology one — the three facts that set the frame, the four currencies value can be paid in, and the five levers it lands on — is argued once, in **the standing case**. None of it changes when the quarter does.

LIVE MODEL · THE RATIO PATH

# What does <55% actually require?

Set annual revenue and expense growth; the blue path recomputes off the Q2 2026 run-rate. The gold line is what the selected delivery scenario produces on the same axes — growth arithmetic is what the bar *requires*, the arc is what the program *delivers*. The arc has no 2026 point on purpose: the plan opens in Q1 2027, and nothing is assumed to ship this year.

Revenue growth: **+5.0% / yr**

−2% +10%

Expense growth: **+3.0% / yr**

−2% +8%

**Growth arithmetic crosses <55% in 2029 · ending 52.1% in 2031**

**Delivery arc on this scenario crosses in 2029 — the two lenses agree.**

Base: ~$99B revenue run-rate

~$57B expense run-rate

57.4% Q2 2026 efficiency

Illustrative arithmetic on annualized Q2 2026 figures — not a forecast. Mix, notables and rate effects will move any real path.

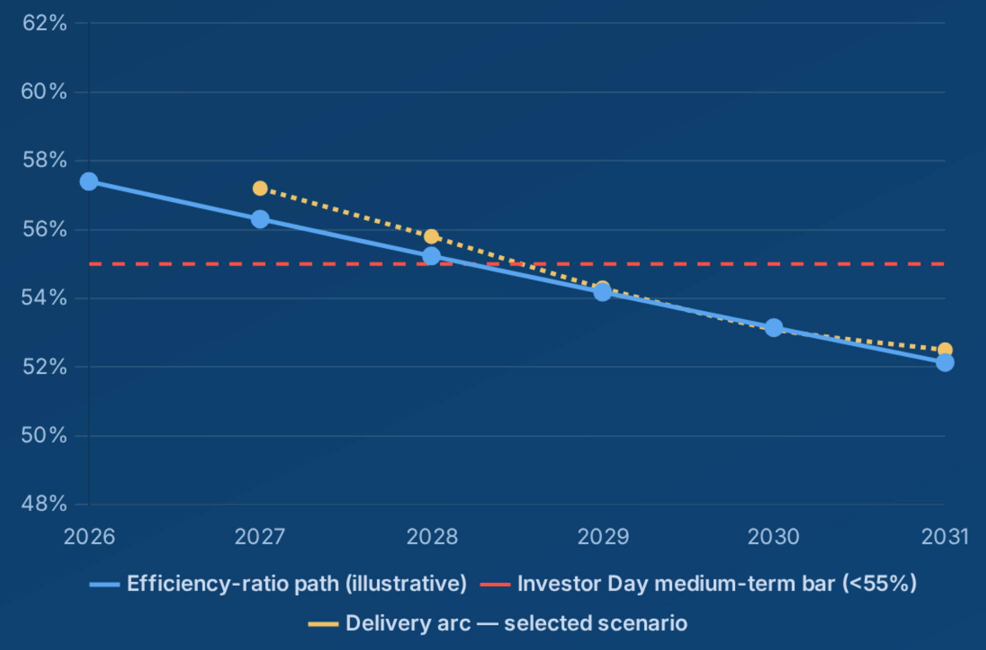


LIVE MODEL · CAPACITY → P&L

# The bridge is a decision, not an outcome.

Take capacity created by AI — hours per week — and split it: released to the expense line, or redeployed into growth work. Both are value; only one is "savings." The split is the management decision this whole map turns on.

**Which line the released side lands on — and who owns it — is set by the five levers.**

**Capacity created: 100K hours / week**

10K 600K

**Share released to the expense line: 0%**

0% · all redeployed 100% · all released

**Average loaded cost per role: $130K**

$80K $200K

Gross and illustrative. Hours ÷ 40 ≈ full-time-equivalents; conversion to P&L assumes released roles are not backfilled — the backfill rate is the honesty check.

Opens seeded off the selected cycle: capacity scales from Citi's disclosed ~100K developer-hours/week toward ~600K (≈15,000 FTE-equivalents, ~7% of a 219K-person firm) as the ~24 convertible flows reach production. The split is then solved to match the run-rate the rest of this page carries. Move either — they are assumptions, not disclosures.

**2,500** FTE-equivalents of capacity — where the slider sends them:

Released → expense reduction (P&L): **$0M** / yr

Redeployed → operating leverage (capacity): **$325M** / yr

> At **Q1 2027** the curve carries **no banked AI run-rate** — the J-curve trough. Capacity already exists; none of it has been released. That is the whole argument for the bridge: the capacity is real long before the saving is.

Scorecard metric this feeds: **hours redeployed vs. released** — the first-class measure that keeps capacity claims honest. For scale: Citi's disclosed ~100K weekly developer-hours ≈ 2,500 FTE-equivalents of capacity — value that banks only when a decision routes it somewhere.

ONE SCORECARD, THREE ALTITUDES

# The artifact that holds it together.

One scorecard read at three altitudes — serving the earnings narrative, the operating review, and the investment committee. The weekly 100-process review is the natural home of the middle altitude; quarterly, the scorecard reconciles capacity claimed with P&L booked.

### Impact — the Street's view

*Is the ratio path credible?*

- Efficiency-ratio trajectory toward <55%
- Operating-leverage points — split revenue-led vs. productivity-led
- RoTCE bridge contribution

This cycle: efficiency **58.5%↓** · RoTCE **11.2%↑**

### Results — the operating view

*Are the flows structurally cheaper and faster?*

- Unit cost-to-serve by flow · STP rate
- Onboarding time-to-revenue
- Hours redeployed vs. released
- Validated-milestone throughput

This cycle: **0** flows live · run from **a desk of ~6**

### Resources — the portfolio view

*Is the envelope earning its keep?*

- Investment drawn vs. the ~$5B envelope
- Adoption-in-workflow — beyond nine-in-ten access
- Model-risk approval cycle time
- Severance-conversion ratio

This cycle: AI run-rate **net-investing** · **= disclosed glidepath** · sub-55% ≈ Q3 2029

citi

**Business Transformation & Operating-Leverage Map — the Operating-Leverage Review · an output of the delivery loop**

**Illustrative concept — independent analysis based on Citi's public disclosures. Not official Citi material; figures are directional, used to frame discussion, and are not Citi projections.**

Sources: Citi Q2 2026 earnings release & presentation (July 14, 2026) · reported Q2 2026 earnings-call remarks · Citi Investor Day 2026 materials (May 7, 2026) · Citi press releases (CFO transition, senior appointments) · Citi Stylus Workspaces (agentic AI) press coverage.