GSTRIDE AI

WORKFORCE CAPACITY PLANNING SPREADSHEET · 2026 EDITION

Workforce Capacity Planning Spreadsheet.

A four-variable capacity model plus quarterly cadence template (Excel + Google Sheets + printable PDF). Built for COOs, CFOs, and Heads of Ops at 200-500 seat IT services, BPO, and offshore-dev shops. India lane primary, EU and US lanes secondary.

BUILT FOR
COO · CFO · Head of Ops

RANGE
200 – 500 seats

EDITION
v1.0 · May 2026

Disclaimer. This spreadsheet is an operational planning aid, not legal or financial advice. The Digital Personal Data Protection Act 2023 is in force with DPDP Rules expected late 2025 / 2026; the EU AI Act enters phased application through 2026 with implementing acts still to land; US state laws (CPRA, NY S2628, CT 31-48d, DE 19 Del.C. § 705, NYC Local Law 144) continue to evolve — verify the latest position with counsel before any compliance assertion. Penalty exposure under the EU AI Act can run up to EUR 35 million or 7 percent of worldwide annual turnover for prohibited-AI infringements, subject to revision; DPDP penalties run up to INR 250 crore per category in the headline tier, also subject to revision under final Rules — verify with counsel before quoting. Margin and leakage ranges quoted are typical for 200-500 seat IT services shops in our review set and should be validated against your own data before use in board materials. gStride is not certified by any data protection authority, the Data Protection Board, the Ministry of Electronics and Information Technology, or the European AI Office.

How to use this spreadsheet

This is a **four-variable capacity planning spreadsheet** built for a COO, CFO, or Head of Operations running a 200-500 seat IT services, BPO, or offshore-dev shop. India lane primary, EU and US lanes secondary. It replaces the spreadsheet most shops keep in their head — utilisation as a single number, attrition as a buffer line, leakage as a yearly write-off discussion.

The model rests on four explicit variables that move quarterly and need named owners. The template carries a quarterly cadence table, a 12-month grid, a variance-check checklist, a roll-up dashboard, a DPDP and EU AI Act overlay for when capacity data feeds AI inference, and a worked example on an anonymised 350-seat shop.

> **How the pages fit together.** Section 1 explains each of the four variables. Section 2 is the quarterly cadence (5 phases). Section 3 is the 12-month grid for fill-in. Section 4 is the 8-item variance-check checklist. Section 5 is the 6-KPI roll-up dashboard. Section 6 is the DPDP + AI Act overlay for India-lane shops feeding AI inference. Section 7 is the worked example.

> **Recommended customisation workflow.** One pass with the COO and CFO to set the Q-start model. One pass with the Head of Delivery to break down account-level targets. One pass with the HR Head to align the attrition replacement pipeline. Run weekly variance with the Delivery PMO. Re-baseline at mid-Q if cumulative variance exceeds 5% on any variable. Retrospective at Q-end.

SECTION 1

The four capacity variables

Intent. *Define what each variable measures, the healthy target, and the leakage zone that flags an upstream issue.*

1 Billable hours captured

HRS / IC / WEEK

How it is measured. Hours that actually land on the timesheet system honestly — not assumed utilisation. Measures the gap between hours worked and hours captured. The first input to recoverable capacity.

Healthy target. 32 – 36 hrs / IC / week (200-500 seat range)

Leakage zone. When captured hours drop below 30 / IC / week without a holiday or leave reason, investigate timesheet discipline before assuming an attrition or bench problem.

2 Recovered leakage

% OF REVENUE

How it is measured. The share of the gap between hours worked and hours billed that is recoverable in the next quarter. Recovery actions: scope-creep re-billing, write-off discipline, rate-card refresh on rolling contracts, account-level under-billed conversation.

Healthy target. 4 – 8% of revenue typical recoverable

Leakage zone. Leakage above 8% is usually structural (rate-card outdated, scope drift unmanaged) not behavioural — flag to the CFO and revisit the contract architecture rather than the timesheet discipline.

3 Bench-to-bill conversion lead time

DAYS MEAN + LONG TAIL

How it is measured. Days from a bench resource being available to ramping to full billable. Track mean and the long tail (90-day legacy-stack bench) separately. Segment by skill (frontend / backend / data / ML / cloud / mobile) and account type.

Healthy target. 30-day mean, 90-day long tail

Leakage zone. If the long tail exceeds 90 days for more than 12% of bench, you have a re-skilling problem masquerading as a sales pipeline problem.

4 Attrition replacement lead time

DAYS END-TO-END

How it is measured. Notice → hiring → onboarding → ramp to full billable. The number runs longer than most CFOs assume — typical end-to-end for mid-senior engineering in India IT services is 90 – 120 days when measured honestly.

Healthy target. 90 – 120 days mid-senior engineering

Leakage zone. If the attrition replacement pipeline is treated as a buffer line rather than an explicit input, every quarterly plan starts with a 5 – 8% optimism error baked in.

SECTION 2

Quarterly cadence template

Intent. *Quarterly cycle sets the model; weekly variance catches drift early; mid-Q re-baseline saves the quarter when drift exceeds 5%.*

PHASE	ACTIVITY	OWNER
Q-start, week 1-2	Set the model — billable target, leakage recovery target, bench conversion plan, attrition pipeline	COO + CFO + Head of Delivery
Q-start, week 3-4	Account-level breakdown — which accounts carry which targets, named recovery actions per account	Account leads + Delivery managers
Weekly through Q	Variance dashboard — actuals vs model, drift flags, recovery action status, blocker log review	Delivery PMO
Mid-Q checkpoint	Re-baseline if cumulative variance > 5% on any of the four variables — re-cast the second half	COO + CFO
Q-end	Variance retrospective — model assumption review, next-Q input refresh, capacity learnings logged	COO + CFO + HR Head

> The discipline that matters is the **mid-Q checkpoint**. Shops that defer re-baselining to the next quarter end up planning the next quarter on broken assumptions. The 5% trigger is approximate — for some shops 3% is the right line, for some 7% — but the principle of an explicit re-baseline trigger is universal.

SECTION 3

12-month variable grid

Intent. *Fill row-by-row each month with the four variables and a variance flag. Print or transcribe into your Excel / Google Sheets file.*

MONTH	V1 BILLABLE HRS	V2 RECOVERED LEAKAGE	V3 BENCH MEAN (DAYS)	V4 ATTRITION (DAYS)	VARIANCE FLAG
Jan	_____	_____	_____	_____	_____
Feb	_____	_____	_____	_____	_____
Mar	_____	_____	_____	_____	_____
Apr	_____	_____	_____	_____	_____
May	_____	_____	_____	_____	_____
Jun	_____	_____	_____	_____	_____
Jul	_____	_____	_____	_____	_____
Aug	_____	_____	_____	_____	_____
Sep	_____	_____	_____	_____	_____
Oct	_____	_____	_____	_____	_____
Nov	_____	_____	_____	_____	_____
Dec	_____	_____	_____	_____	_____

> Variance flag column convention: **0** = within 5% of model, **1** = 5 – 10% drift (watchlist), **2** = > 10% drift (re-baseline trigger). Track over rolling-3-month windows to filter single-month noise from structural drift.

SECTION 4

Variance-check checklist

Intent. *8 fail-fast questions every PMO and CFO should run at the stated cadence. Each maps to a named owner.*

#	CHECK	OWNER	CADENCE	DONE
C1	Are weekly captured hours within 5% of the quarterly model assumption?	Delivery PMO	Weekly	☐
C2	Is recovered leakage tracking to the Q-start recovery target by week 6?	CFO + Account leads	Mid-Q	☐
C3	Has the bench long-tail (>90 days) grown by more than 2 percentage points?	Head of Delivery	Monthly	☐
C4	Is the attrition replacement pipeline carrying the assumed mid-senior backfill?	HR Head + COO	Monthly	☐
C5	Are any of the four variables off by more than 5% cumulative? (Re-baseline trigger)	COO + CFO	Mid-Q	☐
C6	Are timesheet capture rates honest? (Sample-audit one account per quarter)	Delivery PMO	Quarterly	☐
C7	Is the AI inference layer (if any) audit-logged for capacity-related decisions?	DPO / Data Privacy lead	Quarterly	☐
C8	Has the capacity model been stress-tested against a 10% adverse attrition shock?	CFO	Half-yearly	☐

SECTION 5

Roll-up dashboard — 6 KPIs

Intent. *A single sheet the COO and CFO read together. Six derived KPIs that compress the four variables into a board-readable view.*

KPI	HOW TO COMPUTE	HEALTHY BAND	THIS QUARTER
Realisation %	(Billed hrs / Capturable hrs) × 100	≥ 85%	_____
Effective utilisation %	Realisation × Utilisation	≥ 64%	_____
Recoverable leakage %	(Recoverable gap / Revenue) × 100	4 – 8%	_____
Bench mean conversion (days)	Mean(bench → billable days), all skills	≤ 30 days	_____
Bench long-tail share %	(Bench > 90 days / Total bench) × 100	≤ 12%	_____
Attrition replacement (days)	Mean(notice → full billable days), mid-senior engineering	≤ 120 days	_____

> The two KPIs that change board conversations are **recoverable leakage %** (because it usually exceeds operating margin) and **bench long-tail share %** (because the long tail compounds quarter-on-quarter when ignored).

SECTION 6

DPDP and EU AI Act overlay

Intent. *When capacity-planning data feeds an AI inference layer (allocation, attrition risk, utilisation flagging), specific obligations attach. Build these in from the start rather than retrofit.*

India lane — DPDP Act 2023

Section 4 lawful basis. Most capacity-planning processing falls under recognised legitimate use in the employment context. AI-driven evaluation may need explicit consent depending on how implementing Rules land — subject to revision under final DPDP Rules.

Section 8 DPIA. Trigger DPIA when capacity data feeds AI inference making significant evaluation, allocation, or retention decisions. Examples: AI-driven bench allocation, AI-driven attrition risk scoring, AI-driven utilisation flagging.

Section 10 SDF obligations. If designated Significant Data Fiduciary (typically by volume or sensitivity), periodic-audit and DPO obligations attach — build the audit trail into the capacity-planning system from day one.

EU lane — EU AI Act

High-risk classification. AI systems used in employment for performance evaluation, task allocation, and termination are classified high-risk under Annex III — subject to revision under final implementing acts.

Required controls. Risk management, data governance, transparency, human oversight, accuracy, and post-market monitoring. The annual re-baseline cycle of this template can be aligned with the post-market monitoring obligation where appropriately scoped.

Penalty exposure. Up to EUR 35 million or 7 percent of worldwide annual turnover for prohibited-AI infringements — subject to revision. Verify the current position with counsel before any compliance assertion.

US lane — state laws

NYC Local Law 144. Bias audits required annually for Automated Employment Decision Tools used in hiring or promotion. Where capacity-driven AI feeds those decisions, the bias-audit obligation applies.

State notice statutes. CT 31-48d, DE 19 Del.C. § 705, NY S2628 require written notice to employees before electronic monitoring begins. Capacity signals sourced from screen capture or keystroke logging carry the strictest notice obligation — gStride's position is to not collect those signals at all.

CPRA + similar state laws. Notice-at-collection, purpose limitation, retention disclosure — apply the most protective standard where multiple state laws apply to one workforce.

This overlay describes the architecture and operating choices; it does not constitute legal advice. Verify the latest position with counsel before any compliance assertion — penalty schedules, implementing acts, and state legislation continue to evolve through 2026.

SECTION 7

Worked example — 350-seat shop

Intent. *Anchor the four-variable model against a representative baseline. Numbers below are illustrative typical patterns, not a customer reference.*

Illustrative 350-seat India IT services shop with mixed onshore + offshore engagement model. Used here as a worked example only — not a real customer reference.

VARIABLE	BASELINE	Q1 TARGET	NAMED ACTIONS
V1 Billable hrs captured	31.4 hrs / IC / week	34 hrs / IC / week	Move 4 accounts to fortnightly timesheet review; close-out scope-drift on 2 long-running contracts.
V2 Recovered leakage	6.8% of revenue	Recover 2.5 of 6.8	Re-paper 3 rate-card-stale contracts; write-off discipline on 2 accounts where TL hours stay unbilled.
V3 Bench-to-bill mean	42 days mean	34 days mean	Pre-stage skill-up plan for backend Java pool; align with named Q2 pipeline.
V3 Bench long-tail	17% > 90 days	12% > 90 days	Decision on legacy-stack bench: re-skill, redeploy, or planned attrition.
V4 Attrition replacement	128 days mid-senior	110 days mid-senior	Pre-screen pool, sign-on-notice incentives for 60-day notice acceptance, named onboarding owner.

Outcome pattern. On this profile, a disciplined four-variable re-baseline at Q-start delivers a recoverable margin lift in the 2.5 – 4.0 percentage-point range over the quarter. The number is illustrative — the model is the discipline, not the percentage.

Want a second pair of eyes on your Q-start model?

If you are about to set the next quarter's capacity model and want an independent read on the four-variable baseline, the recovery target architecture, or the AI Act / DPDP overlay before adoption, book a 30-minute readiness audit with the gStride team. We will walk through your inputs, flag the assumptions most likely to break, and share the deployer-side dashboards we ship to our own customers in the 200-500 seat range.

Book a 30-min readiness audit · cal.com/gstrideai/30min

Interactive online previewer at **gstride.ai/assets/audits/workforce-capacity-planning-spreadsheet.html** — input your headcount, billable %, attrition %, and ramp time to see capacity gap and bench-to-bill alerts per quarter.