albert

The SWP Big Beautiful Benchmark

2026

Governance | Data | Drivers

In partnership with

WPI THE WORKFORCE PLANNING INSTITUTE

EMPOWERING TOMORROW'S WORKFORCE, TODAY

Foreword

There is a famous French saying: *quand je me regarde je me désole, quand je me compare je me console*. Which basically translates to “I suck but the others are worse” in a rhyming form. This is why I have always struggled with how useful benchmarks can be. It tends to drive practices towards a “good enough” state which is far from the excellence we strive for.

I feel different about this one. The reason is because 2025 really seemed like a pivotal moment for our line of work: the uncertainty and volatility of, well, everything, drove the need for better and ongoing planning. And despite having been around for ages, SWP as a process is still in a rather immature state in most organizations. That is why we collabed with Nick Kennedy and the Workforce Planning Institute on this one.

Nick already runs the WPI State of the Profession: Strategic Workforce Planning, a study that covers the discipline at the institutional level. This survey is deliberately narrower and more operational. The methodology mirrors the chapters of our SWP Cookbook (2025), and for each step of the process we look at what 100+ practitioners actually do, what stands out in the answers, and what good looks like.

Nick and I have shared a stage at every SWP Conference London for the past few years, and I have been a recurring guest on his podcast. We have two passions in common: triathlon and SWP, which says a lot about our masochist tendencies. But we are also committed to bringing you whatever is useful to push the topic forward.

Enjoy the read!

Vincent Barat - CEO - Albert

The Workforce Planning Institute's State of the Profession: Strategic Workforce Planning Report tells you how the world sees SWP.
This tells you what SWP practitioners actually do.
Those are two different things. And the distance between them is, in many ways, what this study is about.

When Vincent and I started talking about this, we kept coming back to the same question: we know the profession is growing. We know investment is rising. We know organisations are recognising SWP faster than almost any other discipline. What we didn't have, with any precision, is a picture of what practitioners are actually doing when they sit down at their desk.
This benchmark answers that. Not what the frameworks say. What happens in practice.

The numbers are not always comfortable. An average maturity of 2.43 out of 5. Seventy percent still running on spreadsheets. One in seven with SWP fully wired into budget and headcount planning. A third of skills framework owners who don't know how many skills they have.
But discomfort is useful. And what this study also shows, through the top quartile data, through the practitioners doing this well, is that the ceiling is visible. The gap between where most organisations sit and where the good practice examples land is not mysterious. It is measurable. And it is closeable.

That is the point. Not to benchmark against average. To know, with evidence, what good actually looks like.

That is why we built this.

Nick Kennedy - Chief Executive Workforce Planning Institute

Table of contents

About the panel

109 organizations responded to the survey. 68 completed it in full. 41 stopped along the way, and we kept their partial answers wherever they did reach the question. So the answered-N varies by question, and we report it on every chart.

The response volume was beyond what we expected. There is a great community at WPI, and Albert has a community of its own, and the overlap of those two communities is exactly the audience this benchmark was built for. The volume itself is a signal that practitioners want this conversation.

Industry mix

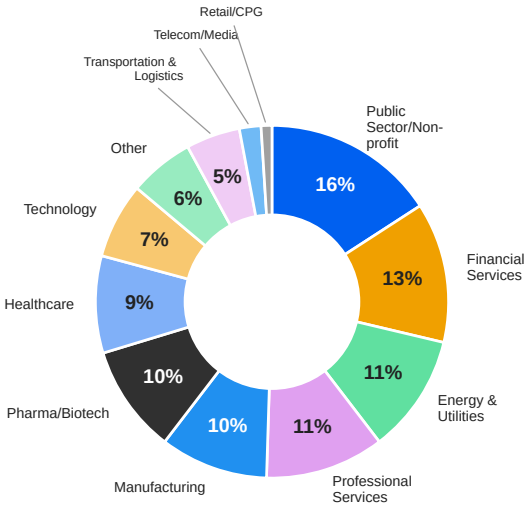

The industry mix is broad. Public Sector and Non-profit is the largest single segment at 16%, followed by Financial Services (13%) and a tightly packed group of Energy & Utilities, Professional Services, Pharma/Biotech and Manufacturing each at 9 to 11%. Technology is slightly under-represented at 7%. Retail and Telecom register barely above the noise floor.

Region mix

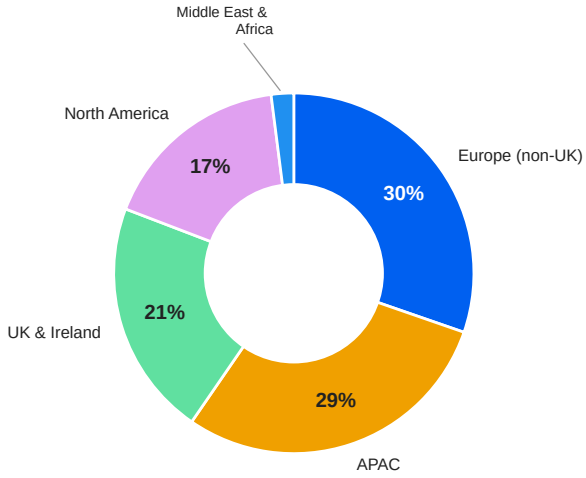

The geographic balance is healthy across the four major regions we hoped to reach. Europe (non-UK) at 30%, APAC at 29%, the UK and Ireland at 21%, North America at 17%. Middle East and Africa at 2%. LATAM did not register a respondent in this edition.

Headcount mix (FTE)

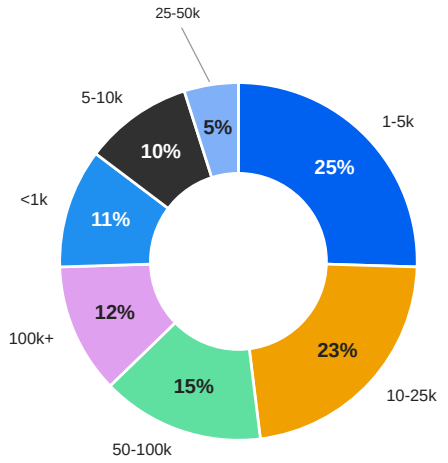

The size distribution covers the full range. The largest single cohort is the 1,000 to 5,000 segment at 25%, followed by the 10,000 to 25,000 segment at 23%. 50,000+ organizations represent 27% of the panel combined, which is a healthier tail than most workforce surveys ever reach.

Respondent Role

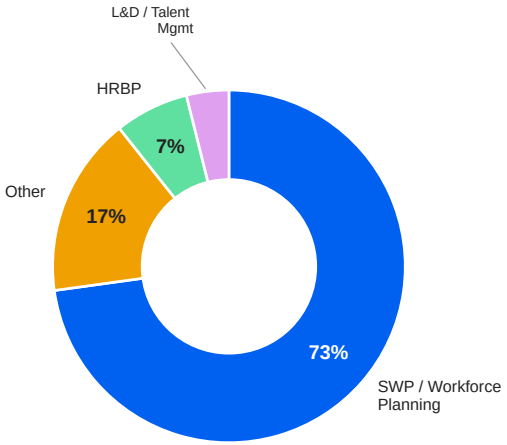

Almost three-quarters of respondents identify their role as SWP or Workforce Planning. The rest are HRBPs, L&D, Talent Management, and other HR roles. This is a practitioner panel.

The bias we have to name

The panel is drawn from the WPI and Albert communities. Anyone reaching this survey through those channels is already aware that SWP exists as a discipline, already curious enough to read about it, and almost certainly more advanced than the median large organization in the broader market. The 8% of the panel who answered "we do not do SWP" is probably the lower bound of that figure outside our communities.

In practice, this means two things. The numbers in this study describe what an engaged practitioner audience looks like, which is exactly the audience the study is for. And those numbers are not representative of the entire universe of large organizations. Read the percentiles with that bias in mind.

The top-line snapshot

Six numbers frame the whole study.

Self-assessed SWP maturity (0-5)

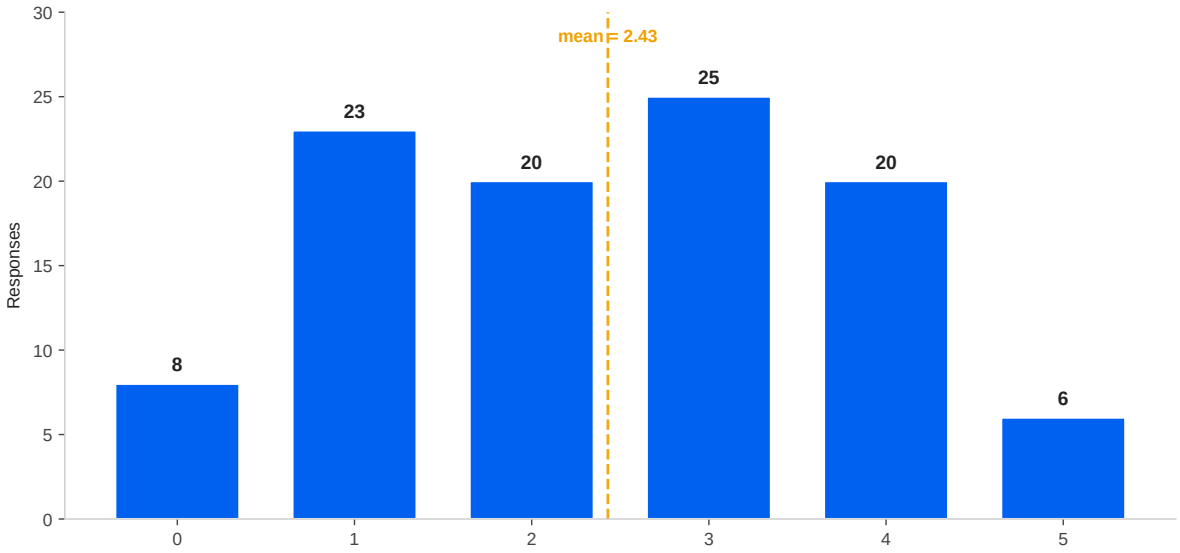

1. Average maturity is 2.43 out of 5. The shape is a bell with a slight left lean and the mode at 3. Most practitioners are doing parts of the methodology well and no part of it consistently. Partial industrialization is the modal experience.

2. Adoption is broad. Consistency is what varies. More than half the panel uses driver-based forecasting. 64% have a skills framework in some form. 82% model attrition. The methodology has arrived as a standard. What separates the top quartile from the rest is whether the standard is applied every cycle, the same way, with the same tooling.

3. 70% of practitioners still run SWP on spreadsheets. Dedicated SWP solutions are the primary tool for 18%. The rest is HRIS modules, BI platforms and custom builds. The tooling layer is the single most striking lag in the data.

4. Half the panel models one scenario. 51% report a single-base forecast. Scenario thinking is broadly underdeveloped, despite being the cheapest place to add value once the demographic baseline is in place.

5. Only 15% have wired SWP fully into budget and headcount planning. Another 51% report partial integration. A third sits in fully separate processes. This is the gap with the most direct line to executive attention, and whether this is intentional remains an open question

6. Buy and Build dominate the action plan. Bridge does not. 84% of the panel deploys Buy. 78% deploys Build. Internal mobility, Bridge, sits at 21%, well behind every other lever in the 8B framework. Closing gaps through internal movement is the most under-used route in the panel.

Governance & Process

SWP is a business process even though HR coordinates it. Of course, Finance is a critical partner and BU leaders provide the operating reality. And People Analytics keeps the data honest. None of those statements are controversial in principle, but the interesting question is what shape they take inside a real organization, and where the friction actually sits.

Who coordinates the process

Who formally owns SWP

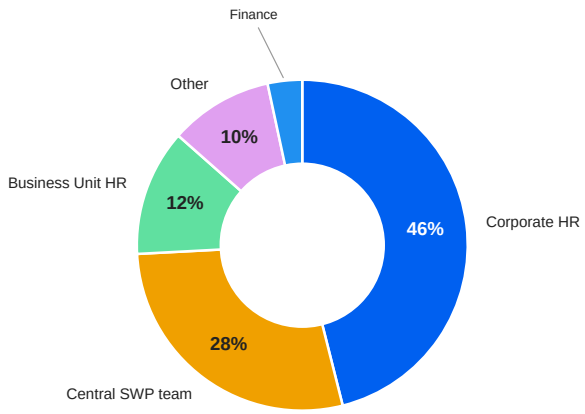

Coordination of SWP sits in HR in 86% of cases. Corporate HR for 46%, a Central SWP team for 28%, BU HR for 12%. Finance coordinates the process for 3% of respondents. The Finance-coordinated model is a curiosity in the data, and the wider practitioner experience supports the same conclusion: SWP run from Finance drifts toward budgetary consolidation, with the workforce question reduced to a cost line.

A point worth being assertive about, because it comes up in every operating-model discussion. Finance is of course a critical partner in SWP:

- They validate the figures.
- They own the hand-off to budget.
- They are an explicit persona that any well-designed SWP operating model has to address with respect and with clarity.

But none of that means they call the shots on the exercise itself. SWP is a workforce question, coordinated from HR, in service of the business. Finance is in the room, and they should be. The chair belongs somewhere else.

The split between "Corporate HR" and "Central SWP team" is the more interesting structural question inside the HR-coordinated majority. A Central SWP team is a structural commitment. It usually means a dedicated calendar, a dedicated set of tooling, and an operating model that survives the departure of the SWP champion (we have seen that situation happen very often...). Organizations large enough to justify a dedicated team should consider it. The 28% in this panel who report one are, in our experience, disproportionately concentrated in the more industrialized end of the maturity scale.

Who is in the room

Stakeholders actively involved

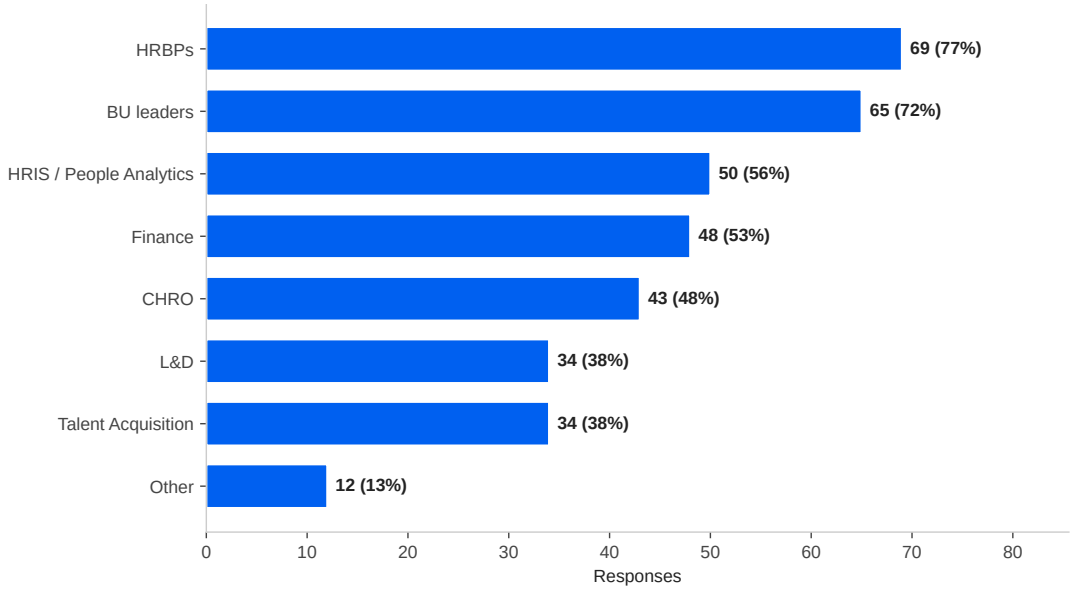

Coordination is one signal. Active participation is another, and the lattice here is wide. HRBPs are involved in 77% of cycles. BU leaders in 72%. HRIS or People Analytics in 56%. Finance in 53%. CHROs in 48%. Talent Acquisition and L&D both at 38%.

Two observations matter.

BU leaders are in the room more often than CHROs. That tells you which conversation matters at the working level:

- The C-suite reviews.
- The middle designs.

That is consistent with how a business process behaves: the work is done at the operational level, the result is approved at the executive level.

Talent Acquisition and L&D show up in only 38% of cycles. They are the two functions that consume SWP outputs most directly. A workforce plan that is not co-designed with the people who will close the gaps will produce gaps that nobody closes. We come back to this in Chapter 7.

How often

SWP Cadence

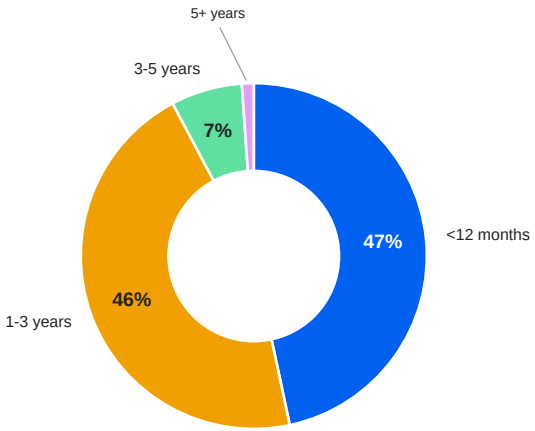

47% of the panel runs or updates the SWP exercise within a 12-month window, which covers everything from quarterly through annual cycles. 46% updates on a 1 to 3 year cycle. 7% on a 3 to 5 year cycle. 1% less often than that.

Two readings sit on top of this distribution.

- **The encouraging read:** half the panel is on at least an annual cadence. That is the minimum required for SWP to function as a recurring process rather than as a periodic study. The mature organizations within this group are typically running rolling updates, where the same 3-year or 5-year forecast gets refreshed every quarter or every six months instead of being rebuilt from scratch each year.
- **The sobering read:** the other half updates only every 1 to 3 years or less. At that cadence, SWP is closer to an episodic study than a process. The snag is that hypotheses age, and drivers drift... The exercise loses its connective tissue with the business decisions it was meant to inform.

A note on terminology that we will keep tight throughout this study:

- **Cadence** is how often the exercise is run or updated (every quarter, every year, every two years).
- **Horizon** is how far ahead the forecast looks (1 year, 3 years, 5 years). The two are independent.

A rolling 3-year forecast updated every quarter is a high-cadence, medium-horizon practice, and it is the shape most industrialized organizations end up with.

How far ahead

Planning horizon

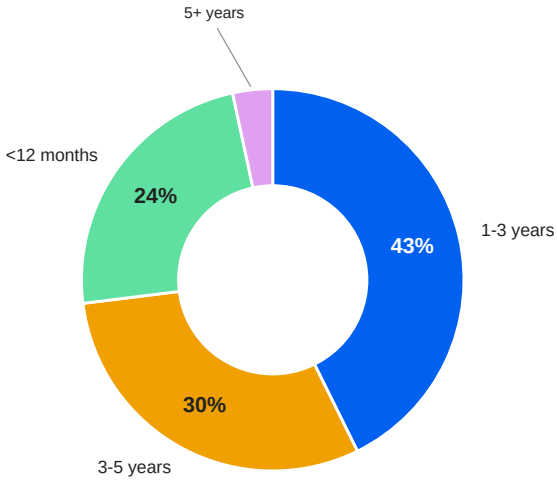

The panel distributes itself across SWP horizons. 24% plan less than 12 months out. 43% plan 1 to 3 years out. 30% plan 3 to 5 years out. Just over 3% plan further.

A planning horizon shorter than 12 months is, by definition, Operational Workforce Planning. The fact that nearly a quarter of the panel sits there suggests either that some respondents are doing OWP and answering the SWP survey because their organization labels it that way, or that they are doing SWP at a horizon too short to act on. The practical implication is the same in both cases.

Recruitment campaigns, reskilling programs, internal mobility paths and retention plans all take quarters to design and years to deliver. A 12-month horizon will run into that wall every cycle.

What it covers

Scope currently covered

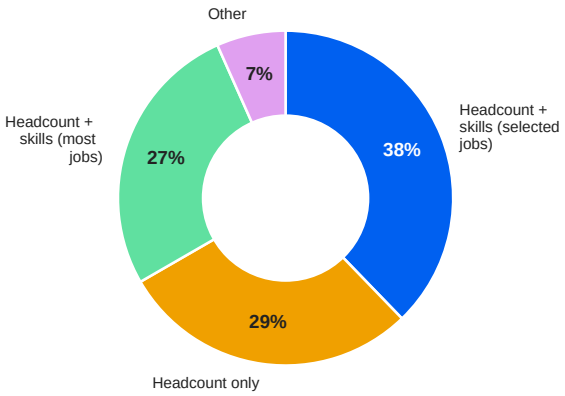

65% of the panel includes skills in their scope, in some form. 29% remain headcount-only. The skills layer has arrived as a methodological standard. What varies is the depth at which it is applied: 38% on selected jobs only, 27% on most jobs.

For organizations of any meaningful size in any evolving industry, the headcount-only model cannot answer the questions that actually keep CHROs and BU leaders awake. How many of the data analysts you have today will still be useful in three years. Whether the engineering organization you plan to grow is the right one for the products you plan to ship. Whether the workforce you are paying for today is the one you will need tomorrow. Those are skills questions wearing headcount clothing.

What good looks like

A healthy governance and cadence arrangement for SWP looks like this:

- A **clear coordinator in HR**. Ideally a Central SWP team for organizations large enough to support one, Corporate HR otherwise. The exercise is owned and driven from HR.
- **HRBPs and BU Leaders** as the main contributors in the field.
- **Finance** addressed assertively as a critical partner. They validate the figures, own the budget hand-off, and have a legitimate seat at the table. The operating model is clear from the start that they are a partner, not the driver, and that distinction is made openly rather than left to ambiguity.
- A **stakeholder lattice** that includes HRIS or People Analytics for data discipline, and Talent Acquisition or L&D for downstream action. Four or five active groups in the room is the right size. Three is too narrow. Eight starts to dilute the conversation.
- A **cadence** of at least **annual**, with rolling updates wherever the organization can support them. The same 3-year forecast updated every quarter is materially more useful than a fresh 5-year forecast rebuilt every other year.
- A **planning horizon** of at **least three years**. Shorter is OWP. Longer is appropriate where the lead times of the business require it.
- A **scope** that includes **skills**, even if only for the 20 to 30 critical job families that drive the strategy. Headcount-only is no longer where the field is.

How to get there

If your scope is headcount-only, pick three to five critical job families and add a skills dimension for those families before extending. A partial skills layer applied where it matters most will outperform a comprehensive skills layer applied at low resolution everywhere.

If your cadence is longer than annual, commit to a rolling refresh. Pick the quarter, lock the calendar, and refresh hypotheses on the same drumbeat. A rolling cadence does not require a full rebuild every cycle.

If your horizon is shorter than 12 months, either rename what you do (call it OWP and design it accordingly) or extend the horizon. There is no useful middle ground, and words matter.

If your coordinator currently sits in Finance, the move is to address the operating model explicitly. Bring Finance closer as a validating partner, put the chair in HR or in a Central SWP team, and reframe the exercise around workforce questions. Finance keeps its rightful seat at the budget hand-off.

Data & Segmentation

Before any forecasting can happen, the panel has to agree with itself on what the workforce is. That comes down to two questions:

- Where the source of truth lives.
- How the workforce is cut into segments that the planning exercise can address.

The panel's answers describe a picture that is broadly competent, with one fragility worth naming early.

Where the source of truth lives

Primary people data source for SWP

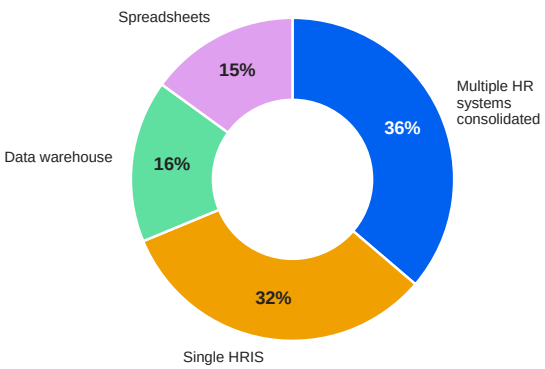

49% of the panel runs SWP from a single HRIS or a data warehouse, 36% consolidate multiple HR systems into a working dataset, which usually means real engineering effort every cycle but a fundamentally reliable answer at the end of it. A whopping number of 15% report spreadsheets as the primary source of truth.

That last number deserves attention. Spreadsheet-as-tool is well known and we return to it in Chapter 8. Spreadsheet-as-source is the more structural problem: when the data the forecast is built on lives in a spreadsheet, versioning is informal, joins are manual, and audit trails are almost always retrospective. A spreadsheet tool sitting on top of a clean HRIS source is recoverable. A spreadsheet tool sitting on top of a spreadsheet source is not.

How the panel slices itself

Segmentation dimensions used

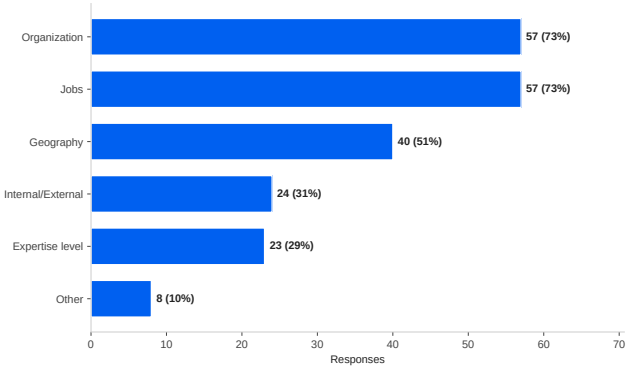

Jobs and Organization are the dominant segmentation axes, both at 73%. Geography follows at 51%. That trio is the structural minimum for any meaningful SWP exercise in a multi-region organization.

The two less common cuts are worth attention:

- **Internal/External** at 31% says only a minority of the panel models the contingent workforce alongside the permanent one, despite contingent labor being increasingly strategic in tech, professional services and pharma.
- **Expertise** level at 29% says fewer than a third explicitly segment by seniority or proficiency, which becomes a problem the moment skills enter the picture.

Job architecture granularity

Granularity of job architecture

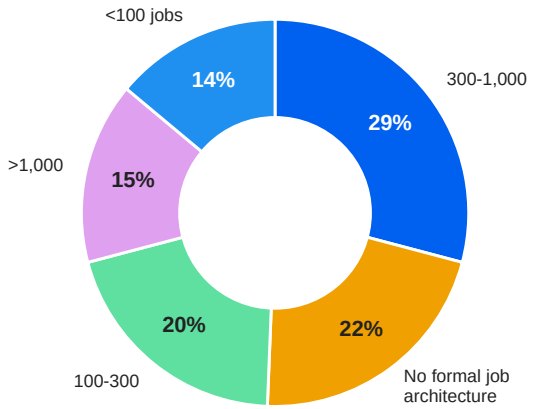

78% of the panel has a formal job architecture. Among those, the modal granularity is 300 to 1,000 jobs (29%), which is the pragmatic sweet spot for a large enterprise: granular enough to find the bottlenecks, coarse enough to govern.

Below 100 jobs, the architecture is usually a job-family layer used as a proxy for jobs proper, which works in early SWP and breaks once skills enter the picture.

Above 1,000 jobs, the architecture starts to need its own owner, because the maintenance work alone can absorb a full-time role.

What good looks like

- **A single source of truth**, or a consolidated layer that masquerades as one. The HRIS or data warehouse is the working assumption. Spreadsheets stay downstream as tools, not upstream as sources.
- **A segmentation that covers Jobs, Organization, and Geography at minimum,** with Internal/External added wherever contingent labor is a meaningful slice of the workforce.
- **A job architecture** in the **300 to 1,000** range for enterprises, with explicit governance (an owner, an update cadence, a versioning convention) once the count crosses 500.

How to get there

If your source is a spreadsheet, the next investment is obviously consolidating to whatever HRIS or warehouse can hold the canonical workforce dataset. The work to do this once is meaningful. The work to keep doing SWP on a spreadsheet source is heavier on a longer horizon.

If you are segmenting on fewer than three dimensions, the most valuable one to add depends on the business. For multi-region operations, Geography. For consulting, professional services and technology, Internal/External. For organizations with significant seniority heterogeneity inside the same job, grading could make sense.

If you have no formal job architecture, start with the 30 to 50 critical job families that actually drive the strategy. Expand from there as the methodology earns the right to more granularity.

Skills Frameworks

The skills layer is where the methodology becomes most uneven. Broadly adopted, very unevenly governed. The panel's answers tell that story in four steps: whether a framework exists, whether it is unified, how big it is, and whether anyone is officially responsible for keeping it that way.

Framework coverage

Skills framework status

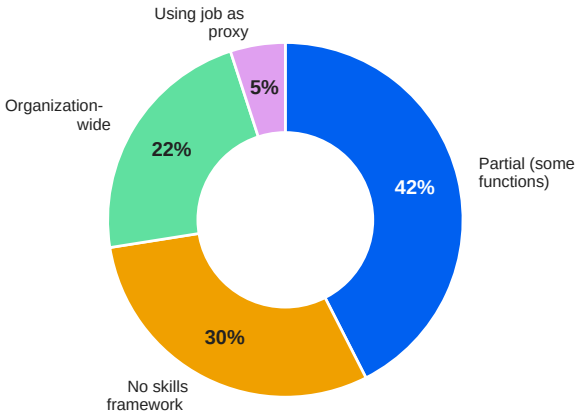

64% of the panel has a skills framework in some form. The modal state is partial coverage, with a framework in place for some functions but not others (42%). Organization-wide frameworks reach 22%. The "no framework" group is 30%, and includes both organizations that have not invested yet. And there is also a small subset (5%) using job titles as a working proxy for skills, which is a defensible interim solution and not a substitute for the real thing.

Unification

Unification across BUs/regions

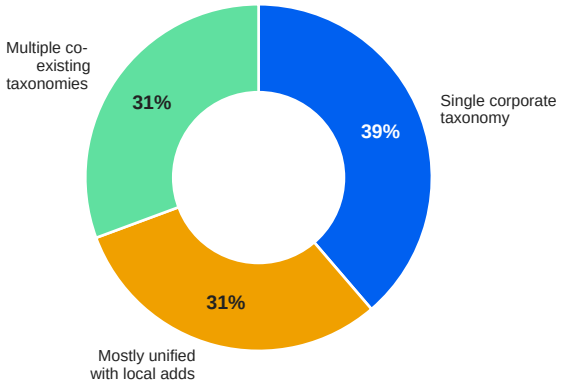

Among the framework-havers, only 39% have a single corporate taxonomy. 31% have a mostly unified framework with local adds, which is the pragmatic middle ground for groups with regional autonomy. 31% have multiple co-existing taxonomies, which is almost always a legacy of mergers, acquisitions, or business-unit-level skill initiatives that never converged. Multiple taxonomies do not stop SWP from happening. They do stop cross-organization skills analytics from happening, which is exactly the analysis CHROs and CEOs increasingly want.

Size

Number of skills in framework

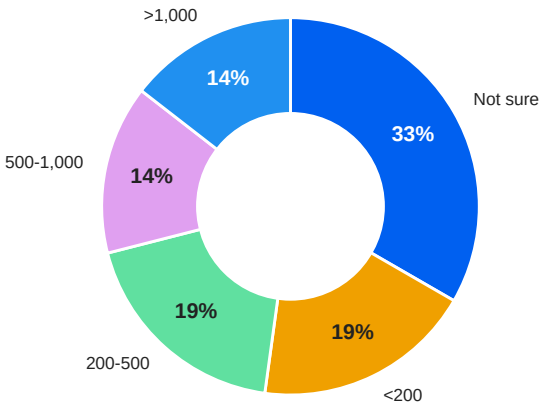

A third of framework-havers (33%) do not know how many skills they have. This is the single sharpest tell of skills governance immaturity in the entire study. Among those who do know, the distribution is roughly even across size bands: 19% in <200 skills, 19% in 200 to 500, 14% in 500 to 1,000, 14% above 1,000. The field has not converged on what the right number is, and probably never will. The useful framing is that the right number is whatever you can keep current with the governance you actually have.

Governance

Skills governance practices

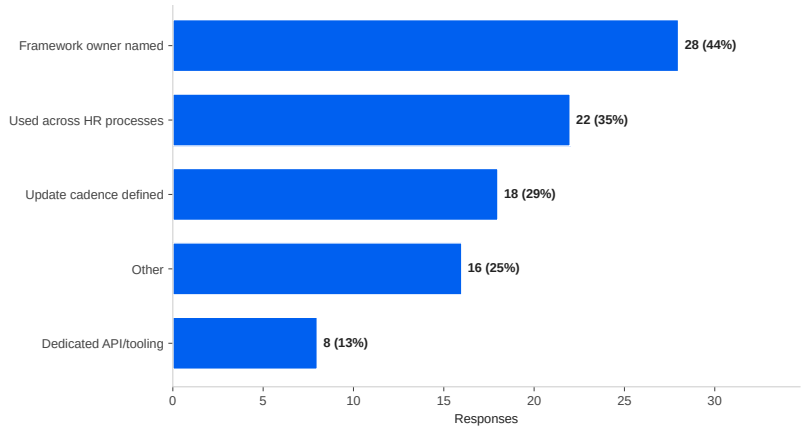

Even among framework-havers, less than half have a named owner (44%). Update cadence defined drops to 29%. Cross-process use (the framework actually showing up in performance, learning, and mobility) is at 35%. Dedicated API or tooling for the framework sits at 13%. The four practices read as a maturity ladder. The framework exists for most of the panel. The framework's maintenance and consumption infrastructure usually does not.

What good looks like

- A framework with **explicit coverage**, at minimum the 30 critical job families that drive the strategy, with a stated plan for extending it.
- A **single corporate taxonomy with permitted local adds**. Multiple co-existing taxonomies are an architectural debt, and the right time to consolidate is before the next round of SWP.
- A size that matches **the governance you can sustain**. 200 to 500 skills is more than enough for most strategic SWP work (we have a preference for “frugal” ontologies). 1,000+ is appropriate only when the maintenance organization is in place.
- **A named owner**. A defined update cadence. Use across at least one other HR process, typically performance management or learning.

How to get there

If you have no framework, start with the same critical job families you used for segmentation. Build a basic skills taxonomy bottom-up from those families. Aim for 100 to 200 skills as a starting point. The first round will be imperfect, and that is perfectly fine.

If you have multiple co-existing taxonomies and want to have a consolidated view, pick the one closest to consensus and declare it canonical. Treat the others as legacy with a sunset plan, map as much as you can to the new reference to minimize friction, then run the next SWP cycle on the canonical version regardless of the local pushback. Pain now beats pain compounding.

If you do not know how many skills you have, the question is not the number but the absence of an owner. Assign one, and the audit follows naturally.

If your framework has no defined update cadence, attach it to an existing HR cycle (annual performance review, biannual learning planning) so the cadence does not need its own organizational gravity.

Demographic Hypotheses

Demographic modeling is the most quietly mature part of the methodology. The panel knows how to model who leaves, mostly trusts internal historical data for the assumptions, and increasingly differentiates the assumptions by criteria that actually matter. This chapter is short because the field is in better shape here than anywhere else.

What departures we model

Departures modeled

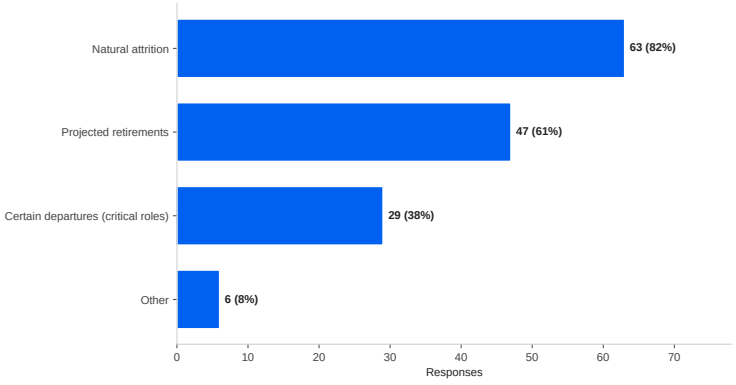

82% of the panel models natural attrition. 61% models projected retirements. 38% models certain departures, which typically means critical-role flight risk and announced exits. The combination of attrition plus retirement covers about three-quarters of the relevant departure dynamics for most large organizations. Modeling certain departures adds the qualitative layer that turns demographic math into actual risk management.

Where the assumptions come from

How attrition/retirement assumptions are set

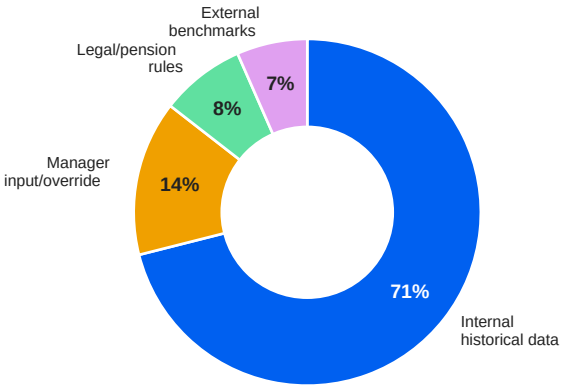

71% of the panel sets attrition and retirement assumptions from internal historical data, typically the last two to three years of HRIS records. 14% allow manager input or override, which is the right safety net for critical roles and small populations where statistical averages get unstable. 8% use legal or pension rules, which is appropriate for retirement projection in jurisdictions with strong statutory frameworks. The striking number is external benchmarks at 7%. That is far below what could be expected and confirms our own empirical data.

Differentiation

Attrition differentiation criteria

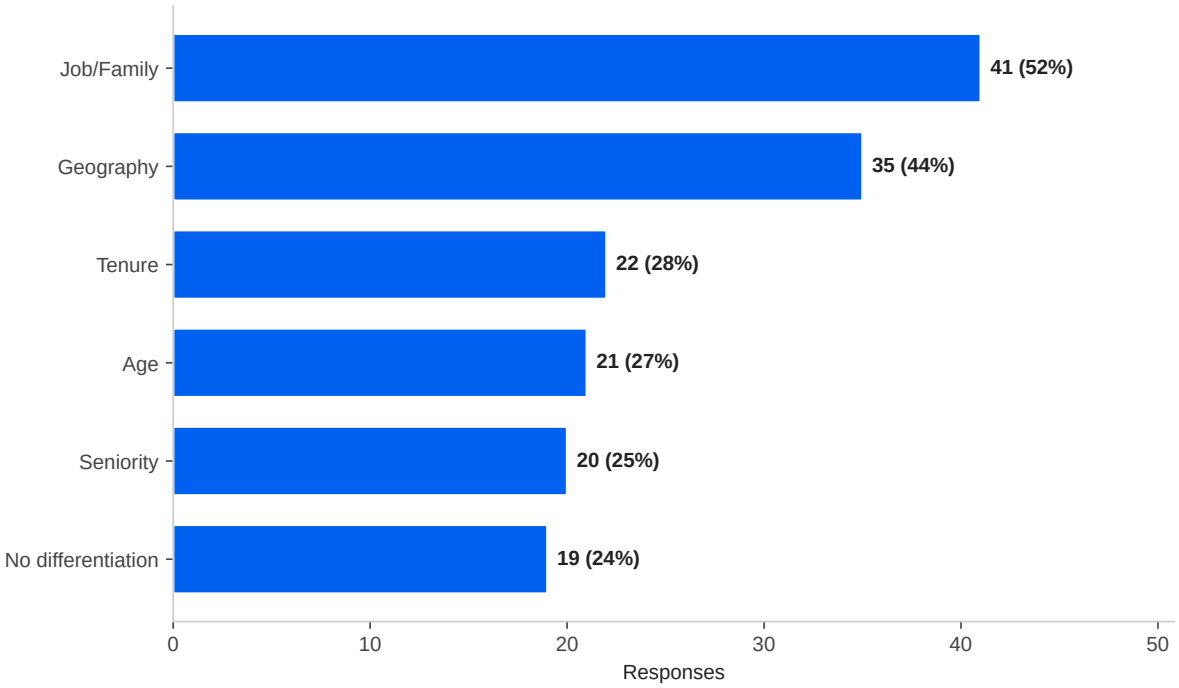

Job or Job Family is the most common differentiation lever (52%), followed by Geography (44%), Tenure (28%), Age (27%), and Seniority (25%). 24% of the panel does not differentiate attrition assumptions at all, which means a single attrition rate is applied across the entire workforce.

For organizations above 5,000 FTEs, that approach almost certainly hides material differences inside the average. The good news is that most of the panel does some form of differentiation. The remaining question is whether the differentiation is the right one for the workforce.

What good looks like

- All **three departure types** need to be modeled: Attrition plus projected retirements plus a critical-role layer for known and likely exits.
- **Internal historical data is the primary source**, with external benchmarks as a second-track cross-check on at least the highest-volume scopes. Manager override available for critical roles and small populations.
- **Differentiation** at minimum by Job Family and Geography. Age, Tenure, or Seniority added where the workforce has meaningful structural variance along those dimensions.

How to get there

If you currently model only attrition, add projected retirements next. Retirement projection is mostly demographic arithmetic and pays back in the most stable forecasting layer you can build.

If you apply a single attrition rate across the workforce, the first differentiation to add is Job Family. Geography follows if you operate across distinct labor markets. The combination of those two will already separate the meaningful signal from the noise.

Bear in mind that using too many variables for attrition will result in very small subsegments, which can skew the rates entirely. A subset with one person only who is leaving naturally results in a 100% attrition for that population, which makes no sense.

Business Drivers

Drivers are the chain link between business strategy and workforce volume. Without them, demand forecasting reduces to extrapolation, which is a competent description of yesterday and a poor model of tomorrow. The panel splits roughly in half on whether drivers exist, and then again on whether the drivers that exist are governed well enough to be useful.

Adoption of driver-based forecasting

Driver-based demand forecasting

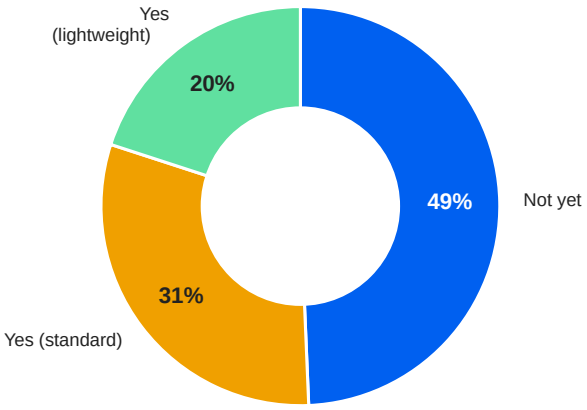

51% of the panel uses a driver-based approach to forecast workforce demand. 31% report it as standard practice and 20% as lightweight. 49% answer "not yet." The split is close to even, which makes drivers the single sharpest dividing line in the methodology.

Organizations that have built a driver model think structurally about what changes their workforce needs. Organizations that have not built one rely on historical trends and BU intuition, which works until a real strategic shift makes both unreliable.

How many drivers

Number of drivers per ~10k FTE scope

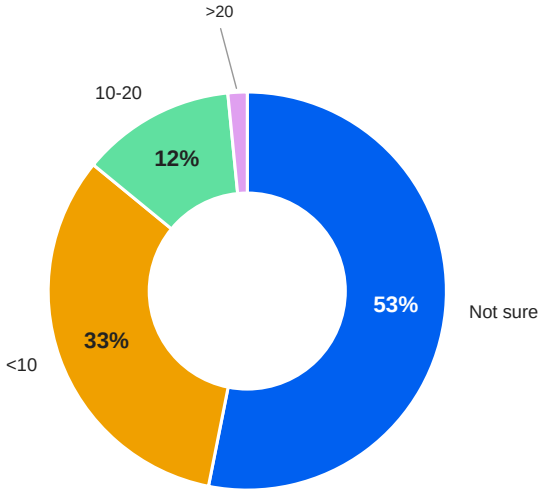

Among the driver-using subset, 53% do not know how many drivers their model has. Of those who do know, the most common range is below 10 (33%), which is the right scale for a focused model. 12% report 10 to 20, which is appropriate for a complex multi-business-unit organization. Above 20 is rare and probably overweight. The high "not sure" percentage matters because it signals the same governance gap we saw with skills: the artifact exists, the inventory does not.

Collection method

How driver impacts are collected

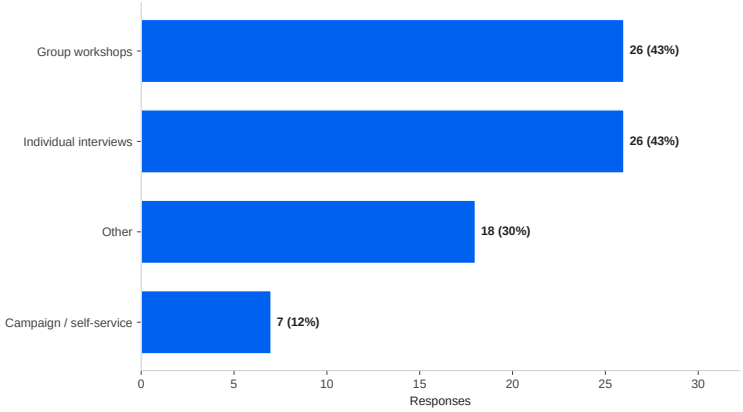

The two dominant methods are individual interviews and group workshops, both at 43%. Campaign or self-service collection is at 12%. The dominant pattern is craft-driven: drivers are surfaced through conversations with BU leaders, not through structured data collection.

That is the right approach for designing drivers, and the wrong approach for scaling them. A mature driver organization typically uses workshops to surface and interviews to validate, with self-service tooling layered on top to keep the driver library current between cycles.

Validation

Who validates final figures

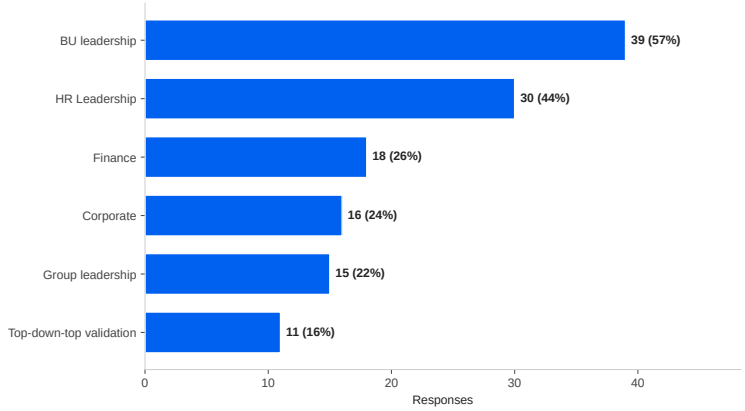

BU leadership is the most common validator (57%), followed by HR Leadership (44%). Finance, Corporate, and Group leadership all sit in the 22 to 26% range. Top-down-top validation, where figures get pushed up from the BU, reviewed at corporate and group level, and then sent back down for final commitment, is at 16%.

This is the gold-standard validation pattern in mature SWP. Its scarcity in the panel is consistent with the broader story that most of the methodology is in place, but the industrialized version of it is not.

What good looks like

- **Driver-based forecasting** as the standard method, even if lightweight at first. The alternative is extrapolation dressed as planning.
- Approximately **10 to 20 drivers** per major scope (roughly 10k FTEs), with a maintained inventory and a named owner. Fewer than 10 is acceptable for focused scopes. More than 20 usually means the model is doing the work that scenarios should do.
- A **collection process** that combines **workshops** (to surface drivers) and **individual interviews** (to validate impact). Self-service tooling layered on for between-cycle maintenance once the driver library is stable.
- **Top-down-top validation** as the working standard. BU leadership owns the impact estimate, corporate and group leadership challenge and consolidate, the final figure is committed jointly.

How to get there

If you have no driver-based approach, start with one BU and one round of workshops. Surface 10 drivers. Estimate their impact on volumes and skills. The output of that first cycle will be imperfect and useful at the same time.

If you do not know how many drivers you have, the audit is small. Catalog them, assign an owner, and prune to the 10 to 20 that materially move the forecast.

If you collect drivers via a single method, add the second. Workshops alone tend to overgenerate. Interviews alone tend to under-surface.

If validation stops at BU leadership, the next step is a corporate or group review pass. The point is not to override BU judgment but to surface inconsistencies across BUs and force them to be reconciled.

Scenarios & Gap Analysis

Scenarios and gap analysis are the hinge between forecasting and action. Drivers and demographics produce numbers. Scenarios stress-test the numbers. Gap analysis turns the comparison into a decision. The panel is most uneven here.

How many scenarios

Number of workforce scenarios modeled

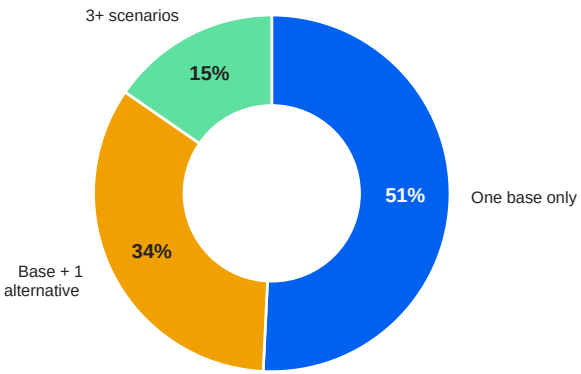

51% of the panel models a single base scenario only. 34% adds one alternative. 15% runs three or more scenarios. This is the most under-developed practice in the methodology relative to its return on effort. Adding a second scenario is the cheapest way to make a forecast more useful, because it forces the assumptions to become explicit and the exposure to become visible. A base scenario alone is a forecast. A base scenario plus an alternative is a planning input.

Gap quantification

Gap quantification (needs vs forecasts)

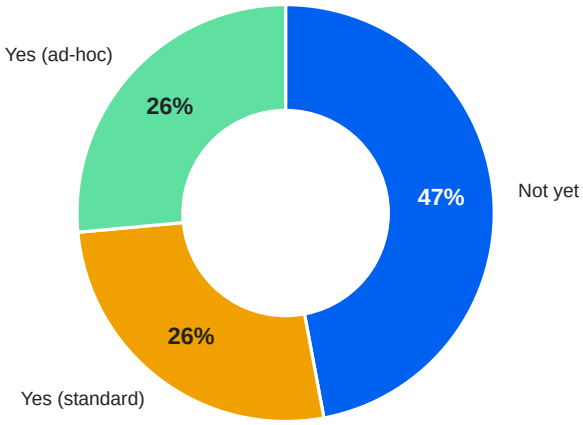

26% of the panel quantifies the gap between scenario needs and demographic forecasts as a standard practice. 26% does it ad-hoc. 47% does not yet. Combined, 53% does some form of gap quantification, which is healthier than the headline number suggests. The split between "standard" and "ad-hoc" is the real distinction. Ad-hoc gap quantification is reactive: someone asks, you do it for that scope. Standard gap quantification is structural: every scope has a gap report every cycle.

Skills gap analysis

Skills gap analysis approach

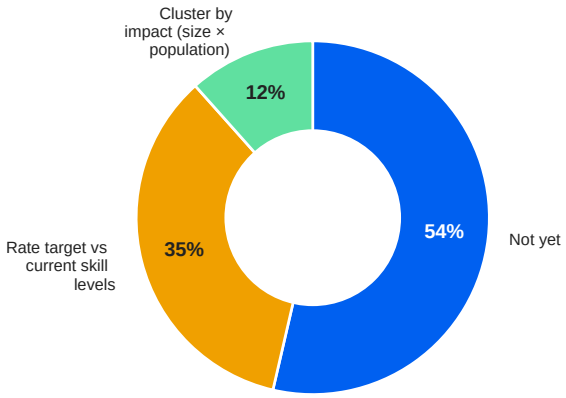

35% of the panel rates target versus current skill levels for the jobs and skills in scope. 12% goes further and clusters skills by impact, which is the right next step for any organization with more than a few dozen critical skills. 54% does not yet do skills gap analysis. The distribution is consistent with what we saw in Chapter 3: the skills framework exists for most of the panel, and the methodology that uses the framework lags the framework itself.

Time to act

Time horizon sufficient to act on gaps (1-5)

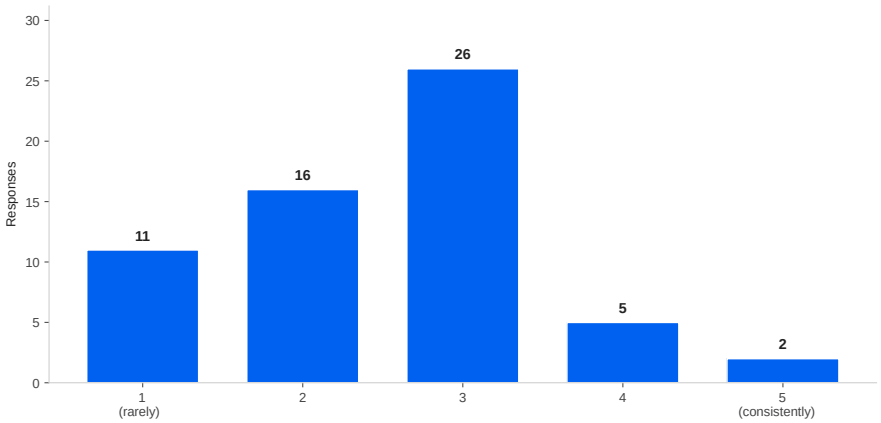

The panel rates its time horizon to act on identified gaps at a mean of 2.52 out of 5, which is structurally similar to the maturity rating in Chapter 1. Most respondents feel short on time. The lever is not analytical, it is structural: act on a longer horizon, or get the cadence and tooling in place to act faster within the horizon you have.

What good looks like

- **Three scenarios** as the working default: downside, base, and upside (or following other "archetypes"). The downside is the scenario that makes you build retention plans. The upside is the scenario that makes you build pipelines. The base is the scenario you commit to.
- **Gap quantification** as a standard step in every cycle, for every scope. The format is the same across BUs, even when the numbers are different.
- **Skills gap analysis** at least at the **target-versus-current** level. Impact clustering (gap size × population) added as the next layer for organizations with more than 100 critical skills.
- A **planning horizon** and cadence combination that gives you **at least 18 months** of runway to act on what the gap analysis surfaces.

How to get there

If you run only a base scenario, the next scenario to add is the downside. It is the one that drives retention and contingency thinking, and it tends to be the most uncomfortable to model, which is exactly why it is the most useful.

If gap quantification is ad-hoc, the move is to industrialize the format. Same template across BUs, run every cycle, archived for audit.

If you do not do skills gap analysis, start with five critical skills in the most strategic job family. Rate target versus current for those five. Use that work as the template for extension.

If your time horizon to act feels too short, see Chapter 1 on horizon and Chapter 8 on tooling. Both contribute, but neither solves it alone.

Action Plans: the 8Bs

The 8Bs framework is the methodology's vocabulary for closing gaps. Eight levers, each a different mechanism for moving the workforce toward the scenario you committed to. The chapter is short because the chart says most of it.

Levers deployed to close gaps (the 8Bs)

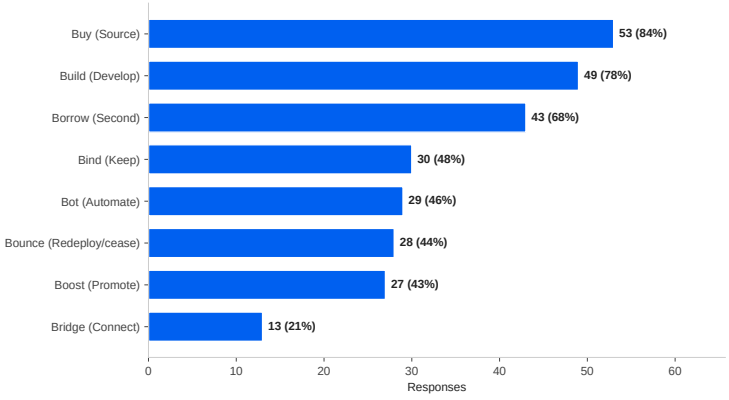

How the levers are actually used

Buy and Build dominate. 84% of the panel sources externally to close gaps. 78% develops internally. Borrow, which covers contracted, seconded, or temporary capacity, follows at 68%. These three together are recognizable as the recruitment-and-development reflex that most HR functions are organized around.

Below the top three, the curve flattens. Bind (retention) at 48%. Bot (automation) at 46%. Bounce (orderly redeployment or downsizing) at 44%. Boost (promotion) at 43%. Then a noticeable drop to Bridge (internal mobility) at 21%.

Two things stand out.

- The first is the gap between adoption levels. The top three Bs are used by two to four times more respondents than the bottom one. The reason has less to do with the bottom levers being less relevant, and more to do with them requiring different organizational muscles. Most HR functions have built the Buy and Build muscles first.
- The second is Bridge specifically. Internal mobility is the lever HR conferences have been talking about for at least a decade. The panel is closing gaps through internal movement in roughly one cycle out of five. That is materially lower than the rhetoric would suggest, and lower than the data on internal mobility's effectiveness justifies. Bridge is the under-used lever the panel should be reaching for first.

What good looks like

- **Conscious** lever choice for each gap, not a default to Buy and Build. The 8Bs are tools, not a checklist.
- **At least four levers** actively deployed across the typical cycle. Three levers is the recruitment-and-development reflex. Four or more starts to look like a real action plan.
- **Bridge** as a first-line consideration before recruitment for any gap that overlaps with internal supply. Internal mobility is almost always cheaper and faster than external sourcing, when the framework exists to support it.
- **Bot** is still a bit of a shady one, as the gaps between theory and reality can be huge.
- **Bind** designed in advance for critical talent. Retention is most expensive when it becomes reactive.

How to get there

If your default is Buy and Build, the next lever to add is Bridge. The prerequisite is the skills framework from Chapter 3 and the segmentation from Chapter 2. If those are in place, internal mobility is the highest-ROI move available.

If you do not consider Bot, the question becomes: what is the most repetitive, rule-based, high-volume work in this scope, and what would it look like to outsource it (Borrow).

If Bind is reactive in your organization, the structural move is to identify critical talent at the time of the gap analysis, not at the time of the resignation letter. The action plan should name names where the workforce model justifies it.

If you deploy fewer than four levers per cycle, the meta-move is to ask which two of the eight you have not seriously considered. Then consider them.

Tooling & Integration

Tooling is the layer where the methodology gets its leverage or loses it. The panel's tooling answers are the most striking single number in the study. The integration answers are not far behind.

The tooling stack

Primary tool(s) for SWP

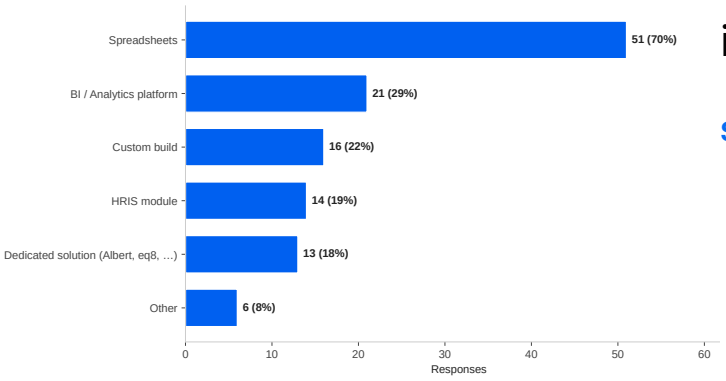

70% of the panel runs SWP with spreadsheets as a primary tool. Dedicated SWP solutions are the primary tool for 18%. HRIS modules at 19%, BI or analytics platforms at 29%, custom builds at 22%. The percentages add to more than 100% because most respondents use multiple tools in combination.

Spreadsheets are useful, and they will continue to be useful at every level of SWP maturity, for ad-hoc analysis, for early pilots, and for the long tail of one-off questions.

What the 70% number really says is that spreadsheets are the workflow itself, not the supporting layer, for most of the panel. That works at small scope. It does not scale through scenarios, versioning, validation workflows, audit trails, and multi-stakeholder review. Past a certain point, the spreadsheet stops being the productivity tool and becomes the bottleneck.

Budget and headcount integration

SWP integrated with budget & HC planning

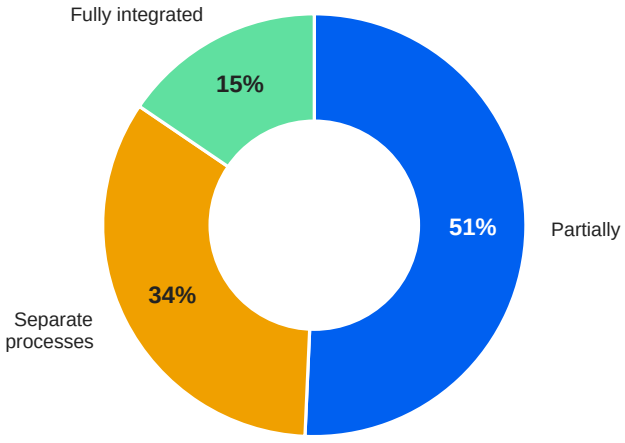

15% of the panel has SWP fully integrated with budgeting and headcount planning. 51% reports partial integration. 34% runs SWP as a separate process. The fully integrated minority is the most operationally connected group in the entire study. They are also the group that finds the easiest path to executive attention, because their SWP outputs are wired into the decisions that the C-suite already makes every quarter.

The partially integrated middle is the largest group, and the most fragile. Partial integration usually means the SWP team shares forecast outputs with the budget team once or twice a year, the budget team uses them with caveats, and neither side has full transparency into the other's assumptions. That works until it does not.

What good looks like

- **A primary SWP environment** that supports scenarios, versioning, validation, and audit. For most organizations of meaningful scale this is a dedicated SWP solution, or a strong consolidated stack on top of a clean data layer. Spreadsheets stay in the toolkit, and not at the center of it.
- **Full integration** with budget and headcount planning as the target state. The same forecast assumptions feed both processes. Mind you: this does not mean it is Finance-driven, quite the opposite. The same scenarios drive both conversations. Disagreements get surfaced and resolved in the planning calendar, not at the budget review.

How to get there

If your tooling stack is primarily spreadsheets, the right time to evaluate alternatives is when you cross 1,000 FTEs in scope, or the moment you need three scenarios run side by side, or the first cycle where versioning becomes a problem you cannot solve with email subject lines.

If your SWP and budget processes are separate, start with a single hand-off. Pick the quarterly budget review, prepare the SWP-side forecast in the same format, and walk through it together once. The process design follows the relationship, and not the other way around.

If you have partial integration, the next layer to wire in is the assumption layer. Both sides should be working off the same drivers, the same demographic baseline, and the same scenarios, even when their commitments differ at the end of it.

Conclusion

THE BENCHMARK, AND THEN WHAT?

109 organizations answered and it seems the methodology has “arrived”. The next mountain is clearly industrialization.

If the chapters have a shared subtext, it is that the panel knows the steps. Governance, data, segmentation, skills, demographics, drivers, scenarios, gap analysis, action plans, tooling... The vocabulary is in place. The hard part now is doing the same steps every cycle, the same way, with the same tooling, and the same validation circle. That is what industrialization means. It is also what separates the top quartile of the panel from the rest.

Three things to take away.

- **Adoption is broad. Consistency is what varies.** If your organization is doing SWP at all, you are almost certainly doing parts of it well. The chapters of this study are a checklist for finding the parts that are not.
- **The weak link is more useful than the strong one.** The chapter where your organization is furthest from "what good looks like" is the highest-ROI place to invest. Industrializing the strong chapters yields diminishing returns. Industrializing the weak ones unlocks the rest.
- **Cadence, tooling, governance.** The fastest path to industrialization is a combination of those three. The methodology rarely fails on its own merits. It fails when one of those three is missing.

We will run this benchmark again in 2027. The questionnaire will possibly change in small ways. The skeleton will not. If you are reading this and we have not yet heard from you, please join the next edition. The conversation gets better with more practitioners in it.

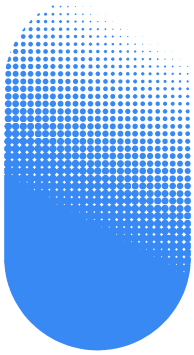

About Albert & The WPI

ALBERT

Albert is revolutionizing Strategic Workforce Planning (SWP) for enterprises by eliminating guesswork and complexity. Our AI-powered SaaS platform simplifies SWP, empowering businesses to plan and adapt with confidence. Unlike traditional spreadsheets or BI tools, Albert seamlessly handles the intricacies of data science and real-time collaboration, helping organizations manage business transformations and optimize their workforce with ease.
Albert is the first challenger to traditional human resources planning practices. Albert was founded in 2021 by Vincent Barat, an engineer who graduated from CentraleSupélec with a ten-year background in SWP, Jérôme Soulard, who graduated from Dauphine with a strong background in HR consulting, and Cyril Siman, our CTO with 20 years of experience in Software Engineering. Our AI-based SaaS platform enables companies to plan their human resource needs, which account for 60% of their expenses. Albert currently has 15 employees and has successfully gained around twenty clients, present in 22 countries, most of which are part of the SBF 120.

THE WORKFORCE PLANNING INSTITUTE

The Workforce Planning Institute (WPI) is a global membership, education and conference organisation dedicated to advancing the practice of Strategic Workforce Planning (SWP). WPI supports professionals and organisations worldwide through industry-leading learning experiences, professional development opportunities and a connected global community focused on building stronger workforce planning capability.
Through memberships, accredited learning programmes, certification pathways and year-round content, WPI helps individuals and organisations develop the skills, knowledge and practical approaches needed to navigate workforce challenges and create future-ready organisations. By bringing together leaders and practitioners from across industries, WPI creates opportunities for learning, collaboration and the sharing of real-world experiences.
As part of its global ecosystem, WPI also delivers SWP Conference, one of the leading global event series for workforce planning professionals. Held across key markets including Chicago, Sydney and London, SWP Conference brings together industry leaders, practitioners and partners to exchange ideas, showcase innovation and shape the future of workforce planning.

albert

LinkedIn
@getalbert

Email Address:
hello@albertapp.com

Website:
www.albertapp.com