MULTIPLEWORKS

Executive Briefing

No. 03

The only generation trained by accident

An executive briefing on building capability by design.
Worth half an hour and a coffee, on a quiet morning.

August 2026

Contents

THE IDEA IN BRIEF

The exposure. Most organisations have never designed how they develop people. They have assumed it. Capable people arrived because the work produced them, gradually and unevenly, as a by-product of doing the job for long enough alongside people who were good at it. That arrangement worked well enough that nobody examined it, and because nobody examined it, nobody funded it.

Why now. The work that quietly did the developing is the work AI does well, and the withdrawal is already measurable. The Stanford Digital Economy Lab, using payroll data through June 2026, finds employment of 22 to 25 year olds in AI-exposed occupations standing 19% below where it would be had it kept pace with their less-exposed peers, with no comparable gap for experienced workers. This is not a forecast about the future of work. It has already happened, and in most organisations there was no pipeline anyone was watching, so nothing has registered as broken.

The cost of waiting. The bill arrives in about five years, as a shortage of people who can be trusted with a decision, and it will not be traced back. No incident report is written when a capable manager fails to emerge. The fields that cannot afford this to fail all train people before the role rather than after it, assess them against a standard, and treat instruction as a specialism. Business has neither their stakes nor their regulator, so it has never been forced to.

The historical position. Formal, structured development is the normal condition of skilled work and always has been: apprenticeships, articles, time-served qualification. The generation shaped by the last three or four decades of white-collar work is the only one trained largely by accident, because the routine cognitive work of that period produced capable people as a side effect. That side effect has now gone.

What to do. Not more courses, and not a programme. The answer is a method for developing people inside real work, managers taught to use it, assessment that runs continuously rather than annually, and a path built around transitions rather than ratings. Most of it costs attention rather than budget. All of it has to be designed, because the thing that used to do it by accident has gone.

Why it rarely gets done. Redesigning a structure is a project, with a completion date and something to hand over at the end. Growing the capability that structure is meant to house has neither, which is why it struggles to reach an agenda assembled out of initiatives, and why an organisation can arrive at a new operating model run by the same people with the same abilities.

PART ONE

The exposure

PART ONE · CHAPTER 1

The pipeline is already thinning

Most arguments about AI and employment are forecasts. This one is not, which is why it belongs on a current agenda rather than a strategic one.

The evidence is concentrated exactly where the theory predicts it should appear first. Brynjolfsson, Chandar and Chen at the Stanford Digital Economy Lab, working from ADP payroll records covering millions of US workers through June 2026, find no sign of widespread economy-wide displacement. What they do find is a gap by age: employment of 22 to 25 year olds in AI-exposed occupations now stands 19% below where it would be had it kept pace with their less-exposed peers, and experienced workers in the same occupations show no comparable gap. That divergence has widened steadily since they first reported it in August 2025, and it operates through reduced hiring rather than through redundancies. The declines concentrate in occupations where AI substitutes for the human task rather than augmenting it, which is the distinction that matters for anyone deciding where to deploy it.

The obvious counter-argument is that AI is being blamed for a contraction driven by interest rates and capital discipline. The authors tested it. AI exposure and interest-rate exposure turn out to be negatively correlated across occupations, and controlling for interest-rate exposure does not meaningfully change the result. Their February 2026 update also found that under the broadest set of controls the declines become statistically significant from 2024, with earlier movements likely driven by other factors. A separate Harvard working paper, covering 66 million workers across more than 280,000 US firms, reports a steeper fall in entry-level hiring at firms adopting generative AI, and the timing there suggests firms adjusted for automation they anticipated rather than automation already in place.

What matters for this briefing is narrower than the employment debate. Whether or not junior roles return, the tasks that used to fill them are being handled differently, and those tasks were carrying something nobody costed. Alongside the shrinking intake sits a second change: employers are loading junior postings with expectations of judgement and stakeholder handling that used to be acquired over several years. The requirement has moved ahead of the mechanism that produced it.

The reason this reaches the executive agenda now rather than in five years is the lag. Capability takes years to build and its absence takes years to appear. An organisation that stops producing capable people in 2026 will not notice until the succession slate is thin, at which point the remedy is another five years away. The decision has to be taken while nothing appears to be wrong.

PART ONE · CHAPTER 2

What the routine work was doing

Every organisation says it develops its people. Few can describe the mechanism. Asked how, the answer is usually a catalogue of courses, a mentoring scheme with variable participation, and a belief that talented people put next to experienced people will absorb what they need.

It is worth noticing how recently that belief started. Capability used to be built formally. Indentured apprenticeships, City and Guilds, articles, chartered training contracts and time-served qualification all exist because trades and professions could not afford to leave competence to chance, and most of them still operate. The undesigned model is specific to white-collar knowledge work over roughly the last three or four decades, when the volume of routine cognitive tasks became large enough to produce capable people incidentally.

What that volume was doing is worth being precise about, because junior work has always been two things at once. The first is low-value repetition: collation, formatting, first-pass review, chasing. Nobody should mourn its automation. The second is apprenticeship, and it was carried inside the first. Reviewing several hundred documents builds a pattern recognition no seminar reproduces, and a scepticism forms only after you have seen a dozen things that looked fine and were not.

But the volume was never set by what learning required. In professional services it was set by the economics of leverage, and the developmental account was laid over the top afterwards. It was true enough often enough that nobody tested how much repetition the learning actually needed.

Most organisations are not structured that way at all, and their position is worse. There is no pyramid, no annual intake, no function whose job is to worry about the pipeline. People join, stay, and absorb the place through the ordinary run of the work. Nobody calls it development and nobody funds it, because it appears to cost nothing and to happen by itself. So when the routine work goes, a professional firm notices. A finance team or a claims operation does not. There was no pipeline to break.

This is not to say nothing was taught. There has always been training, often a great deal of it, and everyone has a record of completed modules to prove it. But a catalogue of self-directed courses, however well made, cannot watch you make a judgement and tell you where it went wrong. What became visible when the routine work went was that the part which could be counted had been standing in for the part that could not.

An undesigned mechanism behaves like an undesigned dependency. You cannot maintain what you have never described, you cannot improve it, and you will not notice it has stopped until the output is missing.

PART ONE · CHAPTER 3

Why freed time does not become development

Reinvesting the hours AI frees into mentoring, exposure and deliberate practice is the intuitive answer, and it does not survive contact with a real organisation. I argued for it myself until recently.

Freed time does not stay free. It is reabsorbed within a quarter by delivery pressure, by utilisation targets, and by the work that is already late. Nobody defends unbilled development hours in a month that is behind, and the manager who tries will lose the argument to the person holding the deadline. A plan that depends on restraint, applied precisely where an organisation has none, is not a plan.

EY's response is an instance of the structural version. Rather than reallocating hours inside existing roles, it has extended its eight-week internship into a paid residency of eight to twelve months, with structured development and entry at a higher grade for those who complete it. Whatever else that is, it changes the shape of the path rather than relying on the discretionary use of time that will be claimed by something else.

Anything that works has to be structural. It has to change the shape of the work or the shape of the path, rather than relying on the discretionary use of hours that will be claimed by something else. That is the test to apply to any proposal in this area, including the ones in this briefing: does it survive a bad quarter, or does it assume a good one?

Structural does not mean a programme, and it is worth being clear about that early, because the reflex in most organisations is to convert anything serious into one. This work has no completion date. There is no target operating model to arrive at, no point at which it is delivered, and nobody who can do it for you and then leave, because what is being built is the ability of your own people and that ability only exists while the conditions producing it hold. It is closer to a maintenance discipline than to an initiative, which is one reason it struggles to reach an agenda assembled out of initiatives, and another reason it has quietly gone missing while the routine work was still covering for it.

The redesign of a structure is a project and can be run as one. Growing the capability that structure is meant to house is not, and the two are routinely confused, which is how an organisation ends up with a new operating model run by the same people with the same abilities, and wonders in eighteen months why so little changed.

What high-consequence fields do differently

Seen against that history, aviation, rail, shipping, medicine and the military are not exotic comparators. They are simply the fields that never stopped. What they have in common is that being wrong kills someone. Aviation, rail, shipping, medicine and the military all train people before the role rather than after it, reassess them periodically rather than once, and treat instruction as a specialism rather than a favour.

It is worth being clear about why, because it is not superior wisdom. The consequence of error in those settings is immediate and physical, and in most cases a regulator requires the training, so it cannot lose an argument with this quarter's workload. The stakes imposed the discipline. Business has neither, and its failures are slower and more diffuse. A manager who was never taught to manage costs the organisation good people across several years, and nobody writes an incident report. The absence of a forcing function is the whole reason development competes with delivery and loses.

That means two things. The parts of the model that depend on a regulator or on non-billable time will not transfer. The parts that depend only on sequencing and attention will, and they are the cheaper parts: training before the role rather than after it, letting the transition rather than the rating decide who gets developed, and matching what is provided to the gap in front of you. Chapter 15 sets out what each involves.

PART TWO

The method

Everything above is diagnosis. What follows is a position on what actually develops people inside commercial work, and it is deliberately narrow. It is not a curriculum. It is how a piece of work is handed over, supervised and closed, and it is available to any manager without a budget line.

Three capabilities come out of it, and they are worth naming first, because they are the base that everything above them depends on.

Initiative. The disposition to act on an ambiguous problem without waiting to be told how.

Lateral thinking. The ability to approach a problem from an angle the obvious route does not offer.

Grit. The capacity to stay with something difficult past the point of comfort.

Each of these resists being taught in a classroom, because each requires real stakes to develop. That is why they were formed by the work rather than by the training, and why their loss is the part of the current shift that will be felt longest. The four chapters that follow are how they are built.

PART TWO · CHAPTER 5

Outcomes, not tasks

The most consequential thing a manager does for someone's development is decide how to hand over a piece of work.

A task carries the thinking already done. Produce this analysis, in this format, using this method, by Thursday. Completing it teaches execution, and very little else, because the route was supplied. An outcome withholds the route. Here is what we need to be true by Thursday, and I have not worked out how.

The construction of the route is the thing that becomes judgement. Working out what the problem actually is, what would count as an answer, which approaches are available and which to abandon: that is the activity the routine work used to force, and it can be reintroduced deliberately at the moment work is delegated.

Managers default to tasks, and the reasons are good ones rather than lazy ones. A task is faster to explain, the output is predictable, and the risk is lower. Under pressure it is the obvious choice. But the manager who hands out tasks is keeping the developmental part of the work for themselves, usually without noticing, and doing so at the precise moment they believe they are being helpful.

The stretch has to be real. An outcome the person already knows how to reach is a task with a longer preamble. It should be beyond what they have done before by enough that they have to think laterally about it, and it should matter enough that the thinking is done seriously.

PART TWO · CHAPTER 6

The rescue window

One rule makes all of this safe: never leave less time on the clock than you would need to do the work yourself.

The objection it answers is that stretching people is a gamble with work that has a client and a deadline attached. It is, unless it is bounded, and this is the boundary.

That single rule makes the practice safe. The stretch is genuine for the person doing it, because they own the outcome and nobody is going to supply the answer. It is bounded for the organisation, because at any point up to the window closing the manager can take it back and still deliver. The risk is transferred to the manager's own capacity, which is where it belongs, and not to the client.

Applying the rule requires the manager to know two things: how long they would need, and where the person has actually got to. The first is experience. The second requires them to be watching, which is the subject of everything that follows.

It also has a limit. In a genuinely overloaded month, the rescue window shrinks until there is no room to stretch anyone, and the correct decision is to hand out tasks and deliver. That is a judgement rather than a failure. What is not acceptable is for every month to be that month, which is the state most organisations have quietly reached.

PART TWO · CHAPTER 7

Productive difficulty

Inside the window, the manager's job is to help someone think without doing the thinking for them.

That is harder than it sounds, and it runs against every instinct a competent person has. You can see the answer. Supplying it would take ninety seconds, would relieve the visible discomfort of someone who is stuck, and would feel like leadership. It teaches nothing.

The discomfort is where the learning lodges. Being stuck on something real, that matters, with your name on it, produces a depth of retention that no exercise reproduces, because the stakes make the thinking serious. The role of the person supervising is to keep the difficulty productive rather than destructive: asking what has been tried, what the problem really amounts to, what would have to be true for an approach to work, and letting the silence do its work.

There is a line, and locating it is the skill. Difficulty that produces thinking is development. Difficulty that produces panic, or that continues past the point where anything further is being learned, is waste and worse. Judging that line is itself an untrained capability in most managers, which is why the method cannot simply be published as guidance and left to good intentions.

Air cover

None of the above is possible unless failure is survivable, and whether it is survivable is decided by the manager rather than by the individual.

If a stretch goes wrong and the consequence lands on the person who attempted it, the organisation has taught everyone watching a precise lesson: do not attempt anything you are not already certain of. Initiative is expensive to build and cheap to destroy, and one visible instance of someone carrying the blame for an honest attempt will undo a great deal of encouragement.

The manager who delegated the outcome owns the outcome. That is the whole of it. Publicly, upward, at the point where it costs something to say so. In exchange the manager gets a team that will attempt things, which is the only way anyone develops beyond what they can already do, and the only way the manager ever gets leverage.

This is also where the method connects to something larger. A team that believes its leader will carry the risk gives back discretionary effort, which is what every difficult programme runs on. It is not a soft consideration. It is the operating condition for everything else in this briefing.

It is also most of what an organisation means when it says culture. Culture is not the values statement; it is the residue of decisions like these, taken a few hundred times a week across the business: whether work arrives as an outcome or a task, whether attempting something is safe, what happens to the person who tried and missed. Which is why culture programmes so often change nothing. They work on the residue rather than on the decisions producing it, and those decisions are still being made the same way on Monday morning.

PART THREE

Making it real

PART THREE · CHAPTER 9

Who teaches the manager

The first competency to define, and the first cohort to teach, is the method itself, and the audience is managers. Nothing in part two is exotic; the obstacle is that commercial managers are rarely taught any of it. I learned it in the military, where developing your subordinates is part of the syllabus of becoming a leader, taught and assessed alongside everything else, because leading and developing are understood to be the same job. In commercial organisations it is at best an optional module and usually a disposition some managers happen to have.

That cannot be delivered by a learning function, however good. The technique is practised in the flow of real work with real consequences, so it can only be taught by someone who does it and assessed by watching someone do it. What is needed is a small cadre of practitioners who can instruct: selected for it, trained in it, and given the standing to do it.

Instruction is a discipline, and it decomposes into components that can be taught.

Command of the subject. Or knowing where to find someone who has it. Without it there is no authority in the room, and people do not learn from someone they do not respect. Note the direction of that: confidence follows from command of the material. Confidence without it is the failure mode that lets a plausible presenter pass for a capable one.

Synthesis. Teaching a subject forces a reconstruction of it that using it does not, because you have to take it apart and rebuild it for someone else. This is the argument that makes an instructor cadre commercially defensible. Taking a senior practitioner out to teach is not time subtracted from their expertise; it deepens it.

A repertoire of methods. Not everyone learns the same way, and the circumstances vary. I spent years as an advanced diving instructor, where a lesson has to be delivered underwater by someone who cannot speak. That constraint forces a separation most corporate teaching never has to make: between what is being taught and the medium it is habitually taught in. Where the deck is the only method available, the deck gets mistaken for the lesson.

Adaptability. The plan does not survive the room. An instructor with one method is finished the moment conditions change, and conditions always change.

Preparation as a discipline

Professional instruction carries a risk discipline. In physical settings it protects people. In commercial settings it protects the outcome, and the method is the same: what has to work for this to succeed, where can it fail, what is the fallback, and at what point do I abandon the plan and still deliver.

An office looks like the easy case, and in one sense it is, since nobody drowns if a session fails. The requirement for craft does not disappear, it stops being obvious. Rooms get reallocated, links fail, half the group is remote, someone is pulled into a client call. An instructor whose only method is the deck is stopped by any of them, and the development quietly does not happen that week.

The useful question, asked before walking in, is what the minimum is that still delivers the outcome if everything else fails. That is the same graded protection this firm has argued for elsewhere in deciding which few things a business cannot fail to do, applied to a session rather than an estate. It also means the instructor demonstrates initiative and adaptability in front of the room, which are among the capabilities the training exists to build.

PART THREE · CHAPTER 11

The job description problem

Every organisation already produces the artefact this requires. It is called a job description, and it is usually the closest thing in the building to a definition of what a person in a role must be able to do.

It is also, in most cases, written for the wrong purpose. A job description exists to attract and filter candidates, so it is drafted by committee, made aspirational to appeal, and padded with everything anyone in the room thought desirable. The result is a laundry list of twenty attributes describing a person who does not exist, phrased in terms nobody can assess: a strong communicator, a strategic thinker, comfortable with ambiguity. It is used once, at hire, and then filed.

Nobody can be developed against that, because none of it is stated as a capability with a standard, and nobody can be assessed against it either, which is why appraisal conversations so often descend into impressions of character.

The consequence worth the executive's attention is a different one. The compromise happens regardless. Nobody can meet a twenty-item list, so the hiring manager sits in the room and privately trades several items away, usually the ones hardest to test in an interview. A candidate is appointed who has, say, six of the ten things that mattered. Which four were traded, and how far short they fell, is not written anywhere. So a person joins with a known gap that nobody knows, and the organisation has quietly converted a development requirement into an unrecorded risk.

That is the useful move, and it costs nothing. Reduce the list to the small number of capabilities the role genuinely turns on, six or eight rather than twenty, each stated as something a person does rather than something they are, described concretely enough that two observers watching the same work would reach the same verdict. Then, at appointment, record which of them the person has and which they do not.

The document that falls out of that is not a job description. It is the first development plan, written on day one, derived from a decision the organisation was making anyway and throwing away. The compromise you make in hiring is the development plan you never wrote.

PART THREE · CHAPTER 12

The case study

The obvious objection to all of that is that the capabilities worth naming are the ones an interview cannot test. A conversation establishes whether someone is articulate about their past work. It says very little about how they think when the ground is unfamiliar.

The instrument I would use is a case study, and it is inexpensive. Three or four sides of A4 describing a real scenario with a good deal of information in it, some of it relevant and some not, followed by open questions with no clean answer. Thirty minutes to prepare, then thirty minutes to present and defend it under questioning.

What it surfaces is close to the list a serious role turns on: whether they can find the problem inside a mass of material, whether they structure an argument or assemble a pile of observations, what they do when they run out of time since everybody does, and whether under questioning they can hold a view, change it on good evidence, and tell the difference between the two. You are watching the capability rather than hearing an account of it. The practised candidate is a smaller problem than it appears, because the questions are open and the assessor sets them live. Preparation carries someone through the presentation. It does not survive being pushed two or three questions past the edge of what they have thought about, and finding that edge is most of the exercise: you are looking for the limit of their ability rather than the extent of their rehearsal. I ran this format for several years when I was responsible for recruitment into a practice, and it reliably got the right people through the door.

Match the instrument to the role, though, because the case study is not universal. It suits work where reading a messy situation, forming a view and defending it is a substantial part of the job, which covers consulting, analysis and most commercial leadership, and it is disproportionate for a role that turns on accuracy, process or craft. The underlying principle generalises even when the format does not: test by sampling the work rather than by discussing it. The case study is what a work sample looks like when the work is thinking.

The limit that does hold is consistency. The exercise produces evidence only if the assessors agree beforehand on what they are looking for, which brings you back to naming the six things.

None of this is a novel insight, and it is worth saying so plainly. Any competent HR function could describe the principle. It is rarely done because the job description belongs to recruitment, the development belongs to the manager, and the join between them belongs to nobody, so the artefact is optimised for the first purpose and abandoned before the second. Someone has to own the join, which in practice means the same person who owns what the role's capabilities are. Without that, the artefact will continue to be written to fill a vacancy rather than to build a person.

PART THREE · CHAPTER 13

Assessment, and the loop that is usually missing

Assessment should run throughout, in the right form at the right time, and the form matters more than the frequency.

Practical application beats a paper test in every case that matters here. A written test establishes that someone can recall the shape of an idea, which tells you very little about judgement. Watching someone apply it to a real problem tells you whether they have it. And assessment has to happen at all, because without it an instructor cannot distinguish comprehension from politeness, and a room full of nodding is the least reliable signal in professional life.

Much of it should be continuous and unannounced, in the sense that a good instructor is always watching rather than periodically testing. This is worth stating carefully. The person is not being secretly graded. They are being observed by someone whose job is to notice, and what is noticed comes back to them as development rather than as an entry in a rating system. That distinction separates this from performance management, which is what most readers will assume is meant.

The discipline that makes it real is what happens next. Seeing a gap, leaving it uncorrected in the moment, then recording it and deliberately creating the situation later where it can be addressed: that is the difference between an instructor and an attentive colleague. It requires the gap to be written down somewhere and revisited, which is a small mechanism and a necessary one.

Which brings us to the loop everyone claims to run. An annual review is not a feedback loop. It is a record. A loop requires the interval between noticing something and acting on it to be short enough that the person can still change the thing being discussed, and twelve months is not short enough by any measure. The work is finished, the context is gone, and what arrives is a verdict rather than a correction. It also runs on the wrong cycle entirely: development happens at the pace of the work, and the review happens at the pace of the compensation calendar.

Where continuous feedback does happen, it is usually because a particular manager makes it happen. That dependence on the individual is exactly what makes it fragile, and exactly what designing it is meant to fix. A loop needs a trigger, something specific to look at, someone accountable for closing it, and an interval measured in weeks.

PART THREE · CHAPTER 14

Appointed, not qualified

In aviation, medicine, shipping and the military, the authority to lead comes with a qualification that was assessed against a standard and can be withheld. In business, the appointment is the credential.

What is assessed at promotion is performance in the previous role, which is a different capability and often the one most likely to be lost. The consequence is visible in the survey data. A Robert Walters study of 2,000 UK professionals found 66% of managers had received no formal training for the role, with a further 22% given responsibility for other people without acknowledgement, title or pay change. Gallup's finding is the other side of the same coin: fewer than half of managers strongly agree they have the skills to be exceptional at the job. These figures come from surveys with obvious commercial interest in the answer, so treat the precise numbers lightly. The pattern they describe will be recognisable to anyone who has watched a promotion round. Nobody establishes that a person can lead before handing them the role, so the first real assessment of their leadership is the damage, or the absence of it, observed years later through attrition and results that have a dozen other causes.

The counter-argument deserves acknowledging. Leadership is more contextual than flying an aircraft, harder to standardise, and formal credentialism degrades into box-ticking as readily as anything else. The answer is not a certificate. It is demonstrated capability against a defined standard, assessed by observation, and partial assessment is a great deal better than none.

For a first-line leadership transition, that standard is not mysterious. Can this person set an outcome rather than issue a task? Can they judge a rescue window? Will they hold the difficulty without rushing in? Will they carry the consequence publicly when an attempt fails? Those are observable, they are the method in part two, and someone who cannot do them is not ready to have people, however good they were at the work.

Two further things belong on that list. The first is what happens to the people themselves. Attrition is the lagging indicator, and by the time it moves the damage is a year old. Morale shifts first and is visible to anyone paying attention: whether problems arrive early or late, whether anyone volunteers for the difficult piece, whether the capable ones have started taking the recruiter's call. A leader can hit every delivery target while the team around them quietly degrades, and an organisation that assesses only the first finds out about the second at the exit interview.

The second is leading without authority, which commercial work demands constantly and assesses almost never. Most consequential work now runs across functions the leader does not control, so progress depends on persuading peers who owe you nothing, and on managing upwards well enough to secure cover for your own people. What that takes is easier to describe as behaviour than as attributes: knowing when to press and when to concede, holding a position without making an enemy of whoever holds the opposite one, and being willing to be wrong in front of others. Written as a list of qualities these become the laundry list of chapter 11. Watched over a month in a matrix organisation, they are perfectly visible, and everyone already knows who has them.

There is a consequence in this for the person at the top, and it is worth stating because it is invisible from where they sit. An organisation that develops people to hold a position under pressure produces people who will say something a senior leader does not want to hear. One that develops compliance produces the opposite. The difficulty is that from the chair the two look identical: silence reads as agreement, and an absence of challenge feels like alignment rather than like a signal that has gone missing. Nobody has ever been told they are surrounded by people who will not disagree with them.

PART THREE · CHAPTER 15

Designing the path

Three changes follow, and they are the cheap parts of what the high-consequence fields do.

Train before the role, not after it. The current pattern is to promote someone and then, some months later, send them on a short course about the job they have been doing badly in the meantime. The material is often very good; the placement wastes it. Preparation belongs before the appointment, and its completion should be part of what makes someone appointable.

Let the transition be the trigger, rather than the rating. Where substantial investment is rationed to those already identified as high performers, it becomes a reward for performance rather than an instrument for producing it. The trigger should be the change in what someone is being asked to do. Entry, first-line leadership, and the step to leading leaders are the three that matter most, because each demands capabilities the previous role did not.

Make it hard, and make it pass or fail. The objection to a universal trigger is that it wastes money on people who do not need it. At scale that objection usually wins, and the result is the mass course: diluted to accommodate everybody, assessed by attendance, and worth little to anyone in the room. Large organisations then conclude that development at scale does not work, when what did not work was the thing they diluted.

The alternative is to go the other way. A promotion course that is demanding enough to be difficult for a capable person is worth the time of the whole cohort, including the strong, because the standard is set high enough that nobody coasts through it. Assess it properly and let people fail. That single change does several things at once: it converts attendance into achievement, which the individual values and can carry; it justifies real investment, because the organisation is buying a verified capability rather than a completed booking; and it supplies the standard chapter 14 argues is missing, so that the appointment stops being its own credential.

The hard part is the consequence, and it should be faced rather than finessed. If the course is genuinely pass or fail and sits before the appointment, then someone who does not pass does not take the role yet. That makes the promotion provisional until the standard is met, which is how it works in every field where the stakes force the issue, and it is a significant change to how an organisation handles promotion rather than a change to its training. It is also the only version that makes the standard mean anything.

Differentiation then applies around the course rather than instead of it. The assessment at the transition tells you who needs preparation beforehand and who needs support afterwards, which is a far cheaper thing to tailor than the core itself.

This is where the record from chapter 11 earns its keep. If the capabilities were named and the gaps written down at appointment, you already know who needs preparation before the course and who will need support after it, and that tailoring costs nothing to apply. Without it, the organisation has no basis for judging either, which is precisely why it defaults to sending everyone through the same diluted event or to rationing the good ones by rating. Both are the same failure: provision decided by something other than what the person needs.

Together these require a different placement of the existing budget rather than a larger one. Fewer, harder, better-attended events cost less than the diffuse provision most organisations already fund, and produce something a person and a promotion panel can both point at.

The moment it lives or dies

Every proposal in this area meets the same objection: it is a good idea for a quieter period. Things never get quieter, so development never resumes.

The practice survives or fails in one recurring moment, and it is identifiable. A deadline is at risk. The manager who handed over an outcome has to decide whether to take the work back early or hold the window open a little longer. Taking it back is always defensible and always safe, and it ends the development. Holding it is the method.

So the executive's lever is not the objectives document. It is what happens to the manager who held the window and delivered a day late. If that manager is backed, the practice spreads. If they are marked down while the one who took everything back early is praised for reliability, every manager in the building learns the real rule within about two instances, and nothing written in an objective survives it. This is the air cover argument from chapter 8 one level up: the leader protects the junior, the executive protects the leader, and without the second the first does not happen.

A rough test helps in deciding what to hand out at all. Complexity, immediacy and stretch: the work's difficulty, the deadline, and how far beyond the person it sits. It is the old scope-time-cost triangle in another suit, and the same logic applies. You can have all three, at real risk to the outcome, which makes it a deliberate choice taken rarely rather than something that happens because nobody thought about it. Complexity is usually a property of the work and immediacy is usually set by the client, so stretch is the dial the manager actually controls. In a pressed quarter it turns down.

That is not a reason to stop. Everything taught in calm conditions is provisional until it is executed under load, and pressure is where you find out whether any of it took. Most organisations waste this. They observe a difficult period, note informally who stepped up, record nothing, and lose the most informative fortnight of their year. Decide beforehand what you are watching for and write down what you saw.

The leadership test is the sharper one. Under load the failure mode is tunnel vision: the leader narrows to their own deliverable, air cover disappears because they are protecting themselves, and delegation reverts to task-handing because tasks feel safer against the clock. Holding the collective outcome in view when everything pulls towards your own is the observable behaviour that separates someone who has been developed from someone who was merely promoted, and it is the thing an appointment never establishes.

One limit worth stating: there are quarters where the correct decision is to suspend the stretch entirely and simply deliver. That is a judgement, not a failure. The problem is the organisation where every quarter is that quarter.

The questions worth asking now

Four questions, and their value lies in not being able to answer them comfortably.

The first. If someone asked you to describe how your organisation develops judgement, what would you say? Not the courses available. The mechanism: what a person will be able to do that they could not before, by when, and who is accountable for getting them there.

The second. In the last month, how much work was handed to your people as outcomes, and how much as tasks? And if the answer is mostly tasks, what were the developmental hours in that work spent on instead?

The third. Of the people who became managers in your organisation in the last two years, how many were assessed against a standard for leading people before they were given any? And if the answer is none, what exactly did the promotion establish?

The fourth is the one only you can answer. The last time a manager in your organisation delivered a day late because they held a stretch open rather than taking the work back, what happened to that manager? And if you cannot recall an instance, that is itself the answer, because it means nobody has judged it safe to try.

An organisation that answers these well has done the work this briefing describes. One that cannot is not necessarily in difficulty. It is carrying a dependency it has never examined, at the moment the thing it was depending on is being removed.

The advantage of asking now is that the answers can still be acted on calmly. Capability takes years to build and the shortage takes years to appear, which means the organisations that start this year will be visibly better at it within two or three, and the ones that wait will not notice what they are missing until the people they needed are not there.

References

EY (2026). *From intern to leader: EY US introduces Career Residency program to transform entry-level professional experience.* Press release, 17 August 2026. Reports the extension of an eight-week internship into a paid eight-to-twelve-month residency, with entry at analyst level. ey.com/en_us/newsroom/2026/08/from-intern-to-leader-ey-us-introduces-career-residency-program

Brynjolfsson, E., Chandar, B. and Chen, R. (2026). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.* Stanford Digital Economy Lab, revised August 2026; originally released August 2025. Uses ADP administrative payroll data covering millions of US workers through June 2026. Finds no widespread economy-wide displacement, and a 19% employment gap for workers aged 22 to 25 in AI-exposed occupations against their less-exposed peers, operating through reduced hiring rather than separations, with no comparable gap for experienced workers. Working paper; not peer reviewed. digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine

Hosseini, S. and Lichtinger, G. (2026). Harvard University working paper on generative AI adoption and entry-level hiring. Analyses résumé and job-posting data covering 66 million workers across more than 280,000 US firms, 2015 to 2025, and reports a sharp fall in entry-level hiring at adopting firms since 2023 while senior employment continued to grow. Working paper; not peer reviewed. Reported in Forbes, 29 May 2026.

Robert Walters (2024). Survey of 2,000 UK white-collar professionals on management training, reported in Fortune, 23 May 2024. Finds 66% of managers received no formal training for the role.

Goodchild, M. (2026d). *Built for a world that's leaving: an executive briefing on technology resilience.* MultipleWorks, Hong Kong. The graded-protection discipline referenced in chapter 10 is developed there at length. Available at multipleworks.com.hk/briefings.

A note on method. *Part two is drawn from military instructional practice and from applying it in commercial settings over twenty-five years. It is offered as a considered position rather than as a survey of what organisations currently do, because the practice argued for here is not yet common enough to survey.*

Mark Goodchild is Founder and Principal of MultipleWorks, a Hong Kong consultancy engaged by executive teams on complex, high-stakes problems across Asia Pacific and internationally, from the first decision through to delivery. He spent eleven years at EY as a Director leading digital and emerging technology consulting across the UK and Asia Pacific, and holds a planning and command background from the British military reserves, where he served as an instructor.

multipleworks.com.hk

SUGGESTED CITATION

Goodchild, M. (2026e). The only generation trained by accident: an executive briefing on building capability by design. MultipleWorks, Hong Kong. Available at multipleworks.com.hk/briefings.

v1.0 · August 2026