

GUIDE 03

THE FOUNDATIONS SERIES

Levels

*How many levels to run, where to draw the lines between them,
and what it costs you to get that wrong.*

What this guide covers

01 The global grade spine
page 3

02 How many levels
page 6

03 Granularity: the worked example
page 8

04 How architectures rot
page 16

05 Spans, layers and rules of thumb
page 18

01

The global grade spine

One scale that every track maps onto. Without it, tracks cannot be compared.

— THE IDEA —

Tracks run in parallel. The spine is what lets you read across them.

A global grade is a single numbered scale — commonly 1 to 15 — that every level in every track maps onto.

Without a spine, P4 and M2 are just two labels. You cannot say whether they are equivalent, and so you cannot compare spans, model workforce cost by seniority, or move someone between tracks without an argument.

The spine does not flatten the tracks. Each track still has its own levels, its own descriptors and its own progression logic. The spine only says how they line up — which is exactly the question every cross-track conversation turns on.

— THE SPINE —

One scale, every track mapped onto it

GRADE	Support	Professional	Technical	Management	Executive
15					E3
14					E2
13					E1
12				M6	
11				M5	
10		P6	T6		
9				M4	
8		P5	T5	M3	
7				M2	
6		P4	T4	M1	
5		P3	T3		
4	S4	P2	T2		
3	S3	P1	T1		
2	S2				
1	S1				

Read across a row to see what is equivalent. P4, T4 and M1 all sit at grade 6 — same size of work, three different kinds of contribution.

02

How many levels

Reference ranges by track and organisation size — starting points, not answers.

— REFERENCE RANGES —

Typical level counts by track and size

These are commonly observed ranges, offered as a sanity check. Your work content decides the answer, not this table.

Track	Under 500 people	500 – 5,000	5,000 – 25,000	25,000+
Executive (E)	0 – 1	1 – 2	2 – 3	3 – 4
Management (M)	2 – 3	3 – 5	5 – 6	6 – 7
Professional (P)	3 – 4	4 – 6	5 – 7	6 – 8
Technical (T)	0 – 3	4 – 6	5 – 7	6 – 8
Support (S)	2 – 3	3 – 4	3 – 4	4 – 5
Operational (O)	2 – 3	3 – 4	4 – 5	4 – 6
Total spine	8 – 10	10 – 13	13 – 15	15 – 18

Ranges reflect commonly observed practice described generically. No survey or vendor data is reproduced.

03

Granularity: the worked example

The same body of work, levelled two ways. This is the centre of the guide.

— THE SETUP —

One discipline, levelled two different ways.

Financial Analysis. Identical work, identical people, identical organisation. Only the number of levels changes.

Model A runs ten professional levels, P1 to P10. Model B runs four, P1 to P4. Neither is wrong. They make opposite trades, and the trade is what this section is about.

The slides that follow show real descriptor language at equivalent career points, then the operational consequences of each choice. Read the descriptors first — the consequences only make sense once you have seen how small the deltas get.

Ten levels, or four

MODEL A

Granular — P1 to P10

Ten steps across the same career span. Each step is a small, carefully defined increment.

- Small deltas between adjacent levels
- Descriptor text for all ten, each distinguishable from its neighbours
- A promotion roughly every 18–24 months for a strong performer
- Heavy calibration load — ten sets of boundaries to police

MODEL B

Broad — P1 to P4

Four steps across the same career span. Each step is a large, obvious change in the work.

- Large deltas between adjacent levels
- Four wide descriptors, each covering a broad range of scope
- A promotion roughly every 4–6 years
- Light calibration load — three boundaries to police

Model A at three equivalent career points

Granular. Note how much work the qualifying words are doing to separate adjacent levels.

	P3 · Early career	P6 · Mid career	P9 · Late career
Knowledge & Experience	Qualified in the discipline with one to two years applying it in context.	Fully proficient; works unsupervised across the normal range of the discipline.	Recognised internally as an authority; sought out on the hardest questions.
Problem Complexity	Well-defined analytical problems with an established method to follow.	Varied problems needing interpretation; precedent exists but rarely fits exactly.	Novel problems with no internal precedent; establishes the method others follow.
Autonomy & Discretion	Work assigned and reviewed in detail; judgement within a set method.	Sets own objectives within agreed direction; decides what warrants escalation.	Wide latitude against strategic intent; reviewed on business outcome only.

Model B at the same three career points

Broad. Each descriptor has to cover the ground that two or three Model A levels covered.

	P1 · Early career	P3 · Mid career	P4 · Late career
Knowledge & Experience	From entry through to full working proficiency in the discipline.	Substantial expertise, including areas where the organisation has limited depth.	Authoritative across the discipline; expertise differentiates the organisation.
Problem Complexity	Defined problems, from routine through to moderately complex analysis.	Ambiguous problems with incomplete information and competing constraints.	Defines which problems are worth solving; works where there is no internal map.
Autonomy & Discretion	From detailed review through to managing own workload against objectives.	Accountable for outcomes rather than activity; sets direction for a body of work.	Effectively self-directing within strategy; accountable to senior leadership.

The same career point, both ways

Mid-career analyst, Autonomy & Discretion. The delta between adjacent levels is the thing to look at.

MODEL A · P5 → P6

P5 — Manages own workload against agreed objectives. Escalates matters of materiality rather than difficulty.



P6 — Sets own objectives within agreed direction. Determines which matters warrant escalation at all.

The delta is a shade of judgement. Two reviewers can read the same analyst into either level and both be defensible.

MODEL B · P2 → P3

P2 — Works to defined objectives with periodic review. Determines the approach to individual tasks.



P3 — Accountable for outcomes rather than activity. Sets direction for a body of work and is reviewed on results.

The delta is a visible change in how the person is managed. Most reviewers land on the same answer without debate.

Neither is wrong. Model A buys recognition and pace; Model B buys defensibility and low maintenance.

What each model costs you to run

This is the comparison that should decide it — not which descriptor set reads better.

	Model A · P1–P10 (granular)	Model B · P1–P4 (broad)
Promotion frequency	Every 18–24 months; frequent visible recognition	Every 4–6 years; long gaps between milestones
Calibration burden	High — nine boundaries, many marginal cases	Low — three boundaries, mostly obvious calls
Title inflation risk	High — ten levels invite ten title variants	Low — few levels, few titles to inflate
Defensibility	Weak at the margins; adjacent levels hard to separate	Strong; differences are visible to any reviewer
System maintenance	More jobs, more codes, more descriptors to keep current	Fewer objects; cheaper to govern over time
Retention signal	Strong short-cycle progression signal	Requires other recognition mechanisms between steps
Best suited to	Large, fast-moving populations with strong HR capacity	Smaller or leaner organisations, and broad-band designs

How to actually decide

Work through these in order. The first one that gives a clear answer usually is the answer.

- Can you write a genuinely distinguishable descriptor for every level? If two adjacent levels need the same words, you have one level too many.
- Can your calibration capacity police the boundaries? Every boundary is an annual argument. Count the arguments you can afford to run well.
- How long is a realistic career in this discipline? Divide by an acceptable time-in-level and you have an upper bound on level count.
- What does the external market recognise? A level structure nobody outside can interpret makes hiring and benchmarking harder.
- When in doubt, choose fewer. Adding a level later is straightforward; removing one means re-levelling people, and that is a change programme.

04

How architectures rot

Two opposite failure modes, both of which destroy comparability.

Compression and proliferation

Both start as small accommodations. Both end with an architecture that no longer means anything.

COMPRESSION



Everyone drifts upward

Levels are awarded to recognise good people rather than to describe bigger work. Over a few years the population bunches at the top.

- Symptom: most of a family sits in the top two levels
- Symptom: the entry levels are empty and nobody is hired into them
- Cause: level used as a reward mechanism, not a description
- Fix: re-anchor with calibration anchors and hold the line at intake

PROLIFERATION



New levels keep appearing

Every retention case or hiring difficulty is solved by inventing a level. Nobody ever removes one.

- Symptom: levels that exist for one person
- Symptom: adjacent descriptors that cannot be told apart
- Cause: no governance gate on creating a level
- Fix: a single owner for the spine, and a real approval to add to it

05

Spans, layers and rules of thumb

How level design collides with org structure, and the heuristics worth keeping.

Level design and org structure constrain each other.

Levels describe work. Layers describe reporting. When the two stop lining up, both become unreliable.

If your management track has six levels but the organisation only has four reporting layers between the front line and the chief executive, two management levels are describing work that structurally cannot exist. People will be promoted into them anyway, and the level will quietly stop meaning anything.

The reverse also happens: eight reporting layers against four management levels produces managers at the same level with wildly different scope. Sense-check the two against each other before publishing, and again at every restructure.

Seven heuristics worth keeping

- One spine for the whole organisation. Local scales defeat the purpose of having a scale.
- Every level needs a distinguishable descriptor. If you cannot write it, the level does not exist.
- Management levels should not exceed reporting layers. Check this at every restructure.
- The top of the technical ladder must reach the same grade as senior management, or the dual ladder is decorative.
- Fewer levels are easier to defend, cheaper to maintain and harder to game.
- Never create a level for one person. That is what an exception log is for.
- Level the job, never the person. Every rot pattern in this guide begins with breaking this rule.

UP NEXT IN THE SERIES

GUIDE 04

Leveling Criteria: The Six Dimensions

The universal criteria model, how to write descriptors that hold up under challenge, and how to calibrate them.

Find the full series and the free mapping workbook at hrdigitalplayground.com

HR Digital Playground